

Systap Product Preview

Bigdata[®] is a pure-Java noSQL parallel database engine, freely available under an open-source license (GPL). The bigdata[®] parallel database engine is a horizontally-scaled, general purpose storage and computing fabric for ordered data (B+Trees), designed to operate on either a single server or a cluster of commodity hardware. Bigdata[®] uses dynamically partitioned key-range shards in order to remove any realistic scaling limits - in principle, bigdata[®] may be deployed on 10s, 100s, or even thousands of machines and new capacity may be added incrementally without requiring the full reload of all data. Bigdata[®] comes packaged with a high-performance RDF database supporting RDFS and OWL Lite reasoning, high-level query (SPARQL), and datum level provenance.

SYSTAP and Nokia have been working in close combination to bring online a High Availability (HA) architecture for bigdata[®]. The HA architecture enables:

- Retention of infinite database history. (Bigdata[®] is an immortal database with the ability to read from historical commit points.)
- Service failover and a high degree of operational continuity.
- Arbitrary compute concurrency via parallel decomposition of complex tasks.
- Partitioning of compute resources so that long running queries do not interfere with normal low-latency query processing.
- Blindingly fast parallel closure of RDFS and OWL entailments.

Information about bigdata will be presented at two sessions. The first is on Wednesday June 23rd at 2:00 PM and will be about work in progress with Nokia on the HA architecture [1]. The second session will be on Friday June 25th at 11:00 AM and will be on Object Design Patterns with RDF [2].