

Dolphin White Paper

***Dolphin Express:* The Most Cost Effective Database Cluster Interconnect Solution**



Dolphin Interconnect Solutions ASA

Reg. No. 984 861 060 in the Norwegian Register of Business Enterprises

Revision 1.1, June 12th, 2007

Companies spend significant resources on choosing and deploying a database solution to meet their business needs. Typical criteria used to quantify the best option for an enterprise include performance, scalability, reliability, availability, and cost. While most would assume Ethernet should provide the most cost effective option available, it turns out, if we compare against Dolphin Express on a cost per transaction rate basis, Gigabit Ethernet is significantly more expensive for even the smallest configurations. In this paper, we will apply recently released performance results to compare database transaction cost for Gigabit Ethernet versus Dolphin Express.

In a recent paper, Mikael Ronström from iClaustron AB, a database and Open Source consultancy, completed a series of benchmark tests using standard database cluster software, the latest systems from Intel and currently shipping Dolphin Express interconnect. The database software runs unmodified through the Dolphin Express SuperSockets™ “socket” interface. His paper clearly shows the performance and scalability advantages enabled by Dolphin Express. Because Gigabit Ethernet is included with the price of most servers, many would assume it to be a more cost effective solution. As we will show in this companion paper, this is not the case. In fact, Dolphin Express is the best cost/performance solution for enterprise databases by providing over 250% transaction rate cost improvement.

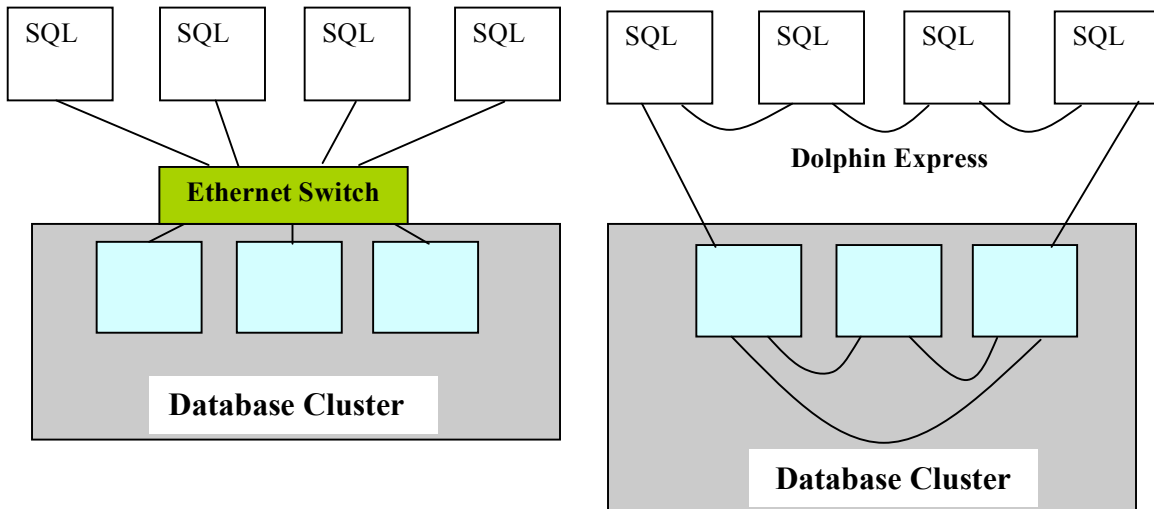
To complete the tests, dual socket, dual core Xeon (5150 processor) based systems¹ were used with 8GB of system memory, two 300GB 15k local drives, dual on-board Gigabit Ethernet, and Dolphin Express HBA. For this paper, we are using the latest pricing available on the internet for each on the system components.

Item	Cost
Server	\$8,770
Gigabit Ethernet Switch	\$318
Dolphin Express HBA ²	\$995

In the first test, seven servers were configured and interconnected with Gigabit Ethernet and Fibre Channel and a second test was executed with Dolphin Express replacing the Gigabit Ethernet between the server nodes.

¹ Systems were 1RU rack servers using 2.67GHz 5150 processors with 8 GB of 667 MHz, dual ranked system memory.

² Dolphin Express HBA cost includes cable(s).



