SQL Server 2012/2014
AlwaysOn Availability Group

Part 2 - Design

v1.0 - 2014 - G.MONVILLE

Summary

Microsoft® SQL Server®

Part 2 - Design

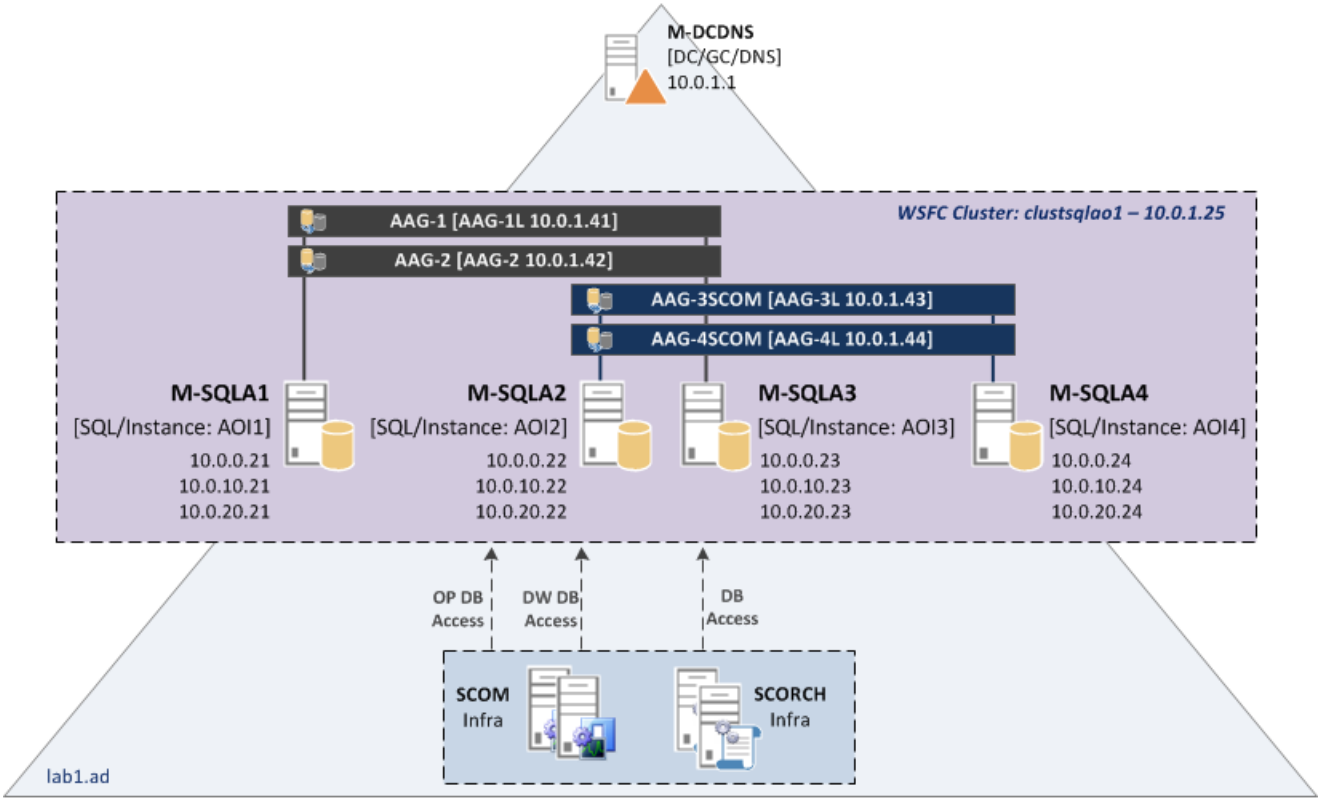

For the tests, I will create an "AlwaysOn Availability Group" cluster with four nodes and four AAG. Each AAG has two SQL Server Instance members, so each SQL node participate to two AAG.

The first two AAG will be used to host only test Databases. The other two will be used to host databases for SCOM 2012 R2 and SCO 2012 R2.

LAB Requirements

Three Networks are required.

vSwitch	Description	Subnet
vSwitch0-Public	Client Access	10.0.1.0 /24
vSwitch1-Cluster	Heartbeat	10.0.10.0 /24
vSwitch2-Replication	AAG Replication	10.0.20.0 /24

Note: I use same subnet for all nodes, I'll write an article for WSFC Cluster Administration/Troubleshoot which also cover cross-subnet configuration.

