

Disaster Recovery Where to Begin



Who Am I?

Tracy Boggiano

- Worked on SQL Server for over 18 years
- Current MCSE Data Platform
- Other hobbies include running and volunteering with foster children (www.casaforchildren.org)
- Current DBA at ChannelAdvisor in Morrisville, NC
- Contact Info
 - Email: Tracy@TracyBoggiano.com
 - Twitter: @TracyBoggiano
 - Blog: www.tracyboggiano.com



Supported Environments

- **119 SQL Servers - 3200+ databases**
- **1.6M+ SQL Server transactions/sec**
- 908 cores / 34TB RAM 100% virtualized
- 100TB storage allocated for SQL Server data - Pure Storage All Flash Arrays
- 80 servers – with at most 1 to 2 databases per server
- 8 GBS per machines/ 2 CPUS per machine
- Max size of DB 200 GBs



Agenda

- Why is Disaster Recovery Important?
- Technologies
 - Replication
 - Log Shipping
 - Failover Clustering
 - Database Mirroring
 - High Availability Groups
- Summary

Why is Disaster Recovery Important?

- The business wants to data to be available after a disaster to continue doing business
- Bad things can happen
- We want to keep our jobs

Defining RTO and RPO

- Recovery Time Objective (RTO)
 - Maximum allowable downtime
 - Must consider how is defined at your company
 - i.e. 24x7, or 9 to 5
 - Zero downtime is not possible
- Recovery Point Objective (RPO)
 - Maximum allowed data loss
 - May be different for different systems
 - Zero data loss is possible, but does the business really need it

Backup Strategy

- Backups only need to plan for time for restore
- Recommend weekly full and nightly differential backups
- Transaction logs depend how much data you can lose
- Take backups to network share and replicate to second data center
- Ola Hallengren's scripts are great to use
 - Add component to generate restore scripts to a file

DBCC CHECKDB

- Run it
- Please run it
- And run it often
- It should at least be able to keep up with your backups
- Recommendation: offload by restoring backups to another system and running CHECKDB

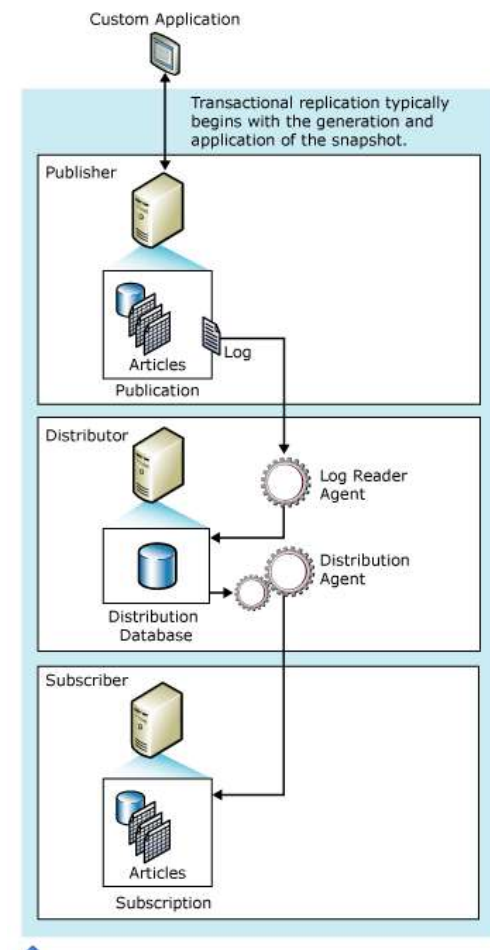
Replication Overview

- Make copies of tables rather than the whole database
- Three types of replication but really only one used for disaster recovery
 - Transactional
 - Mostly used for distributed systems and high availability

Replication Overview

- Publisher pushes out changes from the transaction log
- Distributor distributes the data to the subscribers
- The subscriber receives the changed data from the distributor

<https://msdn.microsoft.com/en-us/library/ms151176.aspx>



Replication Advantages

- Only makes copies of data needed
- Distributes the workload
- Distributes copies of the data to different data centers
- Multiple copies of the data

