



Accelerated Surveillance Analytics

Thank you for downloading this Voltron Data resource. Carahsoft is the distributor for Voltron Data's AI and ML solutions available via NASA SEWP V, ITES-SW2, NASPO ValuePoint, and many more contract vehicles.

To learn how to take the next step toward acquiring Voltron Data's solutions, please check out the following resources and information:



For Voltron Data overview:
carah.io/Voltron-Data



For upcoming events:
carah.io/Voltron-Events



For additional resources:
carah.io/Voltron-Resources



For additional Artificial Intelligence:
carah.io/ai-solutions



To set up a meeting:
VoltronData@carahsoft.com



To purchase, check out the contract vehicles available for procurement:
carah.io/procurement

Accelerated Surveillance Analytics

A surveillance pipeline ingests billions of transactions daily, but the current Spark based workflow struggles to keep up, taking hours to process just 3-4 billion rows and scaling costs linearly with data volume.

To support the new directive, the team requires sub-hour turnaround on datasets 10x larger (30-40 billion rows), a target that remains unattainable even with increased cluster resources.

By adopting Voltron Data's Theseus, a GPU accelerated SQL engine, the team can meet the target SLA for processing 30-40 billion rows while reducing overall compute costs by 50%.

In parallel, refactoring the existing Spark jobs into Ibis, an open source Python framework that provides a portable and backend agnostic interface for writing analytics workflows, enables execution across multiple query engines and helps avoid vendor lock-in.

