Swachh Bharat Mission - Gramin

Ensuring High Availability, Disaster Recovery and Design Implementation

Edited by
MOHAN DAS VISWAM

The eGovernance Solutions including Mobile Apps, developed by NIC, for Swachh Bharat Mission-Gramin (SBM-G) help the Ministry of Drinking Water & Sanitation track the progress of cleanliness campaigns in different states. SBM-G is built upon hybrid deployment, made with the combination of High Availability Solution and Disaster Recovery Solution for SQL Server at a database level.

SEEMANTINEE SENGUPTA
Sr. Technical Director
ssengupta@nic.in

PRAMOD KUMAR
Scientist-B
pramodk.yadav@nic.in

The Swachh Bharat Mission - Gramin (SBM-G) programme, under the Ministry of Drinking Water and Sanitation, is a flagship programme of the Government of India. The e-Governance application used for its monitoring has been developed in-house by NIC under the guidance of the SBM-G programme division. It is an integrated portal having various web-based modules, Mobile Apps and multiple GIS-based dashboards for close and effective monitoring of the progress of implementation at various levels, viz. the PMO, all states and districts across the country and at the Ministry itself. The software deployment has been done on Microsoft Virtual Machines (VMs) on MeghRaj Cloud Infrastructure with SQL Server 2016 as the backend database with a High Availability (HA) and Disaster Recovery (DR) solution. This article describes various steps required to set up an HADR solution, using SQL Server 2016.

Need of HADR Solution

HA and DR strategies strive to address non-functional requirements such as performance, system availability, fault tolerance, data retention, business continuity and user experience. It is imperative that selection of an appropriate HA and DR strategy is driven by business requirements. For HA, all service level agreements expected of the system have to be considered. For defining DR requirements, measurable characteristics such as Recovery Time Objective (RTO) and Recovery Point Objective (RPO) have to be considered.

SQL Server offers solutions for various HADR scenarios and it comes with a set of features and capabilities that can help organisations achieve a wide range of availability/ SLA goals. Typical high availability solutions involve the deployment of costly, redundant and passive servers. However, the aim should be to eliminate idle hardware while improving cost efficiency and performance. The AlwaysOn Availability Groups, a feature of SQL Server, enables the utilisation of secondary database replicas on otherwise passive or idle servers for read-only workloads. The ability to simultaneously utilise both the primary and secondary database replicas, helps improve the performance of all workloads due to better resource balancing across the server hardware investments.

Features such as the Configuration Wizard, support for the Windows PowerShell command-line interface, dashboards, Dynamic Management Views (DMVs), policy-based management and System Centre integration help simplify deployment and management of availability groups. The business goals for Recovery Point Objective (RPO) and Recovery Time Objective (RTO) should be key drivers in selecting a SQL Server technology for High Availability and Disaster Recovery (HADR) solution. The following table indicates the potential data loss and recovery time applicable for SBM-G programme for resuming the services if any disaster occurs. Almost zero data loss is being ensured by using the SQL Server HADR solution.

HADR Deployment Architecture

Hybrid deployment by combining the High Availability Solution with Disaster Recovery Solution for SQL Server at a database level with AlwaysOn Availability Groups has been used as HADR Solution for Swachh Bharat Mission-Gramin.

High Availability Solution

AlwaysOn Availability Groups helps ensure the availability of application databases, and they enable zero data loss through log-based data movement for data protection without shared disks.

Availability Groups provide an integrated set of options including automatic and manual failover of a logical group of databases, support for up to eight secondary replicas, fast application failover and

High Availability and Disaster Recovery SQL Server Solution	Potential Data Loss (RPO)	Potential Recovery Time (RTO)	Automatic Failover	Readable Secondaries
AlwaysOn Availability Group - Synchro-nous-Commit	Zero	Seconds	Yes	1 - 8
AlwaysOn Availability Group - Asynchro-nous-Commit	Seconds	Minutes	No	1 - 8
AlwaysOn Failover -Cluster Instance	NA	Seconds -to- minutes	Yes	NA
Backup, Copy, Restore	Hours	Hours -to- days	No	Not during a restore

