Framework Deployment Guide

Disaster Recovery Using MS Failover Cluster and MS SQL AlwaysOn

4/26/2025

Disaster Recovery Using MS Failover Cluster and MS SQL AlwaysOn

Contents

This section describes a recommended architecture to ensure successful disaster recovery, or business continuity, using MS Failover Cluster with MS SQL AlwaysOn, following a scenario in which the main site was rendered inoperable because of some natural or other disaster. For more information, including configuration details, see Configuring Disaster Recovery Using MS Failover Cluster and SQL AlwaysOn.

Overview

The Genesys system configuration is stored in a single database and can be accessed by only one primary master Configuration Server connection at a time. The Configuration Database, constantly modified by Configuration Server clients, is archived periodically to prevent the loss of data. Database maintenance and periodic backup can cause significant downtime. It cannot prevent partial or whole loss of configuration data if a major disaster occurs, such as one in which the Configuration Database and all updates and modifications made since the last backup is completely lost.

To improve the robustness of the Management Framework solution and to reduce downtime for system maintenance, this architecture replicates a live database to a secondary live standby database. This solution is based on non-shared storage, because each instance of SQL Server in the topology has its own copy of data and does not need to share storage. If a major disaster occurs, that secondary database can be accessed by a secondary Master Configuration Server that is brought online from a dormant state. Network traffic redirection for Configuration Server Proxies residing on each site to a host running the secondary Master Configuration Server will be done by Microsoft Windows Server Failover Cluster. Operations at sites are continued uninterrupted in limited mode without a configuration change until the secondary Master Configuration Server is brought online and restored to normal mode after the Proxy servers reconnect to the secondary Master Configuration Server.

Components

This architecture consists of the following components:

- Main live MS SQL database server at Site 1.
- Secondary live MS SQL database server at Site 2.
- Microsoft Windows Server Failover cluster.
- Microsoft AlwaysOn High Availability Group to replicate the live Configuration Database to a secondary live standby database and log message databases cross sites replication.
- Main live redundant pair master Configuration Server primary/backup pair at Site 1.
- Secondary dormant (not running in normal operation mode) master Configuration Server primary/backup pair at Site 2.
- Main live Solution Control Server in distributed mode to control the main master Configuration Server pair at Site 1.
- Secondary dormant Solution Control Server in distributed mode to control the secondary master Configuration Server pair at Site 2.

- Main Message Server at Site 1 to support communication between Solution Control Servers controlling site components, such as Configuration Server Proxy HA pairs, T-Servers, and Log Message Servers.
- Secondary dormant (not running in normal operation mode) Message Server at Site 2 to support communication between Solution Control Servers controlling site components, such as Configuration Server Proxy HA pairs, T-Servers, and Log Message Servers.
- Live Configuration Server Proxy pair at Site 1.
- Live Configuration Server Proxy pair at Site 1.
- Live Solution Control Server at Site 1.
- Live Solution Control Server at Site 2.
- Live Message Server for network logging at Site 1, connected to the Log Database at Site 1.
- Live Message Server for network logging at Site 2, connected to the Log Database at Site 2.
- Scripts to start and stop the master Configuration Server primary/backup pair and master Solution Control Servers.

Architecture

The following diagram illustrates the disaster recovery architecture for a multi-site configuration under normal conditions.

