ScaleOut Software

ComputeServer

In-Memory Computing for Operational Intelligence

Overview

ScaleOut ComputeServer™ combines a scalable in-memory data grid (IMDG) with an integrated compute engine to deliver immediate results that provide real-time feedback and time-sensitive analytics — "operational intelligence" — in live, production environments. Now you can run continuous, in-memory computations across fast-changing datasets and obtain results with extremely low latency to capture business opportunities before the moment is lost. Object-oriented APIs make it easy to track live systems, eliminating the need for specialized skills in parallel computing.

The Power of Operational Intelligence

In countless applications, from financial trading systems to e-commerce web sites, operational systems continuously create fast-changing data that need to be tracked, correlated, and analyzed to identify patterns and trends — and then generate immediate feedback to steer operations. This is called operational intelligence. Organizations that have it can deliver better results, boost cost-effectiveness, and identify perishable business opportunities that others miss. For example, operational intelligence can help financial trading systems respond to market price fluctuations, create timely and relevant offers for e-commerce shoppers, and alert medical personnel to suspicious patterns in real-time telemetry. Traditional analytics platforms designed for business intelligence, such as Hadoop and Spark, with their focus on offline, batch processing of petabyte datasets, simply cannot match the mission-critical needs of operational systems. ScaleOut ComputeServer *can*.

Parallel Data Analysis

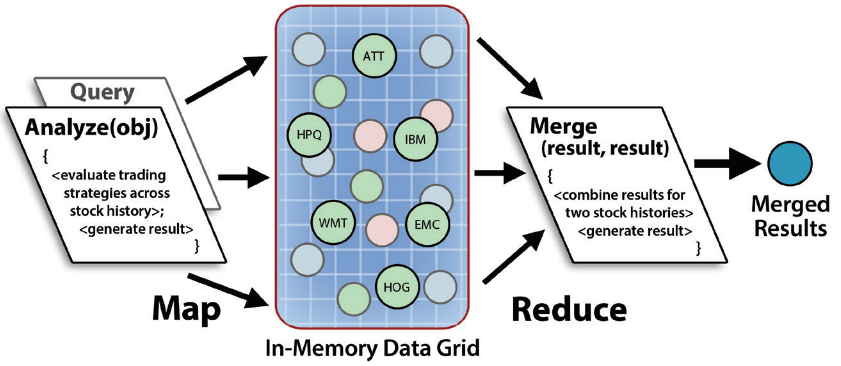

Parallel Method Invocation

ScaleOut ComputeServer lets applications run scalable, data-parallel computations on live, memory-based datasets to deliver fast insights and eliminate bottlenecks in generating real-time feedback. Applications can create large, object-oriented datasets which track the behavior of live systems and perform fast, data-parallel analyses that run user-defined Java, C/C++, or C# methods on selected objects and combine the results with extremely low latency. Called "parallel method invocation," this powerful, patent-pending technology provides intuitive APIs derived from high performance computing. In addition, ScaleOut ComputeServer offers APIs that implement full MapReduce semantics for both Java and C#. The in-memory compute engine delivers fast execution times by minimizing both scheduling overhead and data motion. It also simplifies development by eliminating the need for tuning.

High Performance from In-Memory Computing

ScaleOut ComputeServer performs data-parallel computations on memory-based data stored in its scalable, in-memory data grid. This eliminates delays due to file I/O, minimizes network overhead, and enables analytics to be performed on live data. With today's low memory prices and the proliferation of cloud computing, in-memory data grids have demonstrated the ability to host terabytes of highly available, fast-changing data, which ScaleOut ComputeServer's integrated compute engine can analyze in seconds. Examples of operational data include live shopping or market data, gaming data, logistics information, telemetry from energy or communication systems, and many others. Traditional approaches, such as Hadoop MapReduce, which are designed to analyze large, static datasets, cannot match either its responsiveness or its ability to analyze fast changing data.

