Edge Intelligence IP Overview

Preface

A cornerstone of the Edge Intelligence platform is the autonomy of the store at the edge of the network - without this, the orchestration of a widely distributed network of edge stores becomes untenable.

The conventional approaches to delivering performance from a database are broadly limited to:

- Optimization, such as using indexing or partitioning, to optimize for specific queries.
- Use of hardware scale-out to overcome the performance limitations of the database engine.

But the variety of queries that an edge store will be exposed to is indeterminate and can continually evolve over time. Not only is this a difficult usage pattern to optimize for, but it would require continuous and multiple discrete optimizations to address the variations in data volume and in data distribution patterns between the edge points and their changes over time. Moreover, an edge store may be storing tens to hundreds of terabytes in a very constrained hardware environment and this prohibits a conventional scale-out approach.

This combination of challenging requirements combined with a lack of any individual supervision or direct management at the edge makes the conventional approaches to delivering performance ineffective.

The autonomy and agile performance provided by the Edge Intelligence platform is achieved through unique intellectual property that is applied to how an edge store structures its data. This intellectual property is described below.

Background

The von Neumann architecture currently used in all general purpose computers, which separates the CPU from memory and storage, inevitably incurs data transfer latency when a CPU needs to fetch data from memory or storage - and with the speed of modern CPU clocks, this latency incurs a huge cost to CPU efficiency. Latency is not just about the speed of the storage system but is also about the time taken to transfer data from storage to the CPU. Latency is governed by the laws of physics and the only way to remove latency is to locate the data right next to the CPU – which becomes less and less feasible as data volumes grow. Latency is now the largest performance bottleneck in modern hardware and is a limitation of physics which cannot be overcome by hardware design.

A solution to minimizing the impact of latency is to reduce the number of memory and storage transfers required. This can be achieved by replicating and associating data in a multitude of different ways to provide good spatial locality for a multitude of access paths. Good spatial locality increases the relevance of data in each and every transfer which reduces the number of transfers required overall, thus reducing the performance impact of latency.

While the replication of data increases the overall storage capacity requirements, the cost of storage and memory, including hard disk and flash disk has consistently decreased over time and inevitably will continue to do so. Therefore replication by spatial locality provides a good long term solution for the von Neumann architecture.

Description

The Edge Intelligence IP provides a data storage model that implicitly provides the multi-dimensional replication and association of data to provide good spatial locality for all of the query elements required by any arbitrary query. This is achieved without any user design input and without impacting data acquisition performance and yet only requires 50% storage overhead of the raw data consumed.

This means that the edge store is able to deliver the agility of the logical relational model without all of the physical design baggage that accompanies conventional database engines. These physical aspects include indexing, de-normalization, partitioning and gathering statistics and only serve to assist the database engine in overcoming the limitations of its simplistic storage mechanism. In fact, an index, a de-normalization or a partition are just structures that deliver good spatial locality for specific queries. Whereas the Edge Intelligence IP delivers a storage structure with spatial locality pre-optimized for all query elements in all combinations.

This delivers a number of important implications.

- Administrators are freed from the effort and time overhead of analyzing, designing, building and altering physical design aspects in response to changing data, increasing volumes and evolving requirements. Data can just be acquired with all queries good to go at all times.
- Query performance utilizes the optimized storage structure instead of being dependent on large hardware scale-out deployments.
- The time and cost of joining data between tables is significantly reduced because optimal storage structures are in place to service them. This eliminates both the cost-optimizer search time with its vagaries; and the significant lag from scanning tables often experienced with hash and sort-merge join strategies.
- The edge store becomes immensely agile to on-demand responsive data discovery and analysis because the edge store always delivers good response times regardless of the schema used.

Implementation

The IP has been developed and implemented within PostgreSQL to deliver a high-performance fully ACID relational edge store that can acquire data at rates typical of Hadoop, deliver operational and forensic queries with response times expected of a highly indexed row store and perform analytical queries with response times expected from a column store.