SALIENT® MANAGEMENT COMPANY

UXT® Technology Overview:

A Proven In-Memory Analytical Database

UXT®: A Proven In-Memory Analytical Processing Platform

Executive Summary / Overview:

Salient Corporation's UXT 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 through visualization user interface, SIM, Salient provides an overall analytical solution that is used by executives, analysts and basic users to perform simple through complex analytics – many times faster than previously possible, Salient users are learning more and making better decisions than ever.

Salient has successfully provided Big Data solutions for 10 years

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

Because companies collect and store more and more granular data from a widening number of business operations, databases have grown into huge information silos containing millions, even billions, of data records from business operations and other sources. The compounding effect of database size and number, together with spreading user demand for interactive transactional and managerial decision support, have created a need for analytical systems to integrate information from many processes and scale to very large data volumes without sacrificing ease of use, query performance, or uptime. This paper will discuss how Salient's UXT® technology achieves high performance and high scalability.

The Salient UXT® 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 Enterprise Intelligence Infrastructure Illustration

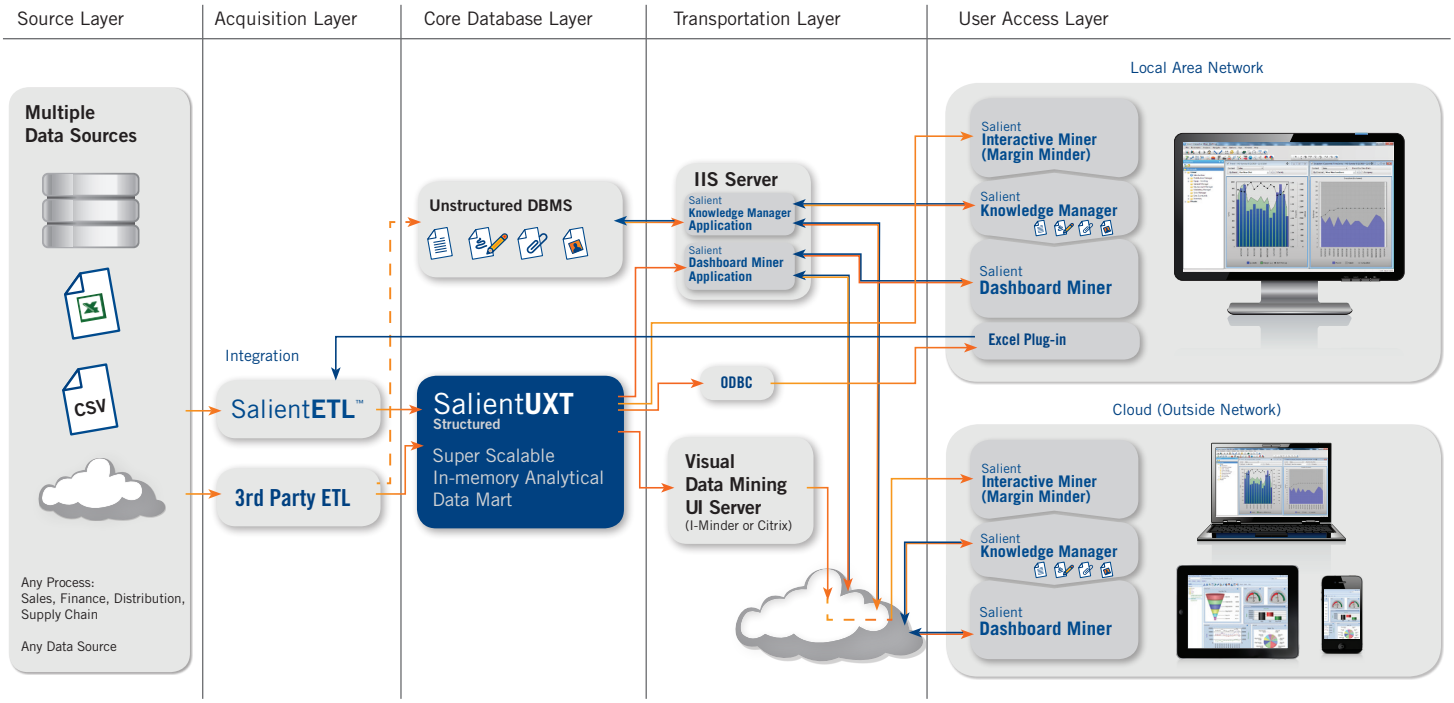

SALIENT® MANAGEMENT COMPANY

Single Machine Scaling

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. Advanced data compression techniques allow UXT to handle very large volumes, and at the same time, take advantage of the speed of in-memory processing. This speed advantage is compounded by Salient's proprietary n-dimensional GRID indexing scheme, which allows a processor to go through only that portion of data most relevant to the query at hand.

