Improve Performance and save Money, Space and Power with Dolphin for MySQL based Web Servers

Topics

- Dolphin Quick Facts
- Dolphin Express for Database Clusters
  - Superior TCO
  - Performance
  - Fault Tolerance
- Q&A

connect better

Quick Facts about Dolphin

- Global provider of interconnect solutions since 1994
  - Main locations Oslo, Norway and Massachusetts, USA
- Leading Products and Technology
  - Dolphin Express
    - Enterprise Clustering and I/O Expansion
    - Embedded OEM

Dolphin Express for Database Clusters

TCO, Performance & Fault Tolerance

connect better

Introducing Dolphin Express

Dolphin Express provides the industry's best solution for maximizing the performance of every database application.

- 4X+ for Oracle
- 4X+ for MySQL
- Dramatically increase concurrent user capacity
- Substantially lower cost than alternatives
- No changes to applications or databases
- Transparent and automatic High Availability

connect better

Dolphin Express™

- Integrated Solution
  - Dolphin Express Hardware
  - SuperSockets™ Software
- Transparent to Applications
  - NO changes to applications necessary
  - ALL Applications using Sockets API (TCP-UDP-RDS) will work
  - Use of LD_PRELOAD
- Scalable
  - Scales to hundreds of nodes with 2-D or 3-D distributed switching
- High Availability - Fault Tolerance
  - NO single point of Failure
  - Automatic rerouting in case of node failures
  - Multi-Channel Support with Automatic Channel Bonding
  - Automatic Fail-Over to Ethernet
- Low TCO with Investment Protection
  - Software compatibility with future hardware products

Lowest Cost and Best Performance

With Dolphin

Qty	Desc	Price
6	Servers	$ 45,000
6	RAC licenses	$120,000
6	Dolphin cards	$ 6,000
1	Support (15/20)	$ 30,750
Total:		$201,750

Without Dolphin

Qty	Desc	Price
12	Servers	$ 90,000
12	RAC licenses	$240,000
12	Ethernet cards	$ 1,800
1	Support (15/20)	$ 61,500
Total:		$393,300

- Dolphin solution is $191,550 (49%) less than competitor's
- Dolphin solution is $65,706 less per year in OpEx (source: innovate.intel.com)
- 3 Year TCO savings: $354,000

Conclusion: Dolphin delivers the industry's fastest solution for applications, while dramatically reducing the Capital and ongoing Operational Expenses associated with those environments.

Dolphin Express for Database Clusters

TCO, **Performance** & Fault Tolerance

connect better

Web Server Requirements

- Adequate Capacity
  - Scale with Traffic
    - Throughput
    - Response Time
- Reliable
  - No Downtime visible to Clients
- Cost Efficient

DBT2 Message Size Distribution

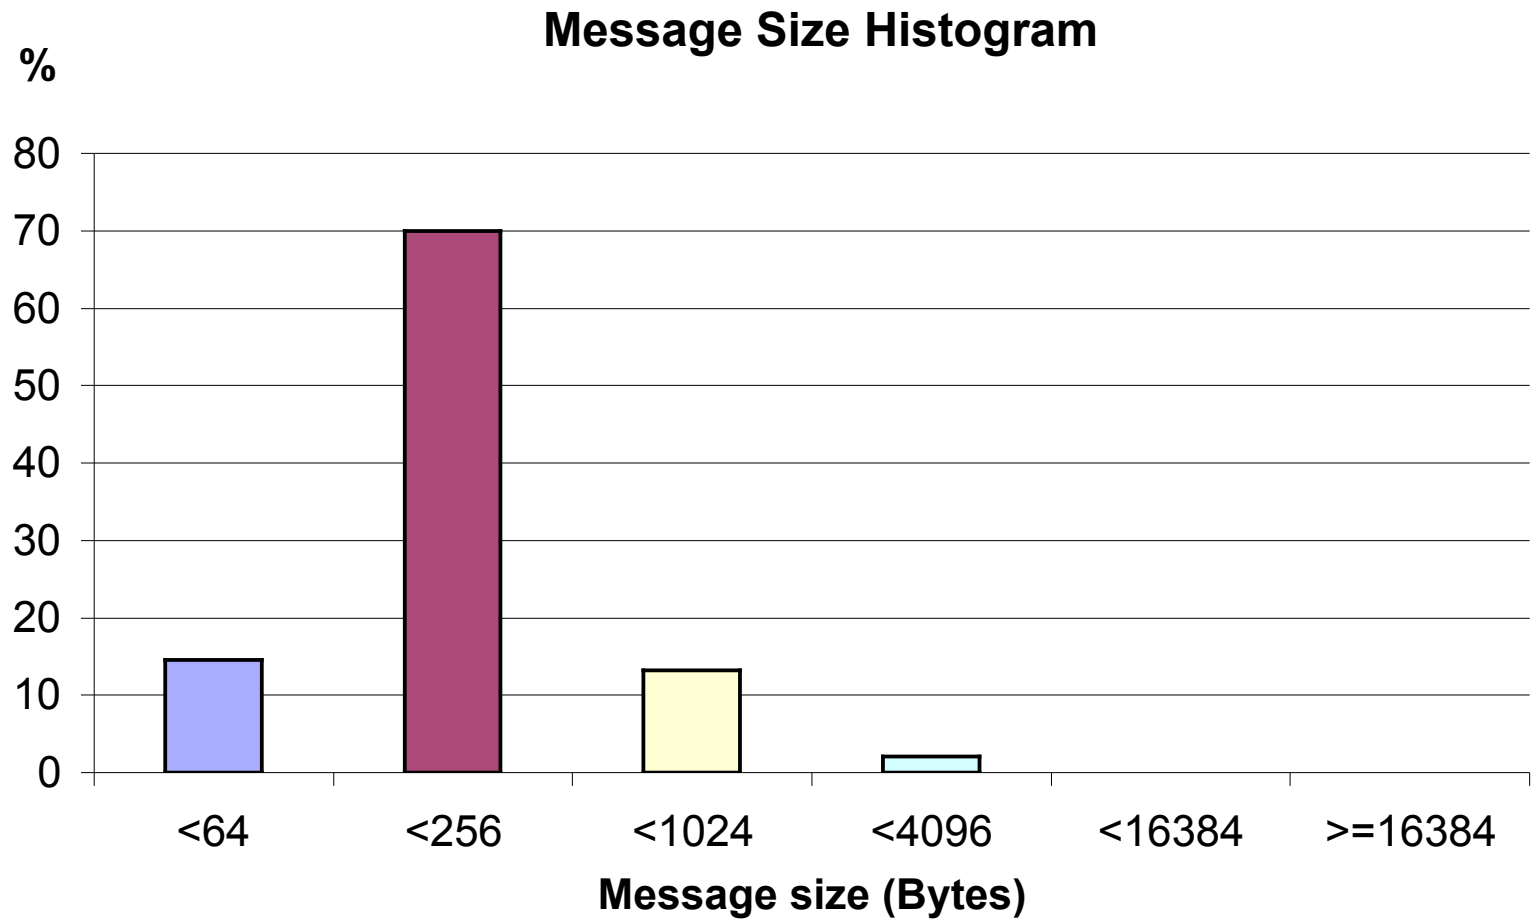

Data from DBT2 Benchmark
by Mikael Ronström

Oracle RAC Message Sizes

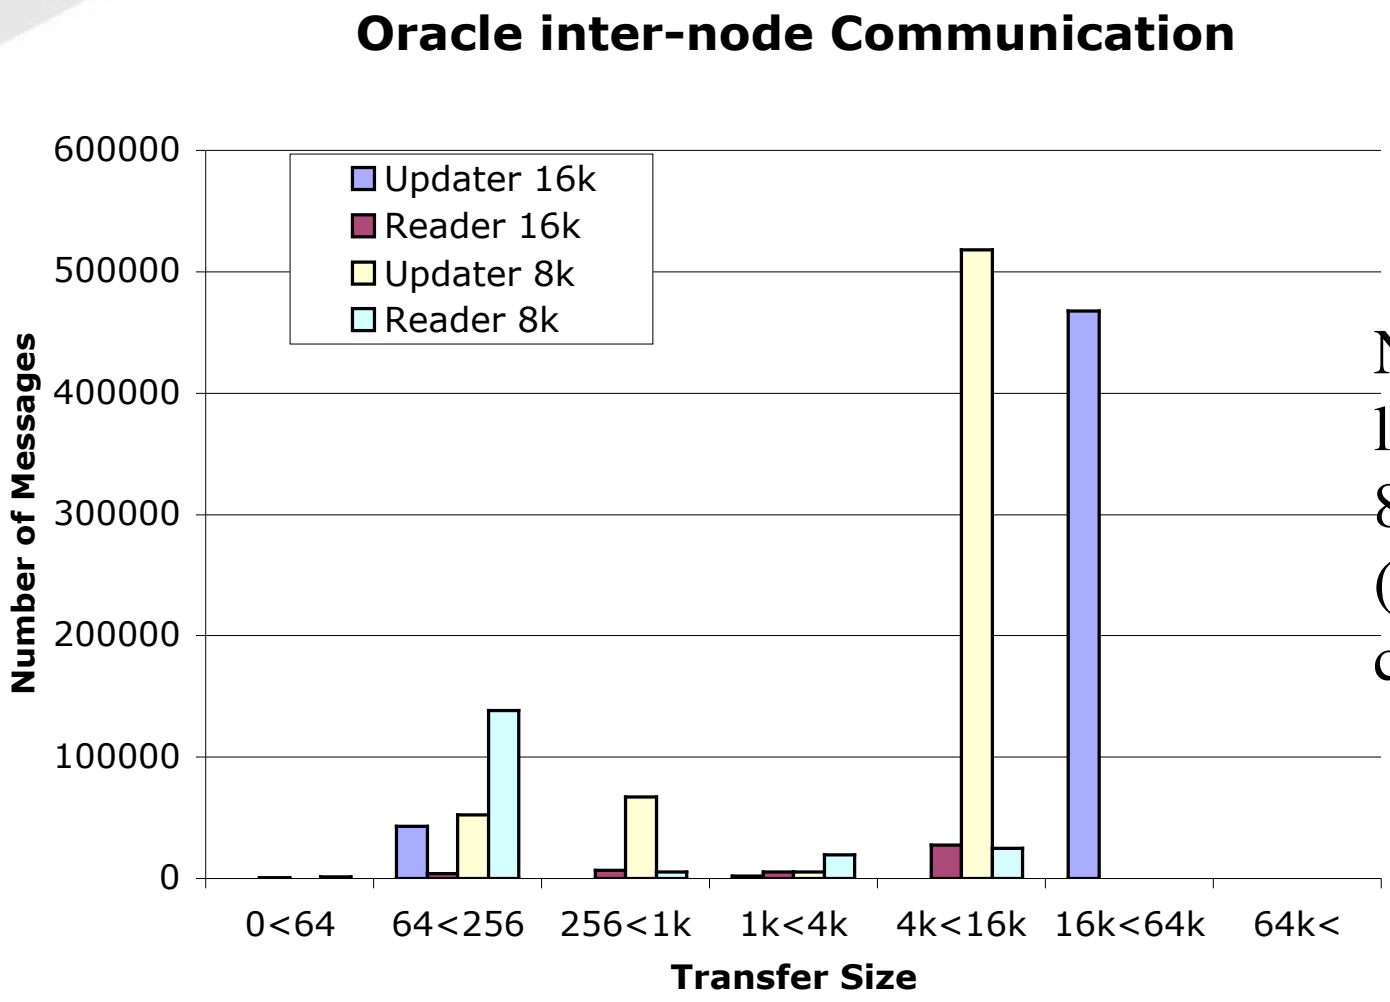

NO transfers larger than 8k/ 16k blocks (16k = max. cache block size)

