DATA+AI SUMMIT
BY databricks

Determining When to Use GPU for Your ETL Pipelines at Scale

Chris Vo, AT&T
Hao Zhu, NVIDIA

Databricks
2023

Agenda

Outline

- NVIDIA Implementation of Spark RAPIDS (GPU) Qualification Tool
- AT&T Usage Pattern for Qualification Tool
- Actual Cost Saving after Conversion
- Key Takeaways
- What’s Next For AT&T

Qualification Tool

Process Spark Event Logs

Spark 2.x or Spark 3.x

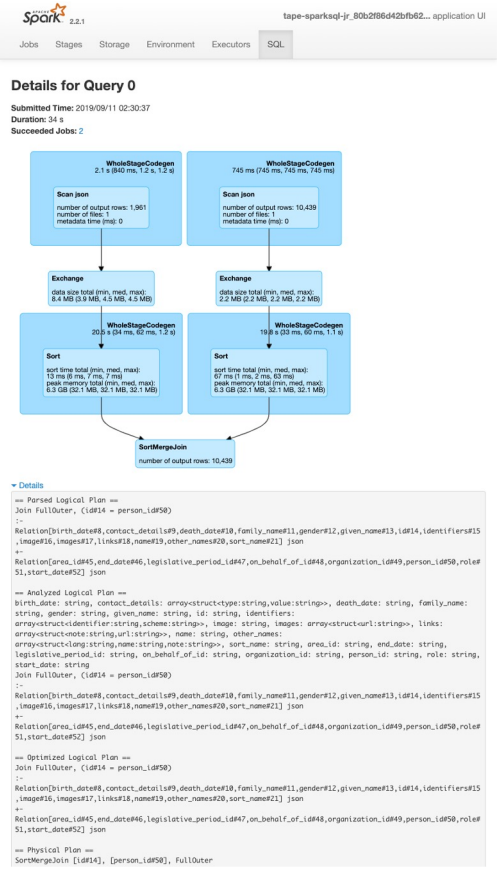

Qualification Tool

Exec and Expression Analysis

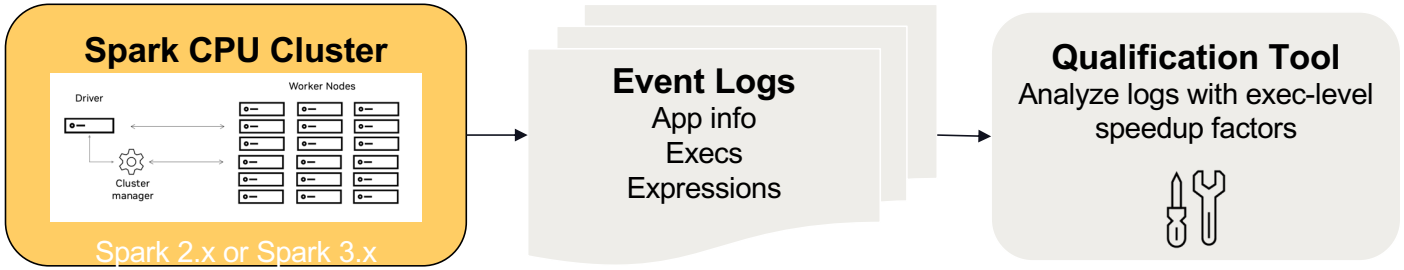

Spark 2.x or Spark 3.x

```
java ${QUALIFICATION_HEAP} \
           -cp rapids-4-spark-tools_2.12-<version>.jar:$SPARK_HOME/jars/* \
           com.nvidia.spark.rapids.tool.qualification.QualificationMain
/usr/logs/app-name1

OR

spark_rapids_user_tools databricks_aws qualification \
   --eventlogs $EVENTLOGS \
   --cpu_cluster $CLUSTER_NAME \
   --remote_folder $REMOTE_FOLDER
```