Replication Disadvantages

- Only makes copies of subset of tables
 - You might want the whole database, if so this is not the solution for you
- Can be fragile
- Manual change of connection after failover
- Can't make schema changes without recreating replication

Replication Summary


- Would only use transactional replication
- Would only use if a small number of tables is needed
- My less favorite disaster recovery solution

Log Shipping Overview

- You have two separate servers
- Setup by taking a full and transaction log backup and restoring to second server WITH NORECOVERY
- Use wizard to setup the rest which sets up each transaction log to be copied to second server and applied the database
- Make sure to use DNS Name to connect instead of server name

Log Shipping Overview

- Primary server runs backup job
- Copy job is run on each secondary
- Restore job is run on each secondary



<https://msdn.microsoft.com/en-us/library/ms187103.aspx>

Log Shipping Advantages

- Two separate copies of the database
- Available in Standard Edition
- Simple setup
- Separate hardware including storage
- Secondary servers located in second data center
- Can read data on secondary server
- Supports multiple targets

Log Shipping Disadvantages

- Data is only current to the most recently applied transaction log
- Some data loss
- Must script out logins and SQL Agent jobs
- Manual failover
 - Always has downtime
- DNS name change required

Log Shipping Monitoring

- Recommend you specify a third server as the monitor server
 - Must be specified when setup initially
- Monitors
 - Last backup
 - Copies to secondary server
 - Last restore
- Be careful to space jobs and alerts apart

Log Shipping Summary


- Simple setup
- Good for systems where some data loss is allowable
- Good for systems where some downtime is allowable
- Has separate hardware components
- Use third server as monitor server

Failover Clustering Overview

- Built on top of Windows Clustering that provides automatic failover
- Standard Edition supports two nodes, Enterprise supports OS maximum
- Supported as VMs in 2005+ and Windows Server 2008+

AlwaysON Failover Clustering Overview

- This in itself does not give you disaster recovery




AlwaysOn Failover Geo-Clustering

- Eliminates storage as single point of failure
- Requires specific hardware on Geographic Cluster Hardware Compatibility list
- Requires stretched VLAN
 - Public and private networks must be non-routed
 - Latency must be under 500ms
- Requires SAN replication technology
- Have to install from command line and skip Cluster_VerifyForErrors option


AlwaysON Failover Clustering Overview

- Typical Geo-Clustering configuration




AlwaysON Failover Clustering Overview

- SAN Failure




AlwaysON Failover Clustering Overview

- Site Failure



AlwaysON Failover Clustering Overview

- Everything comes back online Instance A points to SAN at Site 2
- And SAN replication kicks back in



Failover Clustering Quorum

- Determines number of failures an instance can sustain
- Exists to prevent two nodes from trying to be up at the same time causing a “split-brain” problem
- Determine by votes
 - Each cluster gets 1 vote
 - A disk or file share witness can be configured for 1 vote
 - Configure for odd number of votes

Failover Clustering Advantages

- Automatic failover
- Can have more than one failover server
 - 2 nodes on Standard Edition
- Instance level failover
- Requires Geo-Clustering to avoid single point of failure

Failover Clustering Disadvantages

- Single copy of databases on shared storage
 - Single point of failure
- Must be using SAN replication for disaster recovery
- More complex setup on the Windows and SQL
- Identical hardware required

Failover Clustering Summary

- Automatic Failover
- Single copy of database unless using geo-clustering
- Needs identical hardware
- Complex setup
- Quorum is important to get setup correctly

Database Mirroring Overview


- Continuously sends log records from primary sever to secondary server
- Use same method as log shipping to initialize
 - Backup full and transaction log and restore with NORECOVERY
- Three configurations
 - Asynchronous (Enterprise only)
 - Synchronous not witness
 - Synchronous with a witness

Database Mirroring Overview

- A database snapshot can be created on the mirror to offload read operations
- Each mirroring session is for only on database
- Automatic failover available through connection string parameter
Failover_Partner=<server_name>
- Must script out logins and jobs for second instance