Infra server:

Server	Description	IP
M-DCDNS	AD Root / DNS	10.0.1.1

SQL Servers / Instances Configuration

The lab will be composed on a four node WSFC cluster:

Hostname	OS	IP VLAN Public	IP VLAN CLUSTER	IP VLAN Replication	Note
M-SQLA1	WS2012R2	10.0.1.21	10.0.10.21	10.0.20.21	
M-SQLA2	WS2012R2-CORE	10.0.1.22	10.0.10.22	10.0.20.22	
M-SQLA3	WS2012R2-CORE	10.0.1.23	10.0.10.23	10.0.20.23	
M-SQLA4	WS2012R2-CORE	10.0.1.24	10.0.10.24	10.0.20.24	
clustsqlao1	n/a	10.0.1.25	n/a	n/a	Cluster Resource Name

M-SQLA1 OS will be installed in full GUI mode with the SQL Feature "Management Tools - Complete" (include "Management Studio"; it's not compatible with a Core installation). This server will be used to manage SQL AAG and WSFC cluster.

Note: In a production environment, all servers must be identical (all in core mode, or full/minimal) and a dedicated "management/tools" server with consoles is used for administration.

Best Practices and Recommendations

It's recommended to use the Windows Server Core Installation option for setting up a SQL server environment (especially if it's virtualized).
Advantages of a SQL Core installation:

- reduce the space required on disk.
- reduce the potential attack surface.
- reduce the overhead of updating patches.
- minimize the requirements for servicing and restarting the server.

We need to install one named-instance per SQL Server:

Server	Instance Name	Instance Port	SQL Features
M-SQL1	aoi1	1764	SQL Database Engine Full-Text Search (needed for SCOM)
M-SQL2	aoi2	1764	
M-SQL3	aoi3	1764	
M-SQL4	aoi4	1764	

Note Port Instances/Listener:
For an AAG Environment, you have to choice Ports for instances (here x4) and Ports for AAG-Listener (also x4 in my lab). I choose to use the same port (but not the default 1433) for all instances and all AAG Listeners, but there is no restriction. You can use different ports for each instance, different ports for each Listener, same port for all instances and another port for all Listeners, etc…

Availability Groups Configuration

I will create four Availability Groups:

AAG	Members (Instance)	Default Role	AAG Listener Name	AAG Listener IP	AAG Listener Port	Databases
AAG-1	m-sqla1\aoi1 m-sqla3\aoi3	Primary Secondary	AAG-1L	10.0.1.41	1764	DBTest01
AAG-2	m-sqla1\aoi1 m-sqla3\aoi3	Secondary Primary	AAG-2L	10.0.1.42	1764	DBTest02
AAG-3SCOM	m-sqla2\aoi2 m-sqla4\aoi4	Primary Secondary	AAG-3L	10.0.1.43	1764	SCOM OP
AAG-4SCOM	m-sqla2\aoi2 m-sqla4\aoi4	Secondary Primary	AAG-4L	10.0.1.44	1764	SCOM DW DB Orchestrator

AAG-1 and AAG-2 will serve for tests only. AAG-3SCOM and AAG-4SCOM will be used for my SCOM and Orchestrator Labs.
In this configuration, in nominal mode each instance hosts an "Active" Primary Replica.
The simulation is m-sql1 and m-sql2 in the same room and the two others in another room.
So I can lose one room (all my AAG/Databases remain available)

AAG Listener (VNN - Virtual Network Name)

For reminder, on the WSFC cluster side an AAG is a cluster Resource Group and the VNN is two cluster resources:

- Virtual Name
- Virtual IP

Example:

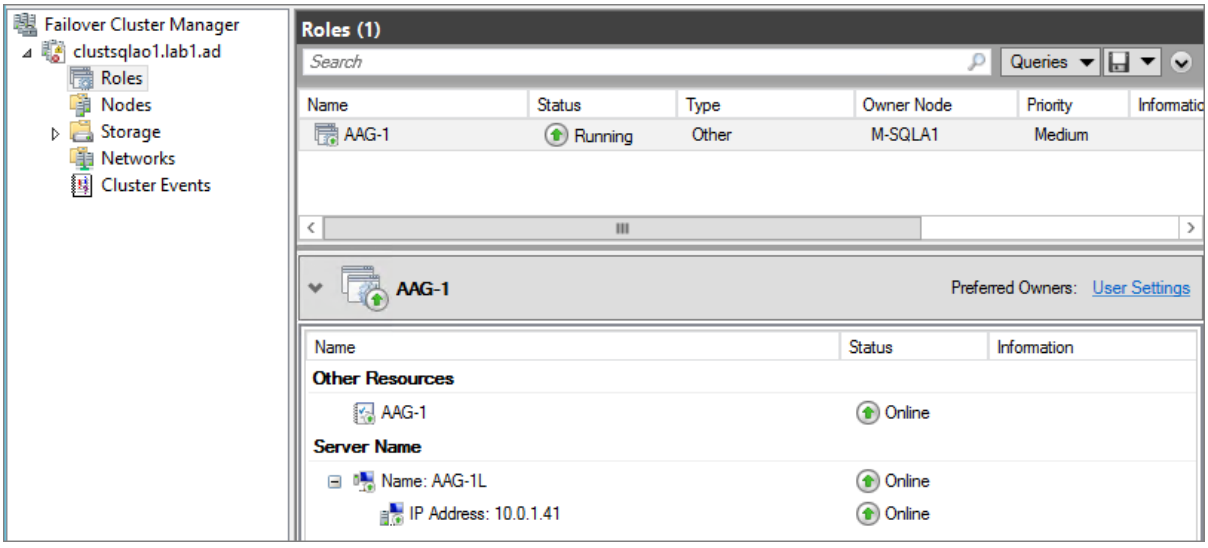

When you configure an application to host its Database on a SQL Availability Group you have to specify the Listener name for the instance name and the Listener port for the Instance port.

AAG Implementation - Version 1

This is the first version that will be configured in the next parts of the article:

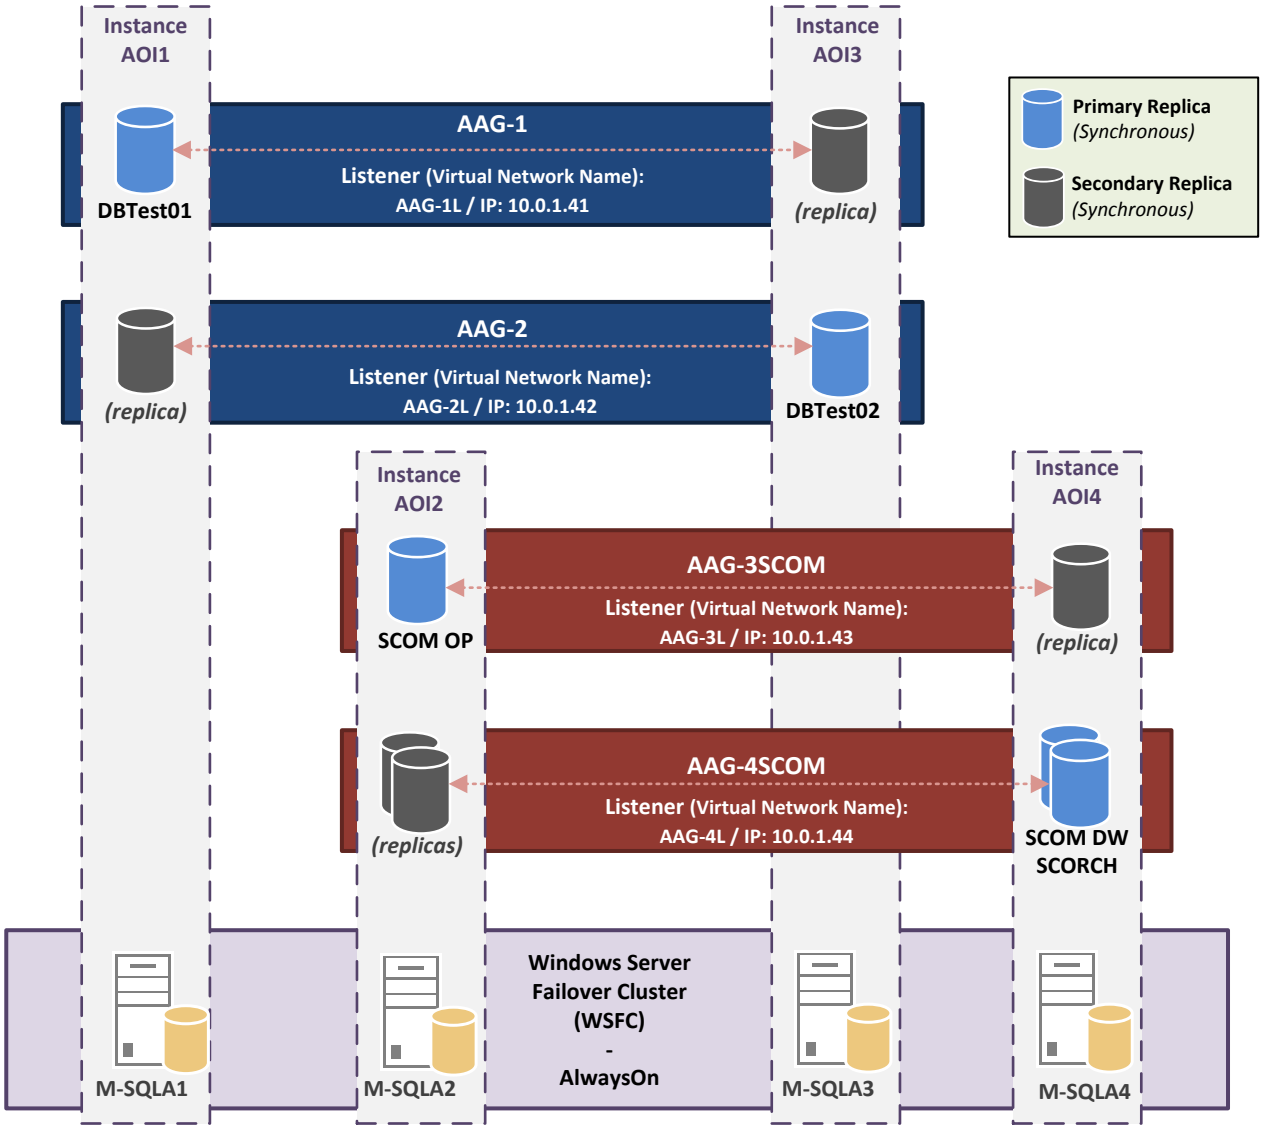

AAG Implementation - Version 2

In another part, to simulate a Remote DRP Site, I will add and additional Instance (with two Replicas in Asynchronous mode on the AAG-1 and the AAG-2):

