

Extending Your Availability Group for Disaster Recovery

Ed Leighton-Dick
Kingfisher Technologies



About Me

- Email: ed@leightondick.com
- Blog: edleightondick.com
- Twitter: [@eLeightondick](https://twitter.com/eLeightondick)



Agenda

- Availability Groups: A Brief Intro
- Multi-Subnet Availability Groups: Differences, When, Why, & How
- Walkthrough
- Troubleshooting



AVAILABILITY GROUPS: A BRIEF INTRODUCTION

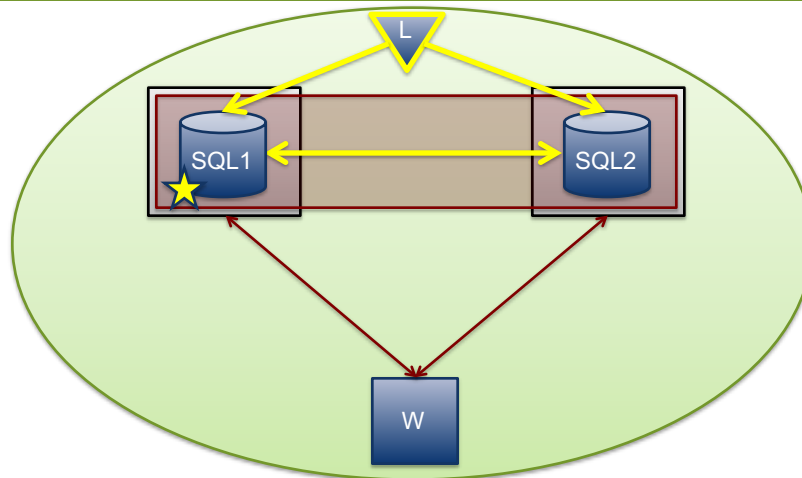


HA/DR Before SQL Server 2012

- Replication
- Log shipping
- Failover cluster
- Database mirroring



Availability Groups



Availability Groups

- Uses Windows Server Failover Cluster
- No shared storage
- Up to 4 secondary replicas (8 in 2014)
- Secondary replicas can be readable
- Backups can be offloaded to secondary replica
- Synchronous or asynchronous
- Automatic or manual failover
- Can be geographically distributed
- Can be combined with other HA/DR methods



Limitations

- Operates at database level
- Practical limit to number of databases per AG
 - Worker Thread Exhaustion
- Licensing
 - Enterprise Edition only
 - Must license all servers
 - ... Except those used strictly for disaster recovery



Limitations: Worker Thread Exhaustion

- Required (per AG)
 - 1 per database in primary replica
 - 1 per database for each secondary replica
 - 1 for each secondary replica
 - +1 if using secondary replica backup
- Limit: Max Worker Threads – 40
- Light Usage: $(D * 2) + 1$
- Heavy Usage: $(D * 5) + 1$



Limitations: Worker Thread Exhaustion

# CPU Cores	32-bit	64-bit
<= 4	256	512
8	288	576
16	352	704
32	480	960
64	736	1472
128	4224	4480
256	8320	8576


Source: MSDN (<http://bit.ly/1DB0nWi>)




MULTI-SUBNET AVAILABILITY GROUPS




Typical Availability Group (HA)



Simple Multi-Subnet Availability Group

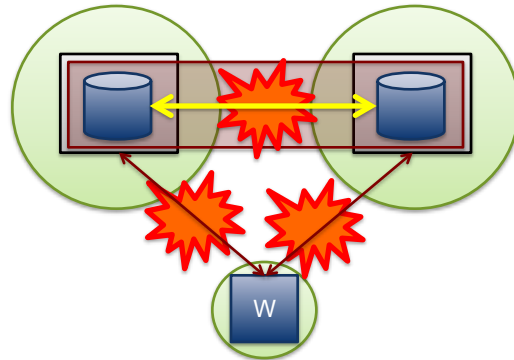


HA+DR – Best of Both Worlds



Potential Pain Points

- It's the network
 - Latency
 - Firewalls
 - Addressing
 - Active Directory
 - DNS



Why Use Multi-Subnet AGs?

- Disaster recovery
- Reporting in remote locations



Prerequisites

- Windows Server Failover Cluster (WSFC)
 - Quorum (File share witness)
 - Node weight (“votes”)



Prerequisites


- Separate subnet/NICs for heartbeat traffic
- IPs on all subnets
 - Public – One per server
 - Heartbeat – One per server
 - Cluster – One per subnet
 - Listener – One per subnet
- FQDNs
 - Cluster
 - Listener



Prerequisites

- Open ports

TCP/UDP	Port	Description
TCP	1433	SQL – AG Listener
TCP	5022	SQL – AG/DBM Endpoint
TCP/UDP	3343	Cluster Network Communication
TCP	135	Cluster DCOM
UDP	137	Authentication
	ICMP	Ping
TCP/UDP	49152-65535	Dynamic Ports (Configurable)




