



Technology and Innovation

In-Memory Computing Summit Europe to Showcase Strategies for Powering the Digital Enterprise with In-Memory Computing

Terry Erisman



Terry Erisman
Vice President of
Marketing
GridGain

Biography

With over 25 years of technology marketing experience, Terry Erisman has initiated and driven high revenue growth for a multitude of award-winning companies in the SaaS, open source, and enterprise software sectors. Prior to joining GridGain (<https://www.gridgain.com>), he was vice president of marketing for Citus Data, where he was responsible for worldwide marketing and established the first annual PGConf Silicon Valley Conference for the PostgreSQL community.

Before Citus Data, Erisman was chief marketing officer at Percona, Inc. Erisman was also instrumental in helping Percona take over, manage, and grow the Percona Live MySQL Conference & Expo held in Silicon Valley, London, and New York. Prior to Percona, Erisman held executive management positions at DotNetNuke, Project Hosts, CaseCentral, Intevac and Atcor. Percona, DotNetNuke, and CaseCentral were all named to the Inc. 500/5000 list during his tenure.

He holds a Master's in Business Administration from the Stanford Graduate School of Business and an A.B. in Chemistry from Harvard University.

Keywords In-memory computing, In-Memory Computing Summit Europe, Data grids, Memory-centric architectures
Paper type Opinion

Abstract

In-memory computing (IMC) can enable companies to deliver real-time, massively scalable applications for today's digital transformation and omnichannel customer experience initiatives. The In-Memory Computing Summit Europe 2019 (<https://www.imcsummit.org/2019/eu/>), scheduled for June 3 and 4, 2019 in London, is the perfect opportunity for you to learn how your organisation can take advantage of IMC platforms for increased application performance and scalability and become familiar with the newest IMC solutions coming to market.

Introduction

Organisations in financial services, fintech, IoT, software, SaaS, ecommerce, healthcare, telecommunications (and more) currently use in-memory computing to solve the challenges presented by complex, data-intensive applications.

- Using a new architecture that allows simultaneous transactional and analytical processing, IMC is powering trading platforms for a variety of European banks,



Technology and Innovation

supporting high speed trading and power FRTB (fundamental review of the trading book) for real-time compliance reporting.

- IMC is delivering real-time performance for data-intensive and compute-intensive applications in the drug discovery and development industry, allowing one company to perform hundreds of thousands of analyses across multiple parameter sets and assumptions. As a result, the company can now complete in just a few hours or even minutes analyses that used to take weeks. It is also able to launch new initiatives that were previously computationally infeasible, opening new approaches to tackling diseases that could be laying the groundwork for new cancer treatments.
- Leveraging dynamically scalable IMC technology, a sports betting platform is able to constantly update huge amounts of data from multiple live events in multiple sports and make this data immediately available to a vast number of clients. The sports betting platform is able to instantly share new betting opportunities and odds for each event, and currently supports over 500 casino and live dealer games and transacts more than £2 million in bets daily, while processing over 700 bets per second.
- IMC is also enabling credit card companies to update their ML models hourly instead of nightly, which helps banks reduce their vulnerability to new fraud vectors.

What Is In-Memory Computing?

In-memory computing platforms are a widely adopted technology that Gartner predicts will be incorporated into the majority of web-scale applications¹ by 2020. They deliver the application speed and scalability organisations require to successfully deploy data-intensive applications to production in response to digital transformation initiatives. Commercially available IMC platforms typically include some or all of the following:

- **IMDG and IMDB** – IMDGs (in-memory data grids) and IMDBs (in-memory databases) are deployed on a cluster of servers on-premises, in a public or private cloud, or in a hybrid environment. The available memory and compute are pooled, allowing all data to be kept in memory, which eliminates the need for applications to constantly read and write data to disks. An IMDG is deployed atop an existing database and keeps the database in sync. An IMDB holds data in memory for processing, with all data written to disk for backup and recovery. The IMDB can also process against the disk-based dataset, enabling fast restarts and allowing a tradeoff between application speed and infrastructure costs by maintaining only some of the data in memory.
- **Streaming data processing** – Streaming data processing solutions manage the complexity around dataflow and event processing. These solutions enable users to rapidly ingest, transact, and analyze high-volume data streams with real-time performance.



- **Machine and Deep Learning Support** – Incorporating ML (machine learning) and DL (deep learning) capabilities into an IMC platform enables what Gartner refers to as “in-process HTAP”. In-process HTAP enables optimal decision making by allowing real-time updates to machine learning models where transactions are occurring. For ML use cases, in-process HTAP enables a continuous learning framework that lets organisations dramatically increase the frequency at which an ML model is updated and deployed to production. For DL use cases, the continuous learning framework can simplify and accelerate moving data into a separate deep learning infrastructure for model training.

In-memory computing can power HTAP (hybrid transactional/analytical processing) or HOAP (hybrid operational/analytical processing) systems which can access a single dataset for simultaneous transactional and analytical processing. This eliminates the need for time-consuming ETL (extract, transform, load) processes, which require moving data to a separate analytical infrastructure and results in higher costs and delays in analytical processing.

The In-Memory Computing Summit Europe 2019

The In-Memory Computing Summit Europe 2019 (<https://www.imcsummit.org/2019/eu/>) is the only place to learn about the full range of in-memory computing-related technologies and solutions. The conference offers multiple tracks, including the latest advances in IMC architectures, new hardware and software solutions for taking advantage of IMC platforms, the latest streaming data use cases, and lessons learned from actual IMC deployments.

Speakers (<https://www.imcsummit.org/2019/eu/agenda/schedule/day-1>) at the event represent some of the most innovative developers and users of in-memory computing technology, including Barclays, CS Financial Group, Capital One, Cerner Corporation, GridGain Systems, GSI Technology, Hazelcast, Oracle, Redis Labs, Sberbank and ScaleOut Software.

The event is attended by hundreds of industry experts and enthusiasts, creating a terrific setting for learning, networking and sharing experiences with in-memory computing.

Reference

¹ <https://www.gartner.com/en/documents/3534832>