

Are Availability Groups A Good Fit For Your Database?

Doug Purnell

About Me

- DBA for Elon University
- Co-leader for Winston-Salem BI User Group
- All things Nikon Photography
- Bring on the BBQ!



Goals

- Understand HA & DR
- Types of HR & DR Solutions
- WSFC/AG Checklists
- Demo
- Resources

CYA

- Management determines the following:
- Recovery Time Objective (RTO)
 - How long can the database be down?
- Recovery Point Objective (RPO)
 - How much data can we lose?
- Recovery Level Objective (RLO)
 - Granularity to recover data (Instance, Database, or Tables)

The answer to these questions helps determine your HA/DR requirements. Read-only file groups might be an option for RLO for historical data.

HA < > DR

High Availability (HA)

- Minimize the impact of downtime (keep the data available!)

A problem has been detected and windows has been shut down to prevent damage to your computer.

DRIVER_IRQL_NOT_LESS_OR_EQUAL

If this is the first time you've seen this stop error screen, restart your computer. If this screen appears again, follow these steps:

Check to make sure any new hardware or software is properly installed. If this is a new installation, ask your hardware or software manufacturer for any windows updates you might need.

If problems continue, disable or remove any newly installed hardware or software. Disable BIOS memory options such as caching or shadowing. If you need to use Safe Mode to remove or disable components, restart your computer, press F8 to select Advanced Startup Options, and then select Safe Mode.

Technical information:

*** STOP: 0x000000D1 (0x0000000C, 0x00000002, 0x00000000, 0xF86B5A89)

*** gv3.sys - Address F86B5A89 base at F86B5000, DateStamp 3dd991eb

Beginning dump of physical memory
Physical memory dump complete.
Contact your system administrator or technical support group for further assistance.

Disaster Recovery (DR)

- Re-establish availability after an outage



DR can be physical (power or storm), HW (SAN or RAID controller), or logical (OS blue screen)

Questions to Ask

- What is the business issue you're trying to solve?
- Do you have experience supporting the HA/DR solution?
- Do you have a test environment to test your solution?
- You are the architect!

Work with your system engineers and SAN administrators to understand your HW capabilities. SAN Replication and VM snapshots other options.

SQL Server HA & DR Options

- Backup and Recovery
- Replication
- Log Shipping
- Mirroring
- AlwaysOn Failover Cluster Instance (FCI)
- AlwaysOn Availability Groups (AG)

Which ones of these are HA and/or DR?

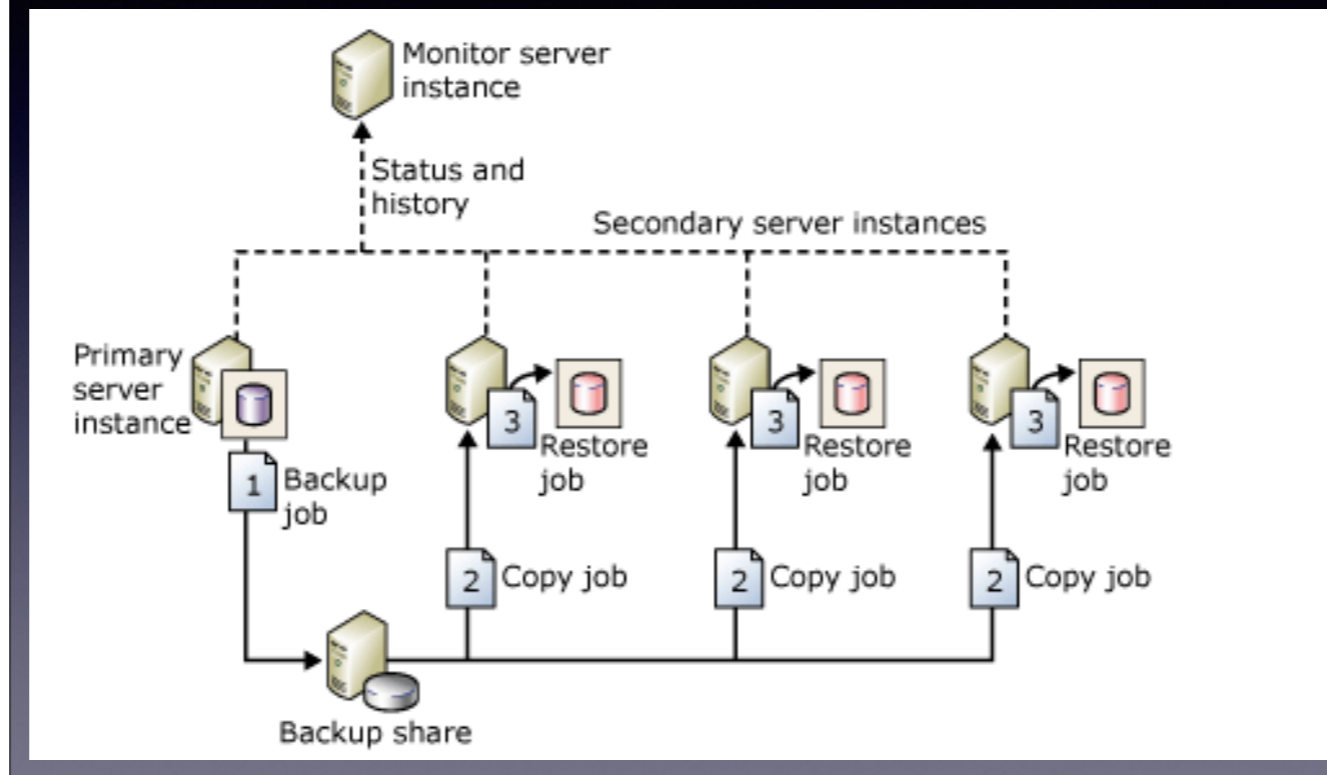
Backup and Recovery is old school. Jobs can be setup to restore database backups and transaction log on another server or offsite.

