

SQL Server 2012/2014 AlwaysOn Availability Group

Part 1 - Introduction

v1.0 - 2014 - G.MONVILLE

Summary

SQL Server 2012 AlwaysOn - Introduction.....	2
AlwaysOn Features	2
AlwaysOn FCI (Failover Cluster Instance).....	2
AlwaysOn Availability Group	3
Use the both features.....	4
License.....	5
Terms and Definitions.....	6
Overview of an AAG	7
Availability Group	7
Availability Replica.....	7
Availability Group listener	9
Database Mirroring Endpoint (SQL Server)	10
Authentication Mode.....	10
Data Encryption.....	11
WSFC.....	12
AAG Health Monitoring	12
Failover.....	13
Restrictions.....	13
TechNet Resources.....	14



SQL Server 2012 AlwaysOn - Introduction

This article describes the new SQL High Availability feature introduced with SQL Server 2012 version: AlwaysOn Availability Group.

AlwaysOn Features

First of all, do not confuse "AlwaysOn Availability Group" and "AlwaysOn FCI". This is two distinct solutions:

AlwaysOn FCI (Failover Cluster Instance)

This is the classic SQL Cluster based on the WSFC (Windows Server Failover Clustering) functionality. FCI provide local high availability through redundancy at the server-instance level (a single instance installed across multiple nodes). In case of failure of a node, the instance will start on another node. Clients connect to this instance from a VNN (Virtual Network Name) which is a Cluster resource.

FCI required Shared Storage (SAN, SMB...) for database data and logs. This storage must be configured on all nodes participating to an FCI. In the FCI solution, Shared Storage is a SPOF (Single Point Of Failure).

SQL Services (Instance) can failover between WSFC nodes (so the loss of a server is covered). If the Storage fails the service stops.

Note:

- For FCI on the same site it's possible to eliminate the Storage SPOF with SAN Replication.
- For FCI across remote sites you need to implement SAN Replication.

Implementing "SAN Replication" Cons:

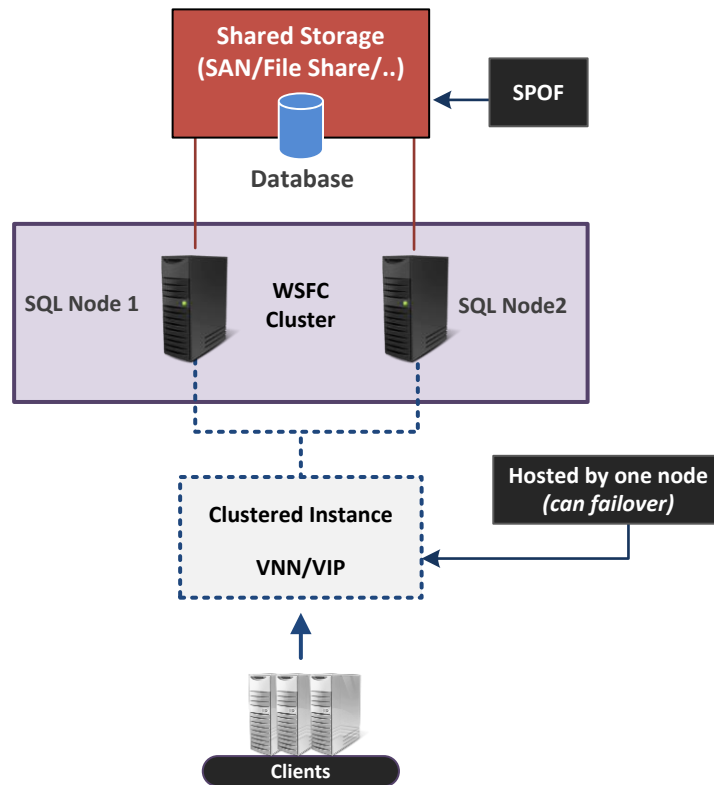
- generates additional costs (this is also the case for AAG; FCI is available on the SQL Standard version, AAG required an SQL Enterprise version) except if it's included in your SAN license.
- it's more complicated to implement on remote geographic sites.
- it's require additional storage skills for the operation team.
- the Availability Group offer more options (like asynchronous replica, backup on secondary replica, ..)