Database Mirroring Overview


- Synchronous with witness



<https://msdn.microsoft.com/en-us/library/ms189852.aspx>

Database Mirroring Overview

- Asynchronous (Enterprise only)
- Synchronous no witness



<https://msdn.microsoft.com/en-us/library/ms189852.aspx>

Database Mirroring Advantages

- Provides minimal to zero data loss
- Provides two copies of the database
- No hardware dependencies
- Transparent client redirect in connection string
- Automatic Page Repair (2008+)
 - Keep eye on msdb.dbo.suspect_pages
- Log Compression

Database Mirroring Disadvantages

- Performance may be impacted in synchronous mode
- Data loss possible in asynchronous mode
- Failover at database level
 - Presents problems for multiple databases needed to failover together
- Marked for deprecation in SQL 2012

Database Mirroring Monitoring

- Unsent log
 - Measured in KBs on primary server
- Unrestored log
 - Measured in KBs on mirror server
 - Effects failover time (<https://msdn.microsoft.com/en-us/library/ms187465.aspx>)
- Oldest unsent transaction
 - Number of minutes on primary server
- Mirror commit overhead
 - Measured in milliseconds
 - Valid only for automatic failover

Database Mirroring Summary

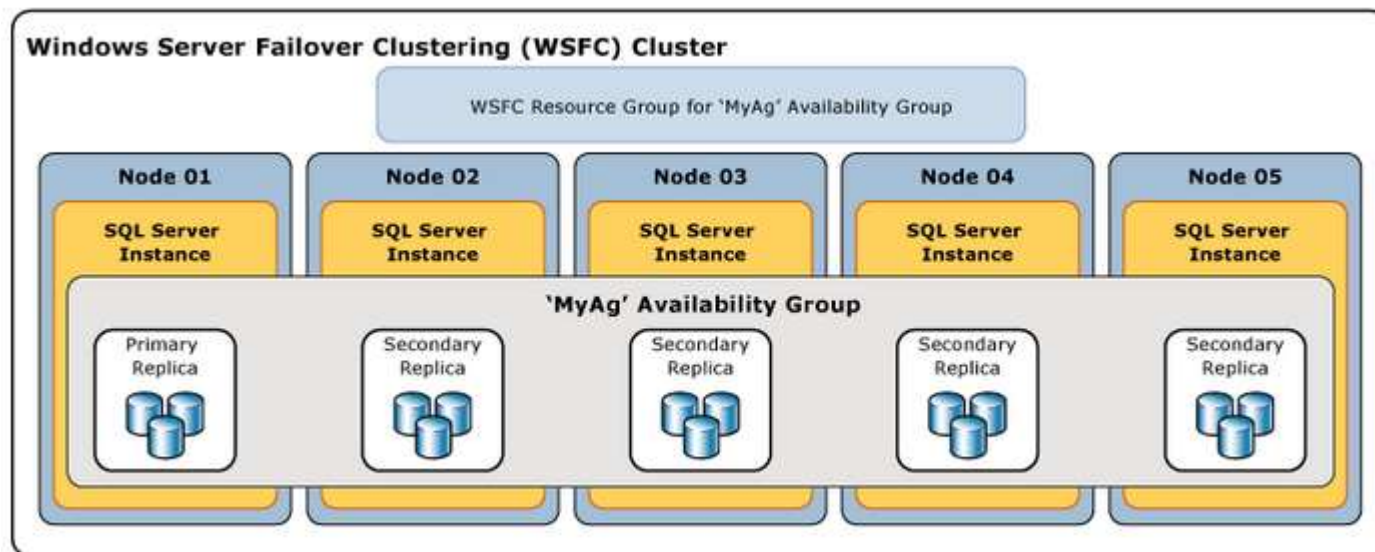
- Simple setup
- Little to no data loss
- Protection at database level
- Possible automatic failover

AlwaysOn High Availability Groups Overview

- Database Mirroring on steroids
- Built in on top of Windows Clustering
 - Can span different subnets but not domains
- Asynchronous and synchronous modes
- Requires Enterprise Edition
- 2012 up to 100 databases in a group, 2014 supports a thousand
- Can be used to offload reads and backups

