

Salient

MPP

Massively Parallel Processing

Big Data Really Fast

A Proven In-Memory Analytical
Processing Platform for Big Data

Executive Summary / Overview:

Big Data can be a big headache for organizations that have outgrown the practicality and usefulness of single-server analytical tools, especially where self-service reporting is a high priority. That’s why successful Big Data users are investing in Massively Parallel Processing (MPP) hardware, a scalable computer architecture that leverages multiple commodity CPUs – potentially hundreds or thousands – to tackle large scale analysis.

Salient MPP is a super scalable, in-memory, multi-dimensional analytical data platform that defeats traditional limitations of speed, granularity, simplicity and flexibility in use. When combined with Salient’s discovery visualization user interface, SIM, it provides an overall analytical solution preferred by executives, analysts and basic users to perform simple through complex analytics much faster than previously possible.

Because companies collect and store more and more granular data from a widening number of business operations and other sources, databases have grown into huge information silos containing millions, even billions, of data records. The compounding effect of database size and number, together with increasing user demand for interactive transactional and managerial decision support, has required analytical systems able to integrate information from many processes, and scale to very large data volumes without sacrificing ease of use, query performance or uptime.

The Salient MPP system represents a new approach to the problem of high- volume information distribution, taking maximum advantage of ever-increasing modern computing power, and using a variety of techniques to achieve new levels of scalability and performance, without sacrificing the informational value contained in the fine grain of everyday business data.

Salient has successfully provided Big Data solutions for 10 years

- *Billions of rows*
- *Many sources*
- *Response under 3 seconds 95% of time*



Large Data In-Memory Processing

Keeping data in random access memory (RAM) allows a system to process data hundreds of times faster than by electro-mechanical input-output (processor to disc) operations. Through advanced data compression techniques, MPP can handle very large volumes, and at the same time, take advantage of the speed of in-memory processing. This speed advantage is enhanced by Salient's proprietary n-dimensional GRID indexing scheme, which enables a processor to go through only that portion of data most relevant to the specific query.

MPP also takes full advantage of both multi-threading platforms and multi-processor machines to accommodate very large numbers of concurrent user queries, without performance degradation. Increasing the number of processors will scale the number of concurrent users in near linear fashion.

Salient MPP is optimized for performance

- *In-memory analytical database*
- *Advanced data compression*
- *N-dimensional GRID indexing*
- *Multi-threading*
- *Multi-processors*

Salient MPP with Partitioned Arrays of Commodity Equipment (PACE)

Salient MPP uses a “divide-and-conquer” methodology to achieve very high scalability. A Salient PACE array consists of a server, known as an Array Master, controlling multiple individual servers, or Nodes. The Array Master is responsible for presenting a coherent view of the data to clients. Users (clients) may access single-server data on any one server, or distributed data through the Array Master.

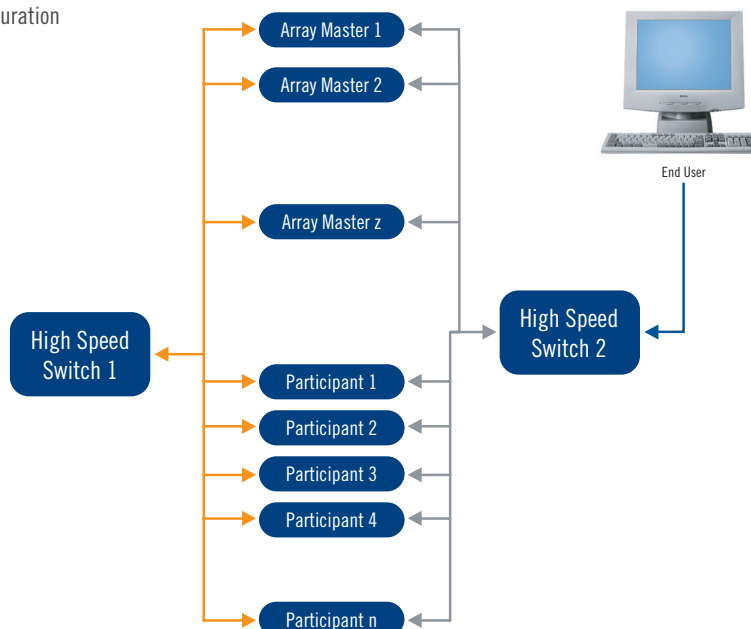
Any of the Nodes in the MPP PACE have complete control over their portion of the data and their computing assets, and use multi-threading to take advantage of multiple CPU's to enhance performance.