Table 1: RTO/RPO for SBM-G HADR solution

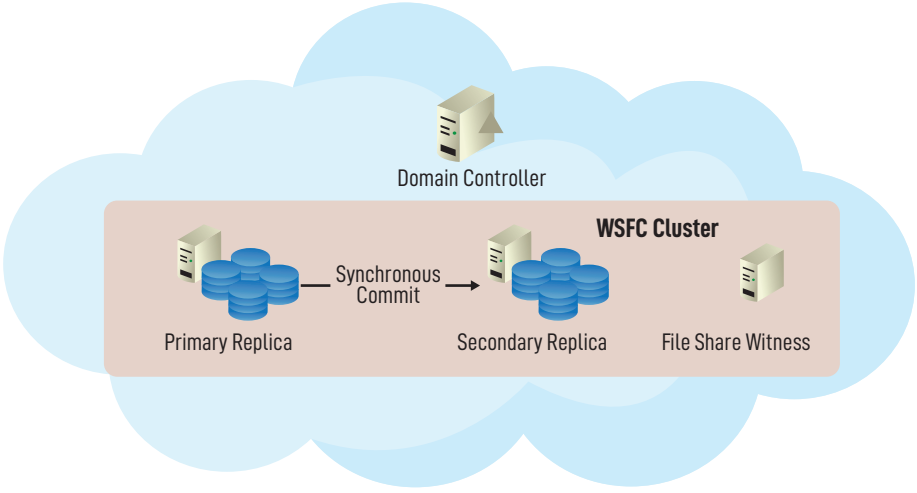

Fig. 1: Primary Domain Controller setup for HADR solution of SBM-G

automatic page repair. Availability replicas running in MeghRaj VMs in the same region (NDCSP-Delhi) provide high availability. Domain Controller and Alternative Domain Controller VMs were created and configured as Forest (Domain). Here it is to be noted that Windows failover clustering requires an Active Directory domain.

Domain Controller: A domain controller is a server that is running a version of the Windows Server® operating system and has Active Directory® Domain Services installed.

Windows Server Failover Clustering (WSFC): A Windows Server Failover Cluster (WSFC) is a group of independent servers that work together to increase the availability of applications and services. SQL Server takes the advantage of WSFC services and capabilities to support AlwaysOn Availability Groups.

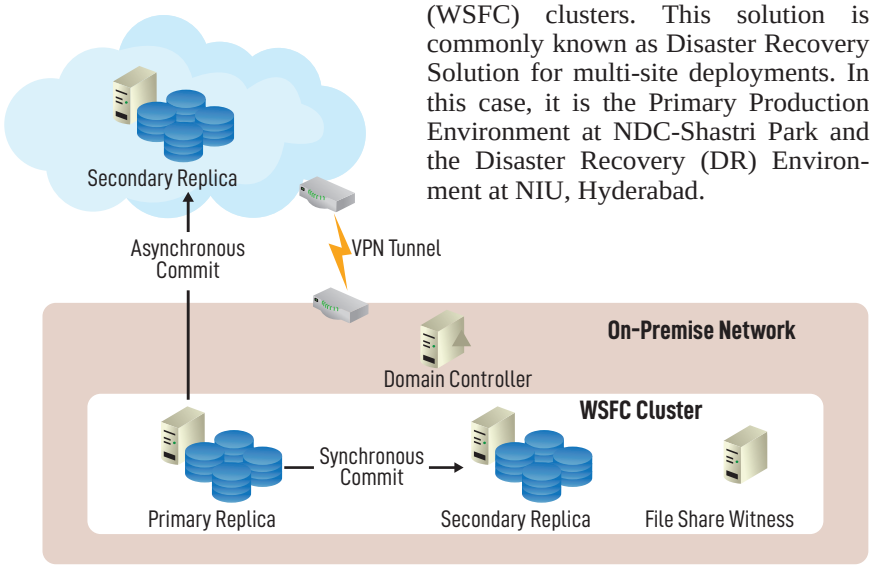

Fig. 2: Async Commit on Secondary Site/Nodes and Sync Commit on Primary Site/ Nodes using AlwaysOn Feature

Deploying a Cloud Witness for a Failover Cluster: SQL Server, AlwaysOn Availability Groups takes the advantage of WSFC as a platform technology. WSFC uses a quorum-based approach to monitor overall cluster health and maximises node-level fault tolerance. A fundamental understanding of WSFC quorum modes and node voting configuration is very important in designing, operating and troubleshooting AlwaysOn Availability Groups solution.

Disaster Recovery Solutions

The Disaster Recovery Solutions for SQL Server using Availability Groups was set up using database mirroring and log shipping, with 'backup-restore' using SAN storage, under MeghRaj Cloud Environment, in NIU-Hyderabad. Distributed Availability Group technology is being used for the DR Solution.

Distributed AG is used ***'when we want the data to continually replicate to the DR site, but don't want a potential network problem or issue at the DR site to bring down the primary site'.***

Any issues with the failover cluster nodes on the DR data centre do not affect the availability of the SQL Server Availability Group on the production data centre. A Distributed Availability Group is a special type of Availability Group that spans two separate Availability Groups. The underlying Availability Groups are configured on two different Windows Server Failover Clustering (WSFC) clusters. This solution is commonly known as Disaster Recovery Solution for multi-site deployments. In this case, it is the Primary Production Environment at NDC-Shastri Park and the Disaster Recovery (DR) Environment at NIU, Hyderabad.

