Using SQL Server AlwaysOn (HADR) Denny Cherry mrdenny@dcac.co twitter.com/mrdenny

## About Me

Denny Cherry & Associates Consulting

Microsoft

CERTIFIED

Master

- 6 books
- Dozens of articles
- Microsoft MVP
- Microsoft Certified Master
- VMware vExpert
- EMC Elect
- Microsoft Certified Trainer

Microsoft

Most Valuable

rofessiona



VEXPERT

vmware<sup>.</sup>

SQL Server 2008

The Real MCTS S Server SECURING

FRTIFIF

Trainer

SERVER



### What is "Always On"?

- AKA "HADR"
- Combines Best of Mirroring and Clustering
- Can allow for backups to be read



### What Parts of Clustering?

- Uses the clustering APIs for failover
- No shared storage, users still connect to normal instance
- Shared IP Address for Availability Group Listener



### What Parts of Mirroring?

- Uses Mirroring Concepts for the data transport
- Uses Mirroring TCP Endpoint



### New Terms to Know

- Availability Groups
- Availability Replicas
- Availability Databases
- Availability Group Listener



### Setup Requirements

- All Instances must be installed on Windows 2008 R2 Enterprise (or higher)
- All servers running as part of HADR much be in a single Windows cluster
- SQL is **not** clustered
- Databases must be in full recovery

### Setup Notes

- Multiple Databases in a single Availability Group
- IO profiles on primary and secondary are different
  - Primary only writes on checkpoint
  - Secondary writes as changes are received
  - (The same as Legacy DB Mirroring)
- All compatibility modes are supported
- No Shared Storage
- Matching hardware is not required

### AlwaysOn – A flexible solution

# AlwaysOn provides the flexibility of different HA configurations



Direct attached storage local, regional and geo target



Shared Storage, regional and geo secondaries





Asynchcronous
Data Movement

### Data Transfer

- Data is transferred via Data Mirroring Endpoint
- Endpoints created via the Wizard automatically if not created



## **Configuration Options...**

- Management Studio
  - Wizard Driven Setup
- PowerShell
  - CmdLets on the next slide
- T/SQL
  - CREATE AVAILABILITY GROUP Setup Primary
  - ALTER AVAILABILITY GROUP Join Secondary



### **PowerShell CmdLets**

- Setup
  - New-SqlAvailabilityGroup
  - New-SqlAvailabilityReplica
  - Add-SqlAvailabilityGroupDatabase
  - Join-SqlAvailabilityGroup
- Breakdown
  - Remove-SqlAvailabilityGroup
  - Remove-SqlAvailabilityGroupDatabase
- Management
  - Suspend-SqlAvailabilityGroup
  - Switch-SqlAvailabilityGroup



## Sample PowerShell Script to

### **Configure HADR**

- # Setup names and TCP addresses of the primary and secondary servers.
- \$primaryServerName = "[PrimaryServerName]";
- \$secondaryServerName = "[SecondaryServerName]";
- \$primaryEndpointUrl = "TCP://" + \$primaryServerName + ".web\_domain:5022";
- \$secondaryEndpointUrl = "TCP://" + \$secondaryServerName + ".web\_domain:5022";
- # Next create the Availability Replicas as templates.
- \$replica = new-sqlavailabilityreplica -Name \$primaryServerName -EndpointUrl \$primaryEndpointUrl -AsTemplate
- \$replica2 = new-sqlavailabilityreplica -Name \$secondaryServerName -EndpointUrl . ssecondaryEndpointUrl -AsTemplate
- # Finally create the Availability Group.
- Sorry About the Eye Chart \$ag = new-sqlavailabilitygroup -Name \$availabilityGroupName -AvailabilityReplica (\$replica,) \$replica2) -Database \$databases
- #Tells a secondary Server to join the Availability Group
- Join-SqlAvailabilityGroup -Path . -Name "My Availability Group"



### Failing over AlwaysOn

- Done via SSMS.
- Switch-SqlAvailabilityGroup PowerShell CmdLet
- ALTER AVAILABILITY GROUP in T-SQL



#### What can we use AlwaysOn for?

- Getting a backup of the database on a remote machine
- Taking a snapshot of a database
- Offloading CheckDB from production server
- Offloading reporting from production server
- Redundancy Dugh



## Big Selling Points - 2012

- Support for up to 4 secondary servers
  - Up to 2 Synchronous
  - Up to 4 Asynchronous
- Backups on the secondary servers
- The ability to fail over without data loss
- Automatic Failover
- Read Only routing to read only servers with just a connection string and driver change on the client

## Big Selling Points - 2014

- Support for up to 8 secondary servers
  - Up to 2 Synchronous
  - Up to 8 Asynchronous
- Backups on the secondary servers
- The ability to fail over without data loss
- Automatic Failover
- Read Only routing to read only servers with just a connection string and driver change on the client

### 2014 Changes

- More Secondaries
- Faster Failover
- Secondary Replicas remain online during failover
- Supports Hekaton
- Supports Clustered ColumnStore

## Licensing

- Enterprise Only Feature
- If Read Only nodes are used by users then licensing is required



#### Lets see how we set this thing up



## **Denny Cherry**

mrdenny@dcac.co http://itke.techtarget.com/sql-server http://twitter.com/mrdenny