Implementing Business Continuity Plans using Active-Active Hosting Strategy

Business Continuity Plan is an essential part of organisation's response planning. It connotes how the business will perform its operation in case of a disaster and how it expects to resume 'business as usual' operations in the quickest possible time following the disaster.

Edited by
MOHAN DAS VISWAM

A Business Continuity Plan is made for critical applications that cater round the clock services to the clients, vendors or general public. They endeavour to ensure that the prioritized critical operation continues to be available even in the event of a disaster.

IMPLEMENTING BUSINESS CONTINUITY PLANS

They can be implemented using hosting strategies wherein both the DC and the DR infrastructure are always active. This approach has some additional advantages:

- Costly investments in the DR infrastructure are put to use.
- In case of a disaster, it was observed that sometimes the DR infrastructure was not found to be working to its full capability.
- Reduced requirement of conducting mock Disaster Recovery Drills

TYPES OF ACTIVE-ACTIVE HOSTING STRATEGY

- **Where both the sites are active:** Transactions can occur on either of the nodes and remains pending until it gets committed on both the ends. Time lag involved in this approach is about 5 ms round trip time for a distance less than 60 kms.
- **When both sites are separately located at relatively far distances:** This offers benefits such as strategic advantages and provides better sustainability, as the chances of a disaster striking at the both the sites simultaneously are less. It can also

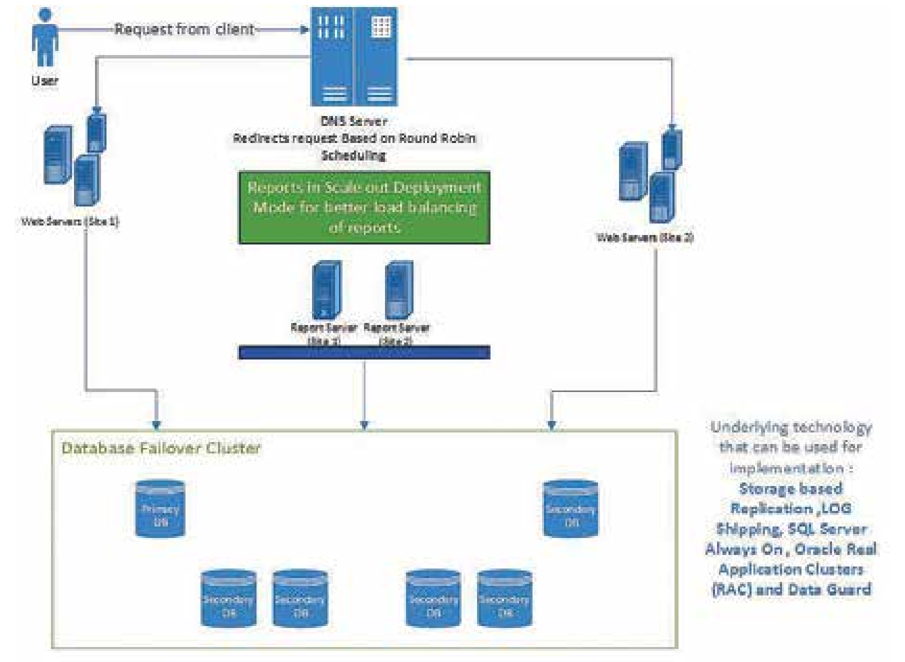
Request from client
User
DNS Server
Redirects request Based on Round Robin Scheduling
Reports in Scale out Deployment Mode for better load balancing of reports
Web Servers (Site 1)
Web Servers (Site 2)
Report Server (Site 1)
Report Server (Site 2)
Database Failover Cluster
Primary DB
Secondary DB
Secondary DB
Secondary DB
Secondary DB
Secondary DB
Underlying technology that can be used for implementation :
Storage based Replication ,LOG Shipping, SQL Server Always On , Oracle Real Application Clusters (RAC) and Data Guard

A Sample Architecture for Active-Active Hosting Strategy

DIPANKAR SENGUPTA
Senior Technical Director
dipankar.s@nic.in

MAYANK PRATIK
Scientist-B
mayank.pratik@nic.in

help in providing faster access times to the clients.