Replication can be setup to copy data from one server to another or offsite. Snapshot, transactional, and peer-to-peer can be used to duplicate data. No automatic failover. Can have as many subscribers as you like. May increase support to keep the publisher, distributor, and subscribers in check. Databases on the subscriber are usable.

Log Shipping

- Available under Standard Edition
- Supports multiple secondaries
- Secondaries can be brought online in case of failure
- Connection string will need to be changed
- No automatic failover
- Secondary databases can be in recovery or standby mode

Log Shipping



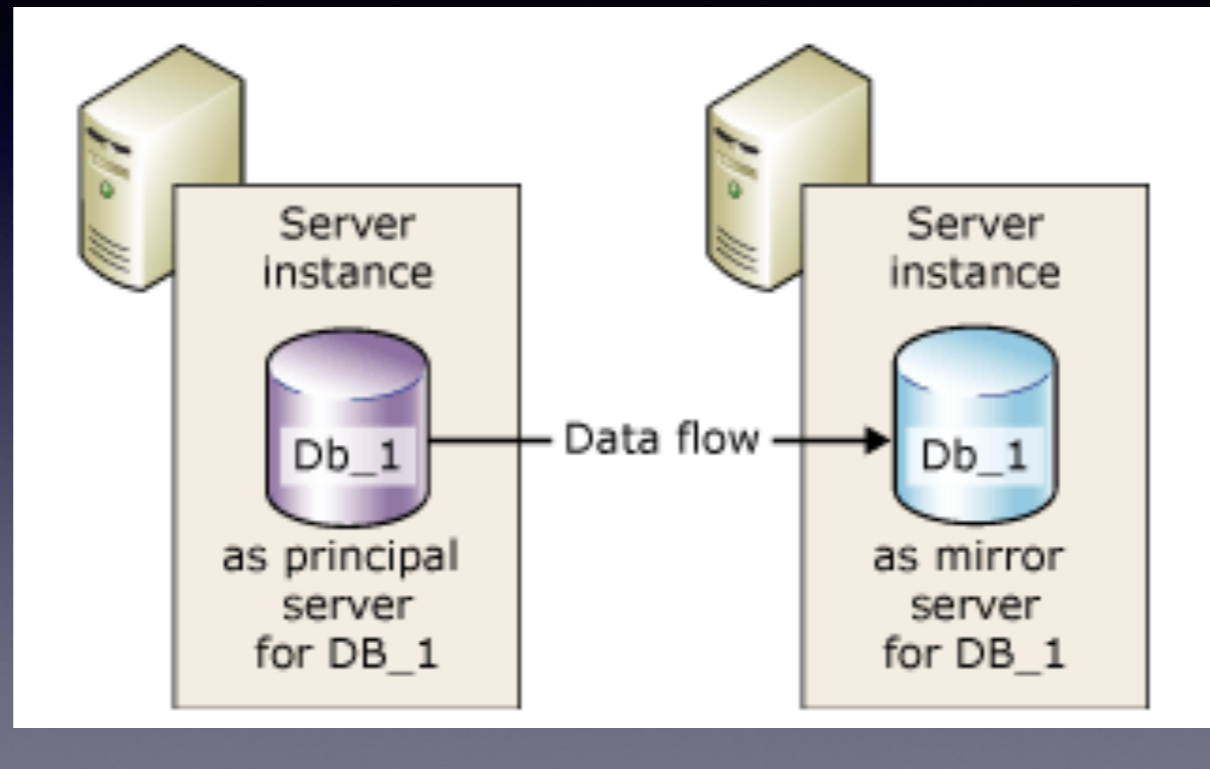
Helps automate Backup and Recovery.

