



Dolphin Express

SuperSockets™

Dolphin SuperSockets μ S

Ultra-low latency

TCP/IP

Dolphin SuperSockets

TCP/IP


μ S

TCP/IP

Maximize your database

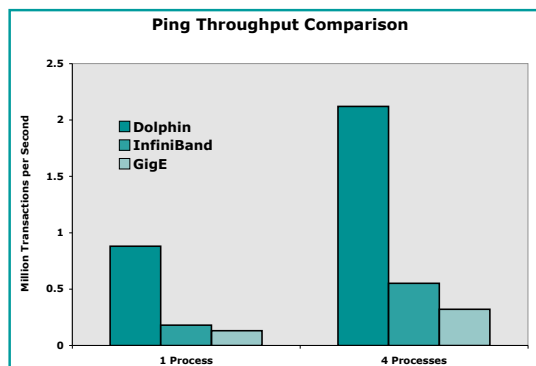
Drive in the fast lane

Dolphin Express SuperSockets™ prevents system bottlenecks and provides super-fast communication for critical applications. It is an innovative implementation of the Berkeley Sockets API designed to run on Dolphin Express hardware. As a plug-in compatible replacement for TCP/UDP/IP protocol stacks, SuperSockets™ complies with the new RDS (Reliable Datagram Sockets) and SDP (Linux Sockets Direct Protocol) standards and delivers unrivalled communication performance for processes run on different individual computers. The combination of Dolphin Express hardware and the SuperSockets™ software provides ultra-low latency, high-bandwidth, low overhead and high availability. It's the ideal solution for applications with significant communication requirements, such as database, client-server and computer-aided engineering.



Eliminates system bottlenecks

The combination of ultra-low latency and high bandwidth increases transactions per second to a level that outstrips alternative solutions and eliminates bottlenecks in the



system. For example, Dolphin Express SuperSockets™ offers users at least ten times better latency than Gigabit Ethernet. In fact, one-byte ping-pong latency is as low as 2.26 microseconds for a user space socket¹.

Dolphin Express hardware provides on-the-fly data integrity checking that eliminates the TCP protocol layer and leaves CPU cycles free for other tasks.

Just plug and play

No special configuration, linking or modifications are required for applications. Just install the Dolphin Express hardware and SuperSockets™ software and run your applications as normal. The solution enables fast and easy upgrades from all existing Ethernet configurations in a matter of minutes. The fast communication ports are selected through configuration files that allow SuperSockets™ to coexist with existing Ethernet links for system management and standard networking purposes.

All applications communicating through TCP/UDP/IP with communication bottlenecks benefit from Dolphin Express SuperSockets™, including older software that has never been supported by high-speed interconnects.

How is ultra-low latency achieved?

The answer lies in the efficiency of the remote memory access mechanism provided by Dolphin Express hardware. Dolphin's technology is based on a combination of RMA (Remote Memory Access) for short transfers and RDMA (Remote Direct Memory Access) for longer transfers.

RMA has clear advantages for short messages, such as lock transfer commands between nodes in a database cluster, as the transfer is completed through a single CPU store operation that moves data from CPU registers into the memory of the remote node².

In most cases, data transfers through SuperSockets™ are completed before alternative technologies have even managed to start their RDMA. SuperSockets™ then further decreases system resources required for sending data and also includes adaptive algorithms to reduce the number

¹ This is measured by running a ping-pong test between two systems. The listed number is half of the full round trip time.

² When the IO system is busy, a four-byte write operation is typically completed in 210 ns. When the IO system is idle, this is completed in one CPU clock cycle, e.g. 300 ps.

e cluster performance



Oracle RAC database solutions benefit significantly from Dolphin Express SuperSockets™

of system interrupts. Quite simply, this shared memory approach is the most reliable, efficient and cost-effective method of performing inter-node communication available today.

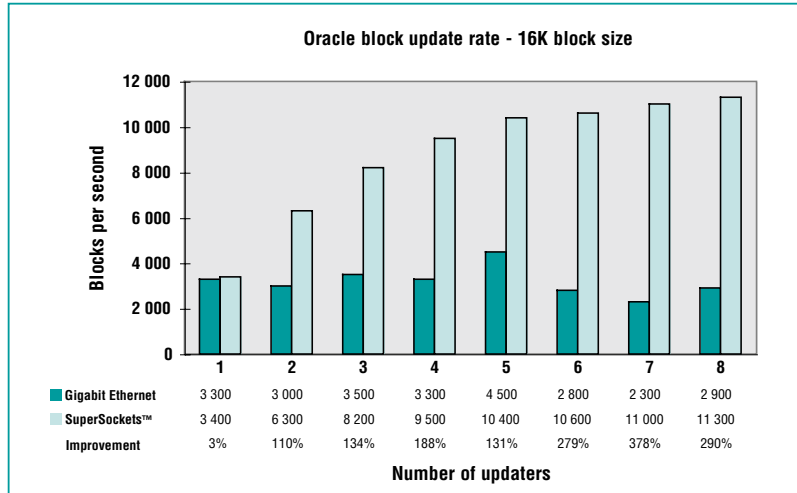
Ideal for a broad range of networking applications

