SQL 2012 AlwaysOn Quickstart

Ryan Adams
Blog - http://ryanjadams.com
Twitter - @ryanjadams
Email - ryan@ryanjadams.com

Objectives

- Discover the AlwaysOn Feature Set
- AlwaysOn Failover Cluster Instances
- AlwaysOn Availability Groups
- Windows Server Failover Clustering
- Multi-database Failover
- Availability Replicas
- Active Secondaries

What is SQL Server AlwaysOn?

- It's an HA and DR solution
- AlwaysOn is a combined term that encompasses
  - AlwaysOn Failover Cluster Instances
    - The old SQL Clustering
  - AlwaysOn Availability Groups
    - The old SQL Mirroring

What does AlwaysOn do for me?

- Server level protection
- SQL Server level protection

- Database level protection
- Client connection level protection/abstraction

AlwaysOn Failover Cluster Instances

What is an AlwaysOn FCI?

- It's SQL Clustering Renamed
- Shared storage requirement

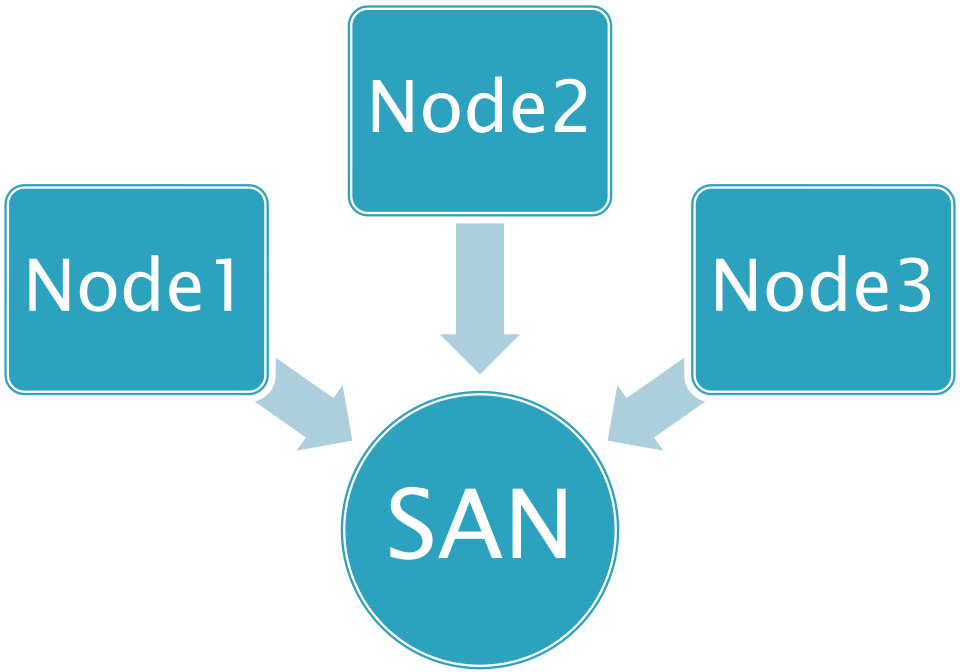

What's New?

- Multi-site Clustering
- TempDB on local storage
- Flexible Failover Policy

Multi-site Clustering

- It's really called Multi-subnet
  - Different subnets in the same datacenter
  - Different subnets in the different datacenters
- Storage level replication required
- Multiple IP addresses that use "OR"
- Network Name has multiple host records in DNS

TempDB

- TempDB is no longer required to be on a SAN or other shared storage
- Install will create the directory and give the SQL Server Service account permissions on the first node
- You must create the directory and add the permissions when adding additional node to the cluster

Flexible Failover Policy

- More options than the old LooksAlive and IsAlive Checks

- 3 Failure Detection Types
  - WSFC monitors the SQL Server service for a stop
  - Response from sp_server_diagnostics must be less than the HealthCheckTimeout setting
  - Conditions as reported by sp_server_diagnostics
    - Level 0 – 5
    - 0 is most restrictive and 5 is least restrictive

AlwaysOn Availability Groups

What is an AlwaysOn AG?