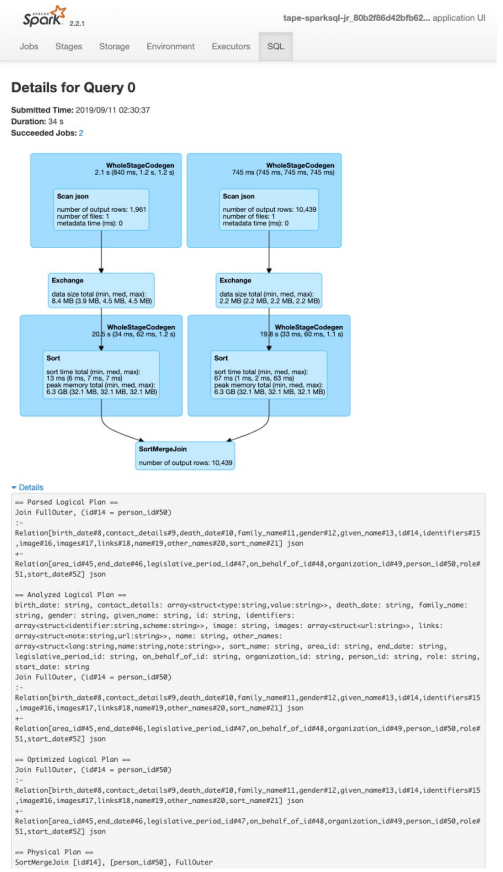

Qualification Tool

Predicting the benefits of Spark + GPUs

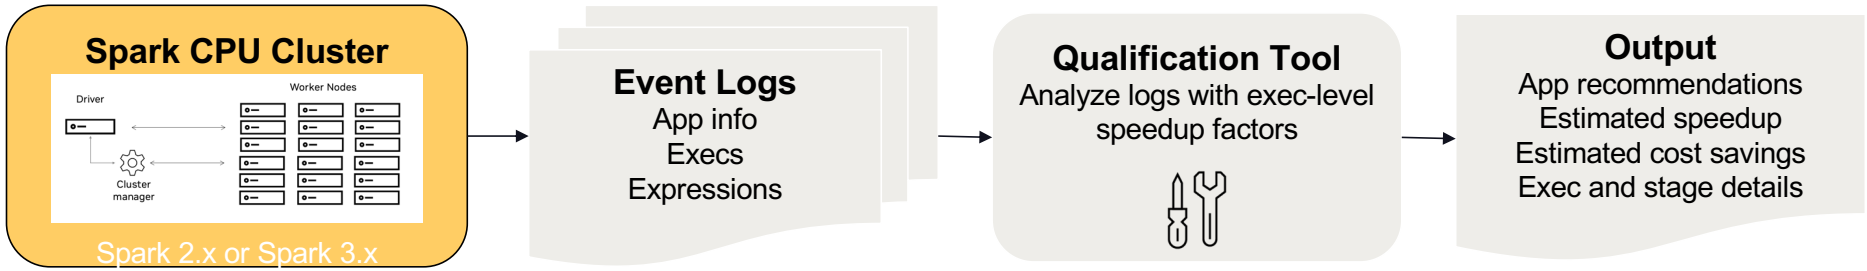

Spark 2.x or Spark 3.x

```
java ${QUALIFICATION_HEAP} \
           -cp rapids-4-spark-tools_2.12-<version>.jar:$SPARK_HOME/jars/* \
           com.nvidia.spark.rapids.tool.qualification.QualificationMain
/usr/logs/app-name1

OR

spark_rapids_user_tools databricks_aws qualification \
   --eventlogs $EVENTLOGS \
   --cpu_cluster $CLUSTER_NAME \
   --remote_folder $REMOTE_FOLDER
```