One of the important steps in deploying this solution was to enable the Domain Controller Replication. For managing Replication between Sites (Production and DR), one has to install a replica Active Directory domain controller in between two data centres (NDCSP-DELHI and NIU-Hyderabad), using NICNET Network to create full trust in all other Database Nodes deployed at DR site.

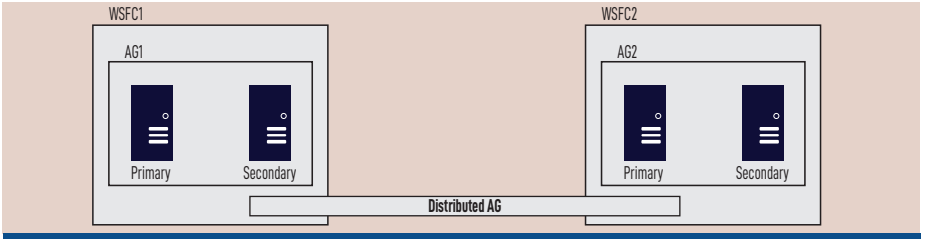

Figure 3: Distributed AlwaysOn Groups (DAG) between Primary (AG 1) and Secondary (AG 2) replicas with different Domain Controllers (WSFC1 and WSFC2)

Another important feature of this solution is that the Distributed Availability Groups allows the creation of an Availability Group even before performing the cutover, thereby ensuring no downtime, in case of hardware malfunction in any of the cluster nodes.

Firewall rules required for Domain Replication

Protocol and Port	AD and AD DS Usage	Type of traffic
TCP and UDP 389	Directory, Replication, User and Computer Authentication, Group Policy, Trusts	LDAP
TCP 636	Directory, Replication, User and Computer Authentication, Group Policy, Trusts	LDAP SSL
TCP 3268	Di rectory, Replication, User and Computer Authentication, Group Policy, Trusts	LDAP GC
TCP 3269	Directory, Replication, User and Computer Authentication, Group Policy, Trusts	LDAP GC SSL
TCP and UDP 88	User and Computer Authentication, Forest Level Trusts	Kerberos
TCP and UDP 53	User and Computer Authentication, Name Resolution, Trusts	DNS
TCP and UDP 445	Replication, User and Computer Authentication, Group Policy, Trusts	SMB,CIFS,SMB2, DFSN, LSARPC, NbtSS, NetLogonR, SamR, SrvSvc
TCP 25	Replication	SMTP
TCP 135	Replication	RPC, EPM
TCP Dynamic 49152 65535	Replication, User and Computer Authentication, Group Policy, Trusts	RPC, DCOM, EPM, DRSUAPI, NetLogonR, SamR, FRS
TCP 5722	File Replication	RPC, DFSR (SYSVOL)
UDP 123	Windows Time, Trusts	Windows Time
TCP and UDP 464	Replication, User and Computer Authentication, Trusts	Kerberos change/set password
UDP Dynamic 49152 65535	Group Policy	DCOM, RPC, EPM
UDP 138	DFS, Group Policy	DFSN, NetLogon, NetBIOS Datagram Service
TCP 9389	AD DS Web Services	SOAP
UDP 137	User and Computer Authentication,	NetLogon, NetBIOS Name Resolution
TCP 139	User and Computer Authentication, Replication	DFSN, NetBIOS Session Service, NetLogon

Summary

HADR solutions are designed and implemented to minimise or mitigate the impact of downtime of any Mission Critical Application like Swachh Bharat Mission – Gramin. High availability is ultimately measured in terms of end user's experience and expectations. The tangible and perceived business impact of downtime may be expressed in terms of information loss, property damage, decreased productivity, opportunity costs, contractual damages, or the loss of goodwill. Data redundancy is a key component of a high availability database solution. Transactional activity on primary SQL Server instance is synchronously or asynchronously applied to one or more secondary instances. When an outage occurs, transactions that were in flight may be rolled back, or they may be lost on the secondary instances due to delays in data propagation. Disaster recovery efforts address what is done to resume the operations. It can provide data and hardware redundancy within and across data centres and improves application failover time to increase the availability of your mission-critical applications. AlwaysOn provides flexibility in configuration and enables reuse of existing hardware investments. ■

For further information, please contact:
SEEMANTINEE SENGUPTA
Senior Technical Director & HoD
Drinking Water/Sanitation Informatics
NIC, A-Block, CGO Complex, Lodhi Road
NEW DELHI- 110003
Email: ssengupta@nic.in
Phone: 011-24362610