Database Mirroring


- Database level protection, only one mirror per database
- Full recovery model
- Principal/Mirror need to be same version
- Mirror state is RECOVERY, but a snapshot can be use for readonly access (Enterprise Edition)
- Witness server can be used for automatic failover with connection string for the application Failover_Partner=<server name>
- Synchronous (Safety) or Asynchronous (Performance) operations.
- Automatic Page Repair is supported between the principal and mirror (two way)

sys.dm_db_mirroring_auto_page_repair

Database Mirroring



Database Mirroring

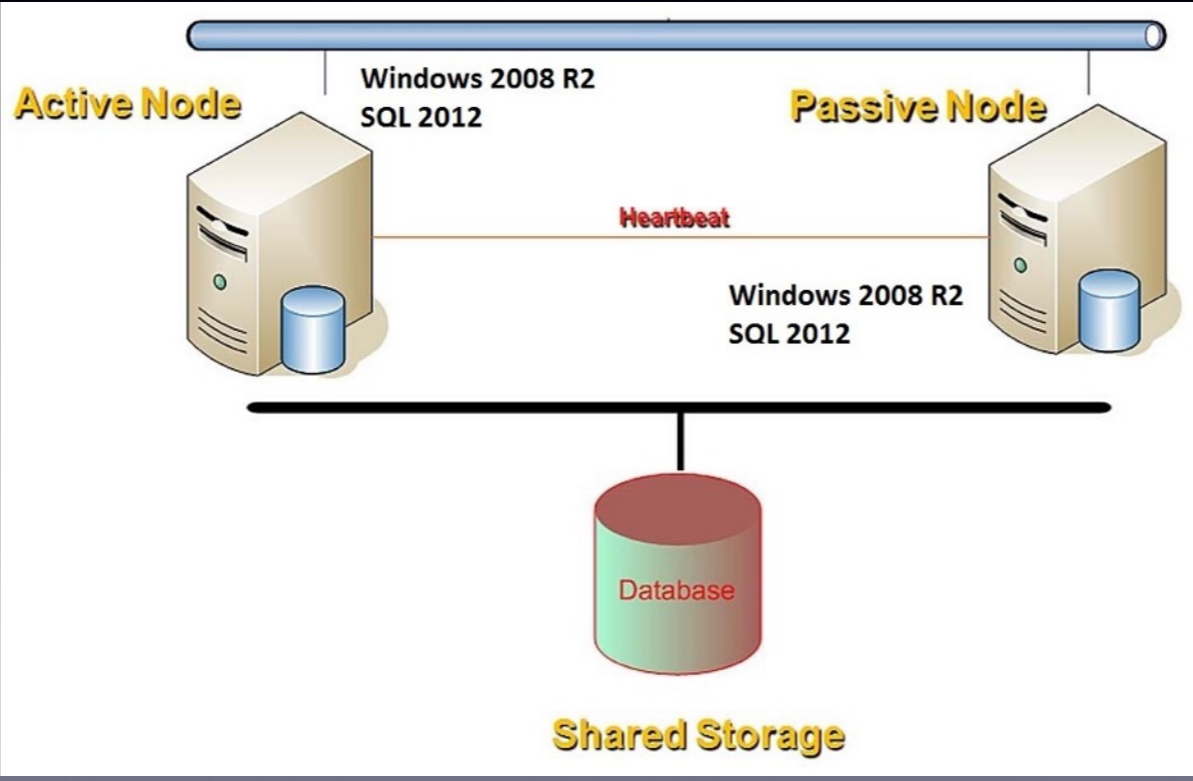


Witness server can be Express Edition.