In the second strategy, the architecture is very easy to implement and forms a part of now highly popular High Availability and Disaster Recovery Strategies. Applications on both the sites have independent access and do not talk to each other. It is at the database level that there is sharing between both the sites. The Database Nodes can be hosted in a failover cluster and any disaster recovery DB solution can be implemented on these nodes.

The implementation of such an architecture previously required the use of Virtual LANs to form an application or database Geo Cluster (Cluster Across Geographical Locations). Advanced technologies available nowadays have eliminated the creation of VLAN for implementing this architecture which might have entailed huge costs in creating a Layer 2 Tunnel for this purpose.

The database nodes forming the part of the failover cluster uses a witness server to determine the quorum. If the quorum is satisfied, the clustering service is up and running and is able to cater to requests from the clients. The database cluster forms a logical group and is assigned a virtual IP / name. The application connects to this name and the cluster decides which of the nodes owns the clusters as of now and is active. The active node caters to the application providing a level of abstraction to the application. The failover of databases can take place but the application is not concerned with these, it just has to make a call to the listener which routes the request to the appropriate node.

The underlying technology that supports this kind of database hosting might involve Storage based replication, SAN-SAN replication, Log shipping or strategies implemented in software viz. SQL Server Always on, ORACLE Real Application Clusters (RAC) and Oracle Data Guard, SAP HANA (not exactly an active-active strategy) etc. Database Mirroring provides greater database availability by providing almost instantaneous failover. It can be used to maintain a standby database or a mirror database to the primary (principal) server.

Log shipping also operates at the database level and use shipping of logs to maintain one or more standby databases for a single primary database (production server).

ALWAYS-ON AVAILABILITY GROUP

An Always On Availability Group allows failover of a group of databases rather than a single database which is marked improvement over log shipping and database mirroring which were configured at the single database level. Availability groups are designed to support a primary database and a set of corresponding secondary database. An availability group fails at the level of an availability replica. Failovers are not caused by database issues such as a database becoming suspect due to a loss of a data file or corruption of a transaction log.

The primary replica makes the primary databases available for read-write connections from clients. Also, in a process known as Data Synchronization, which occurs at the database level, the primary replica sends transaction log records of each primary database to every secondary database. Every secondary replica caches the transaction log records (hardens the log) and then applies them to its corresponding secondary database.

Always On Availability Groups supports two availability modes:

- **Asynchronous-commit Mode:** This mode supports forced failover (manual with possible data loss) wherein the availability replicas are geographically separated. The primary replica sends the transaction confirmation immediately after writing log records to the local log file.
- **Synchronous-commit Mode:** It works similar to two phase commit protocol and supports Automatic and manual failover both.

Client connectivity to the database of a given availability group is provided by creating an Availability Group Listener. An Availability Group Listener is a Virtual Network Name (VNN) to which clients can connect in order to access a database in a primary or secondary replica of an Always On availability group. An Availability Group Listener enables a client to connect to an availability replica without knowing the name of the physical instance of SQL Server to which the client is connecting.

TAKING CARE OF REPORTING LOADS

Read-only Routing refers to the ability of SQL Server to route incoming connections to an availability group listener to a secondary replica that is configured to allow Read-only Workloads.

If Read-only Routing is configured for one or more readable secondary replicas, read-intent client connections to the primary replica are redirected to a readable secondary replica. Once these are configured, reports can be deployed in Scale-out Deployment Mode by which greater scalability and the ability to balance Reporting Loads can be achieved in a better way.

Implementing this strategy for websites will also help our website meet the **resiliency related** needs of critical information Infrastructure protection in line with the National Cyber Security Policy 2013.

For further information:
MAYANK PRATIK
Scientist –B
E-mail: mayank.pratik@nic.in