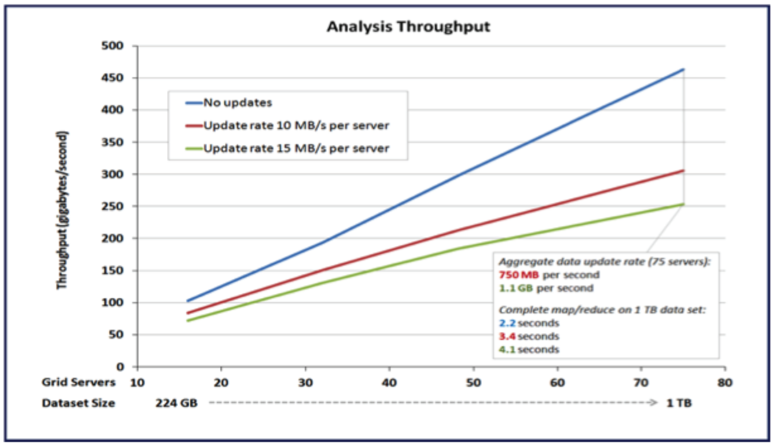

Performance Measurements for a Financial Services Application

The grid's tightly integrated execution engine efficiently schedules computations across multiple cores within each server and across all grid servers to maximize parallelism and minimize data motion within the grid. In addition, it automatically merges analysis results across grid nodes to eliminate the need for complex data reduction steps.

ScaleOut ComputeServer can deliver extremely fast results for large, memory-based datasets. For example, measurements of a real-time, financial services application which back tests stock price histories have demonstrated that a one-terabyte dataset can be analyzed every 4.1 seconds while being updated at the rate of more than a gigabyte per second (the green line in the above chart).

ScaleOut Compute Server provides significantly lower latency than traditional, file-based analytics platforms such as Apache Hadoop, demonstrating more than 40X lower latency in a stock trading application by returning results in 350 milliseconds instead of more than 15 seconds for Hadoop.

Fast Development and Results

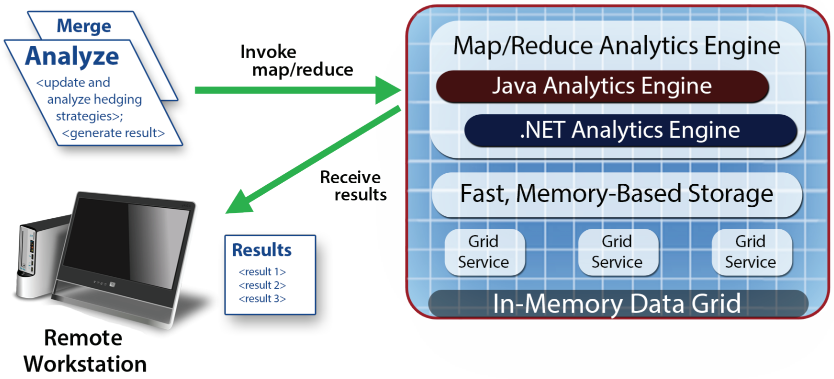

With ScaleOut ComputeServer, there's no need for the user to install and tune a complex MapReduce infrastructure. Applications are written using straightforward, object-oriented techniques that make it easy to track the behavior of live systems and implement data-parallel computations. User-defined methods written in Java, C/C++, or C# operate on in-memory data and avoid explicit grid accesses. Data can be selected for analysis using parallel queries on object properties, avoiding the need for application code to parse files and create key/value pairs. Distributed merging delivers results with minimum latency. Automatic code shipping and pre-staging using multiple Java or .NET "invocation grids" simplify the development cycle and help ensure fast startup. Applications run analysis computations inline without the need for a batch scheduler.

Included with ScaleOut ComputeServer, the ScaleOut Management Pack adds an object browser for visually browsing and managing objects stored in the in-memory data grid and a parallel backup/restore feature for quickly archiving memory-based data in the file system. The object browser's ability to visualize the grid's contents dramatically simplifies application development. The parallel backup and restore feature adds the capability to dynamically save the contents of the in-memory data grid as a snapshot which later can be restored for analysis.

ScaleOut Software
scaleoutsoftware.com

For more information, call us at
(503) 643-3422 or download a
free product trial from our website.