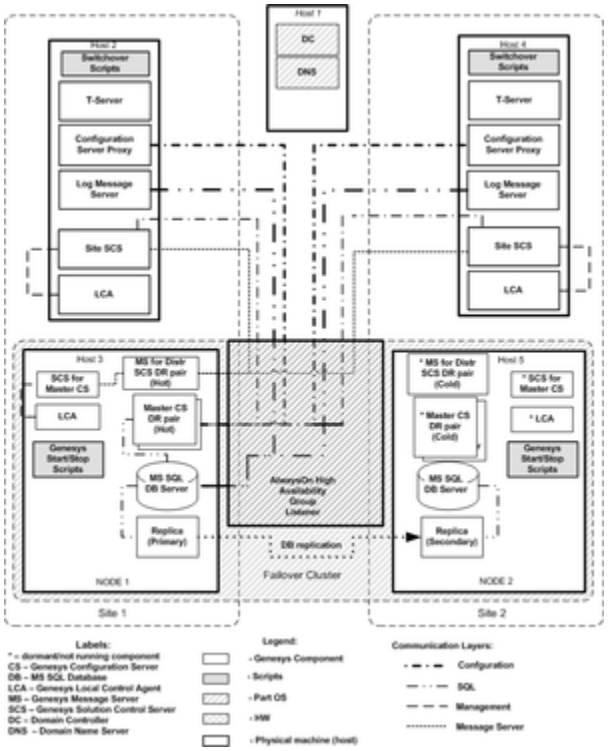

Multi-Site Disaster Recovery Architecture under Normal Operations

Solution Control Server

The Solution Control Servers used in this deployment are configured in Distributed SCS mode. Some or all can also be configured in HA pairs at each site.

At each site, a Solution Control Server is deployed on the management host (Hosts 2 and 4 in the diagram above) and is dedicated to managing Applications on the management hosts, specifically the Configuration Server and the dedicated Message Server for the Distributed Solution Control Servers, described next.

For Distributed Solution Control Servers to communicate with each other, a Message Server dedicated for distributed Solution Control Server use (that is, configured with **[MessageServer]signature=**`scs_distributed`) is also installed on each of the management hosts.

Each site also has a separate Solution Control Server deployed on the Host configured to manage Genesys applications running on each site (that is, the site SCS in the diagram above).

Depending on the number of applications, it is possible to deploy additional Distributed Solution Control Servers for load balancing.

For additional fault tolerance, Solution Control Servers can be deployed in high-availability (HA) pairs.

Message Server

Each site has its own instance of a Log Message Server to be used for network logging by applications running on the same site. The Message Servers are installed on the application host and managed by the site Solution Control Server. A Log Database is used at each site.

As mentioned in the previous section, a Message Server is also dedicated to communications between the distributed Solution Control Servers. This requires two instances (or two HA pairs) of Message Servers to be deployed, one of which is dormant. Each Message Server must be added to the Connections of each distributed Solution Control Server.

DBMS Solution Replication Processes Configuration

The underlying infrastructure of an Availability Group is Windows Server Failover Cluster (WSFC). The Availability Group listener is a virtual network name (VNN) that you create for use with a specific Availability Group. This name is bound to one or more TCP/IP addresses and listener ports and is used to automatically connect to the primary replica hosted at the time. The VNN eliminates the need to specify a failover partner attribute. For example, if an availability group fails over to Node 2 from Node 1, new connections to the availability group listener automatically connect to the replica currently hosting the primary replica. The current Configuration Server/MSSQL client implementation requires a stretched VLAN when cluster nodes are residing in different subnets.

An AlwaysOn High Availability Group has three database replication processes configured to replicate main to secondary for Configuration and Log Message databases.

> ### Important
> This is a non-shared solution. The nodes do not share any storage with another node.

> Each node is running an instance of SQL Server and has its own copy of the data.

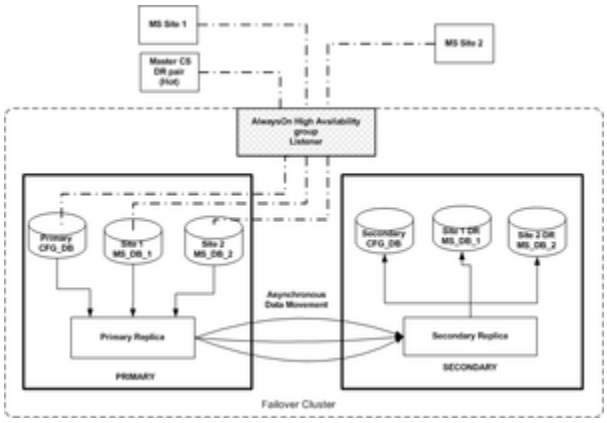

Multi-Site Database Replication