AlwaysOn Failover Cluster Instance

- Provides redundancy at the server-instance level.
- Requires Windows Server Failover Clustering. Provides protection from hardware and software issues.
- A shared disk is required between both servers to host the database files.
- Application does not need to be reconfigured after failover.
- The shared storage has the potential of being the single point of failure.

AlwaysOn FCI



AG Overview

- SQL Server Enterprise Edition*
- Licensing requirements if you're using the secondary replica.
- Active Directory Domain
- Windows Server Failover Cluster (WSFC)
- WSFC can span subnets but not different domains.
- Supports multiple readable secondaries (great for off loading reports).
- Applications connect to each AG via virtual listener name.

SQL Server 2016 CTP2 Standard Edition support 2 nodes (Basic mode)

AG Overview

- Synchronous Synchronization (No data loss)
 - Allows for automatic failures.
 - Transactions are applied to all replicas before control is passed back to application.
- Asynchronous Synchronization (Possible data loss)
 - Transactions are committed on the primary replica and then control is passed back to application. Secondary replicas receive changes at a later time.
- Up to 10 databases per AG and up to 10 AGs

SQL Server 2012 supports primary plus 4 secondary replicas, 3 of them synchronous (Primary plus two)

SQL Server 2014 supports primary plus 8 secondary replicas, 3 of them synchronous (Primary plus two)

SQL Server 2016 supports primary plus 8 secondary replicas, 3 of them synchronous (Primary plus two)

<https://msdn.microsoft.com/en-us/library/ff878487.aspx>

AG Overview

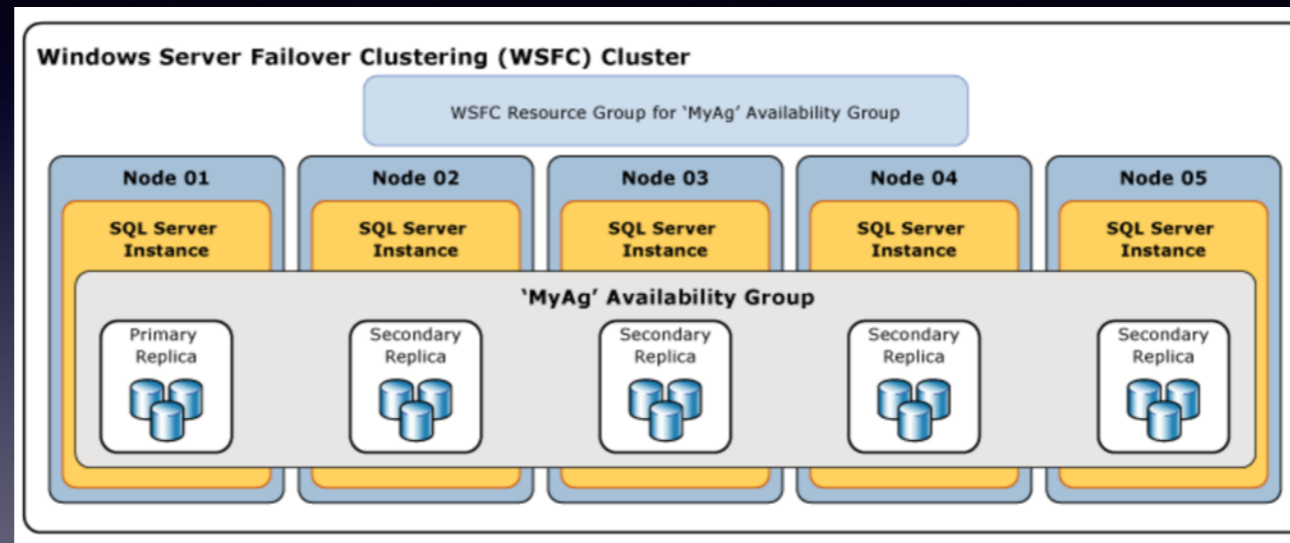
- One or more databases can reside within an availability group.
- The 'listener' is the virtual name for the Availability Group that your application connects to.
- The listener name will follow the primary replica as a failover event occurs.
- Automatic Page Repair is supported between the primary and secondary replica (two way).
- Session-timeout period determines how long connection with another availability replica can remain inactive before the connection is closed.