Salient avoids common analytical database limitations

- *Limits to dimensionality*
- *Limits to historical data permitted*
- *Limits to cube size*

Figure2: PACE Configuration

Array Masters and Participants share a 'private' network segment for communication performance



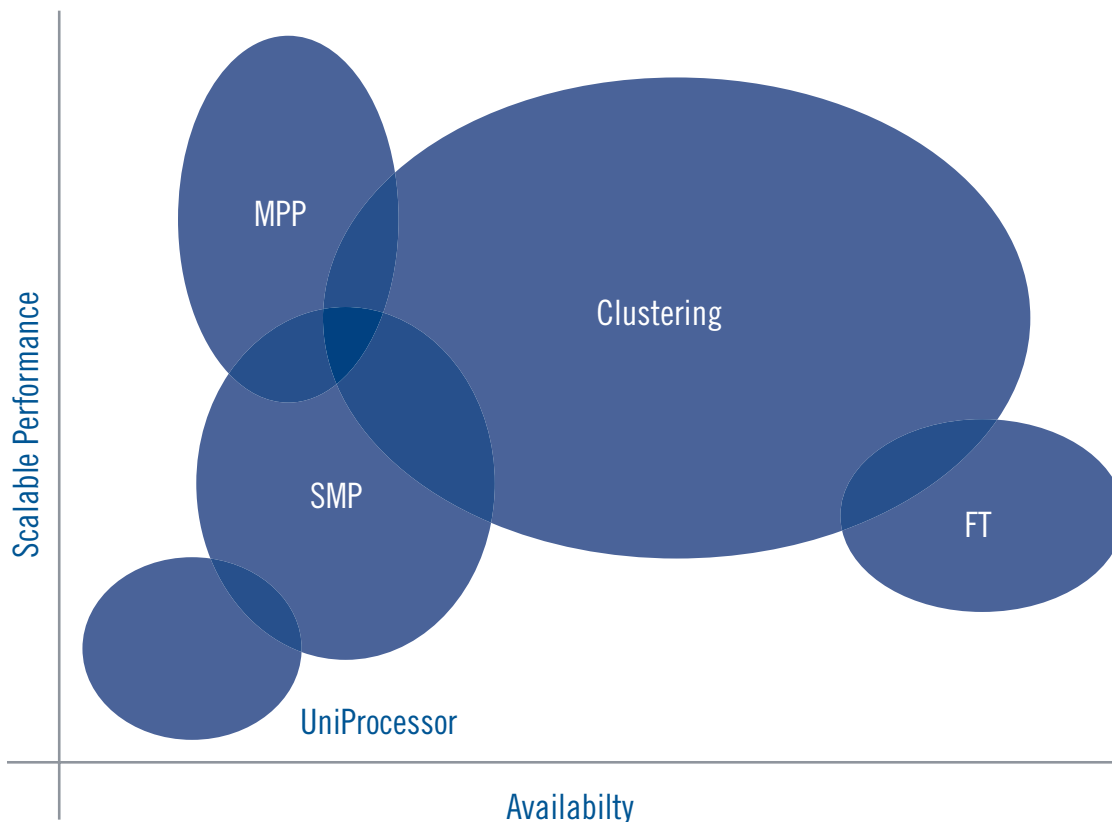
This is Massively Parallel Processing (MPP) in its most efficient form. As shown in Figure 3, MPP is the most scalable form of multiprocessing. The Salient MPP PACE also allows for each Node to have its own “hot” backup to push availability to its limit.

