

The High Availability Features in SQL Server 2016 Standard Edition

MARCH 2016

WHITE PAPER BY EDWIN SARMIENTO

Microsoft SQL Server MVP and Microsoft Certified Master

Trademarks

“StarWind”, “StarWind Software” and the StarWind and the StarWind Software logos are registered trademarks of StarWind Software. “StarWind LSFS” is a trademark of StarWind Software which may be registered in some jurisdictions. All other trademarks are owned by their respective owners.

Changes

The material in this document is for information only and is subject to change without notice. While reasonable efforts have been made in the preparation of this document to assure its accuracy, StarWind Software assumes no liability resulting from errors or omissions in this document, or from the use of the information contained herein. StarWind Software reserves the right to make changes in the product design without reservation and without notification to its users.

Technical Support and Services

If you have questions about installing or using this software, check this and other documents first - you will find answers to most of your questions on the [Technical Papers](#) webpage or in [StarWind Forum](#). If you need further assistance, please [contact us](#).

Copyright ©2009-2015 StarWind Software Inc.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written consent of StarWind Software.

In 2016, Gartner named StarWind “Cool Vendor for Compute Platforms”.

Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

About StarWind

StarWind is a pioneer in virtualization and a company that participated in the development of this technology from its earliest days. Now the company is among the leading vendors of software and hardware hyper-converged solutions. The company's core product is the years-proven StarWind Virtual SAN, which allows SMB and ROBO to benefit from cost-efficient hyperconverged IT infrastructure. Having earned a reputation of reliability, StarWind created a hardware product line and is actively tapping into hyperconverged and storage appliances market. In 2016, Gartner named StarWind “Cool Vendor for Compute Platforms” following the success and popularity of StarWind HyperConverged Appliance. StarWind partners with world-known companies: Microsoft, VMware, Veeam, Intel, Dell, Mellanox, Citrix, Western Digital, etc.

Contents

Introduction	4
Support for Basic Availability Groups	5
Support for Failover Clustered Instances	7
Redundant Copies of the Database: The Case Against SQL Server Failover Clustered Instances	8
StarWind Virtual SAN's Hyper-Converged Storage for SQL Server Failover Clustered Instances	8
Summary	8
Contacts.....	9

Introduction

This technical paper is intended for SQL Server users who are thinking of upgrading their SQL Server 2005. As there are two options available, namely SQL Server 2014 and SQL Server 2016, this document is all you need to choose the right version based on your current high availability and disaster recovery requirements.

As SQL Server 2005 approaches the end of its support next year, customers are already thinking about whether to upgrade to SQL Server 2014 or jump in on the upcoming version – SQL Server 2016. And with Microsoft increasing the pace of releasing SQL Server versions, it adds up to the challenge of making the right decision. This document aims to help you make that decision based on your current high availability needs.

Business requirements change overtime and with those changes come the challenges of addressing them. High availability and disaster recovery (HA/DR) is just one of those business requirements. As the business grow, the amount of data accumulated grows with it. And with data becoming more and more critical to the business operations, it's no surprise that high availability and disaster recovery has become a standard requirement. In fact, as of writing this document, I am working with a customer from five years ago who started off with a relatively small database that has now grown into a terabyte. They are currently evaluating several high availability features to make sure that their applications have increased availability to serve their business needs.

HA/DR solutions are expensive primarily because they are meant to protect digital assets. Most stakeholders don't even want to consider HA/DR because of their very nature – being on standby. They exist to address the “what if” scenarios. That is why I am constantly evaluating technology solutions and features that will address HA/DR requirements without breaking the bank. Below are features made available in SQL Server 2016 that customers and businesses need to consider – primarily because they exist in Standard Edition.

Support for Basic Availability Groups

SQL Server Availability Groups was introduced in SQL Server 2012. It was originally meant to be a replacement for database mirroring. Availability Groups brought SQL Server high availability and disaster recovery to a whole new level by allowing multiple copies of the database be highly available and potentially using them for read-only workloads and offloading management tasks such as backups. Availability Groups allowed you to fail over a group of databases as a single entity, unlike database mirroring where you can only do so one database at a time. This is very useful for applications that access multiple databases in a single SQL Server instance like SharePoint Server 2013.

The biggest concern that many customers have with Availability Groups is cost – it is only available in Enterprise Edition. The SQL Server 2012 release was also the time that Microsoft announced the deprecation of database mirroring in future versions. And with the licensing changes that came with SQL Server 2012, it became an even more expensive option. A lot of customers that use database mirroring were on Standard Edition. They still want to keep their high availability solution while minimizing costs associated with both licensing and administration. Cost is the primary reason why I still recommend database mirroring even on SQL Server 2014 because, even though it is marked for deprecation, it is still technically supported.

With SQL Server 2016 comes support for Availability Groups in Standard Edition. It first appeared in the Community Technology Preview (CTP) 2 release and was called Basic Availability Group.

While the feature is not yet fully baked into the product, it is worth noting that customers now have a viable replacement for database mirroring in Standard Edition. While you are limited to just having two replicas and a single database in an Availability Group, several features make this a great cost-effective high availability solution. Note that some of these may change until the final release.