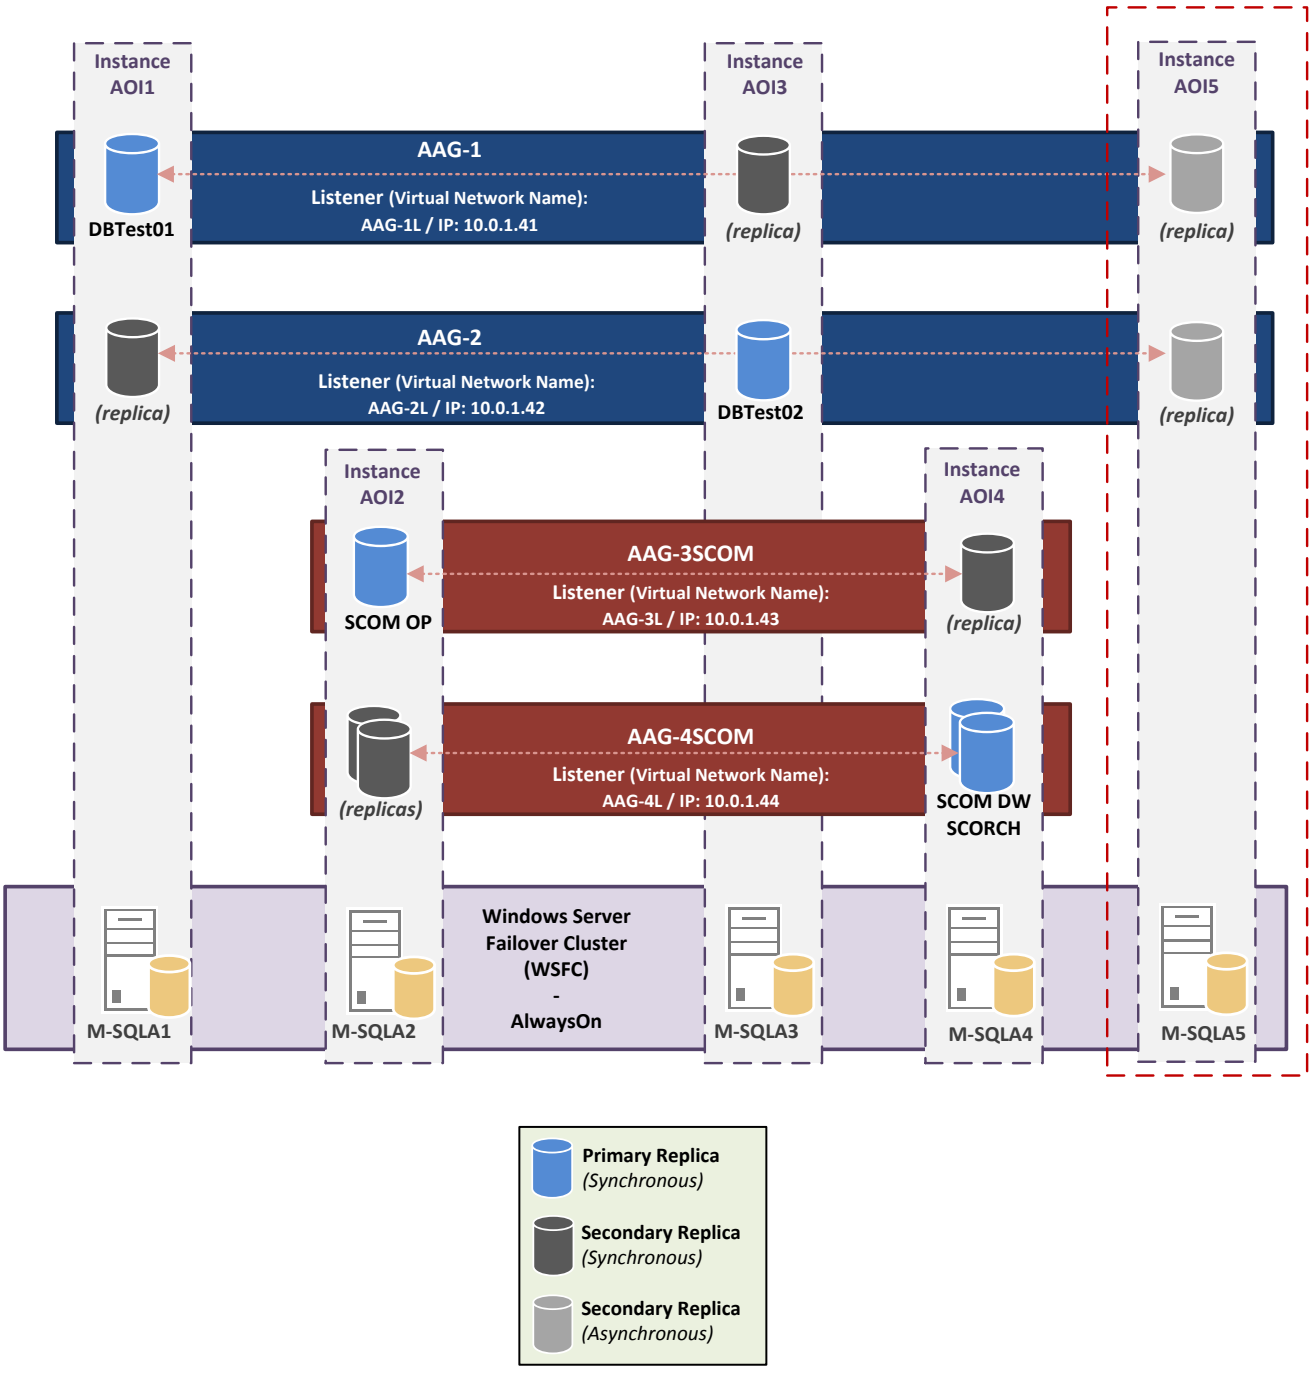

Availability Replicas Configuration

The next part is to specify the detailed availability replica (two per AAG) configuration:

AAG	Server Instance	Initial Role	Automatic Failover	Synchronous Commit	Allow Readable Secondary
AAG-1	m-sqla1\AOI1	Primary	Yes	Yes	Yes
	m-sqla3\AOI3	Secondary	Yes	Yes	Yes
AAG-2	m-sqla1\AOI1	Secondary	Yes	Yes	Yes
	m-sqla3\AOI3	Primary	Yes	Yes	Yes
AAG-3SCOM	m-sqla2\AOI2	Primary	Yes	Yes	Yes
	m-sqla4\AOI4	Secondary	Yes	Yes	Yes
AAG-4SCOM	m-sqla2\AOI2	Secondary	Yes	Yes	Yes
	m-sqla4\AOI4	Primary	Yes	Yes	Yes

All replicas will be configured in "Automatic" failover mode and so in "Synchronous" availability mode.

For more information see TechNet: **Failover and Failover Modes (AlwaysOn Availability Groups)** - http://technet.microsoft.com/en-us/library/hh213151

Readable Secondary Option:

For future tests, I enable Readable Secondary option.

Option	Description
No	No user connections are allowed to secondary databases of this replica. They are not available for read access. This is the default setting.
Read-intent only	Only read-only connections are allowed to secondary databases of this replica. The secondary database(s) are all available for read access.
Yes	All connections are allowed to secondary databases of this replica, but only for read access. The secondary database(s) are all available for read access.

Primary Role Connections:

I use the default settings (Allow all connections).

Option	Description
Allow all connections	All connections are allowed to the databases in the primary replica. This is the default setting.
Allow read/write connections	When the Application Intent property is set to ReadWrite or the Application Intent connection property is not set, the connection is allowed. Connections where the Application Intent connection property is set to ReadOnly are not allowed. This can help prevent customers from connecting a read-intent work load to the primary replica by mistake.

Endpoints Configuration

There is one Endpoint per SQL Server Instance.

During AAG Creation (via Wizard), Endpoint URL is configured with the SQL Instance FQDN. With this default option, instances will communicate over the Public Network (for reminder: 10.0.1.0/24).

So to configured instance communication on the Replication Network (10.0.20.0/24) I have to set my endpoint to: TCP://10.0.20.x:5022.