Salient MPP’s architecture eliminates barriers commonly found in other analytical databases, such as limits to dimensionality, data history and cube sizes.

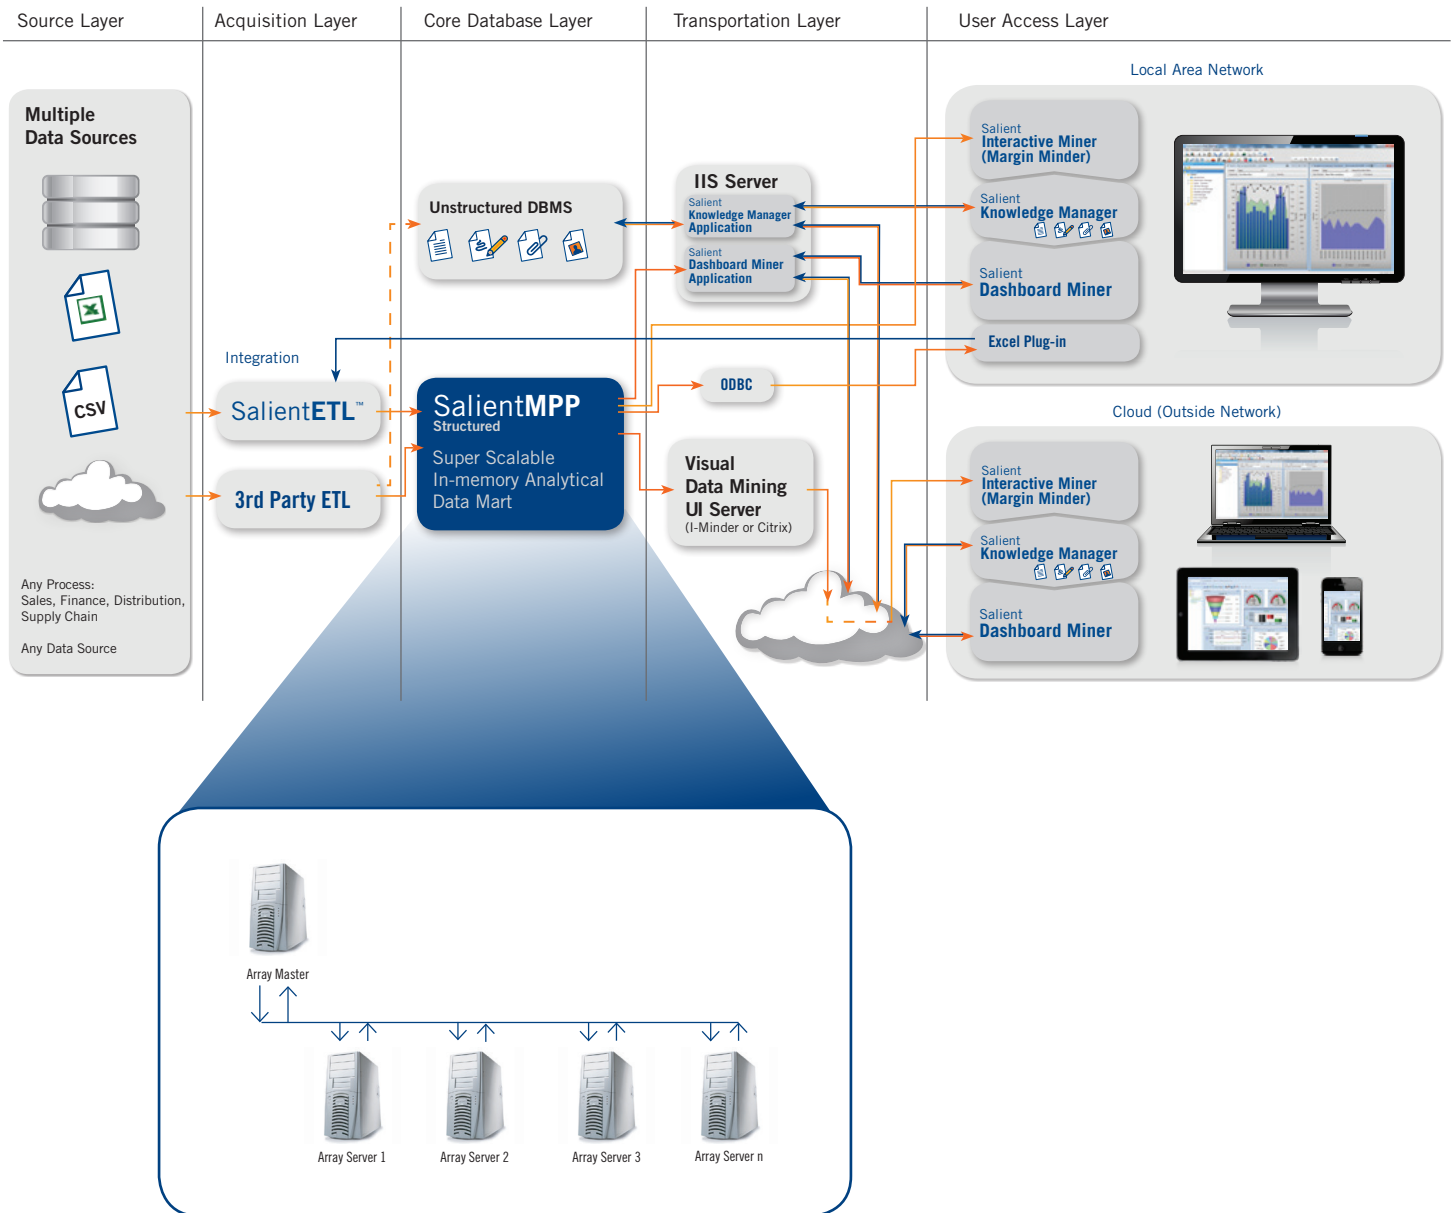
The New York State Office of the Medicaid Inspector General, which uses the Salient MPP solution to identify potential instances of waste, fraud and abuse in its \$52 billion dollar program, is a great example of a leading Big Data user:

- 10 million recipients
- 6 years of data (8 terabytes)
- \$250 million-plus in claims
- Over 4 billion transactions
- 6.9 billion total records

Figure 3: Scalable Performance (Y) vs. Availability (X) for Monolithic, SMP, MPP, FT and Clustering Systems [1]



Salient Enterprise Intelligence Infrastructure Illustration



Putting it all Together

Salient's MPP PACE technology provides consistent throughput with solution scalability to many billions of records. The diagram below depicts the life cycle of a user request utilizing an Array Master server and many Nodes. The Array Master server adds little overhead to the ultimate throughput of the client request.

Typical Request Elapsed Time
 approximately
.5-5 seconds


~ 0 sec - approx. 0 secs



MPP Client



Array Master



MPP request made by client ~ 0 sec

Latency time request to Meta Server ~ 0 sec

Meta turns client request into a request per appropriate Array Server ~ 0 sec

Latency time Meta request to appropriate Array Servers ~ 0 sec

Results returned to client ~ 0 sec

Latency time Meta results to client ~ 0-.5 sec (based on result set size)

Meta server consolidates and aggregates results from Array Servers ~ .1-.5 sec (based on result set size)

Latency time Array Server results returned to the Meta Server ~ 0-.5 sec (based on result set size)

Array servers process individual requests ~ .1-.3 sec (slowest of requested Array servers)

Assumptions:

- Client connects to Meta Server minimum of 100 Base -T
- Meta Server connection to Array Servers 1000 Base-T