AlwaysOn High Availability Groups Overview

- Each node could reside in a different data center



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

AlwaysOn High Availability Groups Overview

- Readable Secondary Options
 - `ALLOW_CONNECTIONS = "ALL"`
 - All connects accepted
 - Client checks to see if is readable secondary
 - `ALLOW_CONNECTIONS = "READ_ONLY"`
 - Only connections with `ApplicationIntent=ReadOnly` are accepted
 - Any write operation will fail
- During failover read only connections are rerouted if the current node becomes primary



AwaysOn High Availability Groups READ_ONLY Routing

- Must be using Listener to connect

```
ALTER AVAILABILITY GROUP [AG1]
  MODIFY REPLICA ON
  N'COMPUTER01' WITH
  (SECONDARY_ROLE (ALLOW_CONNECTIONS = READ_ONLY));
ALTER AVAILABILITY GROUP [AG1]
  MODIFY REPLICA ON
  N'COMPUTER01' WITH
  (SECONDARY_ROLE (READ_ONLY_ROUTING_URL = N'TCP://COMPUTER01.contoso.com:1433'));

ALTER AVAILABILITY GROUP [AG1]
  MODIFY REPLICA ON
  N'COMPUTER02' WITH
  (SECONDARY_ROLE (ALLOW_CONNECTIONS = READ_ONLY));
ALTER AVAILABILITY GROUP [AG1]
  MODIFY REPLICA ON
  N'COMPUTER02' WITH
  (SECONDARY_ROLE (READ_ONLY_ROUTING_URL = N'TCP://COMPUTER02.contoso.com:1433'));

ALTER AVAILABILITY GROUP [AG1]
  MODIFY REPLICA ON
  N'COMPUTER01' WITH
  (PRIMARY_ROLE (READ_ONLY_ROUTING_LIST=('COMPUTER02','COMPUTER01')));

ALTER AVAILABILITY GROUP [AG1]
  MODIFY REPLICA ON
  N'COMPUTER02' WITH
  (PRIMARY_ROLE (READ_ONLY_ROUTING_LIST=('COMPUTER01','COMPUTER02')));
GO
|
```

AlwaysOn High Availability Groups Quorum

- Best practice to have odd number of votes
- Each node has a vote by default
- Any potential primary node should have a vote
- DR sites may not need a vote
- There is a File Share Witness to help break ties
 - Note possibility of File Share Witness going offline breaking cluster

AlwaysOn High Availability Groups Advantages

- Can have automatic failover
- Up to 8 replica copies in 2014
- Replicas can be read
- Offload backups to replica node
- Databases failover as a group
- Partial Database Containment
- Read-Only Routing
- Automatic Page Repair
 - Keep eye on msdb.dbo.suspect_pages

AlwaysOn High Availability Groups

Disadvantages

- Performance may be impacted in synchronous mode
- Data loss possible in asynchronous mode
- Involves Windows Clustering which adds complexity
- Lots of bugs and hotfixes released
- Requires Enterprise Edition

AlwaysOn High Availability Groups Monitoring

- Setup alerts for the following error numbers
 - Detect failover setup - 1480
 - Detect AG data movement suspended - 35264
 - Detect AG data movement resumed - 35265
- Log Send Queue Size (KB)
- Log Send Rate (KB/sec)
- Redo Queue Size (KB)
- Redo Rate (KB/sec)
- Estimated Data Loss (time) for asynchronous