Network and IO applications with a high volume of short messages have the most to gain from Dolphin Express SuperSockets™ – especially where bottlenecks cause systems to slow up. Typical uses are:

- Databases, e.g. Oracle, MySQL Cluster
- Remote file access systems, e.g. NFS, iSCSI
- Cluster file systems, e.g. Lustre, PVFS, GFS
- HPC applications, e.g. Fluent, STAR-CD
- Communication libraries, e.g. MPICH-2, PVM, LAM-MPI
- Web servers

Boost Oracle Grid and RAC performance

Dolphin Express SuperSockets™ speeds up and removes bottlenecks in systems running Oracle and Real Application Cluster (RAC) database solutions. Network transactions in such systems consist mainly of relatively short data transfers in the order of single bytes for inter-process locking, a few more bytes for queries and some 8 – 16 KB for Oracle cache blocks transferred between nodes in the cluster. This concentration of short transfers means that latency is a key network performance indicator – exactly where SuperSockets™ come out on top.



Cost-effective HA clustering

Dolphin's unique distributed-switching technology with automatic instant re-routing capabilities for node or network errors is the most cost-effective way of building high-availability (HA) clusters. The absence of centralized switches eliminates the cost of duplicate switches required with other solutions. As the low-power Dolphin cards insert as standard I/O cards, they use regular node power only. No additional supply is required. In turn, this reduces the number of fans and their risk of failure.

Channel Bonding/Multi-rail

When more than one Dolphin single or multi-channel adapter card is present, Dolphin Express SuperSockets™ provides bandwidth-multiplied performance by automatically using channel bonding to connect to multiple adapters. This feature also includes application transparent instant fail-over between channels.

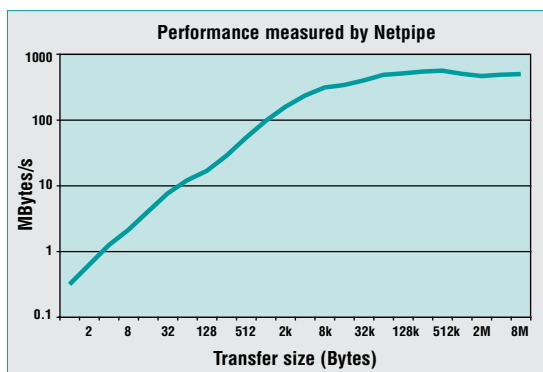


Highly scalable

The Dolphin Express solution is scalable from two to several hundred compute nodes connected with rings, switches or Dolphin's innovative two- or three-dimensional Torus topologies. Having no single point of failure, the Torus topologies are especially well suited for high-availability applications.

Socket bandwidth

Benchmarks using netperf show more than 625 MBytes/s sustained (5,005 Mb/s throughput over one single D350 adapter card).



When it comes to bandwidth, the important thing is bottom-line application performance, not the maximum theoretical bandwidth for the transfer of a very large message. Most applications send only relatively short messages because the data objects to transfer are small. In a database cluster, most messages are very short - in the order of a few bytes only - and none are more than a few kilobytes. Therefore, the most important measure for the cluster interconnect is the bandwidth for short transfers. This is where the Dolphin Express technology excels due to the direct remote memory access (RMA) used for short messages where only one single CPU store operation is needed to send 8 bytes of data. No DMA set-up overhead and no protocol stack processing is required.

Availability

Dolphin SuperSockets™ is currently available for Linux for AMD and Intel x86 and x86_64 processors. Future versions will support Solaris (x86_64 and SPARC) and Windows. Dolphin supplies solutions for any combination of PCI Express/PCI/PCI-X/PMC/cPCI.

Dolphin SuperSockets™ software has been developed by support from the European IST projects GRIA and HPC4U, and is free open source software for Dolphin Express hardware, copyright Dolphin Interconnect Solutions ASA. Please visit <http://www.dolphinics.com> for more information.

About Dolphin

Dolphin Interconnect Solutions of Oslo, Norway and Marlborough, Massachusetts, USA provides high-speed, high-bandwidth interconnect products. Dolphin's unique technology enables customers to build cost-effective, highly scalable, enterprise-class clustering solutions with standards-based hardware and software. Dolphin is listed on the Oslo Stock Exchange with the ticker DOLP.



Dolphin Interconnect Solutions ASA
Olaf Helsetts vei 6
NO-0619 Oslo, Norway
Phone: (47) 23 16 70 00
Fax: (47) 23 16 71 80
E-mail: info@dolphinics.com

Dolphin, Inc.
225 Cedar Hill Street
Marlborough, MA 01752
Phone: (1) 508-786-9950
Fax: (1) 508-786-9785
E-mail: info@dolphinics.com


Dolphin
INTERCONNECT SOLUTIONS
www.dolphinics.com

Note: Dolphin Interconnect Solutions and Dolphin SuperSockets are trademarks of Dolphin Interconnect Solutions ASA in the United States and other countries. Other product names are for informational purposes only and may be trademarks of their respective companies.