- It's a group of user databases that failover together
- Groups are often defined by the application they support
- It's Mirroring re-defined

Requirements

- Windows Server 2008 or later
- SQL Server 2012 Enterprise Edition
- Windows Server Failover Clustering
- Cannot be on a Domain Controller
- User databases only
- Databases must be in Full Recovery Model

What's Not Required

- SQL Failover Cluster Instance
- Shared Storage

WSFC vs SQL FCI

- Windows Server Failover Cluster Required
- SQL Failover Cluster Instance not Required
  - Mutually Exclusive
  - Can be used together
- WSFC Provides…
  - Quorum
  - Abstracts IP and Network Name
  - Monitors Server Health Status
- AGs report SQL Health Status to WSFC

Availability Group Architecture

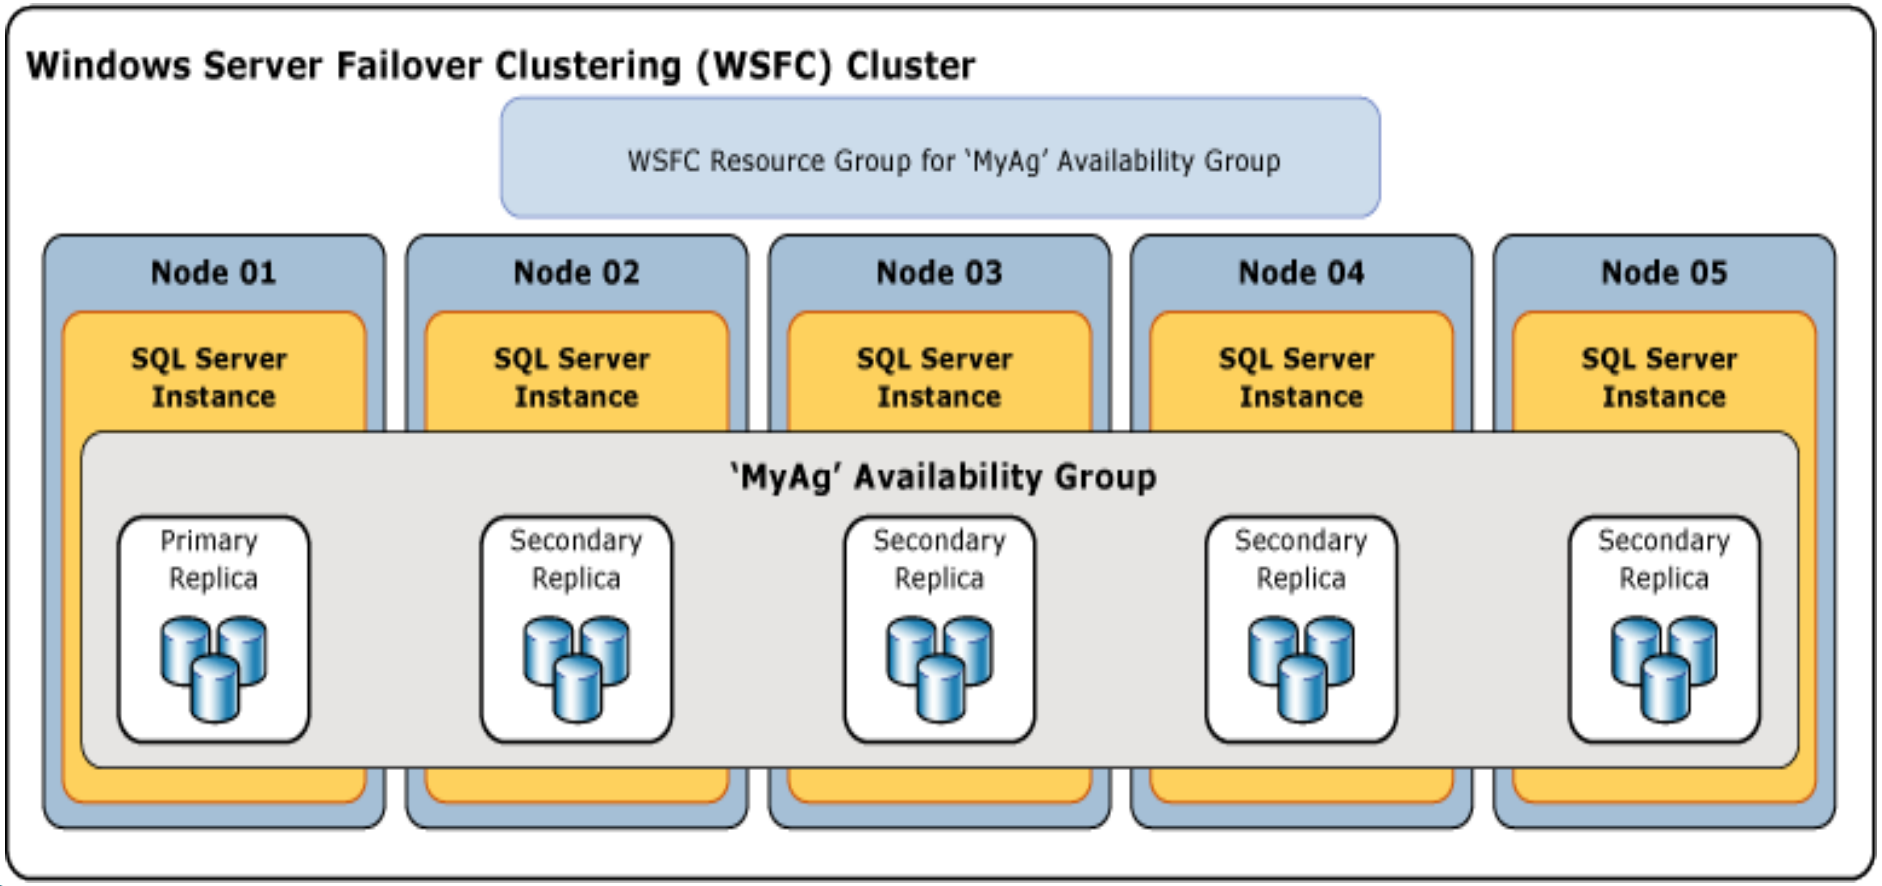

Image taken from SQL Server Books Online
http://msdn.microsoft.com/en-us/library/ff877884.aspx

Mirroring Comparison

What Stays

- Automatic Page Recovery
  - sys.dm_hadr_auto_page_repair
- Compression
- Encryption

What Goes

- Defining the FailoverPartner in the application connection string
- Witness

Don’t we still need that?

- Virtual Network Name
  - Replaced FailoverPartner
- Windows Server Failover Clustering
  - Replaced the witness

What's New

Multi-database Failover

- Group of DBs supporting an app that failover together as a group
- It is defined by an Availability Group

Availability Replicas

- Each instance of SQL Server in an AG is called a replica
  - Each instance can only house one replica for a particular AG
  - Each instance can, however, house a replica for multiple AGs
    - SPECIAL NOTE….A DB can only be part of one AG

Multiple Replicas

- 5 Total Replicas
  - 1 Primary
    - Master read/write copy
  - 4 Secondary
    - Read only copies

Replica Synchronization Modes

Synchronous

- No data loss
- Possible latency on primary replica
- Max of 2 replicas in an AG can use synchronous mode

