

When can “big data” be “in-memory”?

Eric Sedlar
eric.sedlar@oracle.com

1 Abstract

The monikers of “big data” and “in-memory” are certainly hyped in the database world, but some people might argue that they don’t overlap. The terms are vague enough to be treated as mutually exclusive as well as mostly overlapping. In addition, “big data” often implies “data science” which means “not SQL” (based on some programmable framework like Map-Reduce, Spark, or Flink). How do we see the “in-memory” and “big data” trends for analytics evolving in the future (separately or together) and what is the role of SQL vs. other frameworks?

2 Biography

Eric Sedlar is Vice President and Technical Director of Oracle Labs (formerly Sun Labs). This position entails figuring out how to transfer research results from Labs research into Oracle products and services, as well as setting overall technical direction for new research projects in Oracle Labs. Eric manages over 160 fulltime researchers working on over a dozen separate research areas.

Eric’s own research interests are in domain-specific languages and acceleration of database operations both via new hardware and using JIT compilation.

Previously, he led the effort for XML-native storage inside Oracle, starting with Oracle 9iR2. Eric has held various architecture and development management positions at Oracle since starting there in 1990. He holds over 68 patents, and has served on standards organizations for Oracle in the W3C and IETF. He co-authored the Best Paper at SIGMOD 2010 on architecture-sensitive search trees.