UXT 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.

Today's 64 bit operating systems, such as Windows™ Server 2008, can leverage nearly unlimited memory to keep extremely large data volumes in memory on a single PC server. A fully configured server (with 32 GB of random access memory), for example, will handle more than 500 million typical sales records and accommodate hundreds of concurrent users with query response time, typically, of less than 5 seconds.

Salient UXT is optimized for performance

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

 203 Colonial Drive, Horseheads, NY 14845 USA phone 607.739.4511 / 800.447.1868 fax 607.739.4045 www.salient.com

Massive Scaling With Partitioned Arrays of Commodity Equipment (PACE)

Salient UXT uses a 'divide-and-conquer' methodology to achieve very high scalability. A Salient PACE array consists of a server, known as Array Master, controlling multiple individual servers, referred to as 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 UXT 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.

This is Massively Parallel Processing (MPP) in its most efficient form. As shown in Figure 2, MPP is the most scalable form of multiprocessing. The Salient UXT PACE also allows for each node to have its own hot backup to push availability to its limit.

Salient UXT's architecture doesn't have limitations commonly found in other analytical databases. These limitations include, limits to dimensionality, the needs to drop historical data as new data is added and limiting cube sizes to the capacity of a single server.

Salient avoids common analytical database limitations

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

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

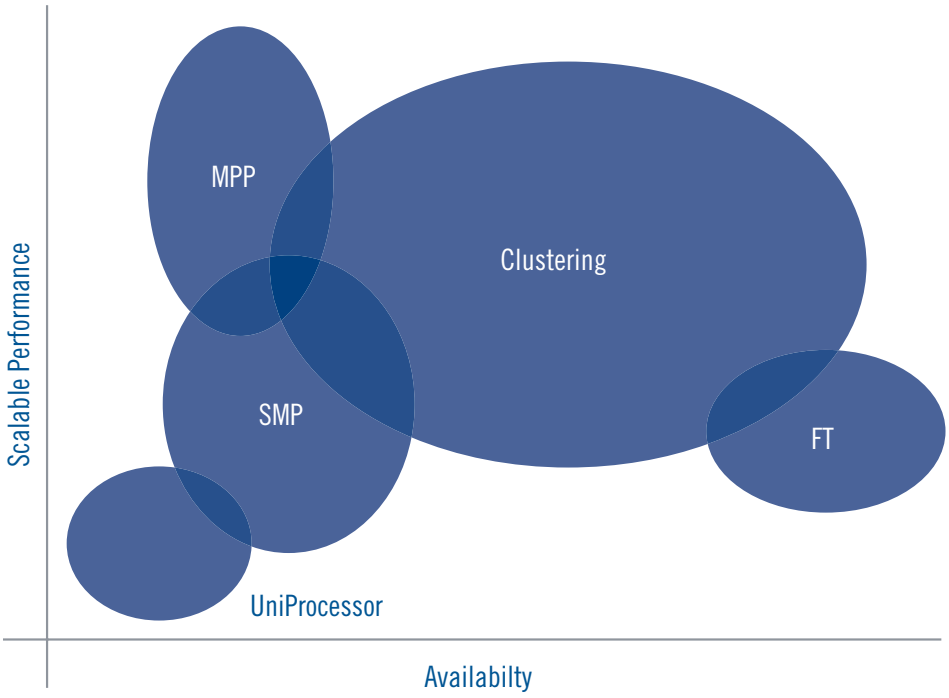

SALIENT® MANAGEMENT COMPANY

Multiple Access Points

UXT PACE array technology scales to thousands of users by providing for multiple access points to a partitioned database. Users who need to access data that is contained entirely on one of the Nodes, will access only that machine. On the other hand, those who need access to data that is spread across more than one participant will gain access through an Array Master.