Latency

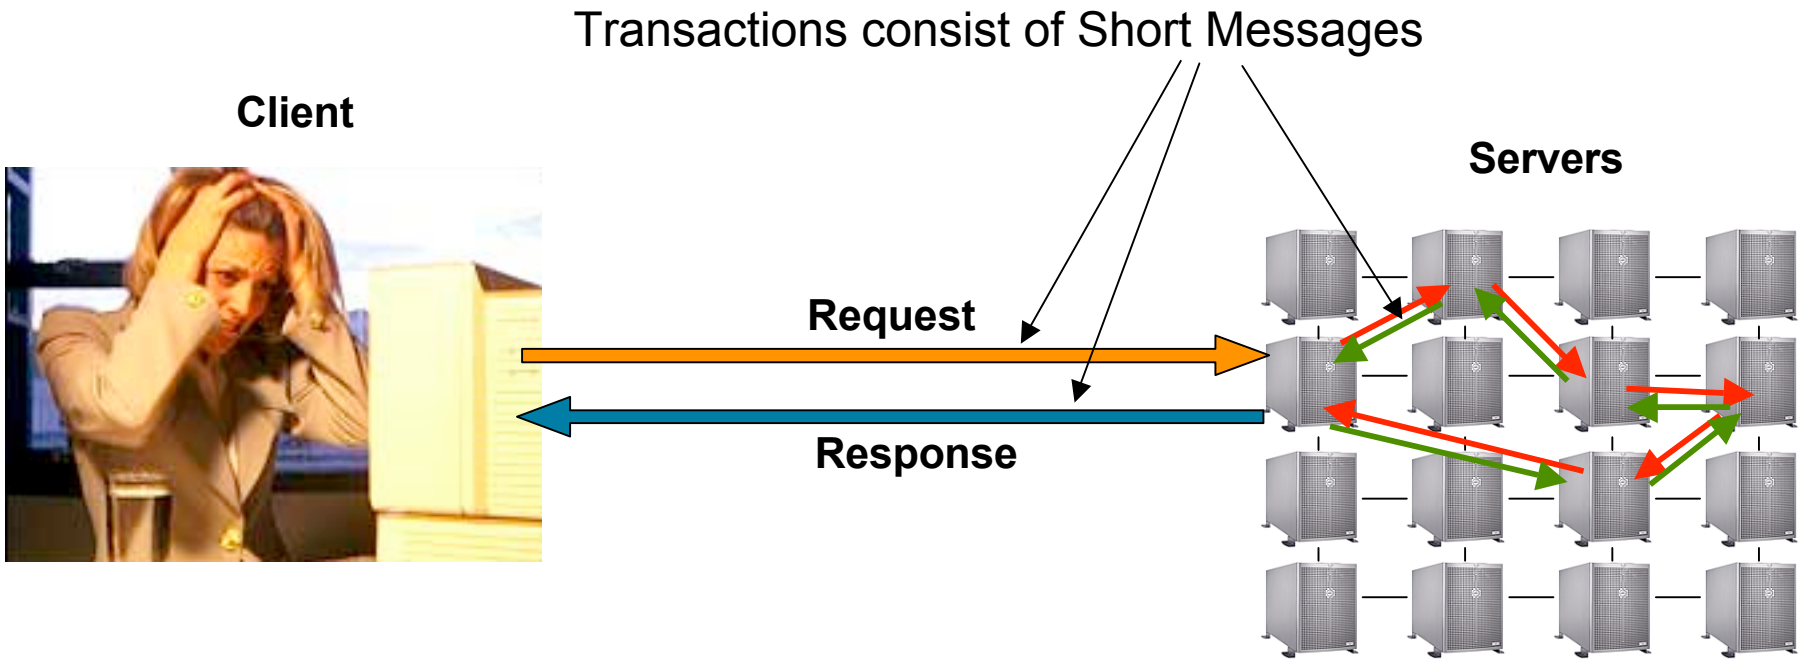

Transaction and other Database applications: Latency is Key for Response Time and Transaction Throughput

Latency and Bandwidth

Low Latency…

- High Bandwidth

Latency Fundamentals

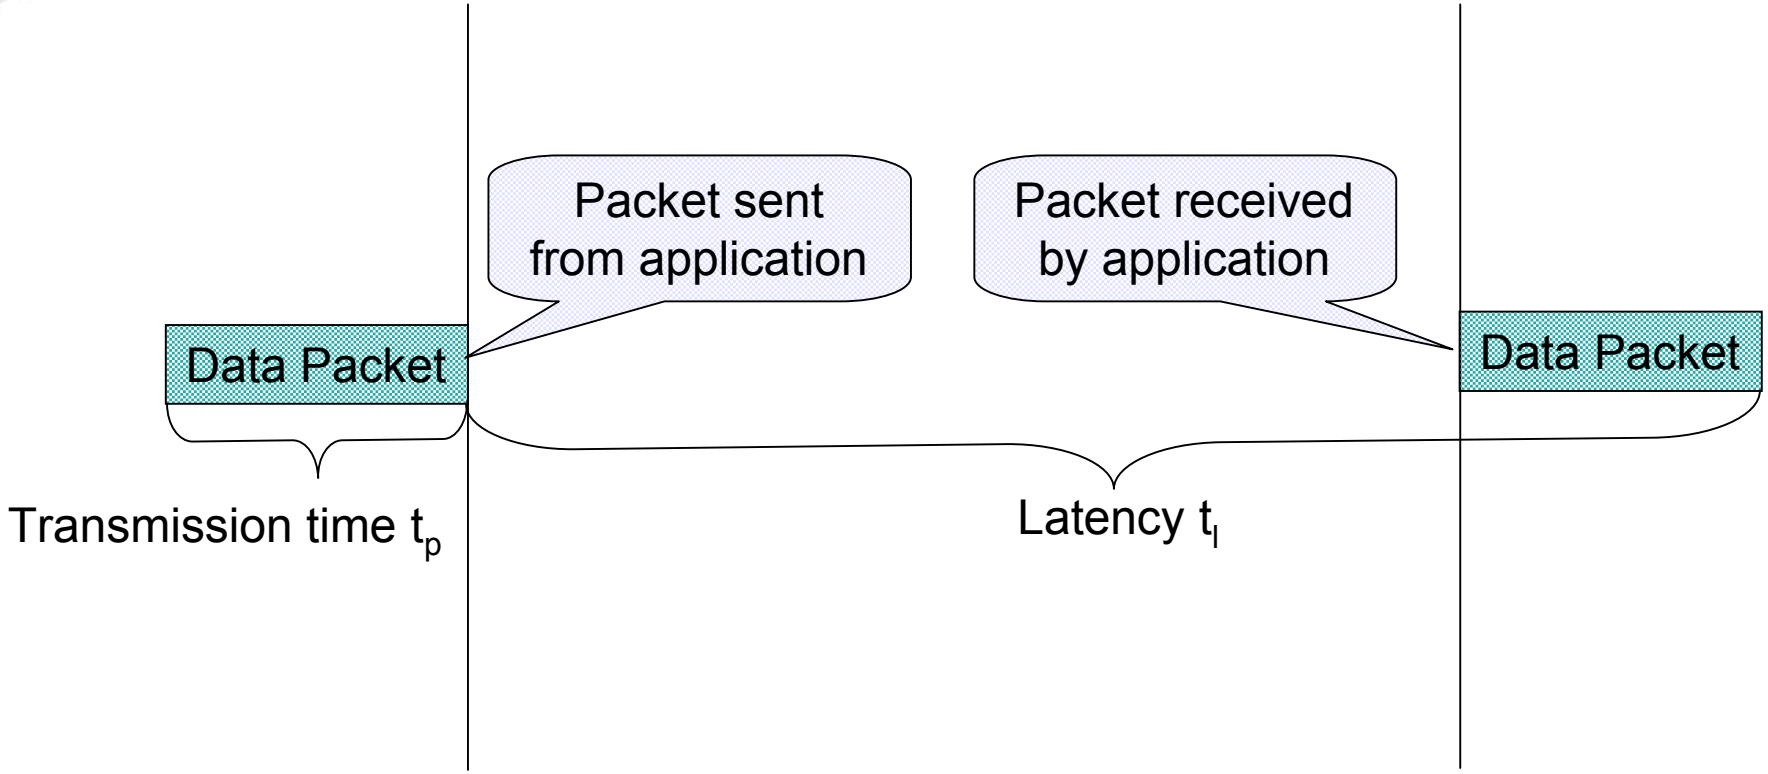

16-bytes Packet Latency

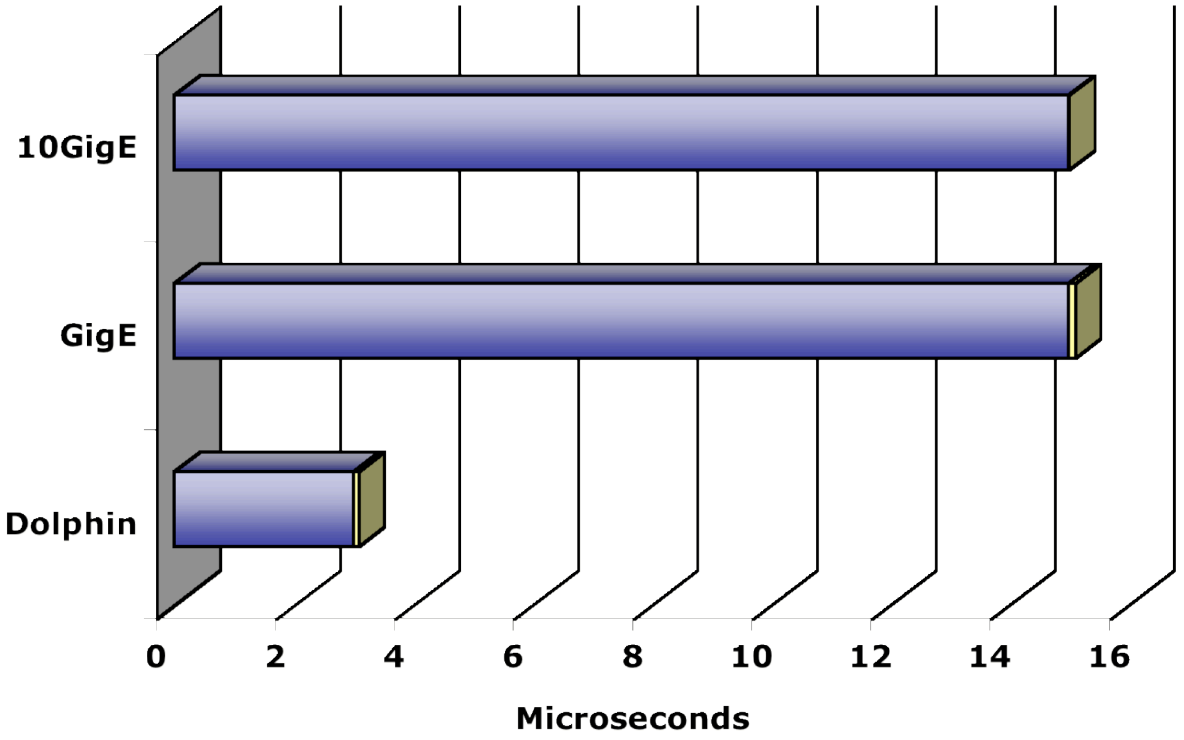

	Dolphin	GigE	10GigE
Packet Flight	0,1	0,1	0,01
Latency	3	15	15

Interconnect for DB Clusters

- Database communication patterns consist of short data packets
  - Performance is sensitive to latency and overhead
    - Response time
    - Throughput
- Batching of messages
  - Improves throughput to a certain degree, but
  - Creates longer response times
  - Introduces more overhead