Synchronous

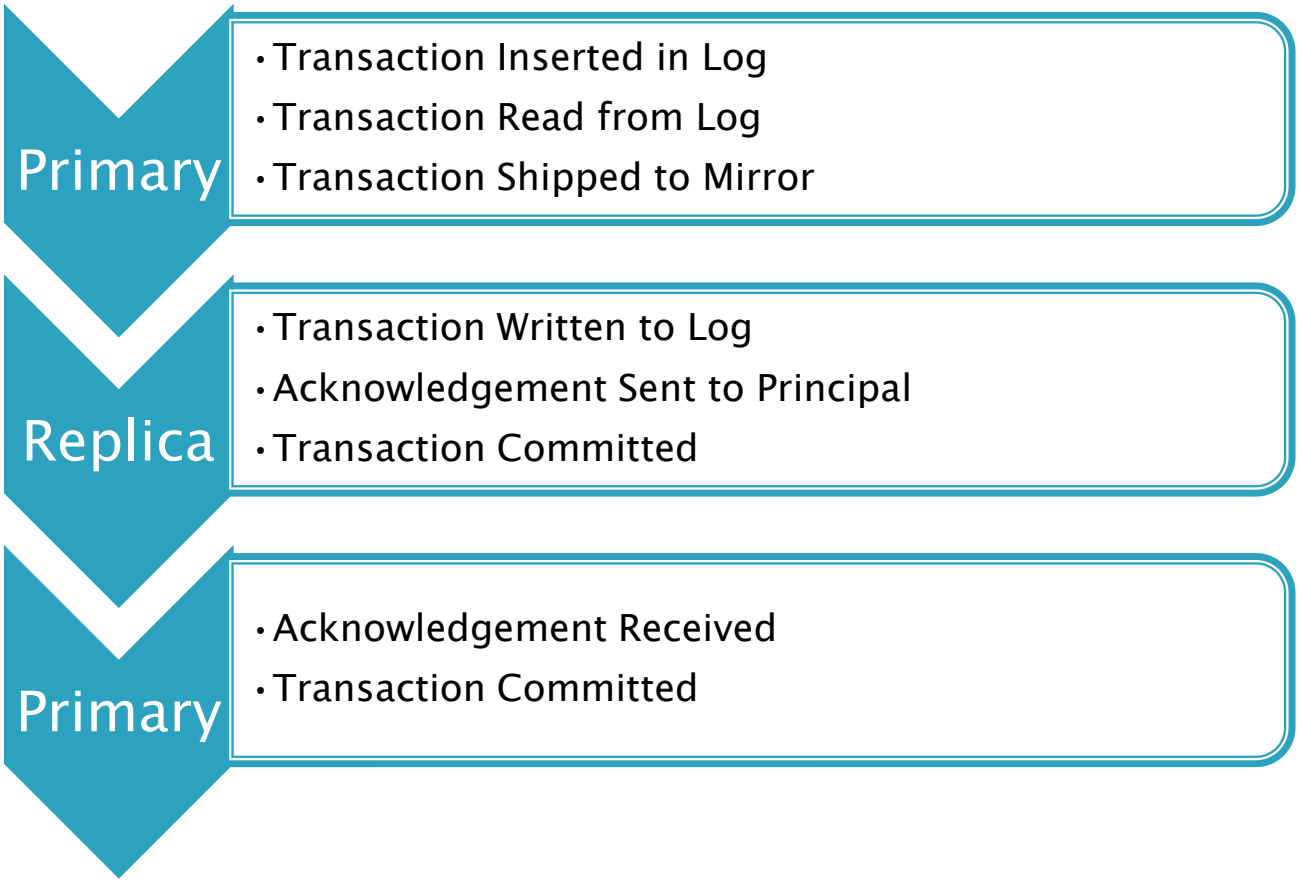

Asynchronous

- Possible data loss
- No latency on primary replica
- Max of 4 replicas in an AG can use Asynchronous mode