The choice of use FCI or AAG depends of your Business requirement (RTO, RPO, DRP, Remote sites ...)

For comparison, AlwaysOn Availability Groups doesn't require Shared Storage.

I will write an article about AlwaysOn FCI with Windows 2012 R2 SOFS Storage (Scaled-Out File Server).

AlwaysOn FCI



AlwaysOn Availability Group

The Availability Group feature is a mix of SQL Clustering and SQL Mirroring (it is also presented by MS as an alternative to SQL Mirroring which is deprecated since SQL Server 2012).

The AG is also based on a WSFC cluster; the difference is that on each node a SQL instance is installed and **active**.

AAG is composed of replicas (a replica is a group of one or more database). There is one Primary Replica and one to four Secondary Replicas (eight for SQL 2014) (a Secondary Replica is a copy of databases from the Primary Replica, when a DB is modified, changes are replicated on all Secondary replica). An AAG is a WSFC cluster resource group (which contains at least the VNN and VIP cluster resources on which the clients will connect). At time T, the AAG is hosted by an instance, so this is the primary replica (database(s) accessible in R/W) and all other replicas synchronize thereon. There are two types of synchronization: synchronous and asynchronous.

If the instance that hosts the primary replica becomes unavailable the AG (the cluster group resource) will switch to another instance and the secondary replica on this instance will become the Primary replica (this failover requires that the secondary replica is set in "synchronous" mode and the status is "Synchronized").

Automatically failover cannot be done to an asynchronous secondary replica. In asynchronous mode only "forced manual failover" (with possible Data loss) is allowed, this action must be performed by a DBA. This mode is generally used for a DRP on one or more remote site or if the maximum number of replica in synchronous mode is reaches - Max for SQL 2012/2014: 3 (1x Primary + 2x Secondary)

A Cluster Resource Group (for SQL the AG) can be hosted only by one instance at a time. So, on a given AG only one instance at a time can be "active" (active = which hosts Database(s) accessible in R/W = the Primary Replica).

On all replicas in an Availability Group only one can be "Primary Replica" at a time. For example, with an AG composed of two nodes and containing two databases, it's impossible to have one Database active on one node and the second DB active on the second node. To achieve this configuration we need to create two AG (on the same two nodes) with one Database in each AG (so there are two Primary Replica, this is the configuration that I will do in this article).

