

From 7 Days to 4 Hours!

Provide Near Real-Time Response to Mobile Web 2.0 Users



Executive Summary

With the exponentially growing number of mobile business users and mobile-oriented applications, IT must adopt modern architectures that provide a near real-time response to hundreds of millions requests every day and to complete the recovery of databases within hours. Reducing the response time and the data recovery time enables the service providers to serve more users over the same equipment and enhance the overall user experience.

Utilizing Mellanox's FDR 56Gb/s InfiniBand solution, a tier-1 Fortune100 Web 2.0 company was able to reduce their database recovery time from 7 days to just 4 hours! With 97% reduction in recovery time, the company is now able to better serve their hundreds of millions worldwide Web 2.0 mobile users and to maintain business continuity.

Overview

In recent years, clustering has become the technology of choice used in database systems. It has also proven to be the most efficient way to meet business requirements, such as high-performance, high-availability and scalability at a lower cost of ownership. In order to achieve those goals, it is critical to choose the right interconnect solution.

One of the most important parameters that a business process must meet is its Recovery Time Objective (RTO), which defines the maximum service downtime allowed, due to failure. Unacceptable consequences could be associated with breaks in business continuity, if optimal RTO is not achieved. As most enterprise applications use databases, RTO must also include the time required to restore the entire database. For example, the expected RTO for databases used in Data Warehouses is 12 hours to a few days. However, this would differ from Online Transaction Processing (OLTP) systems, where the expected

RTO is more stringent and in the timeframe of minutes to a few hours, depending on the size of database storage.

Recently, a tier-1 Fortune 100 Web 2.0 company wanted to build a high-performance database system which would be capable of handling tens-of-millions of mobile user requests per day in real-time, while keeping the Total Cost of Ownership (TCO) intact.



Figure 1. Providing near real-time response for mobile Web 2.0 users at a low TCO is required

Solution

The company's original database system was based on 1GbE network technology, which presented a number of challenges. First, scaling the architecture and technology infrastructure was expensive and time consuming. Second, the company's ability to meet the demands of a near real-time response to its mobile customers was limited. Lastly, in case of database failure, it took them more than 7 days to restore the database system, which was extremely slow compared to their 6-hour goal. As a result, the company considered using a faster and more efficient interconnect technology, that would enable them to meet their demands, while also cutting down the costs associated with maintaining the infrastructure.

By working with Mellanox Technologies, a leading supplier of end-to-end interconnect solutions for data center servers and storage systems, a modular database clustered solution was developed, which utilized Mellanox's end-to-end FDR 56Gb/s InfiniBand technology. The solution supports the deployment of a single database across

a cluster of servers, and provides superior fault tolerance, performance and scalability without any need to change applications. The solution provides a near real-time response for the mobile users and also a 4-hour RTO, 33% faster than the company's target. In addition, the solution offers customer continuous uptime for all of its database applications, on-demand scalability, lowered computing costs, and record-breaking performance.

Applications

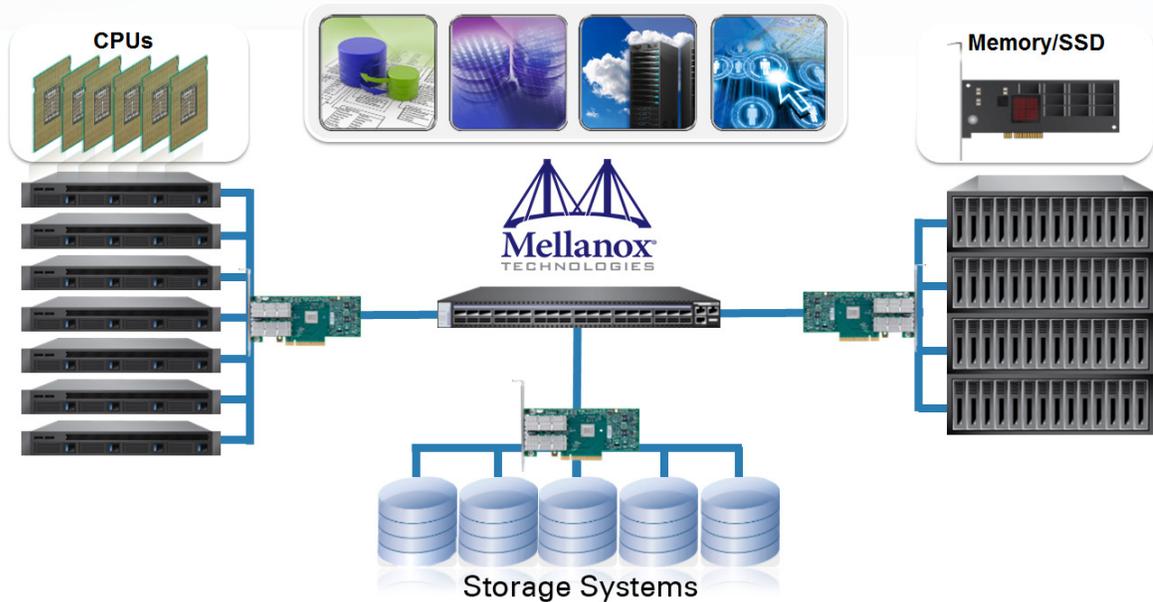


Figure 2. Mellanox's End-to-End FDR 56Gb/s InfiniBand solution

Summary

Modern Web 2.0 companies need to serve hundreds of millions mobile users in a near real-time response time. Providing such a service requires a high-performance OLTP database. The interconnect is the key bottleneck, and if lower in performance, will lead to degraded application performance, limited scalability and poor ROI.

Mellanox's products and solutions enable higher-performance and efficiency for Web 2.0 infrastructures. With Mellanox, businesses do not need to compromise their application performance, in order to achieve lower TCO.

Mellanox's solutions deliver the following features:

- Fastest interconnect adapters with 10, 40, and 56 Gb/s bandwidth per port and sub 1us latency
- Full OS bypass and Remote Direct Memory Access (RDMA)
- A superior high availability solution based on lossless interconnect technology
- I/O consolidation of server-to-server and server-to-storage over a single wire
- Cost-effective, high-density switches and fabric architecture



350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085
 Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com