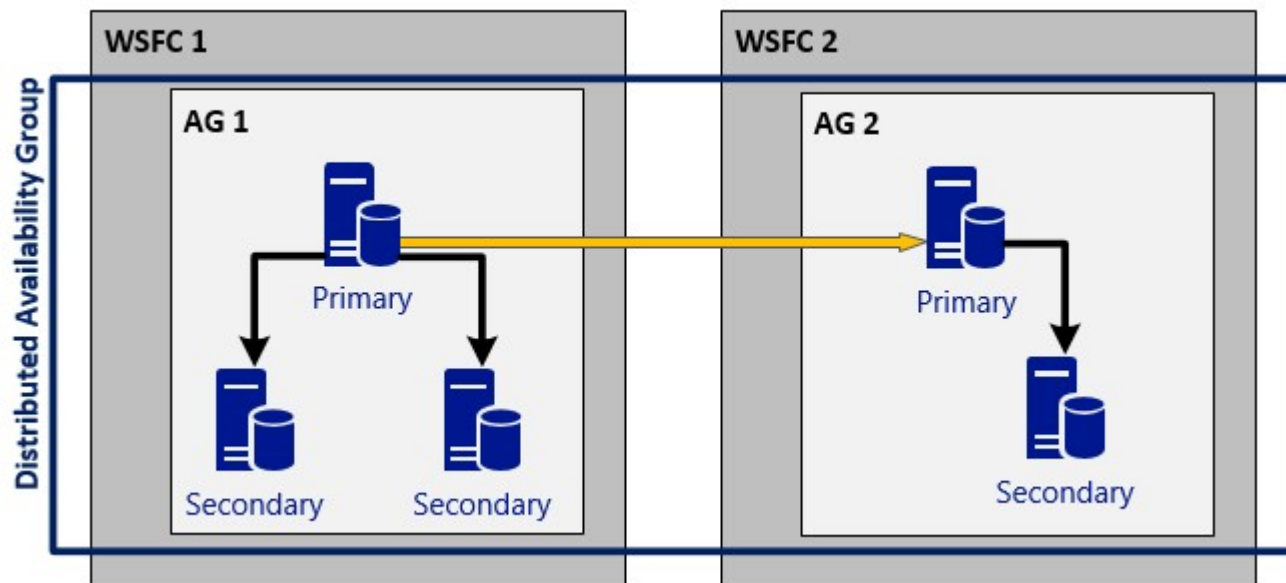
What's New in 2016 for Availability Groups?

- Standard Edition will come with AGs support with one db per group synchronous or asynchronous, not readable (HA/DR only)
- 3 sync replicas, up from 2 in SQL 2014
- Listener will be able to do round-robin load balancing of read-only requests on multiple secondaries
- Microsoft DTC support
- SQL Server AlwaysOn to Azure Virtual Machine



What's New in 2016 for Availability Groups?

- Distributed AGs (SQL 2016 RC0)
 - 2 Different AGs on two different clusters



Distributed AGs

- Helps with avoiding using DR side for quorum votes
- Seeding
 - Automatic
 - Manual
- Second AG is read-only
 - Replicated to primary on second group, then it replicates to it's secondary servers
- Could be different OS (useful for OS upgrades)
- DR only, no automatic failover

AlwaysOn High Availability Groups Summary

- Database Mirroring on steroids
- Asynchronous and synchronous modes
- Servers must be in the same domain
- Read Only Rerouting
- WSFC adds more complexity
- Requires Enterprise Edition
 - But not in 2016

SQL Server vNext as of CTP 1.4


- Will support traditional log shipping, replication, clustering, and AGs
- Cluster-less Availability Groups support added
- Clustering using Pacemaker in Linux
- Availability Groups can now work across Windows-Linux to enable cross-OS migrations and testing

Combining Technologies for DR and HA

- Combining mirroring with log shipping
 - Site 1 has mirroring with a witness providing HA
 - Site 2 has log shipping providing DR
- Combine clustering with log shipping
 - Site 1 has clustering providing HA
 - Site 2 receives providing DR

