


## Introduction

Clustering is a concept that allows an aggregation of servers to achieve a specific goal such as availability, scalability or performance. By combining a number of servers together with the right network and software components, a high performance server cluster can be built. To build high performance storage to connect to the high performance cluster, a storage area network (SAN) is employed. Just as a high performance cluster uses multiple paths to increase performance, a SAN provides multiple high performance paths for storage.


It is now possible to build high performance clusters with the low-cost Windows-based software and Intel processors, bringing new price/performance to the market. One such cluster has been shown on the transaction procession performance council (TPC) website. This was made possible by the introduction of a number of high performance parallel storage paths through a Gadzoox® Networks Capellix™ SAN switch.

## High Performance Clusters

The differences between a cluster built for performance and a cluster built for availability are significant. In the diagrams below these principles are shown in a simplified manner. The high availability cluster provides redundancy of all key components and data paths. For instance, there are two SAN switches so in the event of failure, the data is always available via an alternate path to the servers. Usually, the server operating system or clustering application will be ready to fail-over to another server in the event of a server failure, so the application is always available. The external storage also has dual pathways to allow access to the data which is usually protected either through mirroring or a RAID 5 protection mechanism against disk drive failure.




**Figure 1** *High Availability Cluster*



**Figure 2** *High Performance Cluster*

For performance, parallelism creates higher aggregate bandwidth by combining the capacity of multiple pathways. This usually means a more expensive configuration than that of a high availability configuration. Instead of just redundant components, there are several parallel components to create multiple parallel pathways to reduce data performance constraints. Because mechanical storage devices are so much slower than the electronics in the server, a storage area network can provide performance benefits by providing the servers multiple paths to multiple storage devices to reduce the storage bottleneck. In order to provide the appropriate storage infrastructure, a low overhead SAN switching environment is required. Additionally, a non-blocking or unconstrained pathway for the data is necessary. These pathways in a SAN are provided with Fibre Channel links of 100MBytes/sec through a SAN switch such as the Capellix from Gadzoox Networks.



**Figure 3** *High Performance Cluster with a SAN*

## The Capellix Advantage

In order to support a SAN in a high performance environment, Gadzoox Capellix SAN switches have been designed to bring a low-overhead, high-bandwidth solution to market. The Capellix switch provides a circuit switched technology which provides a dedicated pipeline for unconstrained performance in demanding SAN applications like these. Figure 3 shows a theoretical configuration with a SAN for maximum performance.

One source for proven high performance clustering configurations is found at the transaction processing performance council or TPC. TPC is famous for their industry-standard benchmarking work. A notable example is the experience that Compaq Computer had with their TPC benchmark test. They were able to generate in excess of 100,000 TPC-C throughput benchmark score using a Microsoft Windows NT configuration and a highly parallel configuration. It scored very competitively against Unix systems using more expensive components. The Capellix has also been tested in the Unix environment for clusters wanting that environment.

To review the TPC benchmark using Capellix go to [www.tpc.org](http://www.tpc.org) and review the Compaq PDC/02000 configuration. The configuration used six Compaq hosts, six Capellix SAN switches going into 12 storage arrays and 6 clients. The configuration covered over 100,000 transactions per minute at less than \$35/TPC-C. The specifics on the test can be found at [http://www.tpc.org/new\\_result/c-result1.idc?id=100021201](http://www.tpc.org/new_result/c-result1.idc?id=100021201).

Gadzoox has the experience with high performance SAN components such as the Capellix SAN switch product line. By designing a low-latency high-bandwidth FC-AL switch, it is the logical companion for high performance clusters. The Capellix 3000 provides a scalable framework for large installations or the ability to start small and hot-plug additional ports for easy expandability. The Capellix 2000 provides a low-cost entry into SAN switching with the unique ability in an 8-port switch to add three more additional ports for easy expandability. Either Capellix switch provides high capacity bandwidth for unconstrained performance and the low-latency circuit switching ideal for high performance clusters.

For information on other SAN applications or on the Capellix switches please go to [www.gadzoox.com/solutions](http://www.gadzoox.com/solutions).