

Fast, intuitive in-memory computing platform for operational intelligence

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

IN-MEMORY COMPUTING FOR FAST DATA ANALYSIS

Live systems generate streams of fast-changing data that need to be tracked, correlated, and analyzed to identify patterns and trends — and then generate immediate feedback to steer operations.

ScaleOut ComputeServer lets applications run scalable, data-parallel computations on live, memory-based datasets to deliver fast insights. Its integrated in-memory data grid and in-memory compute engine deliver fast execution times by minimizing both scheduling overhead and data motion. Object-oriented data storage and intuitive APIs simplify development and minimize the need for tuning.

Applications run user-defined Java, C/C++, or C# methods on selected objects hosted within a distributed, in-memory data grid, which executes the methods in parallel and combines the results with extremely low latency. Called **Parallel Method Invocation**, this powerful patent-pending technology derived from high performance computing provides an intuitive API for scalable, data-parallel computing.

In addition, ScaleOut ComputeServer offers APIs that implement full MapReduce semantics for both Java and C#, as well as full, Apache-compatible Hadoop MapReduce. For API developers, ScaleOut ComputeServer extends the widely used Parallel ForEach API to transparently scale data-parallel computing across an in-memory data grid.

MOVE COMPUTING TO WHERE THE DATA LIVES

Beyond just serving as a fast, scalable repository for live data, in-memory data grids (IMDGs) provide the foundation for data-parallel computing. By harnessing the scalable computing power of the server cluster on which they run, IMDGs enable large, in-memory data sets to be analyzed in place, delivering immediate results and important feedback to live systems. While they can serve distributed queries to select data of interest for client applications, their real power lies in the ability to host data-parallel computations within the grid — moving computing to where the data lives — to eliminate network bottlenecks and deliver high performance. In addition, unlike other in-memory computing platforms, they allow low-latency computing on individual data objects to provide deep introspection in stream processing.



Ex: Portfolio tracking, wire-fraud detection, stock back-testing



Ex: Device tracking for manufacturing, vehicles, mobile devices



Ex: Real-time patient monitoring and alerting, health device tracking



Ex: Real-time inventory reconciliation, manufacturing flow optimization

...and much more





KEY FEATURES AND CAPABILITIES

ScaleOut ComputeServer brings the power of data-parallel computing to distributed, in-memory data grids. Its integrated compute engine minimizes data motion and uses techniques from parallel supercomputing, such as software multicast and global combining, to deliver results as fast as possible. These capabilities are delivered as an intuitive, easy to use SDK that makes application development in Java, C#, and C/C++ simple and straightforward. Built-in, automatic code shipping to the in-memory data grid simplifies applications and helps ensure fast startup times.

| Key Features | Benefits |
|---|---|
| Integrated IMDG & compute engine | Enables object-oriented data access. |
| Computations run within the grid | Eliminates network bottlenecks. |
| Automatic code shipping | Simplifies deployment. |
| Data-parallel APIs for live datasets | Delivers fast results for immediate feedback. |
| Data-parallel APIs for specific objects | Enables streaming analytics on "digital twins." |
| MapReduce APIs for Java and C# | Provides fast, portable MapReduce. |
| Apache Hadoop MapReduce | Runs standard MapReduce on in-memory data. |
| Extended API for .NET Parallel.ForEach | Scales familiar .NET API across an in-memory data grid. |
| Query integrated with parallel APIs | Allows selected processing of large datasets. |
| Object browser | Simplifies development. |

POWERFUL PARALLEL COMPUTING TOOLS FOR .NET DEVELOPERS

Targeted at .NET developers, a new API for data-parallel computing builds on capabilities in .NET's widely used Task Parallel Library. Called **Distributed ForEach**, and modeled after the TPL Parallel.ForEach operator, this new API transparently extends data-parallel computing across ScaleOut's IMDG to handle much larger workloads than otherwise possible. It also makes distributed, data-parallel programming easy to implement for any developer familiar with the .NET programming framework.

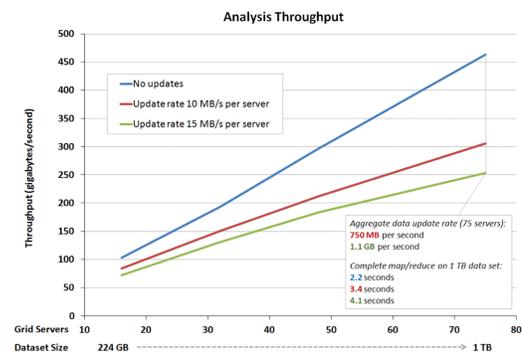
Compatible with the widely used Reactive Extensions (Rx) library, the **Distributed Push-Based Notifications** API for both C# and Java lets applications use the scalable power of ScaleOut's in-memory data grid and compute engine to handle push-based notifications.

Unique Advantages for Live Data

Traditional analytics platforms designed for business intelligence and stream processing, such as Spark, Storm, and Flink, do not offer applications an object-oriented storage model for live data. As a result, they are not well suited to tracking the state of live systems and performing data-parallel computations on this state. Integrating an in-memory data grid with a parallel compute engine gives ScaleOut ComputeServer the unique ability to track live data with high availability and provide operational intelligence for mission-critical applications.

HIGH PERFORMANCE IN-MEMORY COMPUTING PLATFORM

ScaleOut ComputeServer's in-memory compute engine delivers extremely fast results for live datasets. For example, a real-time, financial services application which back-tests stock price histories analyzed a one-terabyte dataset every 4.1 seconds while being updated at the rate of more than a gigabyte per second (the green line in the chart below).



Today's applications need fast access to data for maximum performance. At the same time, the need to quickly obtain feedback from live data has never been greater. ScaleOut's in-memory data grid (IMDG) combines the speed of memory-based storage with an integrated compute engine to enable operational intelligence — the ability to identify and capture business opportunities in the moment, before they disappear.

Whether your data is in a server farm, compute grid, or the cloud, ScaleOut Software's suite of software products eliminates scalability bottlenecks and delivers fast, scalable performance that gives your business a competitive edge.



Try ScaleOut for free.

Experience the power of in-memory computing on Windows or Linux.

www.scaleoutsoftware.com/try-for-free



ScaleOut Software

LEARN MORE

- www.scaleoutsoftware.com
- +1-503-643-3422
- info@scaleoutsoftware.com

- [@ScaleOut_Inc](https://twitter.com/ScaleOut_Inc)
- [linkedin.com/company/scaleout-software](https://www.linkedin.com/company/scaleout-software)
- github.com/scaleoutsoftware
- [youtube.com/user/ScaleOutSoftware1](https://www.youtube.com/user/ScaleOutSoftware1)