GigE/10GigE Network Stack

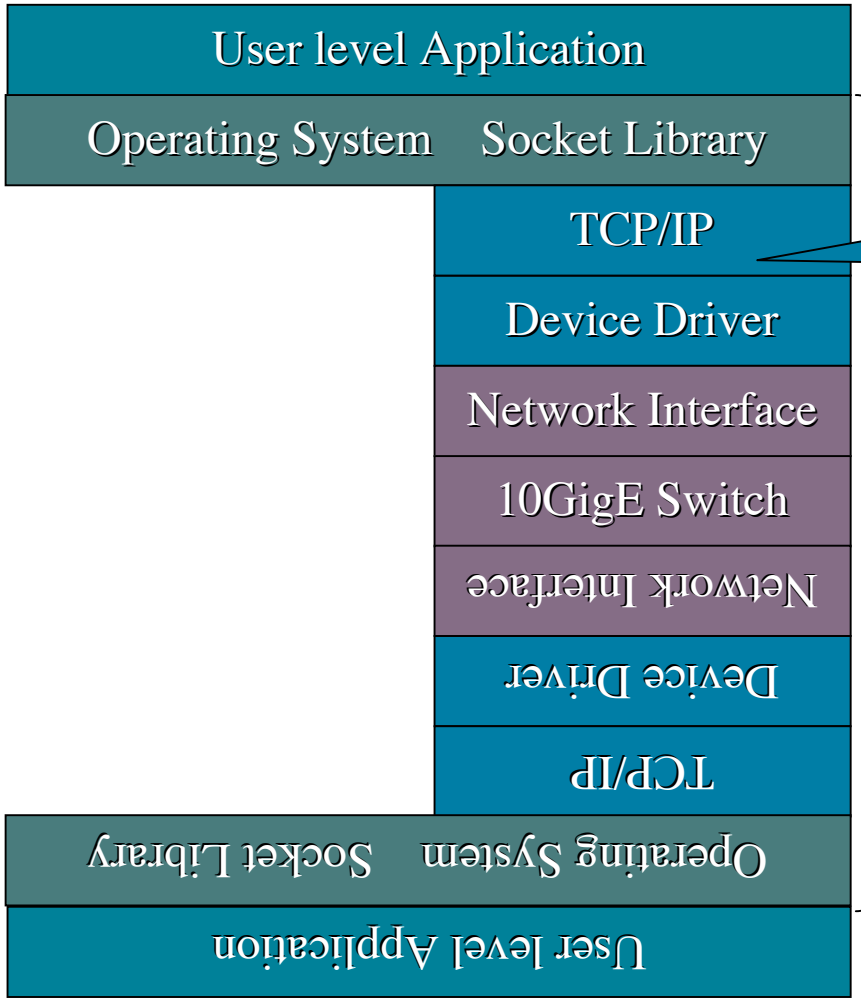

≈1 CPU Cycle
per bit
Requires 1 proc. @10GHz (!)
for full bandwidth

20 - 50μs with 10GigE

Network time @256Bytes payload
GigE : 3.12μs
10GigE: 0.31μs
Diff.: 2.81μs

35μs - 2.81μs = 32.19μs (8%)

Dolphin Express Network Stack

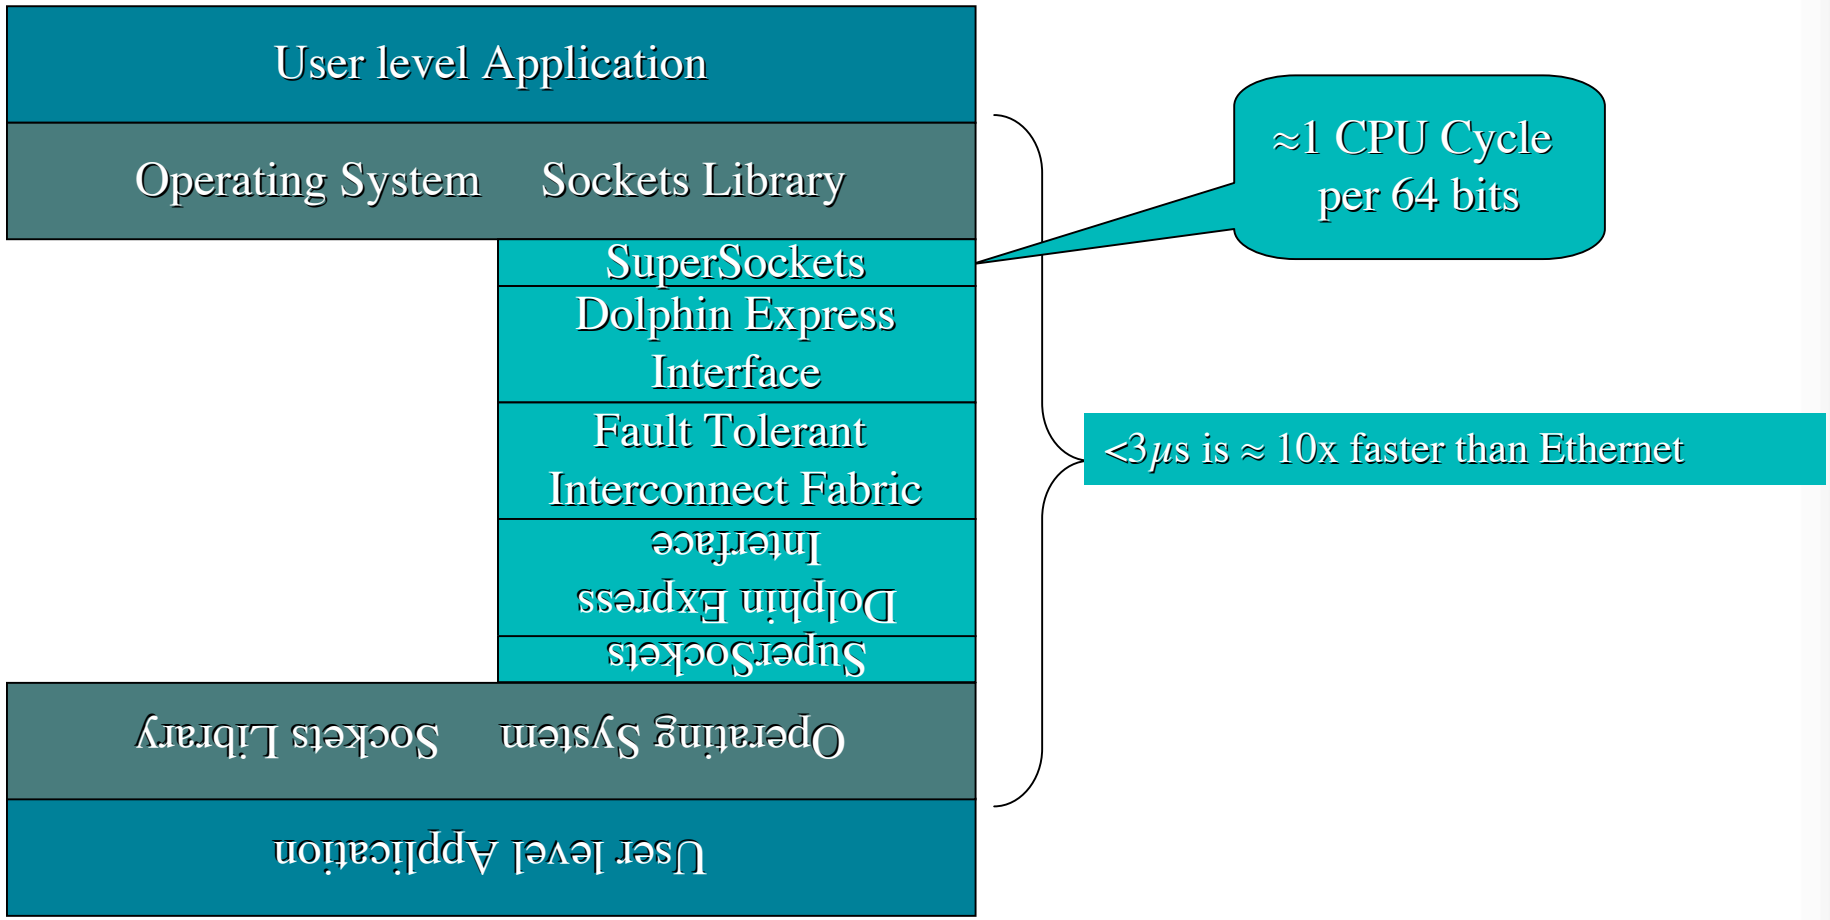

Shared-Nothing Data Transfers (Ether)

This is how other interconnects work:

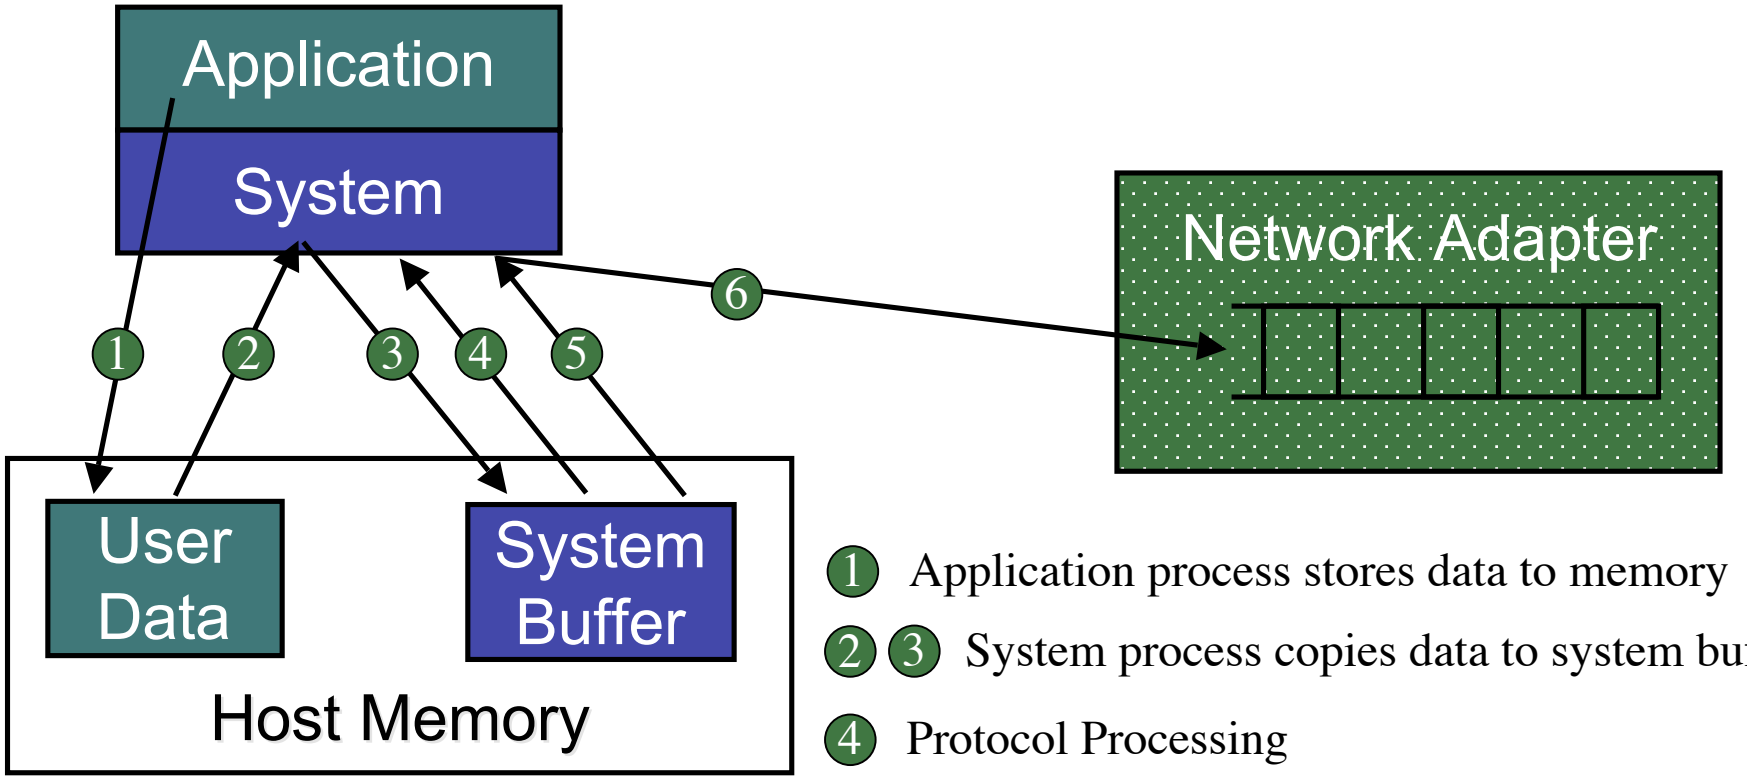

n Bus data transfer

1 Application process stores data to memory

2 3 System process copies data to system buffer

4 Protocol Processing

5 6 Data is copied to network adapter (usually by a DMA engine)

Shared Address Space Data Transfers

This is how Dolphin Express does it:

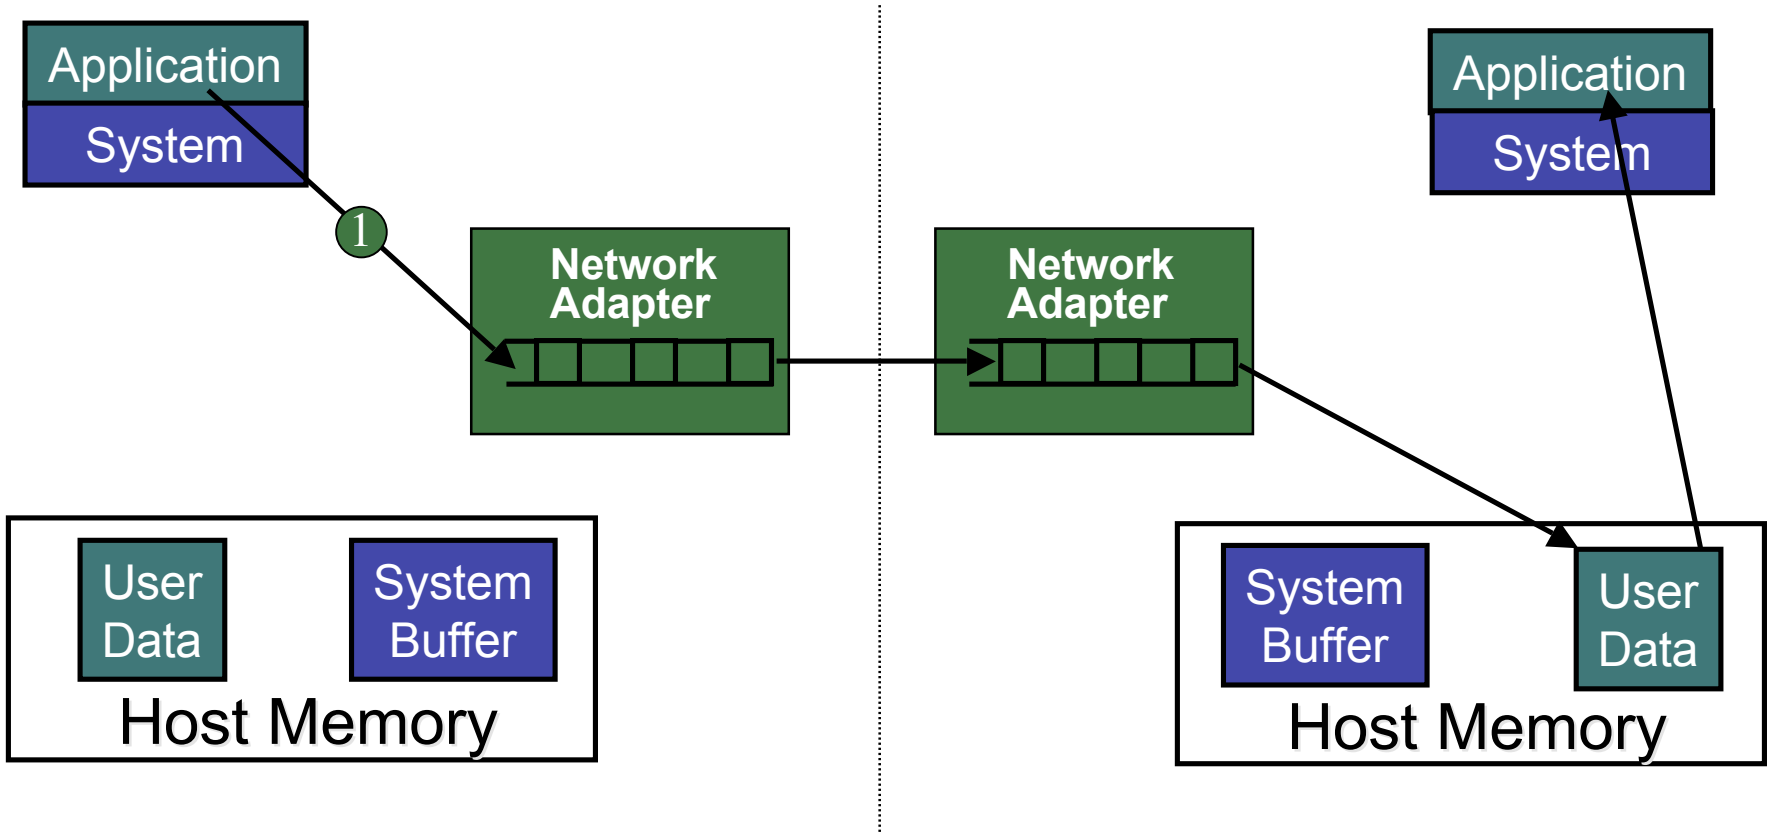

1 Application process stores data directly to remote memory

Dolphin SuperSockets™ Overview

- Berkeley Sockets API on Dolphin Express
- TCP/UDP/RDS compatible
- Data transfer through *remote shared memory*
- Flexibility through Configuration Files
  - Specifying Cluster Nodes
  - Specifying Communication Ports
- IP Address Virtualization

SuperSockets Protocol Handling

- RMA default limits
  - In-line Protocol < 1kBytes
  - Short Protocol 1kBytes < 8KBytes
  - Long Protocol 8kBytes <
- RDMA
  - Selectable
    - Per Port
    - Lower transfer limit selectable per application
- All tuning parameters selectable per application
  - Knowledge based defaults

Removing Database Cluster Bottlenecks

- Available CPU cycles
  - Network Overhead
  - TCP processing eats CPU cycles
  - *Dolphin Express with SuperSockets Removes TCP/IP overhead - Reducing CPU Load*
    - *Dolphin Hardware provides built-in data integrity*
- Network Latency
  - Database Traffic consists of short data packets
- Network Efficiency
  - SuperSockets utilizes Remote Memory Access (RMA)
    - Reduces CPU overhead to transfer data to remote node
    - Important for short transfers
    - Reduces Latency

Removing Database Cluster Bottlenecks (2)

- Interrupt Processing
  - Dolphin Express has integrated Checksums
    - No Interrupt processing for Network Protocol
  - Interrupt only required to wake sleeping process waiting for events on the Dolphin Express socket
  - Per Socket Adaptive cross-over between polling and interrupt
- High Bandwidth with low overhead
  - Dolphin Express provides
    - Remote Direct Memory Access (RDMA)
    - Reduces CPU overhead for large transfers

