

Windows Server 2016 Clustering and SQL Server 2016

Michael Steineke

Director - Data Platform,
Azure Cloud and Datacenter
Concurrency, Inc.



Terminology and Architecture

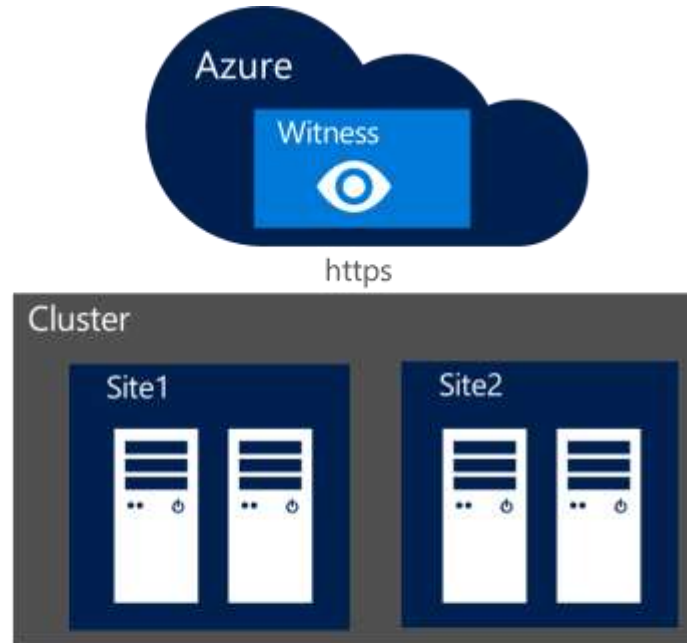
- ▶ AlwaysOn
 - ▶ Marketing name for the high availability features in SQL Server
 - ▶ Includes
 - ▶ Fail over Cluster Instances (FCI)
 - ▶ Availability Groups (AG)
- ▶ Windows Server Failover Clustering (WSFC)
 - ▶ All of AlwaysOn uses WSFC

High Availability vs. Disaster Recovery

- ▶ High Availability
 - ▶ Generally same site
 - ▶ Quick Recovery
- ▶ Disaster Recovery
 - ▶ Separate Site
 - ▶ Slower Recovery

Cloud Witness

 <p>Hybrid Cloud</p>	<p>Leveraging the power of the public cloud to increase resiliency of your private cloud Azure blob storage as an arbitration point</p>
 <p>Flexible Scenarios</p>	<p>Stretched clusters without a 3rd site Multi-site FCI's and AG's SQL Guest Clusters in Azure VM role</p>






Cluster operating system rolling upgrades


- ▶ Enables an administrator to upgrade the operating system of the cluster nodes from without stopping the Hyper-V or the Scale-Out File Server workloads.
- ▶ In-Place OS Upgrade

Failover Clustering Sets for Start Ordering



- ▶ In a private cloud there may be multi-tier applications which are deployed across a set of virtual machines. Such as a database running in one VM, and an application leveraging that database running in another VM. It may be desired to have start ordering of highly available virtual machines which have dependencies.

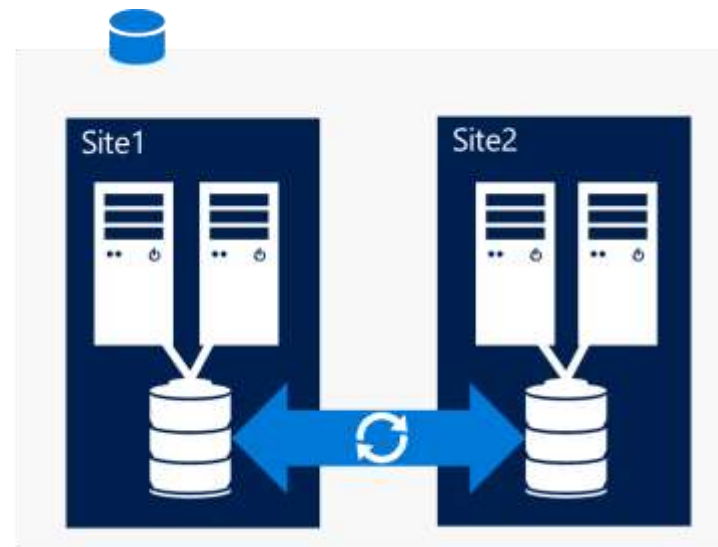
Storage Replica

 <p>Integrated</p>	End-to-end Windows Server disaster recovery solution for SQL FCI's
 <p>Flexible</p>	Volume level software replication between storage of any type Enables 3 rd party / LoB apps as well
 <p>Automatic</p>	Sync or Async replication Automatic cluster failover for low Recovery Time Objective (RTO)







Cluster Site Awareness

 <p>Sites</p>	<p>Define grouping of nodes in a stretched cluster which corresponds to their physical location</p> <p>Impacts placement policies and heartbeating</p>
 <p>Failover Affinity</p>	<p>SQL instances failover to a node within the same site, before failing to a node in a different site</p>



Quarantine of Flapping Nodes




 Protection	<p>Unhealthy nodes are quarantined and are no longer allowed to join the cluster</p> <p>Prevents flapping nodes from negatively affecting other nodes and the overall cluster</p>
 Resiliency	<p>Node is quarantined if it ungracefully leaves the cluster three times within an hour</p> <p>VMs are gracefully drained once quarantined</p>
 Control	<p>No more than 25% of nodes can be quarantined at any given time</p> <p>Nodes prevented from joining the cluster for 2 hours</p>