Qualification Tool

Example 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

Qualification Tool

Example UI Output

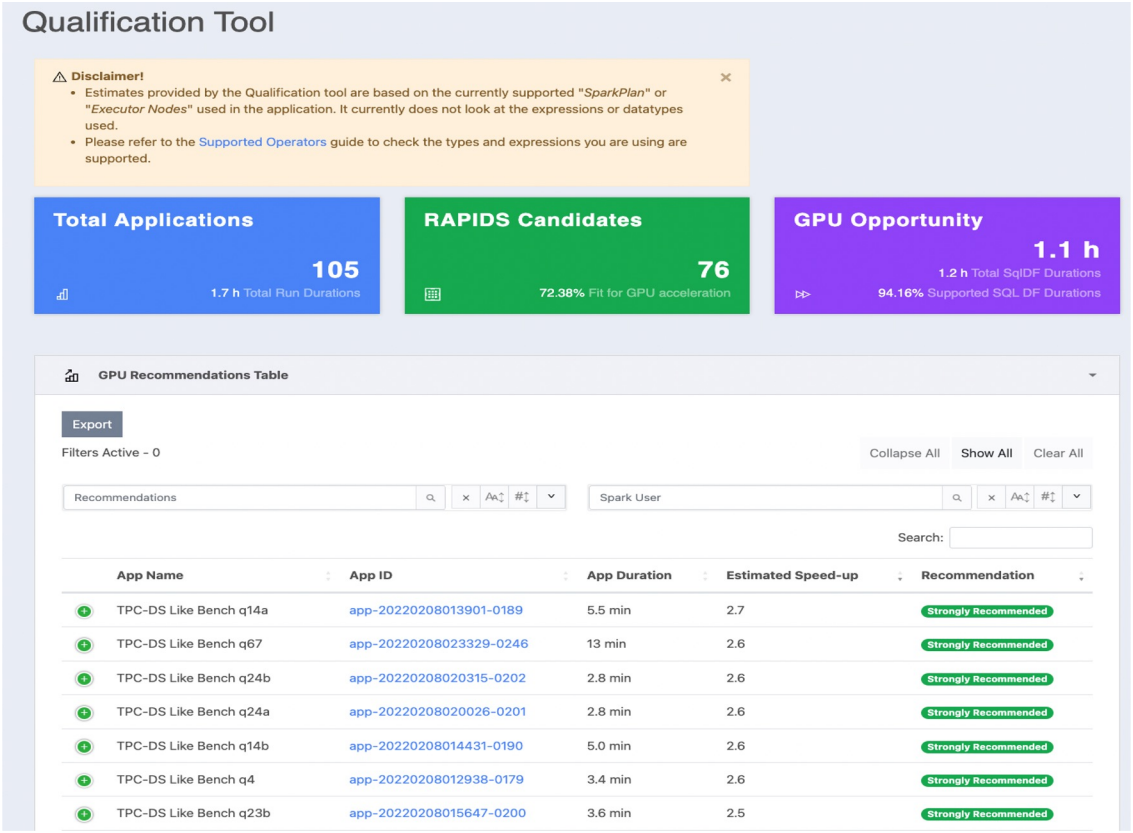

Qualification Tool

Download Links and Documentation

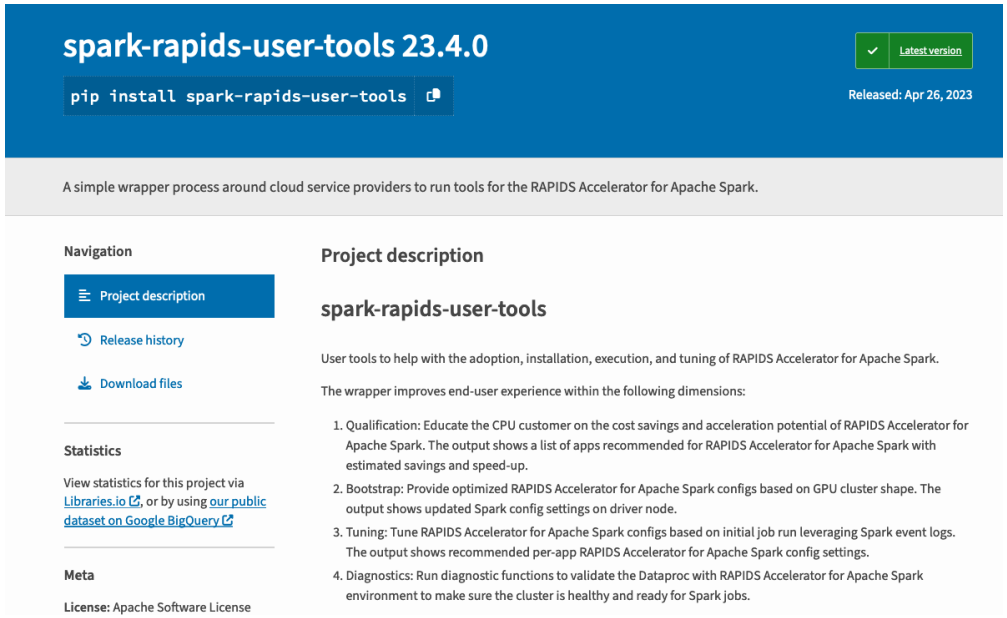

https://pypi.org/project/spark-rapids-user-tools/


databricks

GPU-Qualification for ETL Jobs Scanning Pipeline

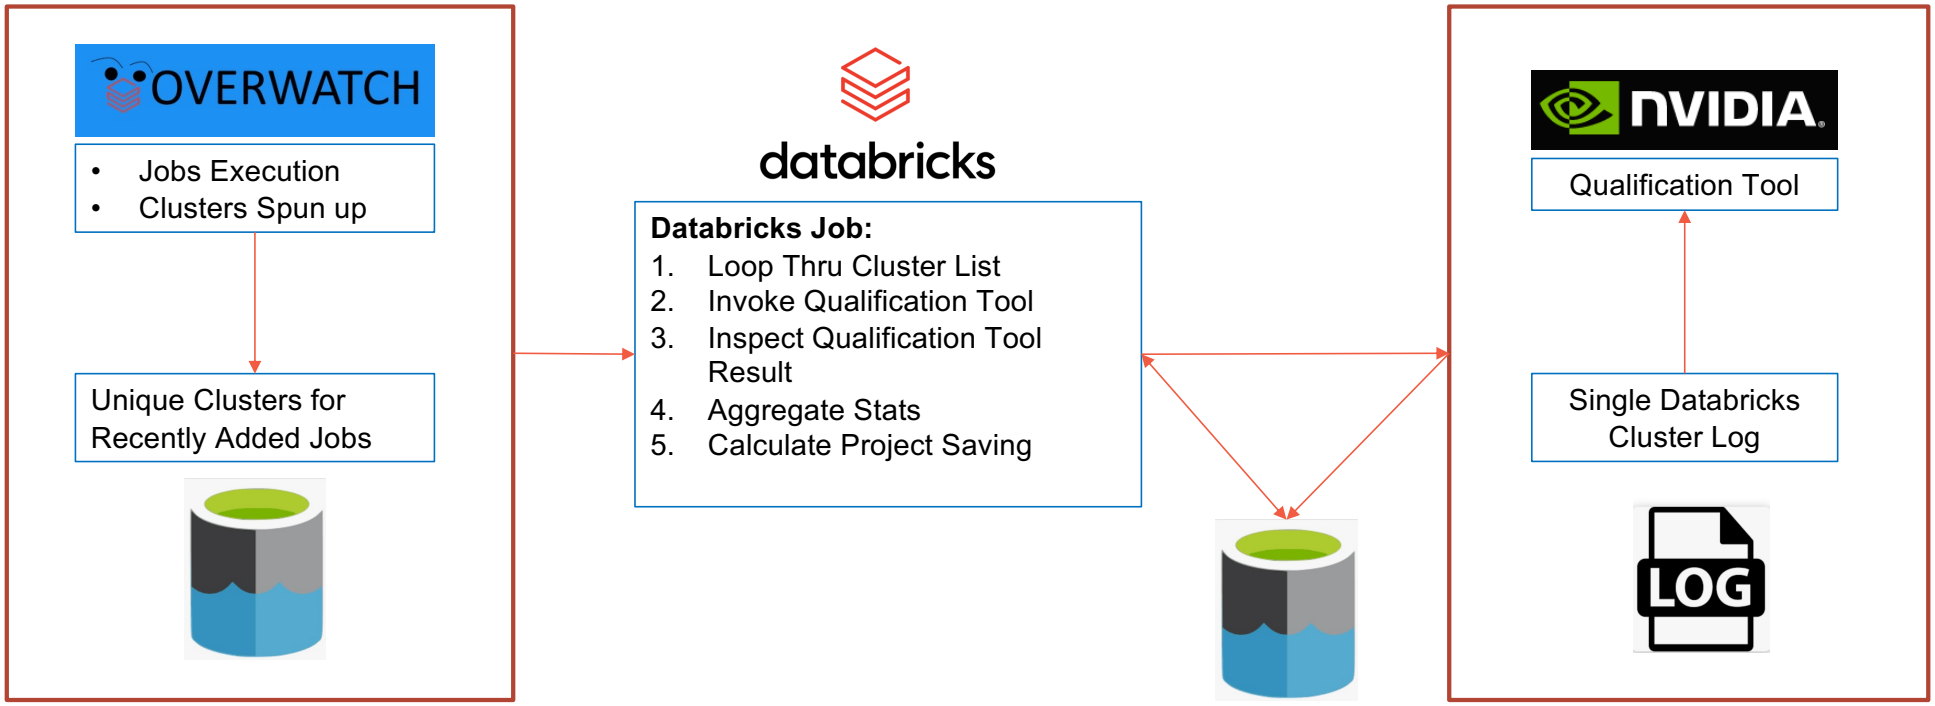

Estimated Saving Projection for Conversion Candidates

Job Name	Contact	Recommendation	Est Hours Saving	Est Unit Saving ($)	Frequency	Est Yearly Saving ($)
Job ABC	xx1234	Recommended	0.5	9.52	Hourly	83,356
Job EFG	xx5678	Recommended	0.4	6.92	Hourly	60,587
…						
Job XXX	xx1111	Strongly Recommended	1.0	28.72	Daily	10,481
…						
Job YYY	xx2222	Strongly Recommended	0.5	15.18	Monthly	182

Databricks Jobs Conversion Checklist

1. Rank Job Candidates List By Highest Saving Projection
2. Convert Job Candidates Where Projected Saving > $10k/yr
3. Record Current CPU Cluster Size and Execution Time
4. Convert Job to Spark RAPIDS GPU
5. Record New GPU Cluster Size and Execution Time
6. Calculate Actual Saving Based on #3 and #5
7. Rollback to CPU If:
    1. Encountered Error
    2. Job Cost Saving is Low (Azure GPU Quota is Limited)

