Virtualization for Big Data Environments – Redefining In-memory Computing

Importance of in-memory computing

The growing performance gap between RAM speeds (measured in >10GBps) and hard disk drive speeds (measured in ~1GBps) is requiring data-hungry applications to utilize RAM rather than slower hard disk drives.

By relying on the main memory for data storage, as opposed to the use of the hard disk drive storage mechanism has several advantages, specifically when it comes to database management system (DBMS). As stated on Wikipedia[1]:

> *Main memory databases are faster than disk-optimized databases since the internal optimization algorithms are simpler and execute fewer CPU instructions. Accessing data in memory reduces the I/O reading activity when querying the data which provides faster and more predictable performance than disk.*

Applications that require faster response time will benefit from in-memory databases.

In-memory databases – limitations

The main limitation for performance with in-memory databases (and any other in-memory applications) is the amount of RAM available in the host system.

Hardware limitations

The cost of RAM for a certain capacity grows with its density. In today's economics, a 16 GB DIMM costs about $7.5/GB while 32 GB DIMMs cost about $30/GB. In addition, number of DIMMs per system is linked with number of processors, starting with dual-socket systems allowing for 8-12 DIMMs per socket and ending with quad-socket systems (or larger) allowing for 16 DIMMs per socket. As a result, to increase number of DIMMs the number of sockets in the system needs to be increased. The cost of processors that fit into systems with higher socket-count systems is higher than those that are installed into lower socket-count systems. Overall, the cost per GB of RAM in a system is increase by over 2.5x (from $15.00 to $39.02) when growing system size.

Sock. / System	Sock. Type	$ / Sock.	DIMM / Sock.	RAM / System	$K / System	$ / GB
Native Systems						
2	Intel E5 (2.4GHz, 8C)	1,440	12	384 - 768	5,760 - 25,920	15.00 - 33.75
4	Intel E5 (2.4GHz, 8C)	2,725	12	768 - 1,536	16,660 - 56,980	21.69 - 37.10
4	Intel E7 (2.4GHz, 10C)	4,394	16	1,024 - 2,048	25,256 - 79,016	24.66 - 38.58
8	Intel E7 (2.4GHz, 10C)	4,616	16	2,048 - 4,096	52,288 - 159,808	25.53 - 39.02
vSMP Foundation						
4	Intel E5 (2.4GHz, 8C)	2,725	N/A	4,096	77,744	18.98
4	Intel E5 (2.4GHz, 8C)	2,725	N/A	6,016	106,224	17.66

Table 1: Cost of RAM as size of system. System cost consider cost of processors (socket) and RAM only.

Software implications

Enterprise software licenses are typically issued on per socket basis, requiring the user to pay for sockets that only drive memory capacity. That raises the overall cost to the user.

[1] http://en.wikipedia.org/wiki/In-memory_database

Oracle Database 11g Enterprise Edition

As an example, Oracle Database 11g Enterprise Edition is priced on per processor basis. Each processor in a system where the software is installed is charged[2] either at $42,750 per processor for a perpetual license, or $8,550 per processor on an annual subscription[3]. In both cases, the first year of support costs $9,405.

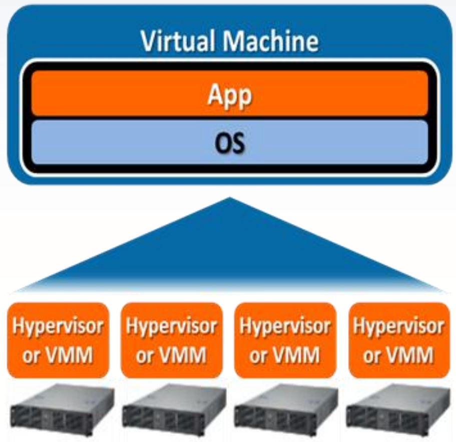

Figure 1: ScaleMP vSMP Foundation

For increased memory capacity for 4 TB, as shown in Table 1, a user may need to grow number of licenses from 2 sockets to 8 sockets resulting in $500K for perpetual license over a 5 year period, or of just over $105K for a one year annual subscription.

In addition, a customer that would like to use a 4 TB coupled with dual-socket system would be able to use Oracle Database 11g Standard Edition (which support maximum of 2 sockets per system), saving an additional 60% of the per-processor cost[4].

ScaleMP vSMP Foundation

ScaleMP has developed the technology to virtualize separate servers to create a powerful shared-memory system (also known as symmetric multi-processor, or SMP). By connecting together different physical servers with a high speed network, systems created with the vSMP Foundation™ software are an aggregation of each of the individual servers CPU, RAM and I/O resources. Very large virtual systems can be created which contain thousands of CPUs, terabytes (TB) of memory and very significant I/O capabilities. Using this same technology within a server that contains both general purpose CPUs and Intel Xeon Phi coprocessor cores and memory, a very powerful virtual SMP can be created, containing many more cores and more memory than previously available in a single system.

