

A Working Prototype of High Availability, High Performance Web Application Platform

(Wintel Platform using all Microsoft Software)

Stefan Stasiewicz

School of Business

Bay of Plenty Polytechnic

stefan.stasiewicz@boppoly.ac.nz

This Web Application platform was setup as a prototype network for demonstration purposes for students undertaking advanced networking papers and also as a basis for the possible future development of a platform for students building web applications using MS SQL Server as a back end.

Since the Clustering and Load Balancing, Distributed File Systems and File Replication features (which are available in the various editions of Microsoft Server products) are relatively mature technologies and readily setup, this presentation focuses more on the newly released MS SQL Server 2005. With its “unsupported” database mirroring feature, MS SQL Server 2005 provides plenty of scope for research on its use in setting up a highly available database platform. All Microsoft software used in this prototype network was obtained and used under the Microsoft Developers Network Academic Alliance (MSDNAA).

1.1 MS SQL Server 2005 - Database Availability Solutions

The three primary solutions for achieving high availability with SQL Server have differing levels of complexity and are appropriate for different situations.

Failover Clustering - is designed for conditions where there is a concern about hardware failure and there is a need to support high transaction rates.

Replication - is designed to increase data availability by distributing the data across multiple database servers. Availability is increased by allowing applications to scale out the SQL Server read workload across databases.

Database Mirroring - relies on a new copy-on-write technology that is included in SQL Server 2005. This availability option requires three servers – principal, mirror and a witness.

1.2 Database Availability Technique Used

Database Mirroring was used because it

- was a completely new feature to the Microsoft SQL Server and provided for new research options

- does not require additional expensive certified hardware as would be needed in the Clustering option (eg. shared storage systems such as fibre channel SCSI RAID5 etc).

The **principal** server writes each transaction to the database and then sends a copy of the transaction to a **mirror** server, which applies the copied transaction as soon as it is received.

The **witness** assists the principal and the mirror to determine whether the other is up and running. If the witness detects that the principal server is no longer available, it instantly designates the mirror server as the new principal server and redirects all client transactions to the new principal server.

1.3 The Prototype Network

For simplicity the network does not include any DMZ or firewalls.

1.3.1 Web Servers (x 3) using Microsoft IIS v6

Setup as a Network Load Balanced Cluster.

The web root directory is setup as a Domain DFS (Distributed File System) with Replication (Mesh Topology).

1.3.2 Database Servers (x 3) using Microsoft SQL Server 2005

The database servers are setup for availability using “**Mirroring**”. The SQL Servers function as – Principal, Mirror and Witness. If the “principal” becomes unavailable then the “mirror” takes over.

References

Microsoft SQL Server 2005 Product Guide, Michelle Dumler. Published: September 2005