CPU Cluster Configuration

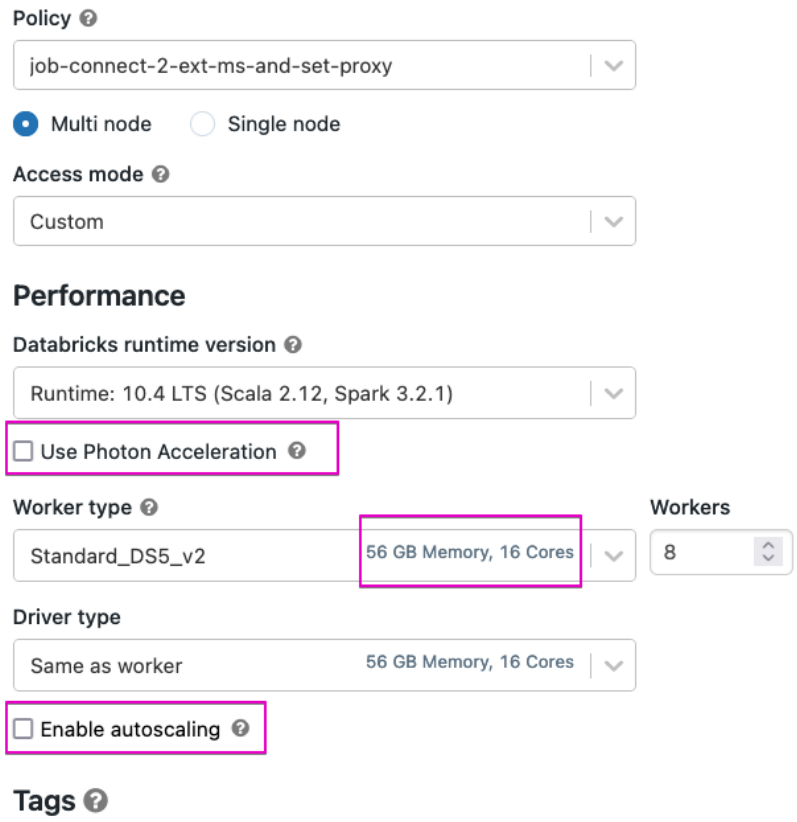

Spark RAPIDS (GPU) Cluster Configuration

Cluster Configuration

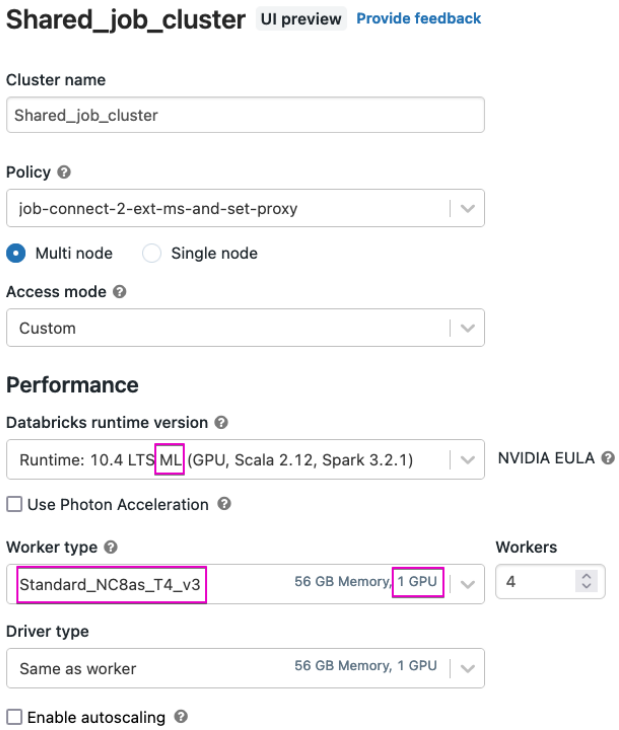

Advanced Options - Spark Config

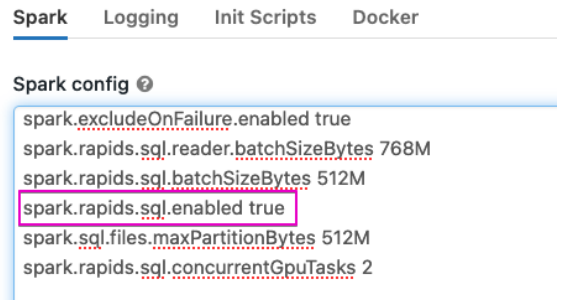

Advanced Options - Init Scripts

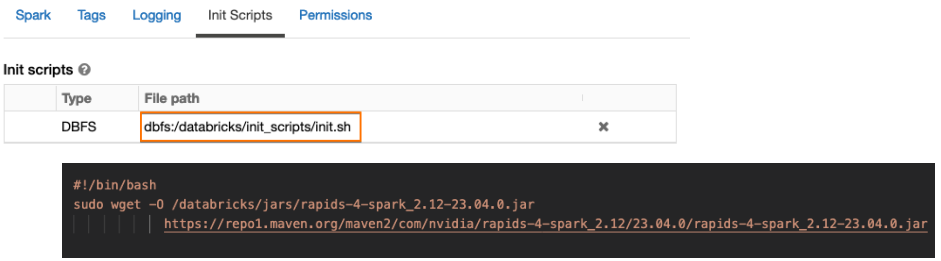

Comparing Query Plans

```
SELECT id1, sum(counter)
FROM BaseView
WHERE SpecialText rlike '{TextSearch}'
GROUP BY id1
```

CPU Plan

```
== Physical Plan ==
AdaptiveSparkPlan (22)
+- == Final Plan ==
   * HashAggregate (13)
   +- ShuffleQueryStage (12)
      +- Exchange (11)
         +- * HashAggregate (10)
            +- * HashAggregate (9)
               +- AQEShuffleRead (8)
                  +- ShuffleQueryStage (7)
                     +- Exchange (6)
                        +- * HashAggregate (5)
                           +- * Project (4)
                              +- * Filter (3)
                                 +- * ColumnarToRow (2)
                                    +- Scan parquet  (1)
```