For tests, I will configure two instances (AOI1 and AOI3) on the Public Network (with FQDN) and the two other instances (AOI2 and AOI4) on the Replication Network.

Server Instance	Endpoint URL	Endpoint Port	Endpoint Name
m-sqla1\AOI1	TCP://M-SQLA1.lab1.ad:5022	5022	Hadr_endpoint
m-sqla2\AOI2	TCP://10.0.20.22:5022	5022	Hadr_endpoint
m-sqla3\AOI3	TCP://M-SQLA3.lab1.ad:5022	5022	Hadr_endpoint
m-sqla4\AOI4	TCP://10.0.20.24:5022	5022	Hadr_endpoint

Note: 5022 is the default port, you can use another port.

Service Accounts Requirement

Isolate Instance Services

Isolating services reduces the risk that one compromised service could be used to compromise others.

At the Instance level, each SQL Service (SQL Server, SQL Agent ...) must be configured with different account.

Isolate Instances

A Security Best Practice is to use different accounts for each instance, but considers these points:

- Microsoft recommends to use the same account for all instances of an AlwaysOn Cluster (it's more simple to assign rights to Endpoints)

- If you want to use Kerberos, instances must use the same account:

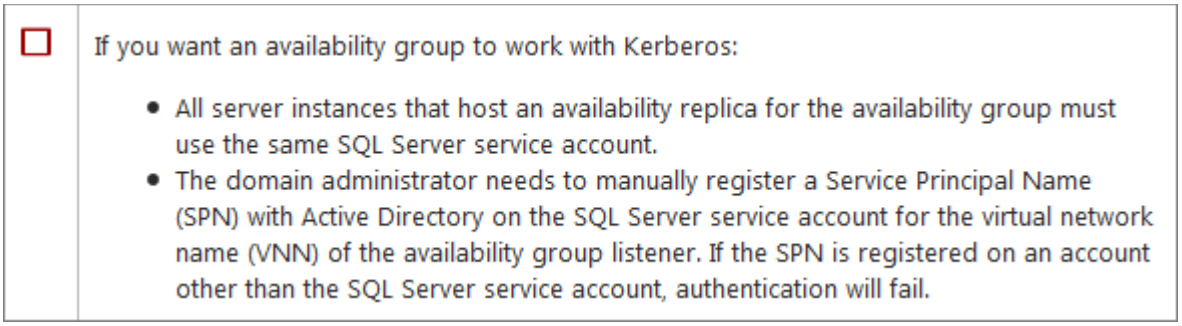

Service Accounts - Solutions

Use the same account for all Instances (enable Kerberos authentication):

- ***gMSA (Group Managed Service Accounts):*** the best solution for the AlwaysOn Availability Group is to use a gMSA (same as a MSA account but available on multiple host). **But it's not supported for the moment on SQL Server...**

  This is article for more information: gMSA/MSA accounts used with SQL
  http://blogs.msdn.com/b/sqlosteam/archive/2014/02/19/msa-accounts-used-with-sql.aspx

  Group Managed Service Accounts Overview
  http://technet.microsoft.com/en-us/library/hh831782.aspx

- ***"Classic" Domain Account:*** you can use the same domain account for all instances (this works), but when you have to change the password account you have to program an interruption of service (all node will be affected at the same time by the password change...)

Use different accounts for all Instances (disable Kerberos authentication):

- ***MSA (Managed Service Account):*** you can use a MSA account per Instance (MSA is a domain account; password is managed automatically by the domain controller; a MSA is assigned to only one host)

- ***Virtual Accounts:*** you can use a virtual account per Instance (the functioning is identical to a MSA except it's a local account managed by the host, not by the DC). This is the default option during a SQL Instance installation.

For more information, see TechNet article: ***Configure Windows Service Accounts and Permissions -*** *http://msdn.microsoft.com/en-us/library/ms143504.aspx*

So actually, there is no possible solution for use Kerberos with AAG in a production environment. I will use MSA account for my lab.

Account	MSA	Description	Member Of / Rights	Instance	Service Mode
lab1\SQLAlwaysOnAdmins	n/a	SQL Administrators Group	Local Administrator of all nodes Sysadmin on all instance	n/a	n/a
lab1\sqlaoinstall	No	Account use for Installation	Member of SQLAOAdmins Group	n/a	n/a
lab1\svc-sqldbe1	Yes	SQL Service - Database Engine	Domain User	aoi1	Automatic
lab1\svc-sqlagt1	Yes	SQL Service - Agent	Domain User		Automatic
lab1\svc-sqldbe2	Yes	SQL Service - Database Engine	Domain User	aoi2	Automatic
lab1\svc-sqlagt2	Yes	SQL Service - Agent	Domain User		Automatic
lab1\svc-sqldbe3	Yes	SQL Service - Database Engine	Domain User	aoi3	Automatic
lab1\svc-sqlagt3	Yes	SQL Service - Agent	Domain User		Automatic
lab1\svc-sqldbe4	Yes	SQL Service - Database Engine	Domain User	aoi4	Automatic
lab1\svc-sqlagt4	Yes	SQL Service - Agent	Domain User		Automatic

Permission needed for Service Account:

Notes: During installation, these permissions are granted by the SQL setup.