Oracle RAC Response Time

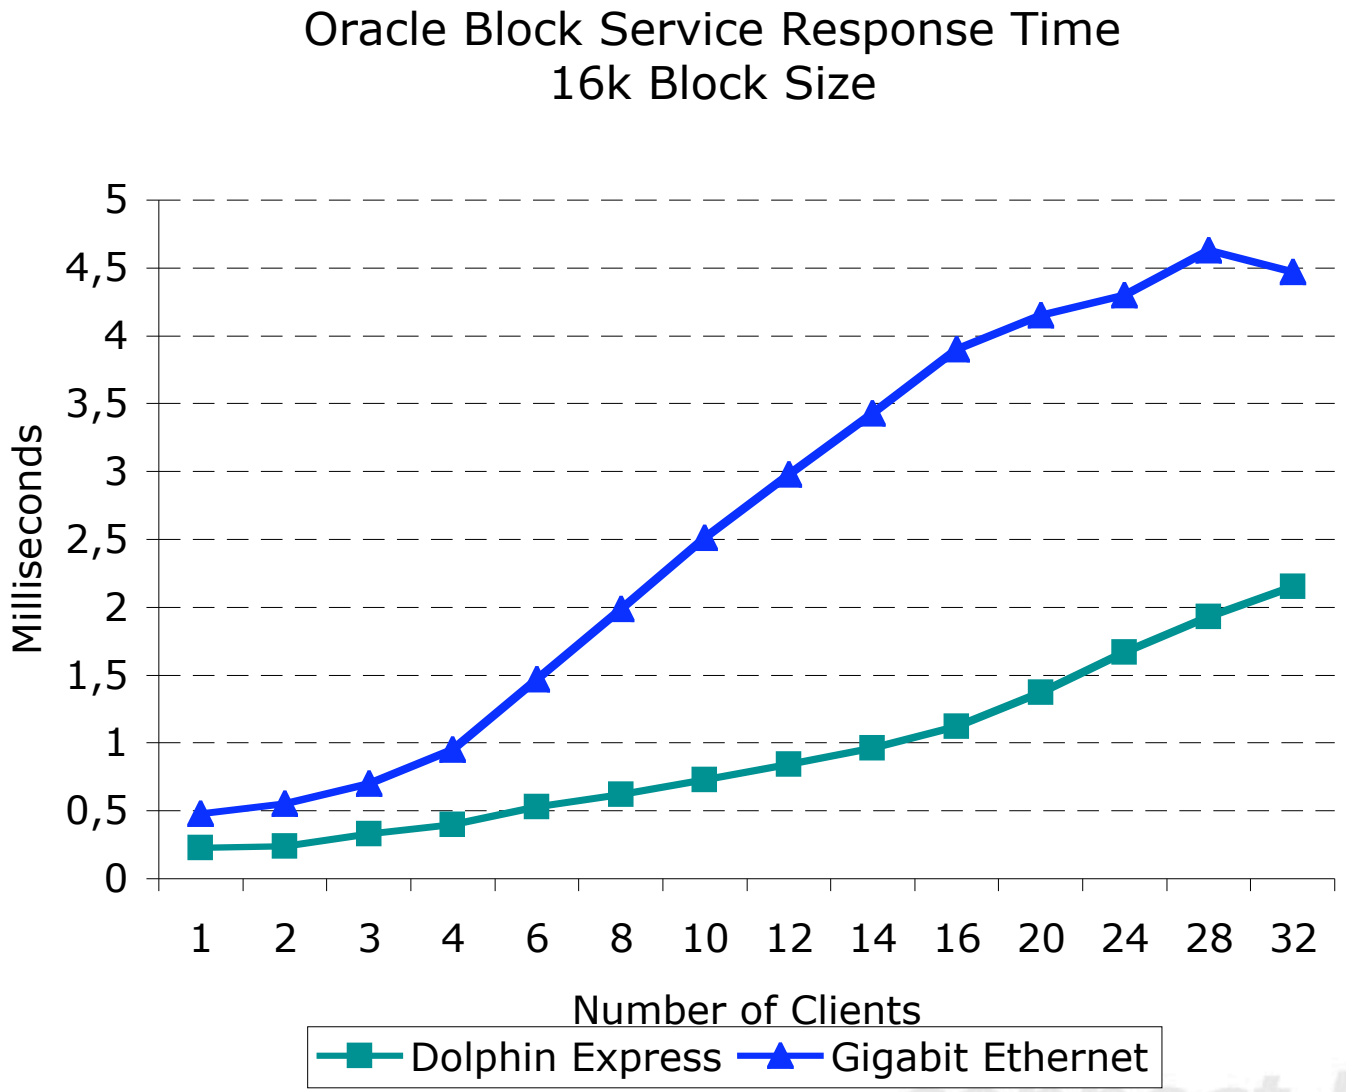

Oracle RAC Block Transfer Rate

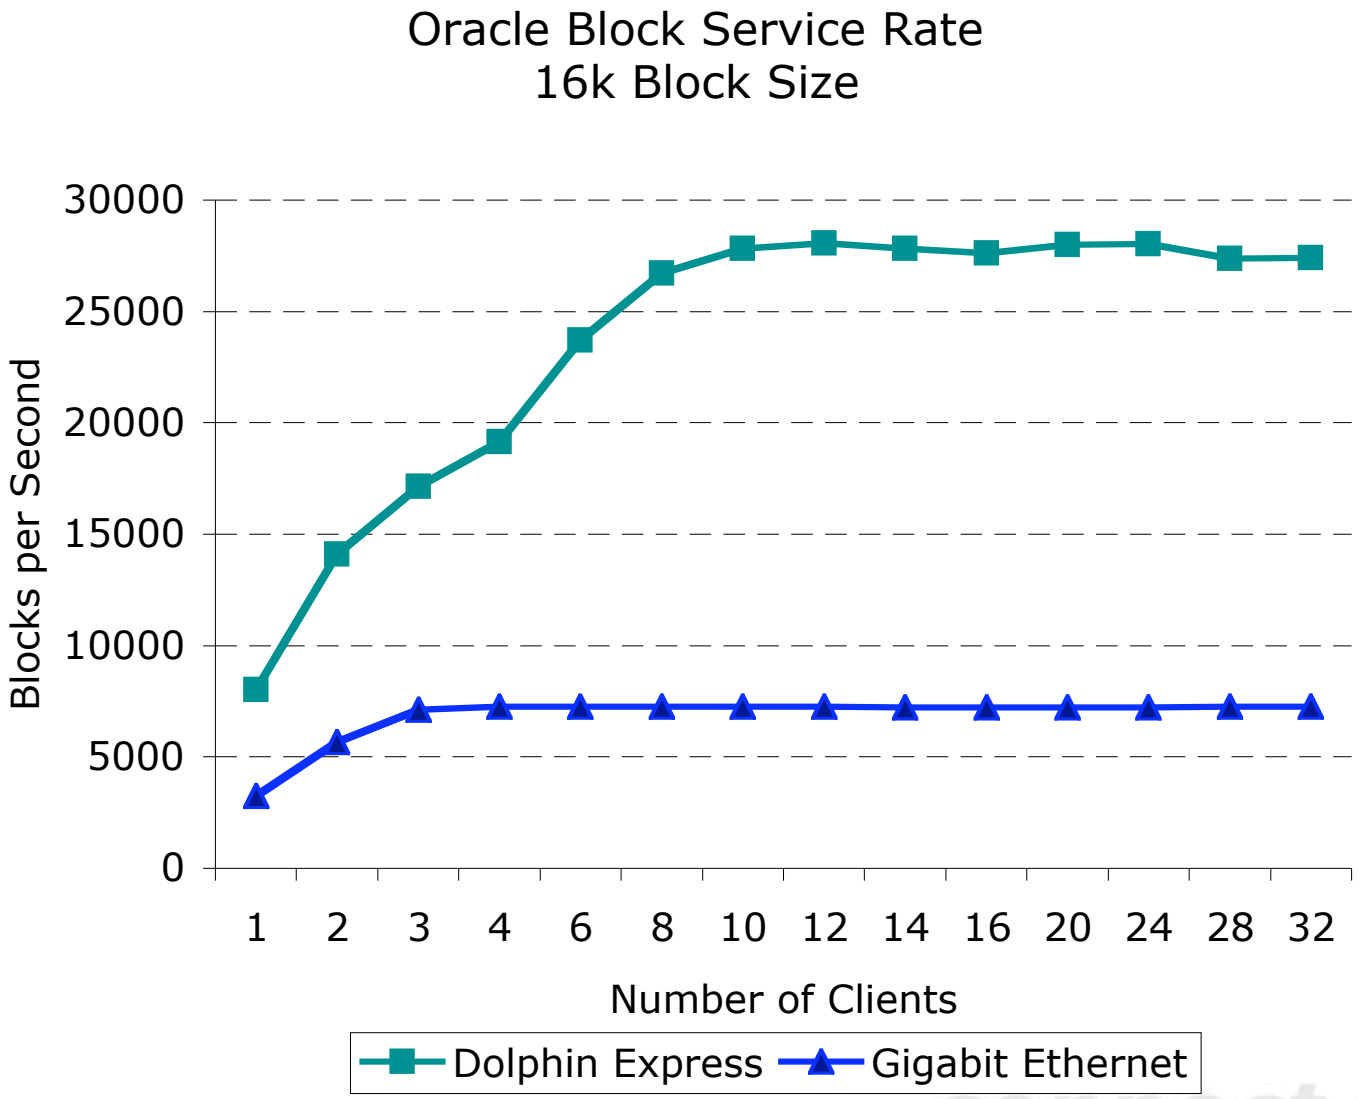

TPC-W Benchmark Set-up

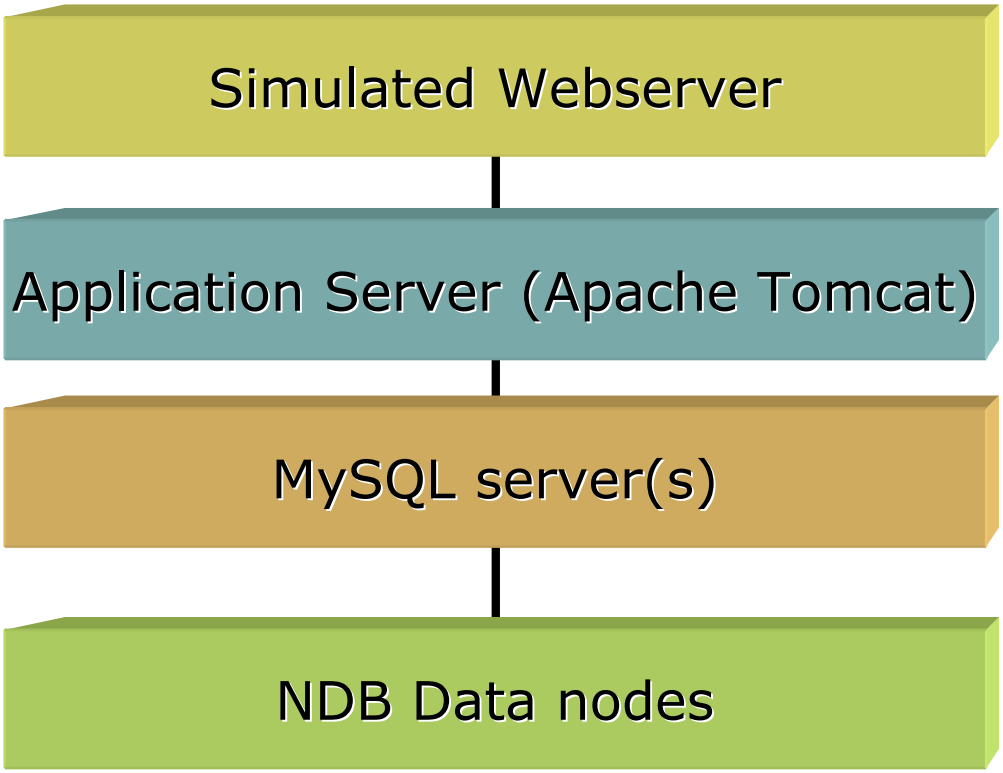

TPC-W Benchmark Cluster Set-up

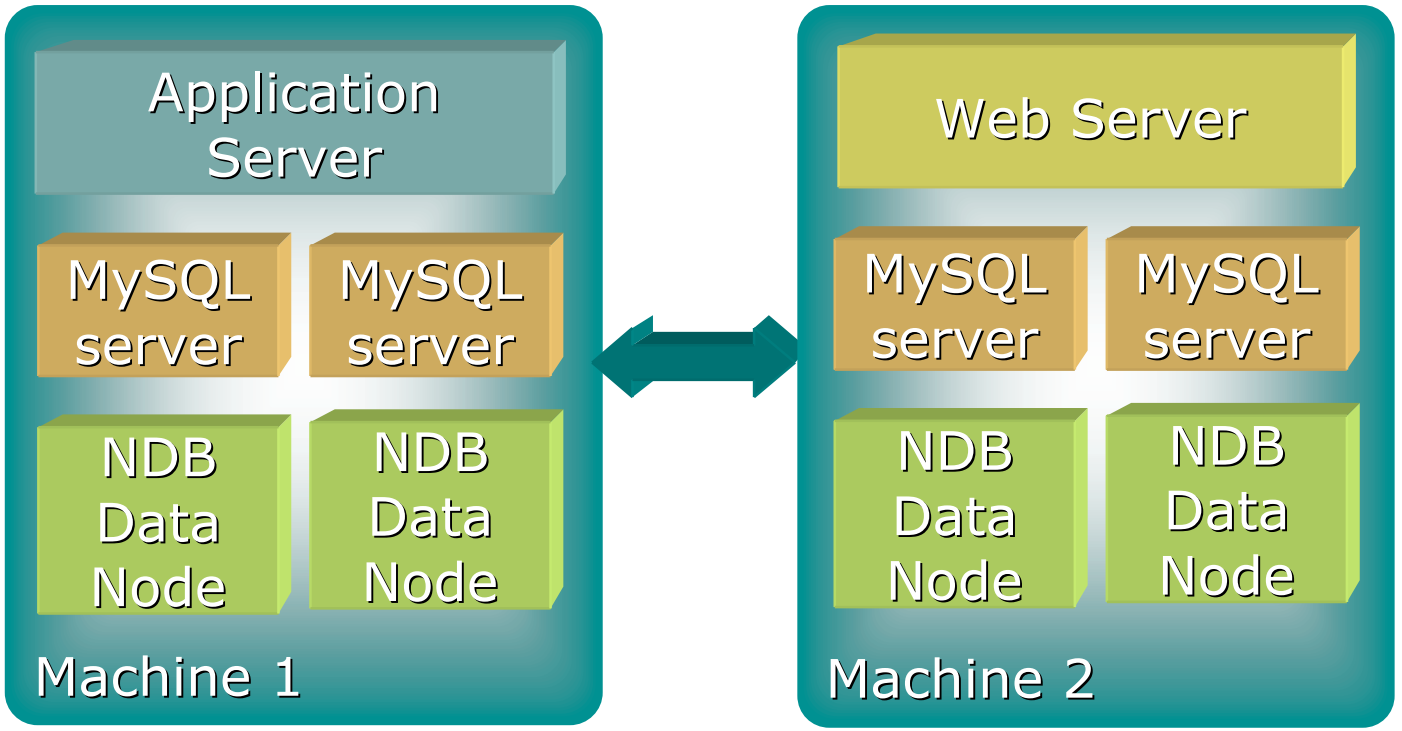

MySQL LAMP Stack TPC-W Response Time