Performance

- ▶ Server Memory to 24 TB & SQL will use it
- ▶ Server Supports 640 Cores & SQL will use them
- ▶ Storage Class Memory
 - ▶ NVDIMM-N

Storage Spaces Direct


 Virtual Storage Array	DAS storage replicated across all nodes SQL FCI Clusters with no shared storage!
 Reliability	No single point of failure Aggregate NVMe to boost throughput and make it highly available
 Reducing CapEx	Reduced hardware costs with no external enclosures and storage cabling SQL FCI or AG with commodity hardware



Shared Disk in a VM

- ▶ VM Shared Disk for FCI guest Clusters
 - ▶ VSS Support to backup shared VHDX with Host backup
 - ▶ Online-resizing of shared VHDX file

Clusters in Workgroups and Across Domains

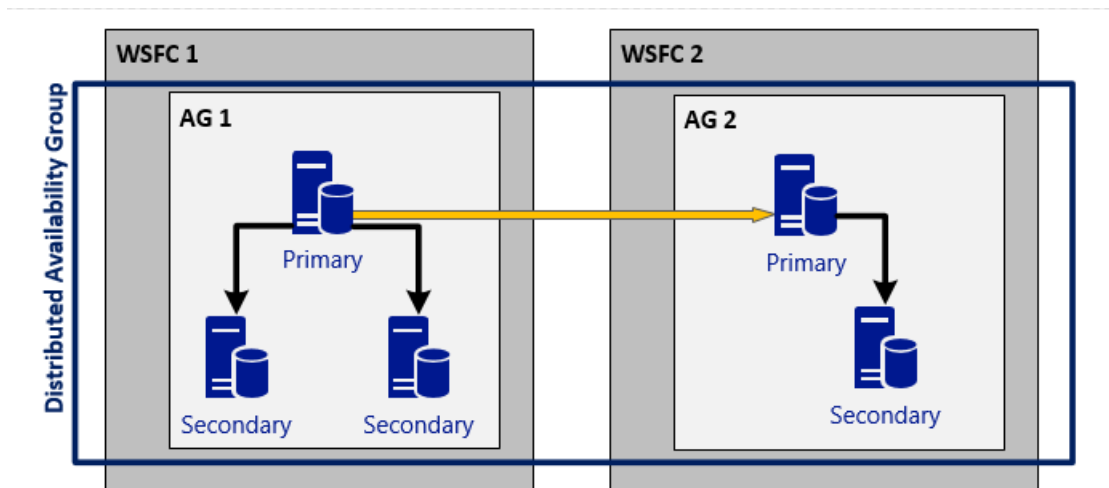


Domain'less



AG's Across Clusters

► Distributed Availability Group



Useful Reference Links

Rolling Upgrade

- ▶ <https://technet.microsoft.com/en-us/windows-server-docs/failover-clustering/cluster-operating-system-rolling-upgrade>

Cloud Witness

- ▶ <https://technet.microsoft.com/en-us/windows-server-docs/failover-clustering/deploy-cloud-witness>

Cluster Sets

- ▶ <https://blogs.msdn.microsoft.com/clustering/2016/10/10/failover-clustering-sets-for-start-ordering/>

Distributed AG

- ▶ <https://msdn.microsoft.com/en-US/library/mt651673.aspx>

Ignite Windows & SQL 2016 Session

- ▶ <https://onedrive.live.com/embed?cid=6B7CF2BD5F47F487&resid=6B7CF2BD5F47F487%21605&authkey=AHjn1A6e5G48jll&em=2&wdAr=1.7777777777777776>

Useful Reference Links

- ▶ SQL AlwaysOn Team Blog
 - ▶ <http://blogs.msdn.com/b/sqlalwayson/>
- ▶ SQL Server Customer Advisory Team
 - ▶ <http://blogs.msdn.com/b/sqlcat/>
- ▶ SQL Server 2014 Technet Reference
 - ▶ <http://technet.microsoft.com/en-US/sqlserver/dn135309>
- ▶ SQL High Availability on MSDN
 - ▶ <http://social.msdn.microsoft.com/Forums/en-US/home?forum=sqlhadr>
- ▶ Failover Cluster Guide to Quorum
 - ▶ [https://technet.microsoft.com/en-us/library/cc770620\(W5.10\).aspx](https://technet.microsoft.com/en-us/library/cc770620(W5.10).aspx)
- ▶ Prestage Computer Objects in AD
 - ▶ <https://technet.microsoft.com/en-us/library/dn466519.aspx>
- ▶ MSDTC
 - ▶ <https://blogs.msdn.microsoft.com/alwaysonpro/2014/01/15/msdtc-recommendations-on-sql-failover-cluster/>
- ▶ AGs on Private Network
 - ▶ <https://blogs.msdn.microsoft.com/alwaysonpro/2013/11/01/configure-availability-group-to-communicate-on-a-dedicated-network/>
- ▶ SQL 2016 Changes to AGs
 - ▶ <https://blogs.technet.microsoft.com/dataplatforminsider/2015/12/15/enhanced-always-on-availability-groups-in-sql-server-2016/>

Thanks!

- ▶ msteineke@concurrency.com
- ▶ www.michaelsteineke.com
- ▶ Twitter @MSteineke