Service	Description	Permissions granted by SQL Server Setup
SQL Server Database Services	The service for the SQL Server relational Database Engine. The executable file is <MSSQLPATH>\MSSQL\Binn\sqlservr.exe.	Log on as a service Replace a process-level token Bypass traverse checking Adjust memory quotas for a process Permission to start SQL Writer Permission to read the Event Log service Permission to read the Remote Procedure Call service
SQL Server Agent	Executes jobs, monitors SQL Server, fires alerts, and enables automation of some administrative tasks. The executable file is <MSSQLPATH>\MSSQL\Binn\sqlagent.exe.	Log on as a service Replace a process-level token Bypass traverse checking Adjust memory quotas for a process
Reporting Services	Manages, executes, creates, schedules, and delivers reports. The executable file is <MSSQLPATH>\Reporting Services\ReportServer\Bin\ReportingServicesService.exe.	Log on as a service
SQL Server Browser	The name resolution service that provides SQL Server connection information for client computers. The executable path is c:\Program Files (x86)\Microsoft SQL Server\90\Shared\sqlbrowser.exe	Log on as a service
Full-text search	Quickly creates full-text indexes on content and properties of structured and semistructured data to provide document filtering and word-breaking for SQL Server.	Log on as a service Adjust memory quotas for a process Bypass traverse checking

Storage

Disk configuration per node:

Disk	Letter	RAID Level	Size	Name	SQL Path	Description
disk0	*c:*	n/a	25GB	System	C:\Program Files\Microsoft SQL Server\ C:\Program Files (x86) \Microsoft SQL Server\ C:\MSSQL\MSSQL11.<instancename>\ C:\MSSQL\MSSQL11.<instancename>\MSSQL\Data	SQL Shared Features SQL Shared Features SQL Server Directory System Databases
disk1	*G:*	n/a	5 GB	SQL_DB	G:\MSSQL\AOREPLICA\Data G:\MSSQL\MSSQL11.<instancename>\MSSQL\TempDB\Data G:\MSSQL\MSSQL11.<instancename>\MSSQL\Backup	Databases TempDB Database Database Backups
disk2	*L:*	n/a	5 GB	SQL_LOG	L:\MSSQL\AOREPLICA\Log L:\MSSQL\MSSQL11.<instancename>\MSSQL\TempDB\Log	DB Transaction Log TempDB Log

Notes about Storage:

If your SQL Servers are virtualized, for production environment you shouldn't use Virtual Disk (except for OS). You have to use Pass-through (via Virtual FC) for Hyper-V, or RDM LUN for VMware.
In addition for better performance you must use a dedicated disk for TempDB.
Install SQL Server (SQL Server Directory) on a separate disk (D:).

You can also add a separate disk for pagefile, but if the SQL server is correctly sized it should not have to swap.

Example of a Production configuration

Disk	Letter	RAID Level	Size	Name	SQL Path	Description
disk0	*C:*	*Raid 1*	xx GB	System	C:\Program Files\Microsoft SQL Server\ C:\Program Files (x86) \Microsoft SQL Server\	SQL Shared Features SQL Shared Features
disk1	*D:*	*Raid 1*	xx GB	*SQL_BIN*	D:\MSSQL\MSSQL11.<instancename>\ D:\MSSQL\MSSQL11.<instancename>\MSSQL\Data	SQL Server Directory System Databases
disk2	*G:*	*Raid 10*	xx GB	SQL_DB	G:\MSSQL\AOREPLICA\Data	Databases
disk3	*K:*	*Raid 5*	xx GB	SQL_BAK	K:\MSSQL\MSSQL11.<instancename>\MSSQL\Backup	Database Backups
disk4	*L:*	*Raid 10*	xx GB	SQL_LOG	L:\MSSQL\AOREPLICA \Log	Transaction Log
disk5	*T:*	*Raid 10*	xx GB	SQL_TEMPDB	T:\MSSQL\MSSQL11.<instancename>\MSSQL\TempDB\Data T:\MSSQL\MSSQL11.<instancename>\MSSQL\TempDB\Log	TempDB Database TempDB Logs
disk6	*R:*	*Raid 5*	xx GB	SQL_SSRS	R:\MSSQL\MSSQL11.<instancename>\MSSQL\Reports	SSRS Feature

Note for Databases/Logs path on AAG:

If you use the default instance path (which contains the instance name) for Databases and Logs, the paths on all the nodes participating to the AAG are different. This has an impact on AlwaysOn AG.