<https://msdn.microsoft.com/en-us/library/hh213417.aspx>

<https://msdn.microsoft.com/en-us/library/bb677167.aspx>

sys.dm_hadr_auto_page_repair

AG Overview



<https://msdn.microsoft.com/en-us/library/ff877884.aspx>

Pre WSFC/AG Checklist

- All WSFC nodes need to be a part of the same domain.
- Validate Configuration and pay attention to errors and warnings.
- AD account needs rights to create computer objects under OU.
 - You can grant the necessary rights or log in with sufficient rights.
- AlwaysOn Availability Groups needs to be enabled for each SQL Server instance (SQL Server Configuration Manager) after the WSFC is established.

Pre WSFC/AG Checklist

- If you are using TDE, you have to implement AGs with T-SQL.
- Databases need to be in FULL recovery model with a recent backup performed.
- You will need IP and Cluster Name for WSFC.
- You will need a name for the AG.
- You will need IP and Listener name for each AG.

AG Wizard does not support TDE, <https://msdn.microsoft.com/en-us/library/hh510178.aspx>

Come up with naming standards for your WSFC/AG

Post AG Checklist

- Backups can be done on primary or secondary.
 - No DIFF backups on the secondary
 - FULL backups on secondary require WITH COPY_ONLY
- sys.fn_hadr_is_preferred_backup_replica function
- DBCC CheckDB can be done on primary or secondary.
- Connection Strings
 - Primary -> Server=Listener;Database=DB_1
 - Secondary -> Server=Listener;Database=DB_1;ApplicationIntent=ReadOnly
 - Requires manual routing (T-SQL or PoSh)

DBCC CHECKDB is recommended on all secondaries that could become primary or where backups are performed because the storage is different across all the nodes.

Post AG Checklist

- Ola Hallengren's scripts are AG aware.
- Logins & Jobs are separate for each node.
- Index maintenance has to be done on primary replica.

Jonathan Kehayias has written a SSMS 2012 add-on to help with logins and jobs.

READ_ONLY Routing

```
ALTER AVAILABILITY GROUP [SQLNikonAG]
  MODIFY REPLICA ON
  N'SQLNikonAG01' WITH
  (SECONDARY_ROLE (ALLOW_CONNECTIONS = READ_ONLY));
ALTER AVAILABILITY GROUP [SQLNikonAG]
  MODIFY REPLICA ON
  N'SQLNikonAG01' WITH
  (SECONDARY_ROLE (READ_ONLY_ROUTING_URL = N'TCP://SQLNikonAG01.sqlnikon.local:1433'));

ALTER AVAILABILITY GROUP [SQLNikonAG]
  MODIFY REPLICA ON
  N'SQLNikonAG02' WITH
  (SECONDARY_ROLE (ALLOW_CONNECTIONS = READ_ONLY));
ALTER AVAILABILITY GROUP [SQLNikonAG]
  MODIFY REPLICA ON
  N'SQLNikonAG02' WITH
  (SECONDARY_ROLE (READ_ONLY_ROUTING_URL = N'TCP://SQLNikonAG02.sqlnikon.local:1433'));

ALTER AVAILABILITY GROUP [SQLNikonAG]
  MODIFY REPLICA ON
  N'SQLNikonAG01' WITH
  (PRIMARY_ROLE (READ_ONLY_ROUTING_LIST=('SQLNikonAG02', 'SQLNikonAG01')));

ALTER AVAILABILITY GROUP [SQLNikonAG]
  MODIFY REPLICA ON
  N'SQLNikonAG02' WITH
  (PRIMARY_ROLE (READ_ONLY_ROUTING_LIST=('SQLNikonAG01', 'SQLNikonAG02')));
GO
```