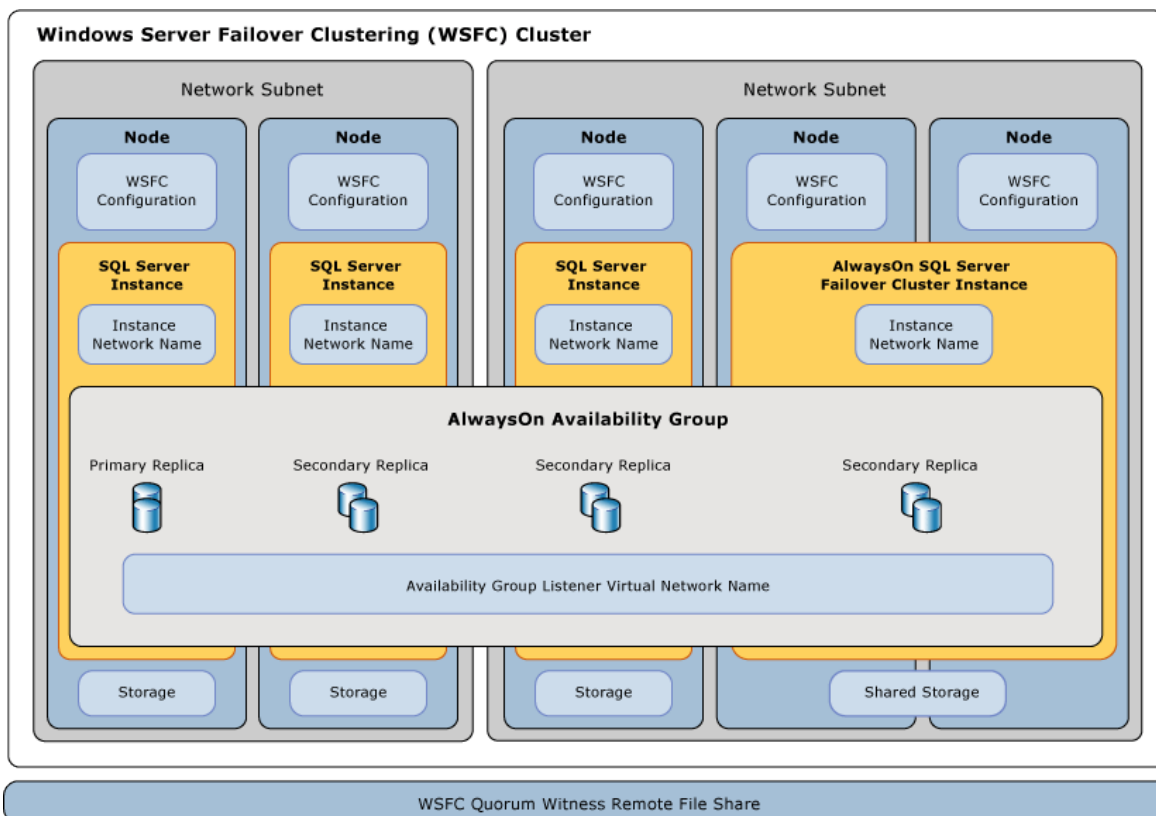
Use the both features

It's possible to use AlwaysOn FCI in an AlwaysOn Availability Group solution. In this case a SQL Failover Clustering Instance can host an Availability Replica for Availability Group

Note for Failover Mode: SQL FCIs do not support automatic failover by availability groups, so any availability replica that is hosted by an FCI can only be configured for manual failover.

The implementation of a FCI in an AG will be covered in another article.

Overview (from TechNet):



<http://technet.microsoft.com/en-us/library/hh270278.aspx>

License

High Availability features for SQL Server 2012 Licenses:

Feature Name	Enterprise	Business Intelligence	Standard	Web	Express
Server Core support	Yes	Yes	Yes	Yes	Yes
Log Shipping	Yes	Yes	Yes	Yes	
Database mirroring	Yes	Yes (Safety Full Only)	Yes (Safety Full Only)	Witness only	Witness only
Backup compression	Yes	Yes	Yes		
Alwayson Failover Cluster Instances	Yes (Node support: OS maximum)	Yes (Node support: 2)	Yes (Node support: 2)		
AlwaysOn Availability Groups	Yes (up to 4 secondary replicas, including 2 synchronous secondary replicas)				
...					

High Availability features for SQL Server 2014 Licenses:

Feature Name	Enterprise	Business Intelligence	Standard	Web	Express
Server Core support	Yes	Yes	Yes	Yes	Yes
Log Shipping	Yes	Yes	Yes	Yes	
Database mirroring	Yes	Yes (Safety Full Only)	Yes (Safety Full Only)	Witness only	Witness only
Backup compression	Yes	Yes	Yes		
Alwayson Failover Cluster Instances	Yes (Node support: Operating system maximum)	Yes (Node support: 2)	Yes (Node support: 2)		
AlwaysOn Availability Groups	Yes (up to 8 secondary replicas, including 2 synchronous secondary replicas)				
...					

For more information on Features/Licenses:

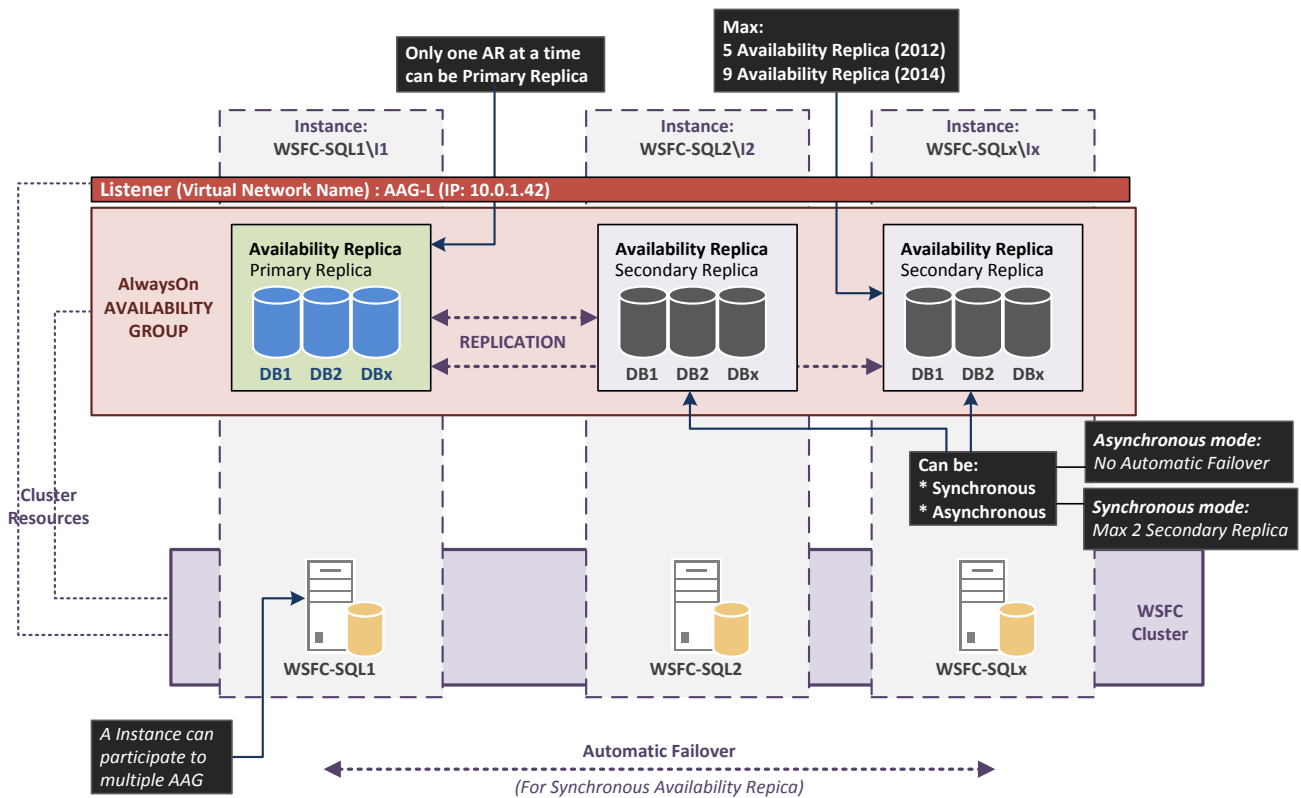
- **SQL Server 2012 Features:**
http://msdn.microsoft.com/en-us/library/cc645993.aspx#High_availability
- **SQL Server 2014 Features:**
http://msdn.microsoft.com/en-us/library/cc645993%28v=sql.120%29.aspx#High_availability

Terms and Definitions

Term	Definition
availability group	A container for a set of databases, availability databases, that fail over together.
availability database	A database that belongs to an availability group. For each availability database, the availability group maintains a single read-write copy (the primary database) and one to four read-only copies (secondary databases).
primary database	The read-write copy of an availability database.
secondary database	A read-only copy of an availability database.
availability replica	An instantiation of an availability group that is hosted by a specific instance of SQL Server and maintains a local copy of each availability database that belongs to the availability group. Two types of availability replicas exist: a single primary replica and one to four secondary replicas.
primary replica	The availability replica that makes the primary databases available for read-write connections from clients and, also, sends transaction log records for each primary database to every secondary replica.
secondary replica	An availability replica that maintains a secondary copy of each availability database, and serves as a potential failover targets for the availability group. Optionally, a secondary replica can support read-only access to secondary databases can support creating backups on secondary databases.
availability group listener	A server name to which clients can connect in order to access a database in a primary or secondary replica of an AlwaysOn availability group. Availability group listeners direct incoming connections to the primary replica or to a read-only secondary replica.