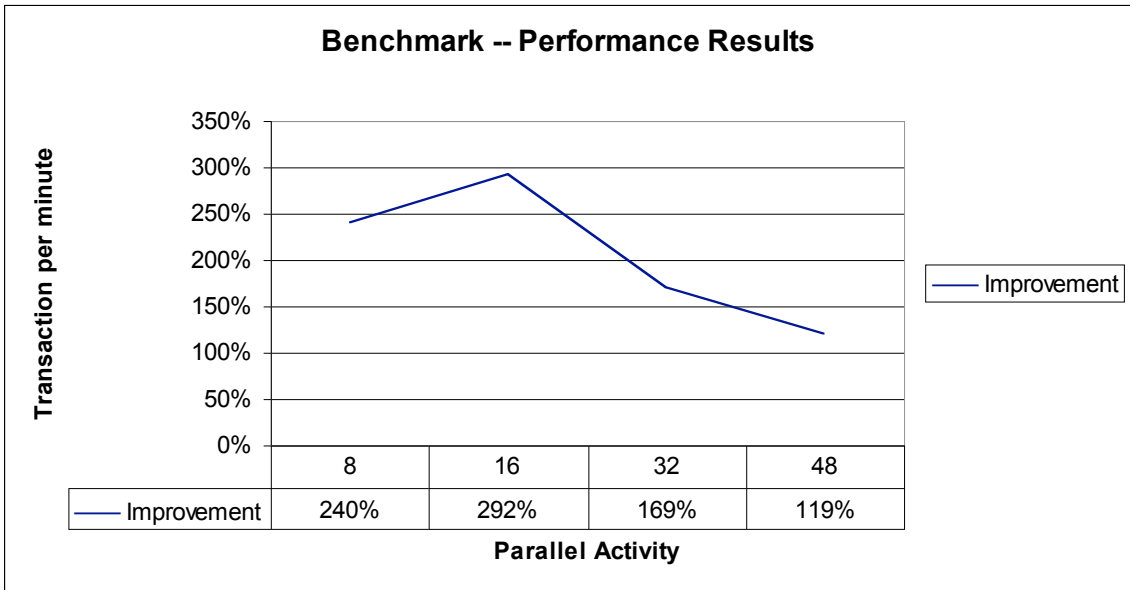
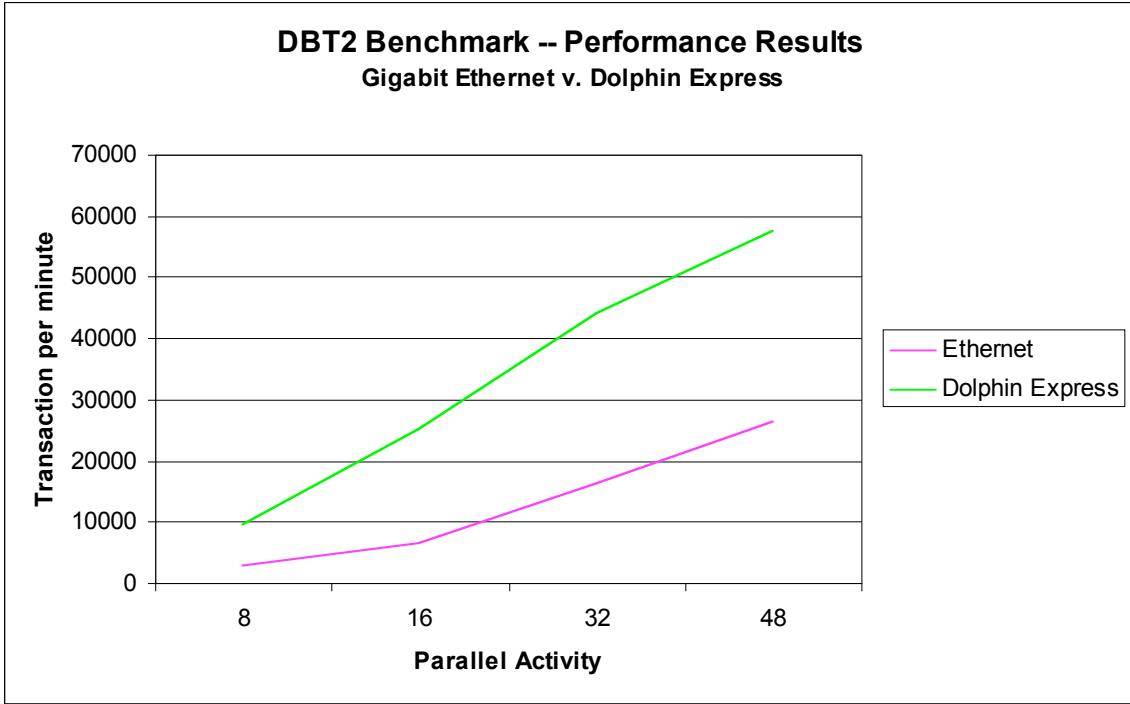
Three of these servers were running the MySQL Cluster³ (NDB blocks), with each CPU operating as one data node (12-data nodes in total). Four servers were executing MySQL server (SQL blocks), with each CPU executing a separate MySQL instance (16 MySQL servers). The cost of purchasing this equipment is shown in the table below:

Gigabit Ethernet		
Item	Quantity	Cost
Server	7	\$61,390
GE Switch	2	\$636
Total		\$62,026

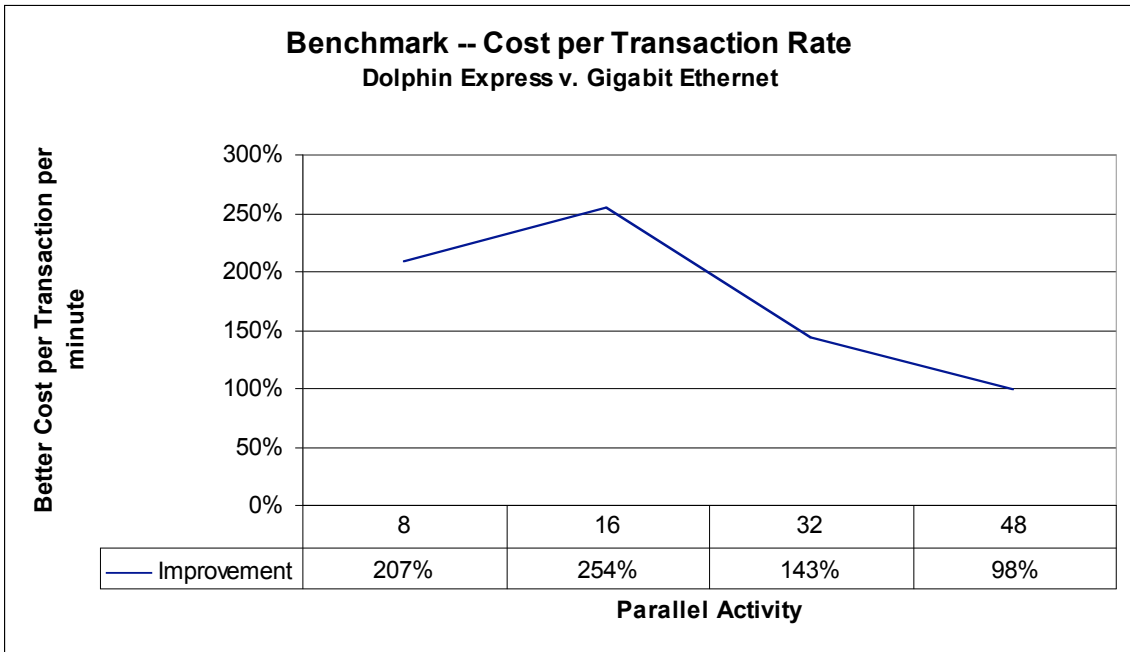
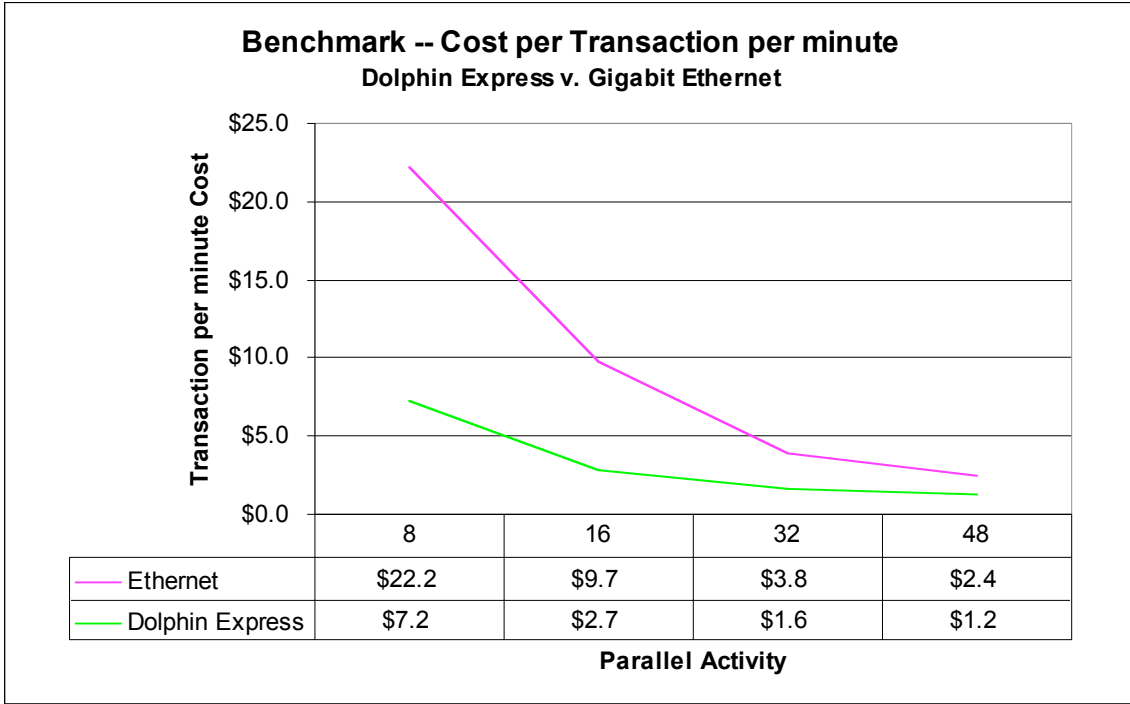
Dolphin Express		
Item	Quantity	Cost
Server	7	\$61,390
Dolphin Express HBA	7	\$6,965
Total		\$68,355