GPU Plan

```
== Physical Plan ==
GpuColumnarToRow (13)
+- GpuHashAggregate (12)
   +- GpuShuffleCoalesce (11)
      +- GpuColumnarExchange (10)
         +- GpuHashAggregate (9)
            +- GpuHashAggregate (8)
               +- GpuShuffleCoalesce (7)
                  +- GpuColumnarExchange (6)
                     +- GpuHashAggregate (5)
                        +- GpuProject (4)
                           +- GpuCoalesceBatches (3)
                              +- GpuFilter (2)
                                 +- GpuScan parquet  (1)
```

Compute Cost Calculation

VM Cost = Cost Per VM x Number of Nodes
DBU Cost = DBU x DBU Price x Number of Nodes

Compute Cost = VM Cost + DBU Cost

Databricks Cluster Creation

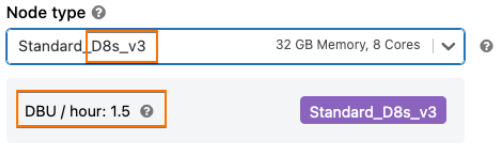

Databricks Website

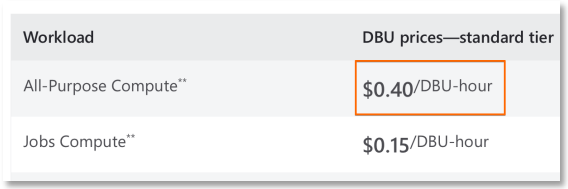

Workload	DBU prices—standard tier
All-Purpose Compute**	$0.40/DBU-hour
Jobs Compute**	$0.15/DBU-hour

Microsoft Azure Website

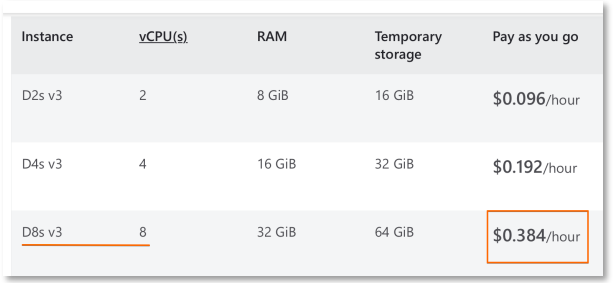

Instance	vCPU(s)	RAM	Temporary storage	Pay as you go
D2s v3	2	8 GiB	16 GiB	$0.096/hour
D4s v3	4	16 GiB	32 GiB	$0.192/hour
D8s v3	8	32 GiB	64 GiB	$0.384/hour

CPU-GPU Conversion Actual Saving

Job Name	Contact	Frequency	CPU Cost ($)	GPU Cost ($)	Saving ($)	Saving (%)	Multiplier	Yearly Saving ($)
Job ABC	xx1234	Hourly	7.64	2.96	4.68	61%	8760	41,039
Job EFG	xx5678	Hourly	5.59	1.34	4.25	76%	8760	37,195
…								
Job XXX	xx1111	Daily	30.57	8.65	21.92	72%	365	8,000
Job ZZZ	zz3333	Daily	31.75	19.90	11.85	37%	365	4,326
…								

Key Takeaways

- Qualification Tool
  - Useful for identifying candidate for Spark RAPIDS
  - Can get estimation for cost saving
- A recommendation status doesn’t mean actual cost saving
- Focus on the jobs with high execution frequency (ie: hourly and daily)

What's Next for AT&T

- **Automate Cluster Configuration, Adjustment, Rollback**
  - Optimize CPU
  - Optimize GPU
  - Optimize Photon
- **Automate Tracking**
  - Compute Costs
  - Savings

NVIDIA RAPIDS Accelerator for Apache Spark

Enterprise Support available with license purchase

5x

Faster Execution Time

- Take advantage of faster analytics
- Accelerate AI pipelines

4x

Lower Costs

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

Full Enterprise Support

NVIDIA AI Enterprise 3.1 offers

- Mission critical support
- Bug fixes
- Professional services

Amazon EMR

Google Cloud Dataproc

databricks

aws

Microsoft Azure

nvidia.com/spark


databricks