Source TechNet: <http://technet.microsoft.com/en-us/library/aa427606-8422-4656-b205-c9e665ddc8c1#TermsAndDefinitions>

Overview of an AAG



Availability Group

An AG consists of:

- Two or more Availability Replica (max for SQL Server 2012: 5 / max for SQL Server 2014: 9)
- One (or more) Availability Group listener

Availability Replica

An Availability Replica contains:

- A set of Databases (at least one) (there are no maximum for a set; this depends on the load and Server Performances).

A Database add to an AG is known as Availability Database

Each availability replica is hosted by an instance of SQL Server residing on different nodes of a WSFC cluster (All nodes members of the same Cluster).

Role:

An Availability Replica can have the role of:

- **Primary Replica** - Database(s) is accessible in Read/Write
- **Secondary Replica** - Database(s) is accessible in Read Only (if it's configured)

On a given AAG, only one availability replica can be "Primary" at a time, others are "Secondary". Secondary Replica contains a "copy" of databases from Primary (during creation a backup is created from the Primary DB and imported on all secondary).

Availability Mode:

An Availability Replica can be configured in mode:

- **Asynchronous-commit mode** - Under asynchronous-commit mode, the primary replica commits transactions without waiting for acknowledgement that an asynchronous-commit secondary replica has hardened the log. Asynchronous-commit mode minimizes transaction latency on the secondary databases but allows them to lag behind the primary databases, making some data loss possible.
- **Synchronous-commit mode** - Under synchronous-commit mode, before committing transactions, a synchronous-commit primary replica waits for a synchronous-commit secondary replica to acknowledge that it has finished hardening the log. Synchronous-commit mode ensures that once a given secondary database is synchronized with the primary database, committed transactions are fully protected. This protection comes at the cost of increased transaction latency.

Allow Readable Secondary:

When you configure the AAG, you can configure this option (Known as "Active Secondary Replica"): With this option you can do the following action:

- **Performing backup operations on secondary replicas** - The secondary replicas support performing log backups and copy-only backups of a full database, file, or filegroup. You can configure the availability group to specify a preference for where backups should be performed.

For more information, see Technet: [Active Secondaries: Backup on Secondary Replicas \(AlwaysOn Availability Groups\)](#).

- **Read-only access to one or more secondary replicas (readable secondary replicas)** - Any availability replica can be configured to allow read-only access to its local databases when performing the secondary role, though some operations are not fully supported. Also, if you would like to prevent read-only workloads from running on the primary replica, you can configure the replicas to allow only read-write access when running under the primary role.

For more information, see [Active Secondaries: Readable Secondary Replicas \(AlwaysOn Availability Groups\)](#).

If an availability group currently possesses an availability group listener and one or more readable secondary replicas, SQL Server can route read-intent connection requests to one of them (read-only routing).

For more information, see [Using a Listener to Connect to a Read-Only Secondary Replica \(Read-Only Routing\)](http://technet.microsoft.com/en-us/library/hh213417.aspx#ConnectToSecondary) : <http://technet.microsoft.com/en-us/library/hh213417.aspx#ConnectToSecondary>

An excellent TechNet article about Read-Only Routing:

Read-Only Routing with SQL Server 2012 Always On Database Availability Groups
<http://social.technet.microsoft.com/wiki/contents/articles/13503.read-only-routing-with-sql-server-2012-always-on-database-availability-groups.aspx?Sort=MostUseful&PageIndex=1>

Availability Group listener

An availability group listener is a Virtual Network Name (VNN) to which clients can connect to access a database (without knowing the name of the instance). Availability group listeners direct incoming connections to the primary replica or to a read-only secondary replica. The listener provides fast application failover after an availability group fails over.

A VNN it's a WSFC cluster resource and consists of:

- A DNS Name (a VNN)
- A Port Number
- An IP Address (a VIP)

The Listener is always owned by the SQL Server instance where the primary replica resides.