Based on the hardware costs, there is 10% premium for Dolphin Express, although after factoring in the cost of the database, operating system, etc. this premium drops to less than 5%. The chart below shows the performance increase by using Dolphin Express running the open source version of the TPC-C benchmark, DBT2.

³ MySQL Cluster version 5.1 was used for this benchmark.



As shown, there is a considerable improvement in performance using Dolphin Express, with the largest benefit achieved when an application needs many parallel transaction processes. If you compare the above result and factor in hardware cost, the following result is observed.



When parallel requests are low, Dolphin Express provides the greatest cost per transaction rate benefit by as much as 250% better than Gigabit Ethernet. As transaction rates increase, Dolphin Express continues to provide the best cost, performance and scalability. The next question to ask is, “Why?”

Dolphin Express and Dolphin SuperSocket software provides three fundamental capabilities that create significant value for our customers with the lowest transaction rate cost available in the industry. These are:

- 1) Low Application Latency
- 2) High Application Bandwidth
- 3) Minimum Processor Overhead

An analysis of the most common message sizes transmitted between nodes during normal database activity was conducted by iClaustron AB, and the results were that 80% of the messages are small in size (<1 KB). Dolphin has optimized its hardware and software so that small messages can be delivered with the lowest latency available in the industry.

The remaining 20% of the messages are medium (>1 KB and < 8 KB) or large (>8 KB) messages, and the analysis shows that most of the remaining messages are of the large category. To address the bandwidth and overhead needs of large messages, Dolphin uses very high performance DMA engines and the latest IO interface technology, PCI Express, to move the data between nodes with a resulting bandwidth as high as 750 MB/s.

Using Ethernet to communicate between servers in a reliable manner requires that the TCP/IP software stack be employed. The TCP/IP stack is necessary because at the hardware level, Ethernet is inherently unreliable and needs software help to ensure that messages accurately arrive at a specified destination. This software is natively integrated into the operating systems and drivers for Ethernet. The standard metric for TCP/IP overhead is 1Hz per bit, so for a Gigabit Ethernet based systems, each node must dedicate 1GHz of a CPU to processing the TCP/IP stack (~37% of CPU in the test configuration). Dolphin Express hardware includes a number of reliability mechanisms to ensure messages are reliably delivered between servers, which allows the SuperSockets software to minimize the CPU overhead consumed for data transmission. The result is that less than 5% of a single CPU is consumed by overhead tasks and as a result, more transactions can be processed by the system.

Dolphin Express products provide the lowest transaction rate cost database solution available. By providing software transparency, optimized data movement, and reduced CPU overhead, Dolphin Express enables database deployments to occur quickly and reliably with headroom to scale. Dolphin Express offers unique capabilities to IT managers deploying database solutions. Dolphin Express is used worldwide by customers in all industries in mission critical and other environments.