Prerequisites

- File share for initial data synchronization
- Service account (SQL)
- Active Directory/DNS security
 - Cluster machine account (*domain\cluster\$*)
- Mechanism for synchronizing server objects
 - Jobs
 - Logins



WALKTHROUGH



SPECIAL CONSIDERATIONS



Special Considerations

- Synchronous or asynchronous?
 - Synchronous
 - High availability
 - No data loss is acceptable
 - WAN latency must be low
 - Asynchronous
 - Disaster recovery
 - Distributed reporting
 - Data loss must be acceptable
 - Hybrid



Special Considerations

- Do you need automatic failover?
 - Latency must be low (<1 sec *roundtrip*)
 - CrossSubnetDelay (Def 1000, Max 5000)
 - CrossSubnetThreshold (Def 5, Max 120)
 - Usually reserved for high availability scenarios



Special Considerations

- IP Handling
 - Works best with SQL Server Native Client v11 or SQL Server ODBC v11
 - Other drivers cannot support MultiSubnetFailover property
 - Manual adjustments required
 - Slower reconnection after failover



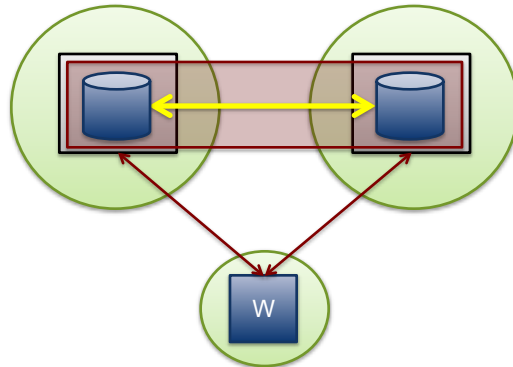
Special Considerations

- IP handling
 - Register all listener IPs
 - SQL driver must be able to support MultiSubnetFailover
 - RegisterAllIPs = 1
 - HostRecordTTL = 300
 - Register only current listener IP
 - RegisterAllIPs = 0
 - HostRecordTTL = as low as possible
 - Cluster must be able to update DNS
 - Much slower failover



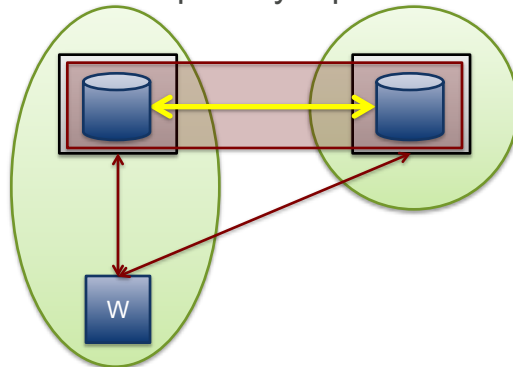
Special Considerations

- Where to put the witness?
 - 3rd datacenter




Special Considerations

- Where to put the witness?
 - Co-located with primary replica



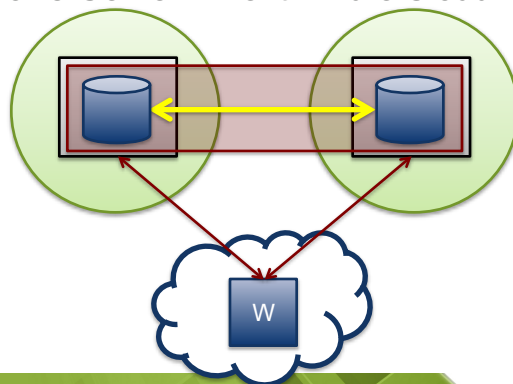
Special Considerations

- Where to put the witness?
 - Co-located with secondary replica



Special Considerations

- Where to put the witness?
 - Windows Server v.Next: Azure Cloud Witness



TROUBLESHOOTING AVAILABILITY GROUPS



Troubleshooting

- SQL Server
 - Availability Group Dashboard
 - Extended Events
 - Special session for availability groups
 - DMVs (sys.dm_hadr_*)
 - Instance log
- Third-party SQL utilities
 - SQL Sentry Performance Advisor
 - Idera Diagnostic Manager



Troubleshooting

- Windows
 - Logs – Application, System
 - Cluster logs
 - Including cluster.log
 - PerfMon
 - SCOM
- Network utilities
 - iperf
 - Packet trace (Wireshark)



Review

- Availability groups can be a great disaster recovery solution!
- Additional considerations when multiple subnets are involved
 - Network latency
 - Firewall configuration
 - Active Directory configuration
 - DNS configuration



Any Questions?

- Email: ed@leightondick.com
- Blog: edleightondick.com
- Twitter: [@eightondick](https://twitter.com/eightondick)

Thank you for attending!

Please remember to
fill out your comment cards