<https://msdn.microsoft.com/en-us/library/hh710054.aspx>

Quorum

- Best practice is to have odd number of votes.
- Each node has a single vote by default.
- Any host that could host a primary replica should have a quorum vote.
- Secondary disaster recovery sites may not need a quorum vote.
- You can use File Share Witness to break a tie.
- More than one-half of the voting nodes/witnesses need to report back for a cluster to be healthy.
- Windows Server 2012 introduced dynamic quorum. Votes are removed if nodes leave the cluster membership. R2 introduced dynamic witness

<https://technet.microsoft.com/en-us/library/jj612870.aspx>

(Get-Cluster).WitnessDynamicWeight

<http://windowsitpro.com/windows-server-2012-r2/dynamic-witness-windows-server-2012-r2-failover-clustering>

AG Demo

Use SQLskills app to do failover events.

Show dashboard

Show AG Properties (Setup and Backup preferences)

Automatic and Manual Failovers

`ALTER DATABASE [DBA] SET HADR RESUME`

Quorum/Dynamic Witness

AG Transaction Test

- The goal of this test was to understand how AG would handle open transactions during a failover event
- .NET Entity Framework application with retry logic
- SP was built to open transaction, insert row, wait one minute, insert a second row, and commit transaction

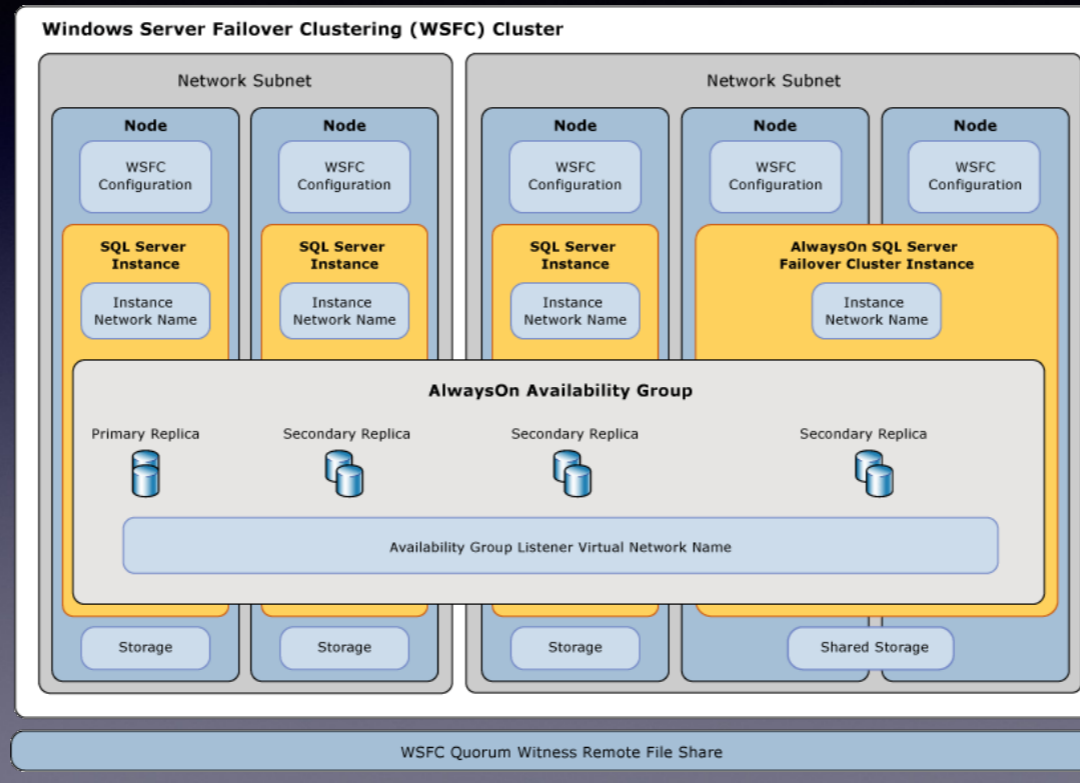
AG Transaction Results

1. .NET app started at 9:23:45 (open trans)
2. AG Failover at 9:24:11
3. Rollback trans from 9:23:45
4. Database down from 9:24:11 through 9:24:20
5. .NET retry at 9:24:11 (unable to open listener)
6. .NET retry at 9:24:15 (unable to open listener)
7. .NET retry at 9:24:22 (listener available)
8. .NET called SP
9. Trans committed at 9:25:22 (2 records inserted)