A useful example is a major tennis event, where many matches are happening at the same time. Some people may hold tickets for only one match, and some may hold tickets for all of the matches. If everyone is required to go through one entrance and exit, then an artificial bottleneck is created, and the spectators may not get seated in time. On the other hand, if each court has its own set of entry and exit points, all fans can get to their seats before the matches start. This corresponds to UXT's direct Node access scheme.

The tennis complex may also have an area where special ticket holders and media reporters can move about more freely from match to match. This corresponds to access through the Array Master.

Figure2: PACE Configuration

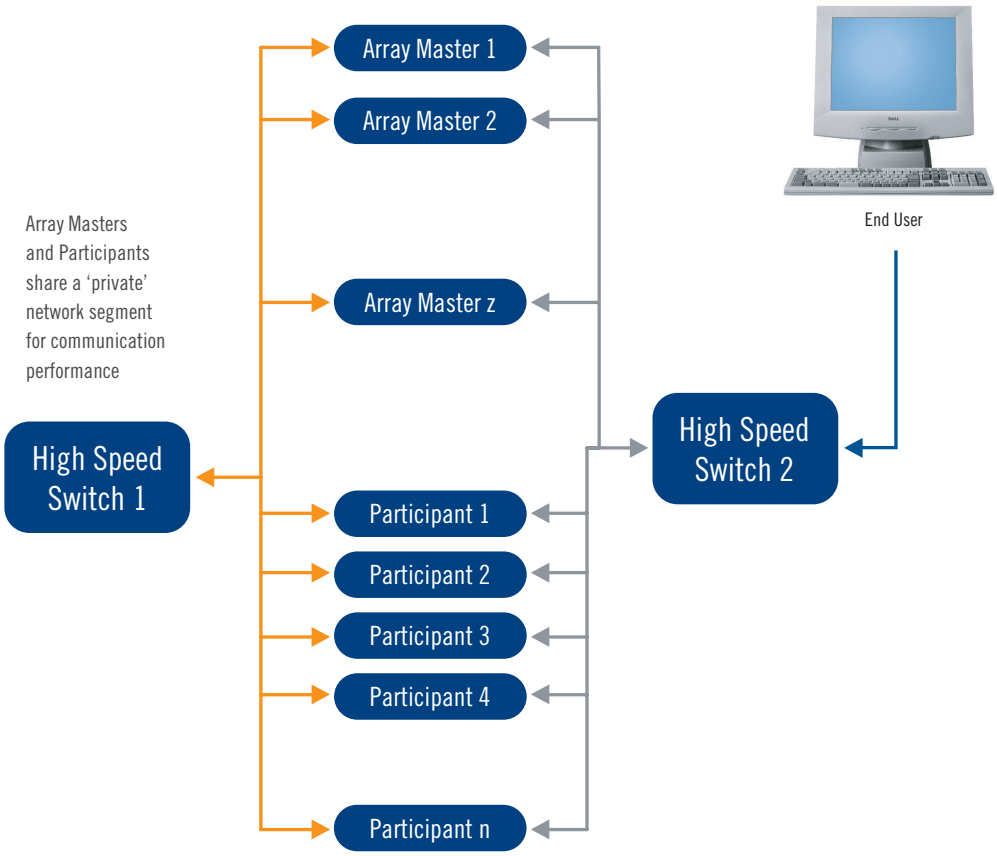

Putting it all Together

Salient's PACE Technology provides comparable throughput to an individual UXT server solution with solution scalability moving into many billions of records. The diagram below depicts the life cycle of a user request utilizing an Array Master server and many Array servers. 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

UXT Client

Array Master

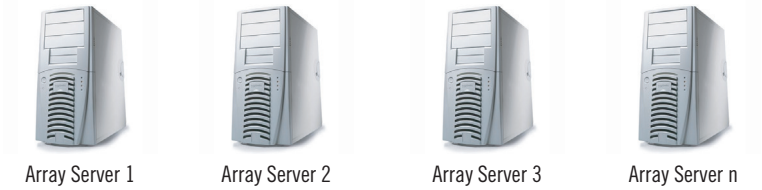

Array Server 1 | Array Server 2 | Array Server 3 | Array Server n

UXT 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

SALIENT MANAGEMENT COMPANY