Asynchronous

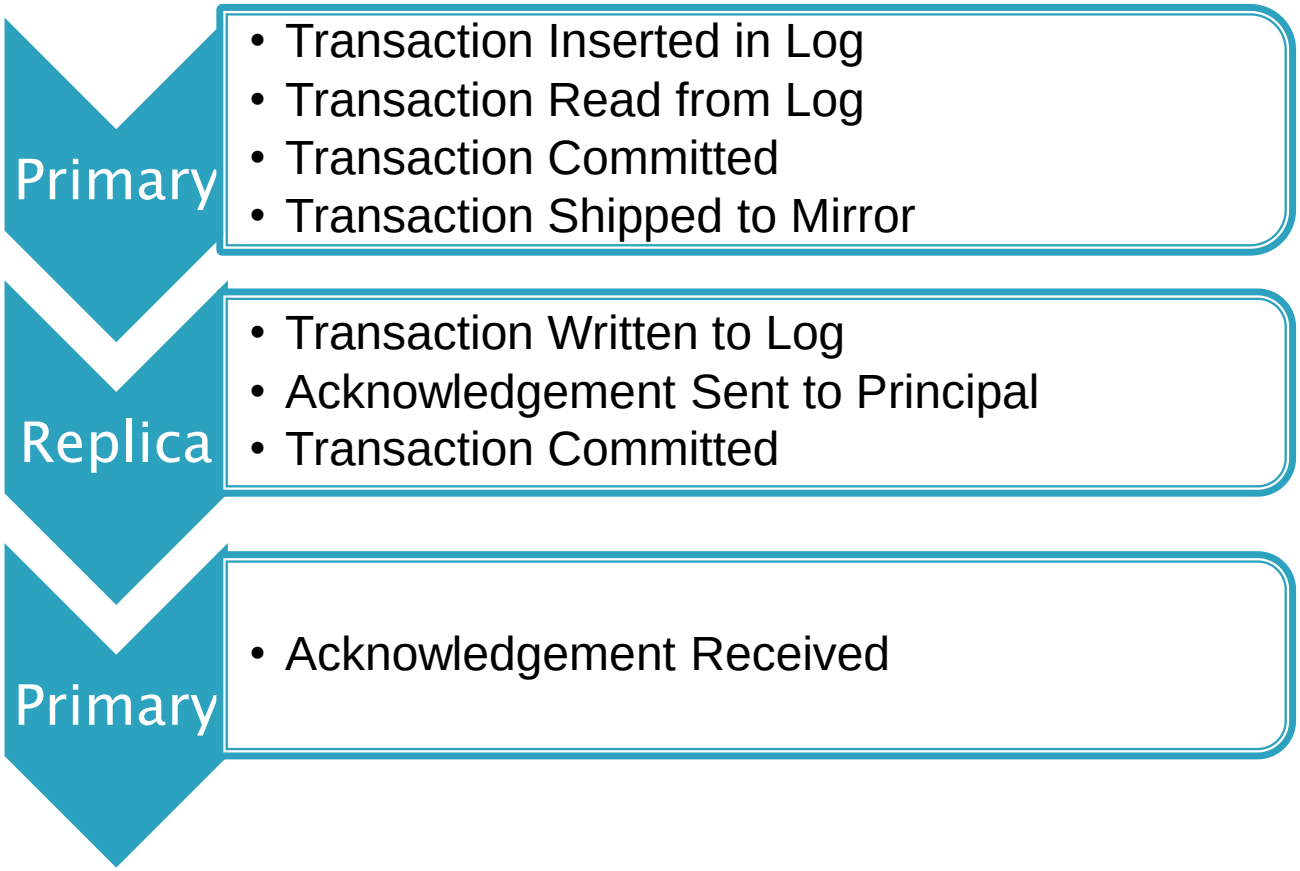

Replica Failover Modes

Automatic

- No data loss
- Synchronous Mode required
- Automatic Failover Enabled
- Not supported if primary or any secondary replica is on a SQL Failover Cluster Instance

Manual

- Only choice if a replica is on a SQL FCI
- Planned
  - Cannot be Asynchronous
  - No data loss
  - Everything is in a synchronized state
- Forced
  - Potential data loss
  - Primary Replica has failed
  - Secondary is not synchronized with the primary

Active Secondaries

What are they?

- Readable Replicas
  - Can be used for reporting
  - Can be used for backups
    - Copy-Only required
    - Differentials not yet supported
  - Can be used for DBCC

Active Secondary Gotchas

- Snapshot Isolation mode used so dirty reads are possible
- Cannot create separate indexes from the primary
- Temporary statistics are stored in TempDB
  - Plan for heavy TempDB usage and storage
  - These stats are removed upon restart

Active Secondary Replica Connection Types

- Read Intent Setting
  - NO
  - YES
  - READ INTENT
    - ApplicationIntent=ReadOnly

Availability Group Listener

- Allows you to connect to an AG without knowing which replica it is being hosted on
- Virtual Network Name with one or more IP Addresses
  - VNN is a Windows Clustering Network Name
  - Clients try each IP until one succeeds
- Can be viewed in Server Cluster Manager