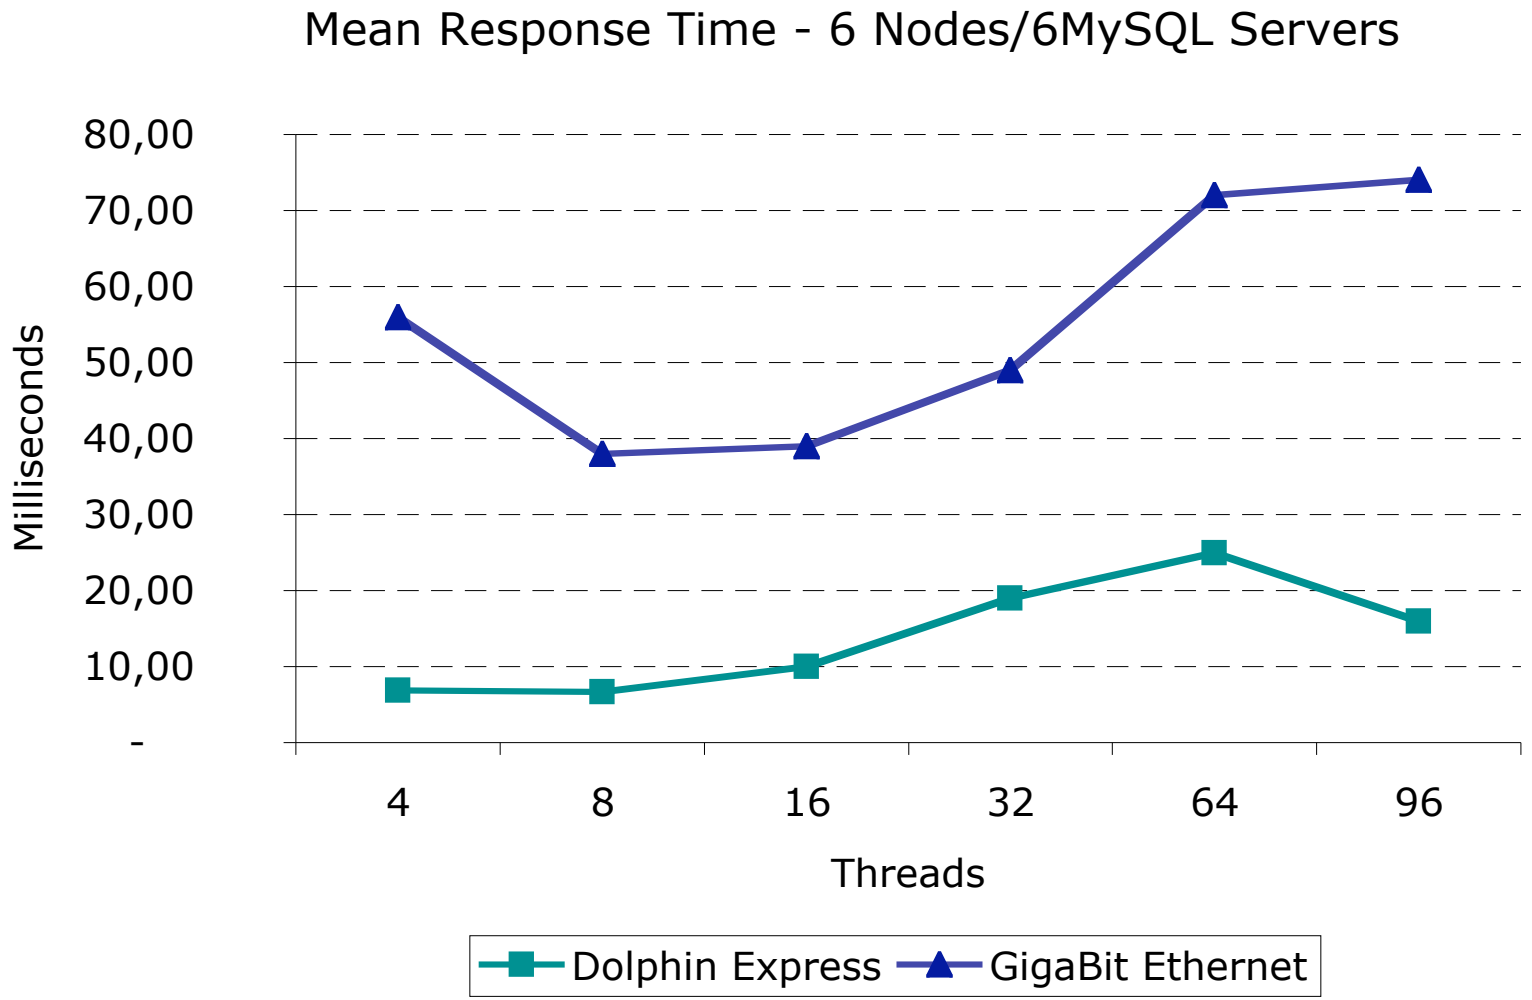

MySQL LAMP Stack TPC-W Throughput

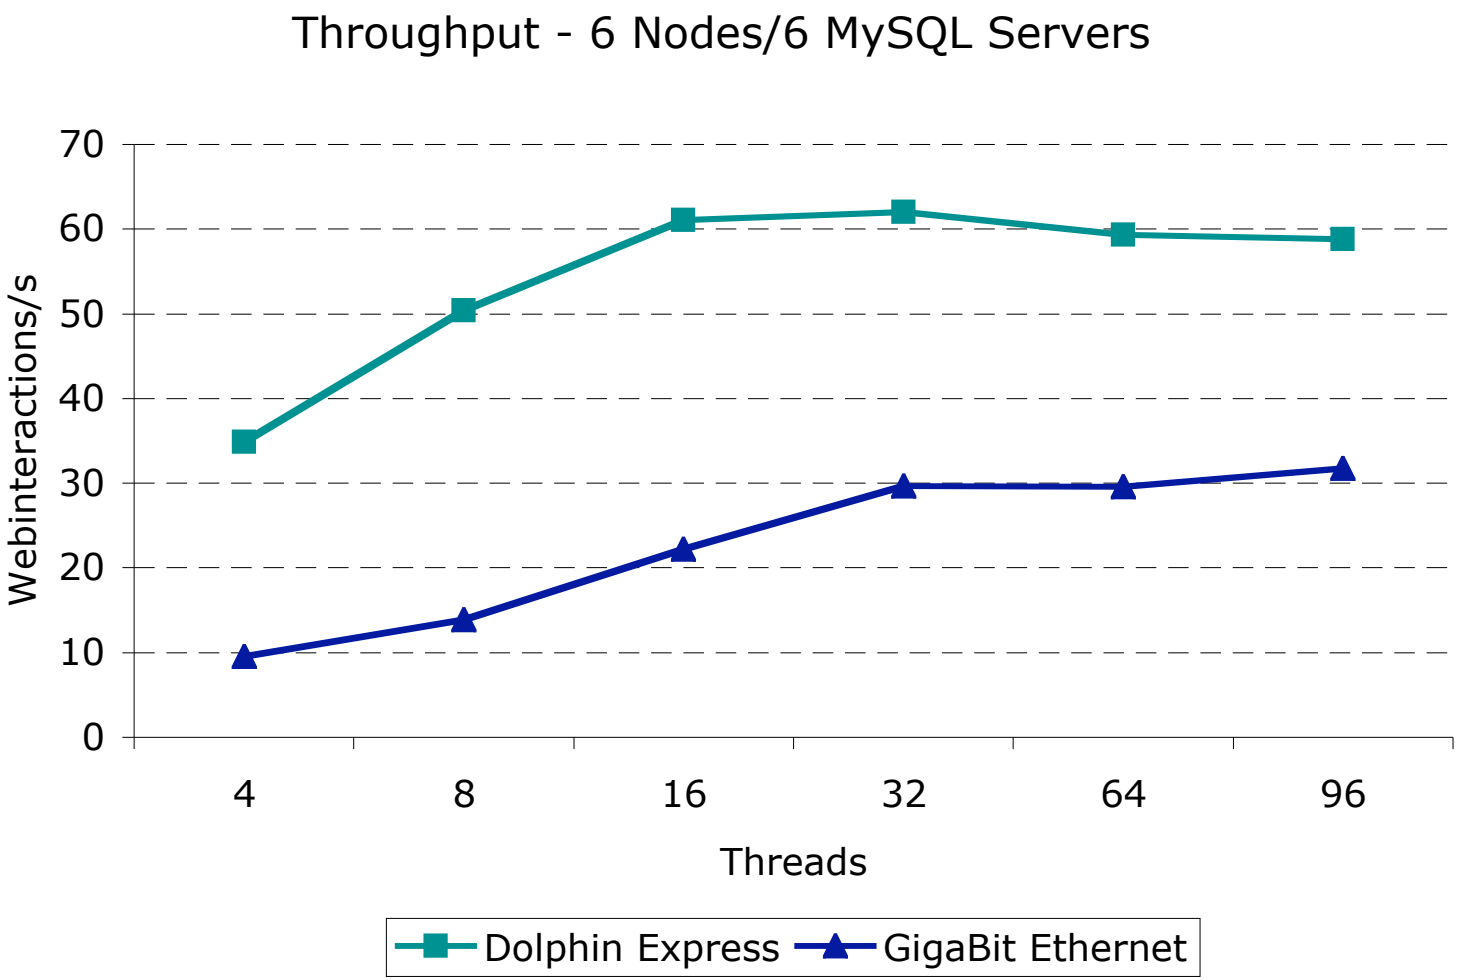

Dolphin Express for Database Clusters

TCO, Performance & **Fault Tolerance**

connect better

Two Interconnect Fabrics - Normal Operation

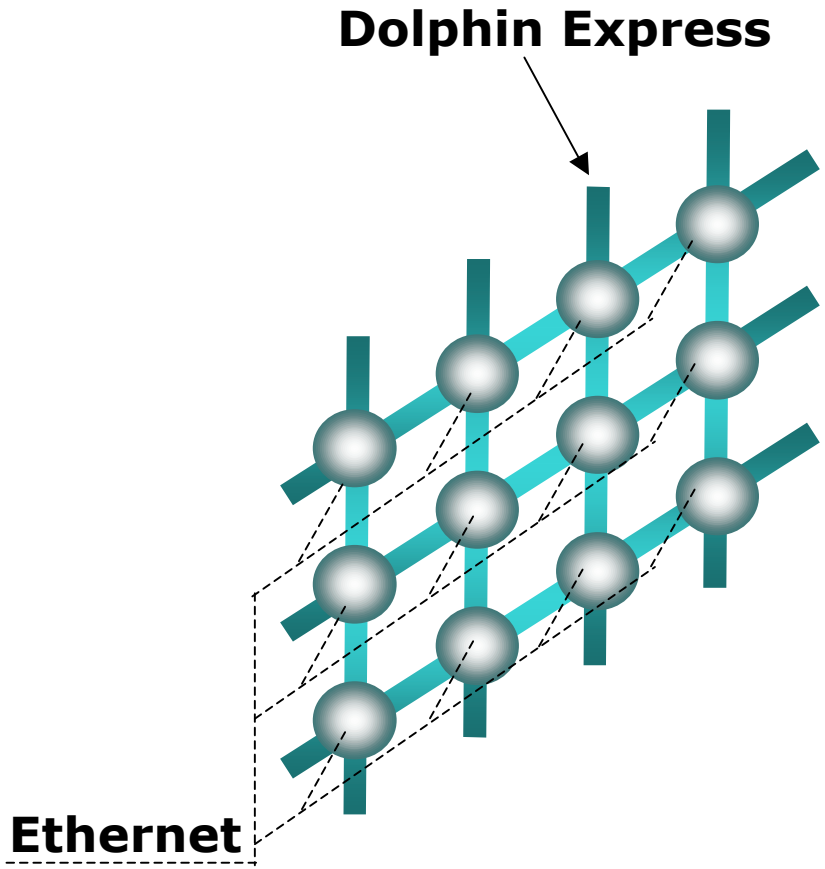

Failover scenarios:

- A node in the cluster fails
- A network card fails

2-D Dataflow

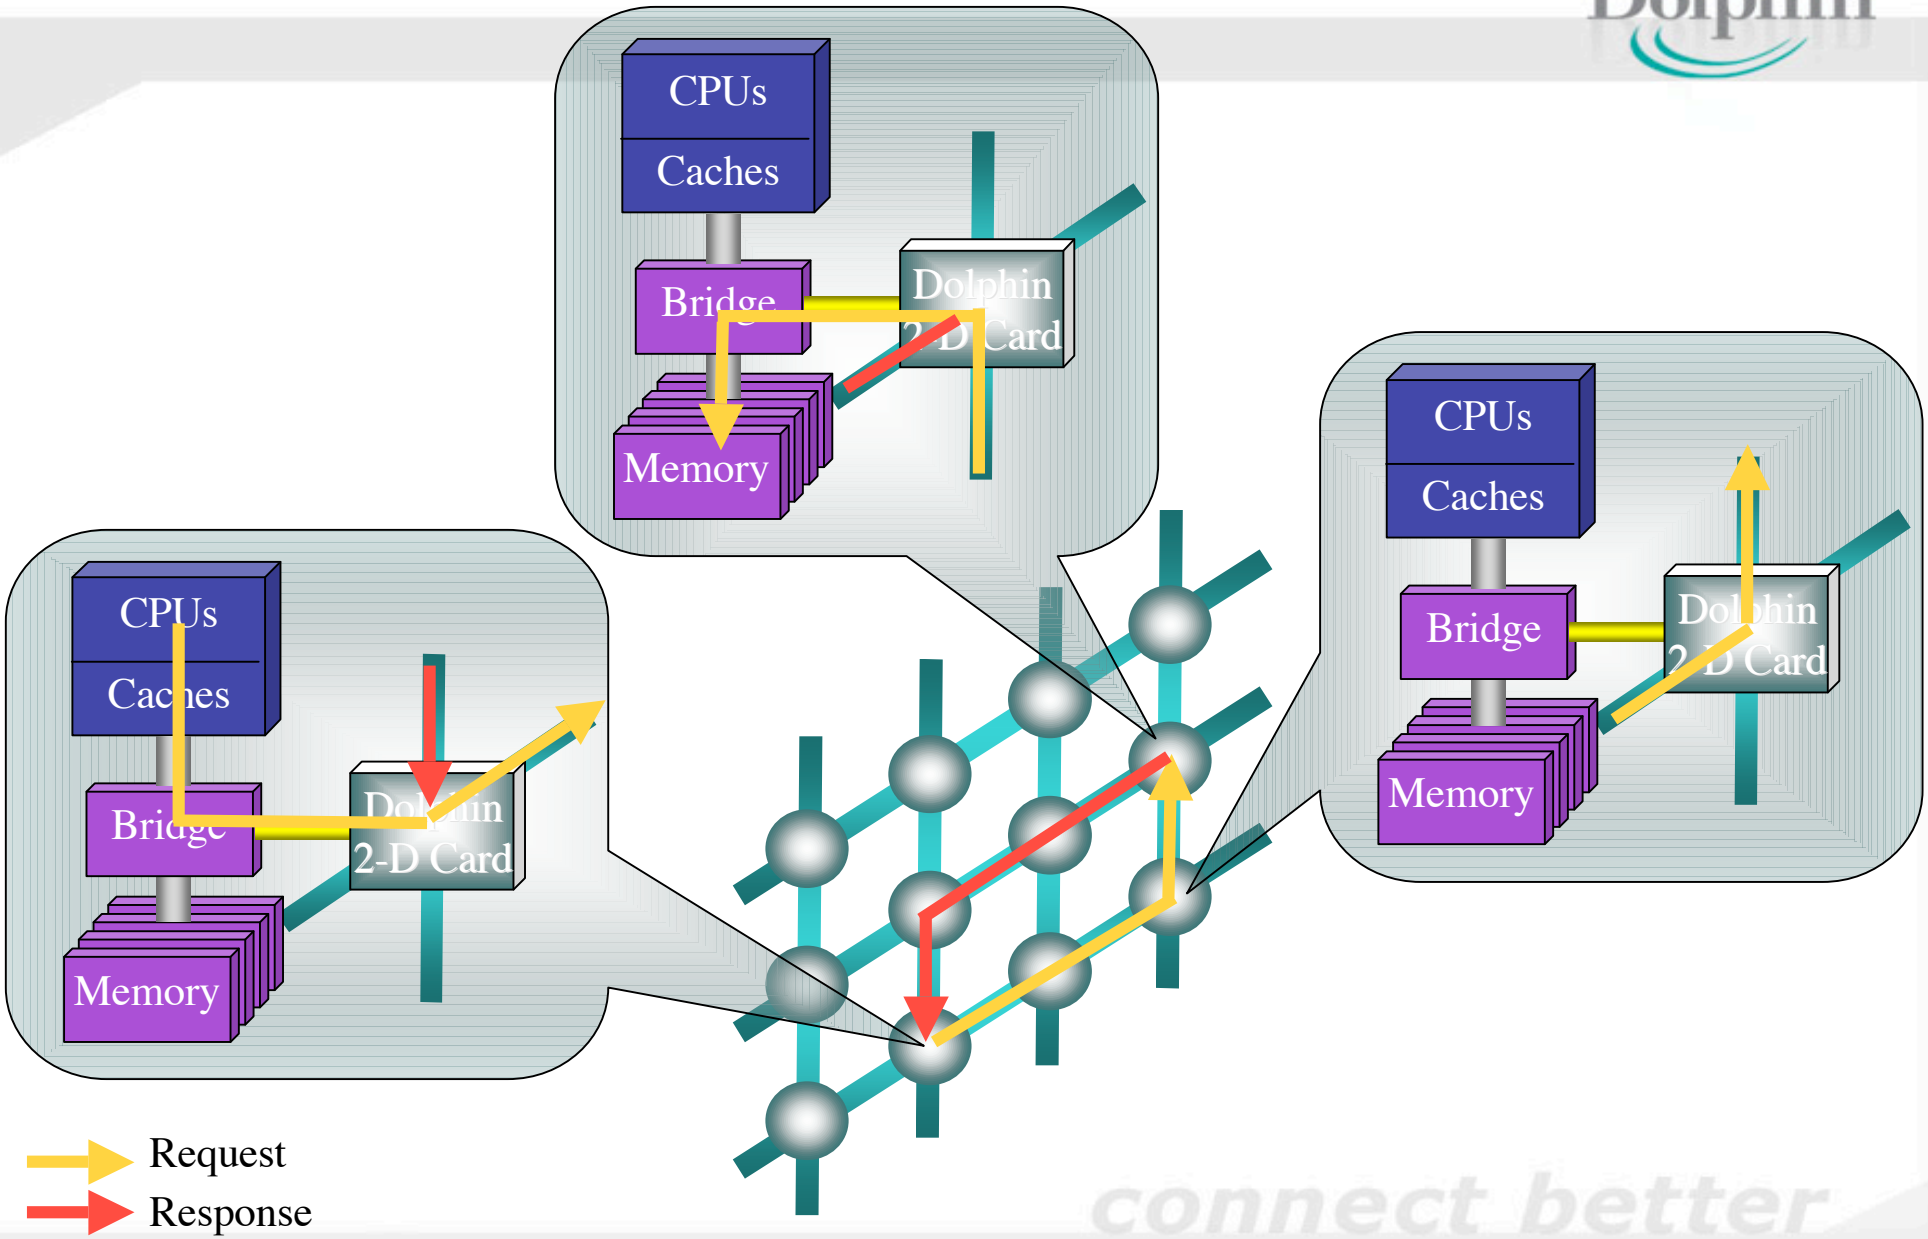

Failing Node

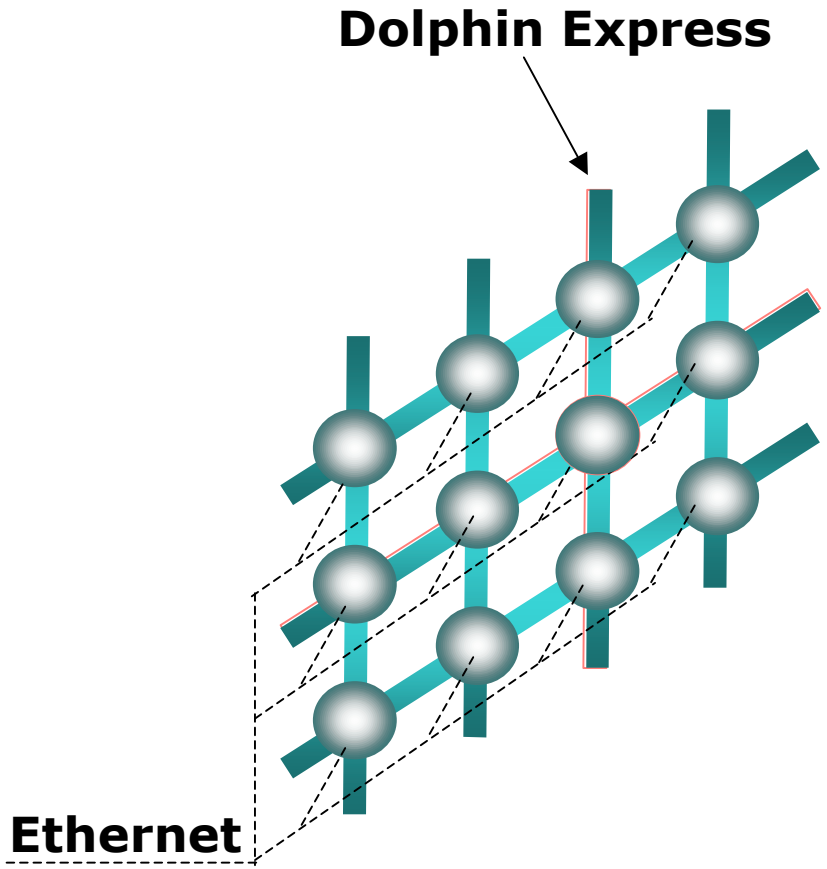

2-D Dataflow w/failing Node

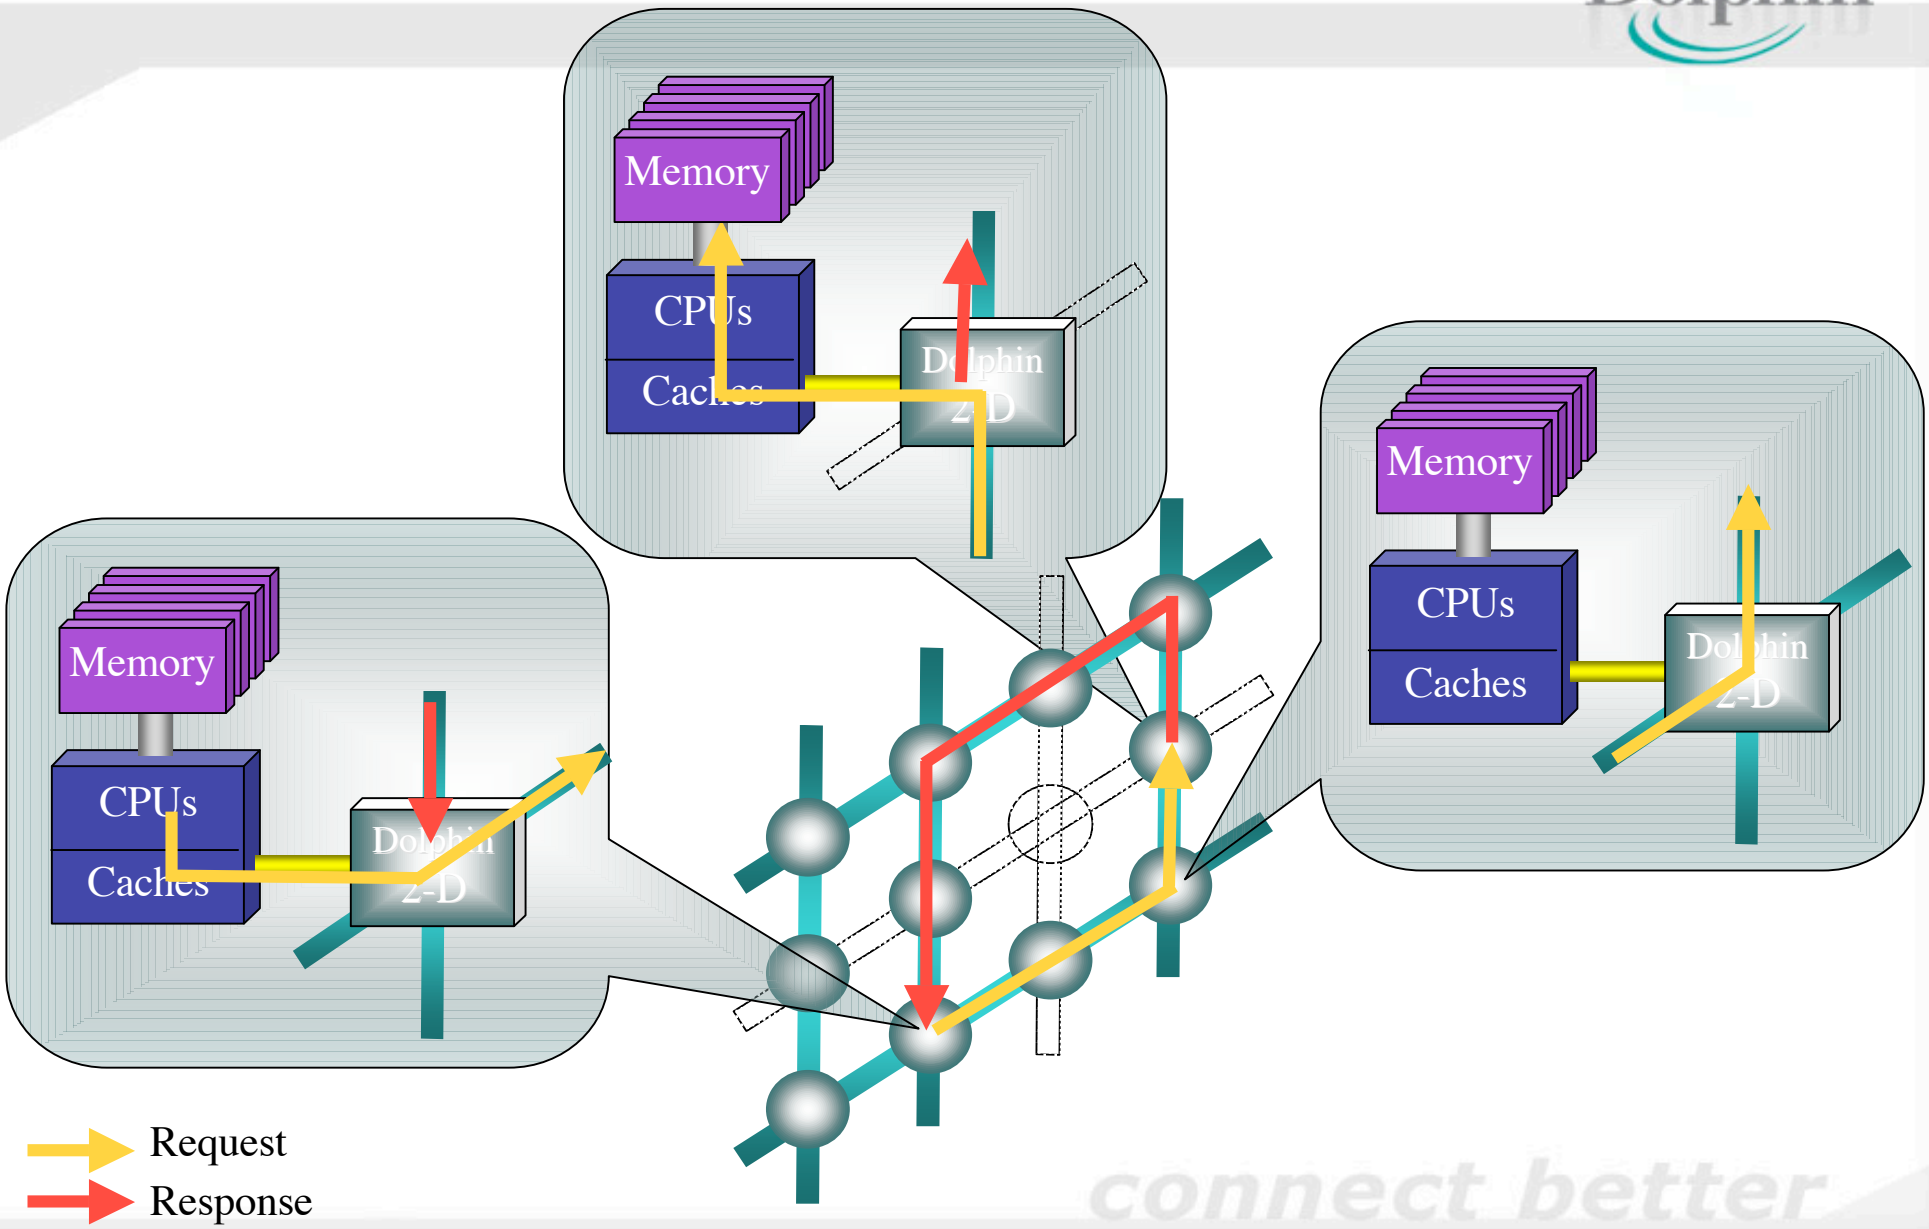

Failing Network Card

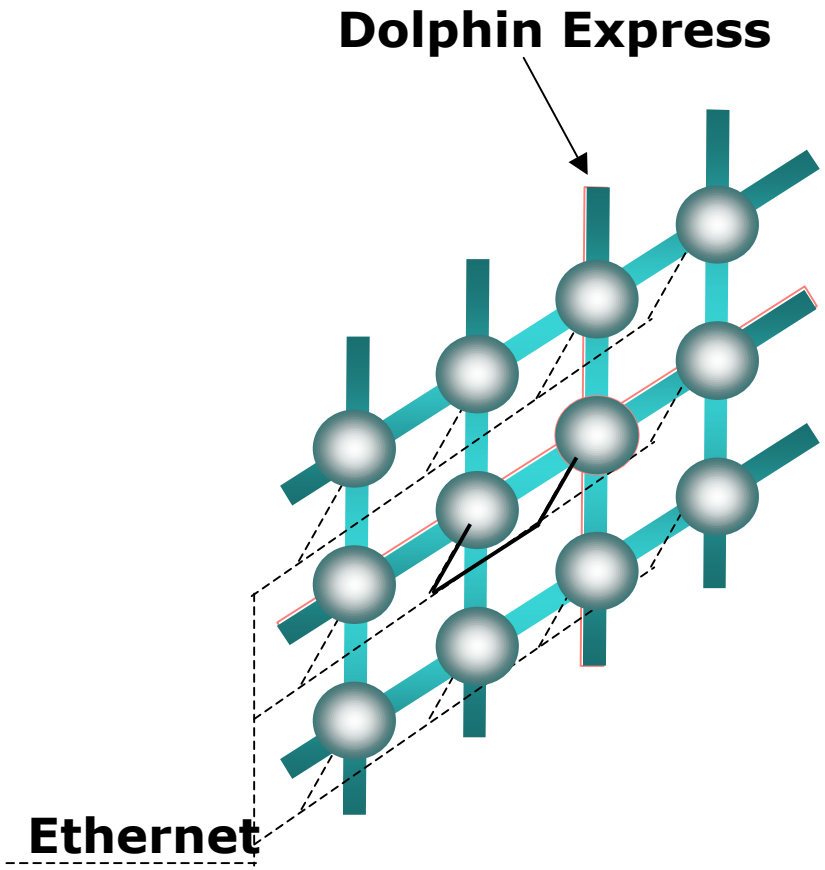

Dolphin SuperSockets Automatically fails-over to Ethernet

Questions & Answers

Thank you for attending!

- Contact Dolphin:
  - Einar Rustad: er@dolphinics.com
  - John Longmire: jl@dolphinics.com
  - Hugo Kohmann: hugo@dolphinics.com

  - US switchboard: 1 508 786 9950
  - www.dolphinics.com

connect better