The advantages of shared-memory systems

Shared-memory systems ease the programming tasks for developers. By relying on compilers to help with the parallelization-effort and using shared memory rather than distributed memory for certain tasks, it has been shown that the programming is much simpler than using a distributed system. A virtual SMP also simplifies the management of a large amount of resources compared to the management of a cluster of individual servers.

vSMP Foundation for Memory Expansion

The vSMP Foundation software from ScaleMP is now capable creating VMs using processors of one system and RAM of other systems. This enables the user to:

- Aggregate the entire RAM from all of the systems, enabling any process to use all of the aggregated memory.
- Use only the main systems processors, which can reduce licensing costs, but allow for use of all of the memory.
- Optimize the system price by using low-cost, lower-performance CPUs where not required. The main server which runs the application should have the highest performing CPUs available. However, the servers which are present and are only used to house the addi-

[2] http://oln.oracle.com/static/opn/BUS_PRAC/Full_Use_Tech/Pricing_Tech/2_3_6_techlicensemodels.htm
[3] https://shop.oracle.com/pls/ostore/product?p1=OracleDatabaseEnterpriseEdition
[4] https://shop.oracle.com/pls/ostore/product?p1=OracleDatabaseStandardEdition

tional memory can be outfitted with lower cost and lower performing CPUs.

The chart below demonstrates the performance of database application handling 100 GB of raw data, which requires about 43GB for runtime processing. Running database that is twice the size cannot fit into a 96 GB system (due to the runtime memory footprint), as the runtime memory consumes almost all system memory leading the application to run out of memory.

With vSMP Foundation and using 3 systems for memory expansion (total cost $25,000), the customer is able to run databases up to 500 GB, in a pay as you grow methodology without needing to add more DBMS software licenses. As the size of the database grows, scaling using vSMP Foundation across multiple physical servers is very good. Without the aggregation made possible by vSMP Foundation, the largest example, of a 500 GB database would not be possible.

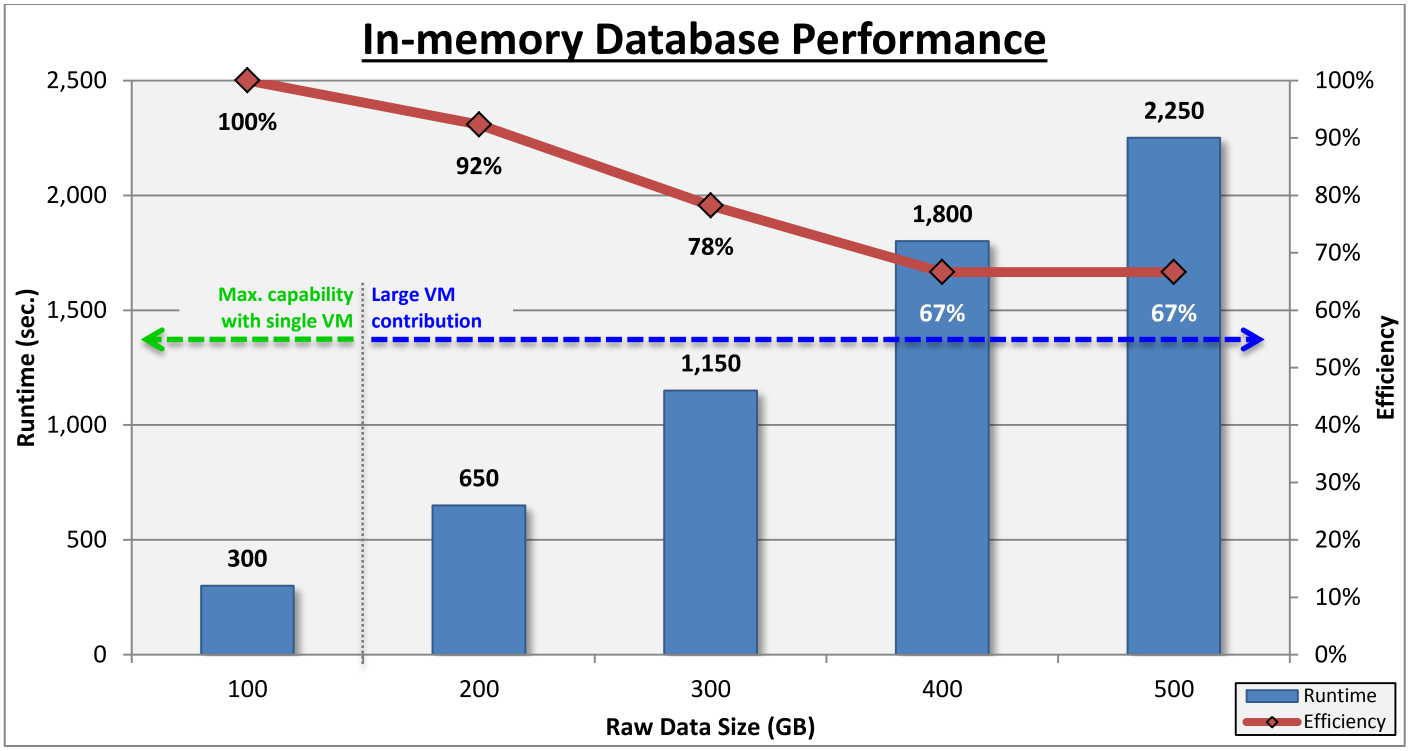

Figure 2: Performance of a database requiring memory footprint that is 4 times larger than a single server can hold, using ScaleMP's vSMP Foundation 4.0 in a Memory Expansion usage model.

Summary

Applications that perform better when using more memory or have the ability to use more memory as the data required grows can benefit from a unique solution from ScaleMP. vSMP Foundation aggregates multiple servers memory to present to the application more memory than is physically available on the system where the application is running. This reduces the cost in both the hardware acquisition as well as the licensing costs when running database oriented software applications. By utilizing industry standard servers, virtual systems can be created that are cost optimized and which also contain large amounts of resources.

Want More Info? Want to Test vSMP Foundation?

Need additional technical information, system requirements or want to know more about testing and implementing vSMP Foundation? Visit our site www.scalemp.com or mail info@scalemp.com