- Support for **Listener Name**. With database mirroring, you are responsible for redirecting client applications to the mirror partner after a failover unless the client application supports the **Failover Partner** connection string attribute. The Listener Name enables a client application to connect to an Availability Group replica without knowing the name of the physical instance of SQL Server to which the client is connecting. The client connection string does not need to be modified to connect to the current location of the current primary replica. When Basic Availability Groups was introduced in SQL Server 2016 CTP2, the Listener Name option was removed. It will be re-introduced in future CTP releases.
- Support for single or multiple Active Directory domains, even workgroups. Database mirroring does not require Windows Server Failover Clustering (WSFC) unlike Availability Groups. Because

Availability Groups rely on the underlying WSFC, the infrastructure requirements are the same between the two. Database mirroring did not require an Active Directory domain to set it up. Availability Groups made it more challenging for database administrators because they are now dependent on Active Directory and WSFC. The limitations of Availability Groups were a by-product of the limitations of WSFC. The upcoming Windows Server 2016 addresses these challenges by allowing you to create a WSFC with cluster nodes being a member of a single domain, different domains and even non-domain-joined servers. Of course, this assumes that you will be deploying SQL Server 2016 Availability Groups on top of Windows Server 2016 WSFC. Authentication between the SQL Server instances will be done via certificates instead of Active Directory.

- **Support for multi-subnet failover.** Database mirroring was typically deployed as a high availability solution with SQL Server instances on different network subnets. While client application connection string still requires a change (or may already have it using the **Failover Partner** attribute) when a failover happens, the DNS entries that represent the SQL Server instances are still unique across the network. And while Availability Group Listener Name supports having multiple virtual IP addresses in different network subnet, you still need to rely on the DNS client cache to successfully connect to the Availability Group replica during a failover. Basic Availability Groups support multi-subnet failover, allowing client applications to seamlessly connect to the Availability Group replica during a failover. And with the latest SQL Server Native Client, using the **MultiSubnetFailover** parameter in the connection string is no longer necessary because it will automatically parallelize connections to the Listener Name's virtual IP address in the different subnets.

*This limited information was made available with the permission from the SQL Server product team responsible for Availability Groups.

Support for Failover Clustered Instances

SQL Server Failover Clustered Instances on Standard Edition has been available since SQL Server 2005. Customers have since implemented two-node SQL Server Failover Clustered Instances on Standard Edition with the use of shared storage.

The features in Availability Groups – support for Active Directory domains or workgroups and multi-subnet failover - are also in Failover Clustered Instances, mainly because both of them rely on the same WSFC architecture. But, unlike Availability Groups that protects a group of databases, Failover Clustered Instances protect an entire SQL Server instance. This is because the system and user databases are stored in the shared storage accessed by the nodes of the cluster. Hence, the need for shared storage in the WSFC to store the SQL Server databases.

Several things make SQL Server 2016 Failover Clustered Instances on Standard Edition a cost-effective high availability solution.

- **Support for Cluster Shared Volumes.** Cluster Shared Volumes (CSV) was introduced in Windows Server 2008 R2 primarily for Hyper-V workloads. However, SQL Server 2014 was the very first version that supported this feature. CSVs improved the use and management of shared storage in SQL Server which also removed the drive-letter limitation. It also increased resiliency of shared storage failover by making the storage an indirect dependency of the SQL Server clustered resource.
- **Support for different storage solutions.** As Failover Clustered Instances rely on shared storage, different storage solutions are made available like the traditional WSFC shared disks as well as server message block (SMB) file shares. The option to run SQL Server on SMB file shares was introduced in SQL Server 2012. This continues to improve with the SMB 3.1.1 enhancements in Windows Server 2016.

Redundant Copies of the Database: The Case Against SQL Server Failover Clustered Instances

One of the reasons why Availability Groups was a preferred option over Failover Clustered Instances is because of the ability to have redundant copies of the database in multiple storage locations. Because Availability Groups involved different SQL Server instances with their own storage subsystem and copies of the database being replicated, it created the ability to have redundant copies of the database should the primary replica and its corresponding storage subsystem becomes inaccessible. This is in contrast to a Failover Clustered Instance where only a single copy of the databases is stored in a shared storage. The underlying storage subsystem becomes a single point of failure

StarWind Virtual SAN's Hyper-Converged Storage for SQL Server Failover Clustered Instances

StarWind Virtual SAN is a pure software defined storage platform which eliminates the need for expensive dedicated storage hardware. StarWind Virtual SAN achieves this by installing the software on the same hardware as the nodes in a SQL Server Failover Clustered Instance. The storage replication feature allows for redundant copies of all of the database files on a per-block-level instead of on a per-file-level like that of Availability Groups. This creates a fault tolerant storage subsystem that a Failover Clustered Instance can take advantage of.

Using StarWind Virtual SAN with SQL Server Failover Clustered Instances, we address the cost concerns of Enterprise Edition license requirement for Availability Group, the single database limitation in Basic Availability Group and the single copy of the database in a traditional Failover Clustered Instance. StarWind Virtual SAN provides a cost-efficient strategy for delivering high availability in a SQL Server environment.

Summary

These are just some of the features that make SQL Server 2016 on Standard Edition a cost-effective high availability solution. And Microsoft isn't done yet. They will continue to improve these features in future CTPs as they approach general availability. Availability of the documentation for these features will be limited as well as their focus is to get the features in the product and perform extensive testing.

Contacts

US Headquarters	EMEA and APAC
 1-617-449-7717	 +44 20 3769 1857 (UK)
 1-617-507-5845	 +49 302 1788 849 (Germany)
	 +33 097 7197 857 (France)
	 +7 495 975 94 39 (Russian Federation and CIS)
	 1-866-790-2646

Customer Support Portal: <https://www.starwind.com/support>

Support Forum: <https://www.starwind.com/forums>

Sales: sales@starwind.com

General Information: info@starwind.com



StarWind Software, Inc. 35 Village Rd., Suite 100, Middleton, MA 01949
©2015, StarWind Software Inc. All rights reserved.