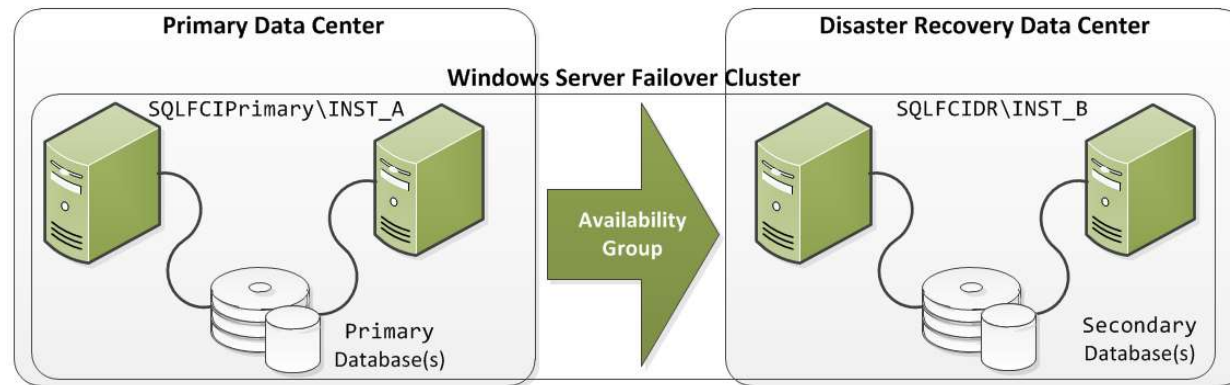
Combining Technologies for DR and HA

- Combine clustering with database mirroring
 - Site 1 has clustering providing HA
 - Asynchronous mirroring to site 2 providing DR



Combining Technologies for DR and HA

- AlwaysOn FCI and AlwaysOn High Availability Groups



Summary

- Only use replication if you meet the use case of only needing a few tables in multiple locations and distributed workload
- For small shops, go less complex Log Shipping or Database Mirroring
- For larger shops, go with AGs if you have the expertise
- AGs is where Microsoft is investing so it's time for us to invest in it as well
 - <http://www.brentozar.com/archive/2015/10/training-plan-for-a-sysadmin-to-manage-availability-groups/>

Questions

- Contact Info
 - Email: Tracy@TracyBoggiano.com
 - Twitter: @TracyBoggiano
 - Blog: www.tracyboggiano.com

Resources

- Transactional Replication
 - <https://msdn.microsoft.com/en-us/library/ms151176.aspx>
- Log Shipping
 - <https://msdn.microsoft.com/en-us/library/ms187103.aspx>
- Whitepaper: High Availability with SQL Server 2008
 - <https://technet.microsoft.com/en-us/library/ee523927%28v=sql.100%29.aspx>



Resources

- AlwaysOn Failover Cluster Instances
 - <https://msdn.microsoft.com/en-us/library/ms189134.aspx>
- Database Mirroring
 - <https://msdn.microsoft.com/en-us/library/ms189852.aspx>
- AlwaysOn Availability Groups
 - <https://msdn.microsoft.com/en-us/library/ff877884.aspx>

Resources

- AlwaysOn Architecture Guide: Building a High Availability and Disaster Recovery Solution by Using Failover Cluster Instances and Availability Groups
 - <https://msdn.microsoft.com/library/jj215886.aspx>
- Microsoft SQL Server AlwaysOn Solutions Guide for High Availability and Disaster Recovery
 - <https://msdn.microsoft.com/library/hh781257.aspx>



Resources

- Log Shipping Monitoring -
<https://msdn.microsoft.com/en-us/library/ms190224%28v=sql.130%29.aspx>
- Database Mirroring Monitoring -
<https://msdn.microsoft.com/en-us/library/ms190030.aspx>
- AlwaysOn High Availability Monitoring
 - <https://msdn.microsoft.com/en-us/library/ff878305.aspx>
 - <http://www.concurrency.com/blog/monitor-availability-groups/>

Resources

- Distributed AGs
 - Allan Hirt - <http://sqlha.com/2016/03/10/distributed-availability-groups-ag-seeding-and-ag-changes-in-rc0/>
 - BOL - <https://msdn.microsoft.com/en-US/library/mt651673.aspx>

Thank You Sponsors

Platinum 

Gold



Microsoft



redgate



TEKsystems
Our people make IT possible.

SentryOne[™]

Silver

I D E R A

Bronze



COZYROC[™]
Go to the next level



revgen partners
DRIVING PROFITABILITY[®]



Swag