AG Failover Test

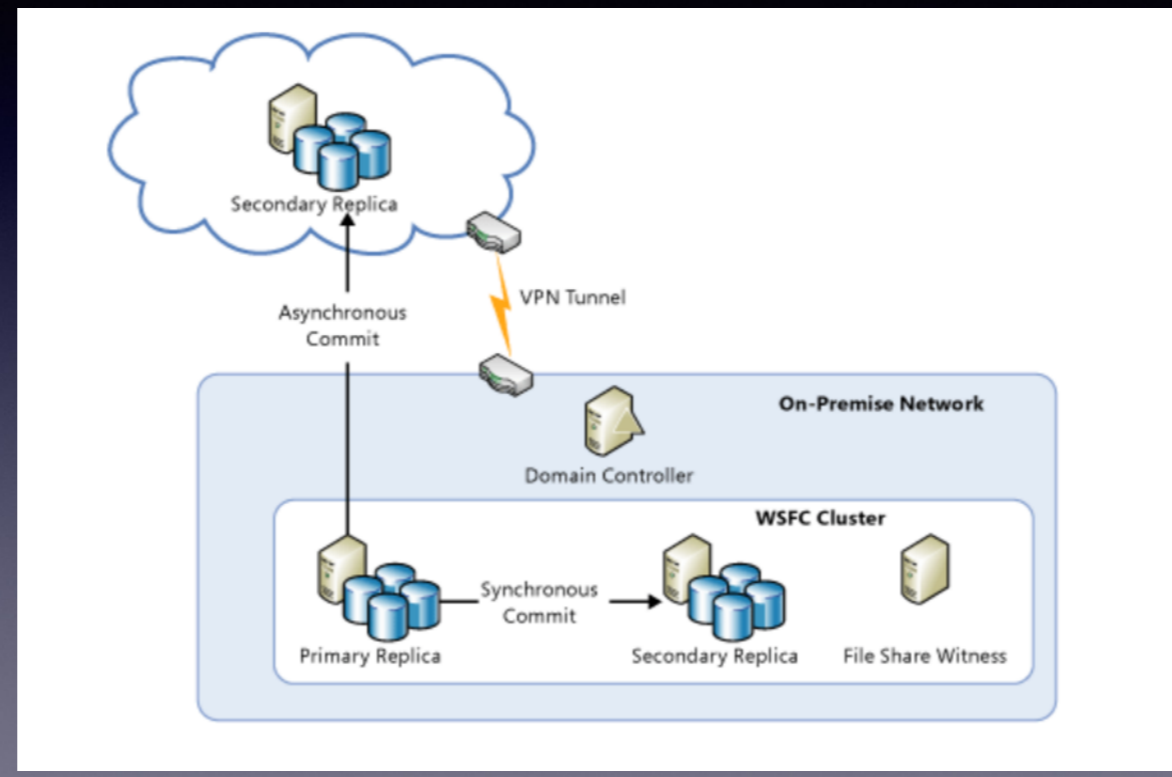
- The goal of this test was to understand how a .NET Entity Framework application would handle an AG failover in between database requests
- Application was written to query SQL Server before and after the AG Failover
- Entity Framework does not keep open connections to the database so failover had no effect on app behavior

Combine FCI & AG



Some of the HA/DR options can be joined together.

Azure Hybrid IT



Here is a possible solution with Azure.

Resources

- [Allan Hirt's Mission Critical SQL Server Book](#)
- [SQL Server Customer Advisory Team Blog](#)
- [SQL Server Blog](#)
- [SQL Release Services Blog](#)
- [SQL Server Builds Blog](#)
- [SQL Server AlwaysOn Support Team](#)

Resources

- [Brent Ozar AG Quiz](#)
- [Kendra Little's Hard Truth About Patching SQL Server AGs](#)
- [HA DR Virtual User Group](#)
- [AG Ports](#)
- [Prerequisites, Restrictions, and Recommendations for AlwaysOn Availability Groups](#)
- [VMware GARP Issue](#)
- [How to setup AGs under SQL Server 2016 Standard](#)

Questions?

Doug Purnell

@SQLNikon

sqlnikon@gmail.com

sqlnikon.wordpress.com

