

## Windjammer EMR Accelerator Q&A

1) What is EMR Accelerator with SNE (Spark Native Engine)?

SNE is an Apache SPARK Native Execution Engine that seamlessly plugs into EMR 6.x clusters during cluster deployment

2) How Does SNE Work?

SNE transparently accelerates query run times, reduces cost of running queries, and adds efficient high availability to EMR jobs. Queries are executed entirely in native code, fully exploiting Data Lake cloud store bandwidth, and utilizing S3-based query checkpointing

3) What Customer ISSUES does Windjammer Accelerator Solve when plugged into EMR?

- Reduces EMR SparkSQL query and job run times by ~3X on same cluster

- Achieves 2.1X query speedup using half the number of slaves

- Achieves 1.5X query speedup using one quarter the number of slaves

- Reduces CPU utilization to  $\frac{1}{4}$ , increases S3 bandwidth utilization by 3X

- Reduces EMR SparkSQL cost/query to  $\frac{1}{4}$  when using  $\frac{1}{4}$  the number of slaves

- Provides rapid query continuation when spot instances are reclaimed, when the cluster is auto-scaled up or down, and from full cluster failures without the need for an external shuffle service

- Reduces query run time variance to  $\frac{1}{4}$

4) How does Windjammer Accelerator work with AWS EMR Spark?

Add a bootstrap-action to EMR Cluster create which deploys SNE as a .JAR to transparently accelerate EMR Spark query execution

5) What Data is available to review SNE results?

TPC-DS benchmarking results are presented in the [Technical White Paper](#) You can quickly run A/B experiments of TPC-DS and your own queries in your own clusters to get your own results as discussed in the [Deployment and POC Guide](#)

6) Is there a FREE Trial version Available?

Yes, a 14-day free Trial is automatically provided when you subscribe to EMR Accelerator on the AWS Marketplace

7) How do I deploy EMR Accelerator in my EMR clusters? Add bootstrap action to your emr 6.2 create-cluster as discussed in [Deployment and POC Guide](#)

8) Is there a technical Whitepaper available?

Yes, click [Technical White Paper](#) for detailed architecture, design and benchmarks

9) Is SNE compliant with Open Source Apache SPARK?

Yes, Windjammer SNE is 100% compatible with Apache SPARK and EMR Spark

10) How Does WJ Accelerator compare with other SQL accelerators?

EMR 6.x+Windjammer Accelerator greatly outperforms stock EMR, Databricks' Photon, Dremio, and Starburst as discussed in [Technical White Paper](#). EMR Accelerator results are all with No caching and No Data Skipping to normalize across databases.

11) Does SNE work with other Cloud native SPARK services?

Yes, today SNE accelerates GCP-DataProc and Apache Spark 3.x. as well as EMR Spark. Technically we can support ANY SPARK environments with SNE.

12) Does WJ-Accelerator work with Alluxio caching services?

Yes, we have validated with Alluxio and can share the results and install instructions on request

13) How is the EMR Windjammer Accelerator Product Supported?

EMR Windjammer Accelerator is supported via Windjammer Technologies SLACK customer channels (created during subscription) and via e-mail. Please email [support@windjammer.io](mailto:support@windjammer.io) with any questions.

14) What is Windjammer's website?

<https://windjammer.io/>