Note about Port:

A Listener can share a same port with an Instance (for example default instance port is used: 1433, and the Listener is also configured with the 1433 port) or use a different port.

For more information:

Availability Group Listeners, Client Connectivity, and Application Failover (SQL Server)
<http://msdn.microsoft.com/en-us/library/hh213417.aspx>

Database Mirroring Endpoint (SQL Server)

Instances participating to an AlwaysOn Availability Group (or a Database Mirroring) use "Database Mirroring Endpoints" to communicate among themselves.

Each Instance must have its own dedicated Endpoint (only one per instance, you cannot create several endpoints for a same instance). After SQL Server Instance installation, the endpoint is not created.

Note: An endpoint is a SQL Server object that enables SQL Server to communicate over the network. "Database Mirroring Endpoint" is a specify endpoint use by AAG or SQL Mirroring for instances communications.

To host an availability replica an instance must have a Database mirroring endpoint (when you create a Replica you have to specify the Endpoint URL of the instance that will host the replica). The instance uses this endpoint to listen messages from availability replicas hosted by other instances.

Database mirroring endpoints use TCP protocol to send and receive messages between the server instances. The Endpoint URL is composed with the Server FQDN (can be also the Server name, an IPv4 or IPv6 address) and a defined port number. The port number uniquely identifies an instance.

Example: Two instances on one server, you can have two endpoints:

TCP://sql-srv1.ad.corp:**5022** => refers to **sql-srv1\i1**

TCP://sql-srv1.ad.corp:**5023** => refers to **sql-srv1\i2**

Authentication Mode

Two types of Authentication are available for database mirroring endpoints:

- **Certificate-based authentication** (If any server instance is running under a built-in account, such as Local System, Local Service, or Network Service, or a non-domain account, you must use certificates for endpoint authentication).

This part is not covered in this article, for more information see:

Use Certificates for a Database Mirroring Endpoint (SQL Server)

<http://technet.microsoft.com/en-us/library/ms191477.aspx>

- **Windows Authentication**

- If all instances use the same domain account, no extra configuration is required. (Microsoft recommends using the same account)
- If instances use different domain accounts, the login of each account must be created in **master** on each of the other server instances, and that login must be granted CONNECT permissions on the endpoint.

Data Encryption


By default encryption is configured to "Required" on a DB Mirroring Endpoint. You can change this (not recommended) to "Disabled" or "Supported"

Encryption algorithms:

You can also change encryption algorithms. By default on an AlwaysOn instance encryption is configured with AES algorithm.

TechNet:

Optionally, you can control the encryption algorithms that can be used by an endpoint, by specifying one of the following values for the ALGORITHM option in a CREATE ENDPOINT statement or ALTER ENDPOINT statement:

ALGORITHM value	Description
RC4	Specifies that the endpoint must use the RC4 algorithm. This is the default. <div data-bbox="347 786 1369 958"> Note The RC4 algorithm is deprecated. This feature will be removed in a future version of Microsoft SQL Server. Do not use this feature in new development work, and modify applications that currently use this feature as soon as possible. We recommend that you use AES.</div>
AES	Specifies that the endpoint must use the AES algorithm.
AES RC4	Specifies that the two endpoints will negotiate for an encryption algorithm with this endpoint giving preference to the AES algorithm.
RC4 AES	Specifies that the two endpoints will negotiate for an encryption algorithm with this endpoint giving preference to the RC4 algorithm.

If connecting endpoints specify both algorithms but in different orders, the endpoint accepting the connection wins.

For more information, see TechNet article "**Choose an Encryption Algorithm**":

<http://technet.microsoft.com/en-us/library/ms345262.aspx>

WSFC

The AlwaysOn Availability Group (and also the AlwaysOn FCI) feature is based on the WSFC service. All SQL Server nodes participating to an AAG are members of a WSFC Cluster.

On each node an SQL Server 2012/2014 **Enterprise** Instance must be deployed.