TechNet:

If the file path (including the drive letter) of a secondary database differs from the path of the corresponding primary database, the following restrictions apply:

- New Availability Group Wizard/Add Database to Availability Group Wizard: The Full option is not supported (on the "Select Initial Data Synchronization" Page),
- RESTORE WITH MOVE: To create the secondary databases, the database files must be RESTORED WITH MOVE on each instance of SQL Server that hosts a secondary replica.
- Impact on add-file operations: A later add-file operation on the primary replica might fail on the secondary databases. This failure could cause the secondary databases to be suspended. This, in turn, causes the secondary replicas to enter the NOT SYNCHRONIZING state.

So it is recommended to use the same path on all instances:

Data	Default Path	New Path
DB	G:\MSSQL\MSSQL11.<instancename>\MSSQL\Data	G:\MSSQL\AOREPLICA\Data
LOG	L:\MSSQL\MSSQL11.<instancename>\MSSQL\Log	L:\MSSQL\AOREPLICA\Log

Note about TempDB

- The TempDB shouldn't be store on the same disk as your Databases
- In Production, autogrow operations can affect performance so preallocate space to allow for the expected workload (autogrow should be used to increase disk space for unplanned exceptions)
- SQL CAT team recommends one file per CPU Core. Microsoft Note:

> Create as many files as needed to maximize disk bandwidth. Using multiple files reduces tempdb storage contention and yields significantly better scalability. However, do not create too many files because this can reduce performance and increase management overhead. As a general guideline, create one data file for each CPU on the server (accounting for any affinity mask settings) and then adjust the number of files up or down as necessary. Note that a dual-core CPU is considered to be two CPUs.

But this recommendation is subject to discussion and depends of your SQL environment (and the TempDB Contention). I'm not going to analyze this in this article, but I invite you to read the great articles of Paul Randal:

A SQL Server DBA myth a day: (12/30) tempdb should always have one data file per processor core:
http://www.sqlskills.com/blogs/paul/a-sql-server-dba-myth-a-day-1230-tempdb-should-always-have-one-data-file-per-processor-core/

The Accidental DBA (Day 27 of 30): Troubleshooting: Tempdb Contention:
http://www.sqlskills.com/blogs/paul/the-accidental-dba-day-27-of-30-troubleshooting-tempdb-contention/

Another "General" Recommendation:

> Last year at PASS 2011 Bob Ward, one the Sr Escalation Engineers for SQL, made the following recommendation which will be updated in the Microsoft references that other people provided on this thread:
>
> *As a general rule, if the number of logical processors is less than 8, use the same number of data files as logical processors. If the number of logical processors is greater than 8, use 8 data files and then if contention continues, increase the number of data files by multiples of 4 (up to the number of logical processors) until the contention is reduced to acceptable levels or make changes to the workload/code.*
>
> http://social.msdn.microsoft.com/Forums/sqlserver/en-US/dceea24c-7a53-4450-94cd-8327b5daa759/what-is-the-best-practice-for-configuring-tempdb

Security

Firewall Rules

Incoming rules:

Service	Protocol	Port	Name	Managed by Windows (*)	Note
SQL	TCP	1764	Instance and VNN Port		
SQL	TCP	5022	Instance SQL Endpoint		User for AAG Replica Communication
WSFC Cluster	TCP	3343	Failover Clusters (TCP-In)	**Yes**	Required during a node join operation
WSFC Cluster	UDP	3343	Failover Clusters (UDP-In)	**Yes**	
WSFC Cluster	TCP	135	Failover Clusters (DCOM-RPC-EPMAP-In)	**Yes**	
WSFC Cluster	TCP	445	Failover Clusters - Named Pipes (NP-In)	**Yes**	
WSFC Cluster	TCP	<Dynamic>	Failover Clusters <RPC Server Programs>	**Yes**	

(*) Rules are automatically created during the feature/role installation

For more information about Microsoft Products Port requirements see MS KB "**Service overview and network port requirements for Windows**" - http://support.microsoft.com/kb/832017/en-us#method70

Antivirus Exclusion

Exclusions for Cluster:

Type	Detail (Path, Extension,...)	Description
Folder	%Systemroot%\Cluster	Cluster folder
Folder	Q:\mscs	Qurom disk

Exclusions for SQL Server:

Type	Detail (Path, Extension,...)	Description
File-name extensions	.mdf .ldf .ndf	SQL Server data files
Process	<installpath>\MSSQL11.<Instance Name>\MSSQL\Binn\SQLServr.exe	SQL process