An AAG is register as a cluster Resource Group in the WSFC Cluster (Resource Group name is the AAG Name). The corresponding Listener (*) of AAG is also register as resource in the Resource Group:

- Resource - AAG-Listener Name (VNN)
- Resource - AAG-Listener IP (VIP)

(*) By default one Listener per AAG is sufficient, but if need, it's possible to create additional Listener (only via PowerShell or WSFC Console)

The WSFC Cluster is responsible of Resources monitoring, in comparison with the SQL Mirroring solution there is not SQL Witness, the WSFC Quorum is used.

It is best practice to always have an odd number of quorum votes in a WSFC cluster.

AAG Health Monitoring

The WSFC Cluster monitors the health of the AlwaysOn Availability Group (physical and logical cluster resources). Few points about WSFC Monitoring:

- The active SQL Server instance periodically reports a set of component diagnostics to the WSFC resource group.
- Issues at the database level, such as a database becoming suspect due to the loss of a data file, deletion of a database, or corruption of a transaction log, do not cause an availability group to failover.

Note: For "Automatic Failover" a Failover Policy is defined and can be customized:

A flexible failover policy provides granular control over the conditions that cause automatic failover for an availability group. By changing the failure conditions that trigger an automatic failover and the frequency of health checks, you can increase or decrease the likelihood of an automatic failover to support your SLA for high availability.

For more information see TechNet article: Flexible Failover Policy for Automatic Failover of an Availability Group (SQL Server): <http://technet.microsoft.com/en-us/library/hh710061.aspx>

Failover

During a failover, a Secondary Replica becomes the Primary Replica.

There are three Failover Mode:

- **Planned manual failover** (without data loss) - DBA action
- **Automatic failover** (without data loss) - In case of failure

For these two modes, the Primary and Secondary Replica must be configured in "Synchronous Availability Mode" and the Secondary must be synchronized.

- **Forced manual failover** (with possible data loss)

This is a Disaster Recovery option; it can only be initiated manually.

It is the only form of failover that is possible when:

- the target secondary replica is not synchronized with the primary replica.
- the target secondary replica is in "Asynchronous availability mode"

Note: All manual failover actions must be done through the SQL Management Studio, PowerShell or Transact-SQL. No action should be done through the WSFC console.

Restrictions

Availability Replica:

- Availability replicas must be hosted by different nodes of one WSFC cluster
- One primary replica and up to four secondary replicas per AAG
- All of the replicas can run under asynchronous-commit mode, or up to three of them can run under synchronous-commit mode.

Maximum number of availability groups and availability databases per computer:

- No Limitation, depends of Server Performances
- Microsoft has extensively tested with 10 AGs and 100 DBs per physical machine

Database (To be eligible to be added to an availability group):

- Use the full recovery model
- Possess at least one full database backup (After setting a database to full recovery mode, a full backup is required to initiate the full-recovery log chain.)
- Be a read-write database. Read-only databases cannot be added to an availability group.
- System databases cannot belong to an availability group.

TDE Protected Databases:

If you use transparent data encryption (TDE), the certificate or asymmetric key for creating and decrypting other keys must be the same on every server instance that hosts an availability replica for the availability group. For more information, see [Move a TDE Protected Database to Another SQL Server](#).

You can find here a FAQ about AlwaysOn capabilities: **SQL Server 2012: Always On FAQs**
<http://blogs.msdn.com/b/srgolla/archive/2012/09/17/sql-server-2012-always-on-faqs.aspx>

TechNet Resources

AlwaysOn Availability Groups (SQL Server)

<http://technet.microsoft.com/en-us/library/aa427606-8422-4656-b205-c9e665ddc8c1>

- **Overview of AlwaysOn Availability Groups (SQL Server)**
<http://technet.microsoft.com/en-us/library/ff877884.aspx>
- **Failover and Failover Modes (AlwaysOn Availability Groups)**
<http://technet.microsoft.com/en-us/library/hh213151>
- **Availability Modes (AlwaysOn Availability Groups)**
<http://technet.microsoft.com/en-us/library/ff877931>
- **Prerequisites, Restrictions, and Recommendations for AlwaysOn Availability Groups (SQL Server)**
<http://technet.microsoft.com/en-us/library/ff878487.aspx#NetworkConnect>