



DocAve® 6 High Availability

User Guide

Service Pack 4, Cumulative Update 2

Revision H

Issued July 2014

Table of Contents

About DocAve High Availability	4
Complementary Products	5
Submitting Documentation Feedback to AvePoint	6
Before You Begin.....	7
Configuration	7
Agents	7
Required Permissions	7
Required Permissions for SQL Mirroring Method.....	7
Required Permissions for AlwaysOn Availability Group Method	8
Required Permissions for Log Shipping Method.....	10
Required Permissions for Platform Backup Log Shipping.....	11
Getting Started.....	13
Launching High Availability	13
User Interface Overview	14
About Throttle Control.....	15
Configuring Throttle Control Profile	16
Configuring SQL Instance Settings	17
Configuring a High Availability Group.....	18
Using Wizard Mode to Create a Single Farm Mode High Availability Group.....	18
Using Wizard Mode to Create a Standby Farm Mode High Availability Group	22
Using Form Mode to Create a High Availability Group.....	28
Adding SharePoint Components to a Group.....	29
SQL Mirroring, Log Shipping, and Platform Backup Log Shipping.....	29
AlwaysOn Availability Group.....	30
Adding Custom Databases to a Group.....	32
Managing Groups.....	33
About the Dashboard.....	35
Distinguishing Failover-Ready Databases on Dashboard.....	36
Configuration File Settings.....	38

Forcing a Full Synchronization on a Database with a Broken LSN	38
Appendix A – Accessing Hot Key Mode	39
Common Access	39
Home Interface	39
Group Manager Interface	40
Appendix B – SharePoint 2010 Components Supported for High Availability.....	41
Appendix C – SharePoint 2013 Components Supported for High Availability.....	43
Appendix D – Checking the Pre-Scan Rules	45
SQL Mirroring.....	45
AlwaysOn Availability Group.....	46
Log Shipping	47
Platform Backup Log Shipping	47
Common Rules from Health Analyzer.....	48
Notices and Copyright Information	50

About DocAve High Availability

DocAve High Availability is a one-switch disaster recovery solution for Microsoft Windows SharePoint Services and Microsoft SharePoint Server. It enables SQL database replication to a standby environment in order to minimize the perceived downtime for users, or allows you to leverage SQL alias to replicate databases within the same farm to separate SQL instance. As of DocAve 6 Service Pack 4, High Availability supports both of the SharePoint 2010 and SharePoint 2013 environments.

Complementary Products

Many products and product suites on the DocAve 6 platform work in conjunction with one another. The following products are recommended for use with High Availability:

- DocAve Report Center for SharePoint – Examine pain points in the SharePoint infrastructure and report on SharePoint user behavior and changes
- DocAve Replicator for SharePoint – Copying SharePoint content within the same SharePoint farm or from one SharePoint farm to another
- DocAve Granular Backup and Restore – Back up all farm content and restore content down to the item level
- DocAve Platform Backup and Restore – Back up the entire SharePoint environment, including farm-level components

Submitting Documentation Feedback to AvePoint

AvePoint encourages customers to provide feedback regarding our product documentation. You can [Submit Your Feedback](#) on our website.

Before You Begin

Refer to the sections for system and farm requirements that must be in place prior to installing and using DocAve High Availability.

Configuration

In order to use DocAve High Availability, the DocAve 6 platform must be installed and configured properly on your farm. High Availability will not function without DocAve 6 present on the farm.

Agents

DocAve Agents are responsible for running DocAve jobs and interacting with the SharePoint object model. DocAve Agent must be installed on all of the servers in the SharePoint farm. For instructions on installing the DocAve Platform, DocAve Manager, and DocAve Agents, see [DocAve 6 Installation Guide](#).

Required Permissions

To install and use High Availability properly, refer to the following sections for detailed information.

Required Permissions for SQL Mirroring Method

Agent account configured on the SharePoint servers that are included in the Agent group.

1. Local System Permissions
 - Member of the Administrator local group
2. SharePoint Permissions:
 - Member of Farm Administrators group
 - Full Control permission to the User Profile Service Application
3. SQL Permissions: These permissions must be manually configured prior to using DocAve 6 High Availability; they are not automatically configured.
 - Database Role of **db_owner** for SharePoint configuration database, and Central Administration content database
 - Database Role of **db_owner** for all of the databases that you want to perform High Availability jobs on
 - Database Role of **public** to SQL Server
 - Database permission of **View server state** to SQL Server
 - Database role of **db_owner** for the master database or the **View Any Definition** permission to the SQL Server

- Server role of **dbcreator** or the **Alter Any Database** permission or **View Any Definition** permission to the SQL Server
- Permission of **Control Server** to the destination SQL instance
- Server role of **securityadmin** to the destination SQL Server

Agent account configured on the SQL Server

1. Local System Permissions
 - Member of the **Administrators** group
2. SQL Server Permissions
 - Database Role of **db_owner** for SQL Server master database
 - Database Role of **db_owner** for all of the databases you want to perform the High Availability jobs on
 - Database Role of **dbcreator** and **securityadmin** to SQL Server
 - Permissions of **Create Endpoint** and **Alter Login** to SQL Server
 - Server Role of **sysadmin** in the SQL instance

SQL Server Service account configured on the SQL Server

The SQL Server Service account configured on the SQL Server must have the following permissions:

- **Read** and **Write** permissions to the **Temporary Buffer**, which is configured in **Control Panel > Agent Monitor > Configure**
- **Read** and **Write** permissions to the directory of *... \AvePoint\DocAve6\Agent\Jobs*.
- **Read** and **Write** permissions to the sparse file location

VSS Writer account configured on the SQL Server

The VSS Writer account configured on the SQL Server must have the **Read** and **Writer** permissions to the database file location (including the path in file share).

Web application pool account configured on the SQL Server

If the standby application pool account does not exist in the standby SQL Server and does not have the **db_owner** role to the production database selected for High Availability, the standby application will not be readable, that is, the **Warm Standby** feature is not applicable. To ensure that the **Warm Standby** feature works, please grant Web application pool account the server role of **sysadmin** in the standby SQL Server.

Required Permissions for AlwaysOn Availability Group Method

Agent account configured on the SharePoint servers that are included in the Agent group.

1. Local System Permissions
 - Member of the **Administrator** local group
2. SharePoint Permissions:
 - Member of **Farm Administrators** group
 - Full Control permission to the User Profile Service Application
3. SQL Permissions: These permissions must be manually configured prior to using DocAve 6 High Availability; they are not automatically configured.
 - Database Role of **db_owner** for SharePoint configuration database, and Central Administration content database
 - Database Role of **db_owner** for all of the databases that you want to perform High Availability jobs on
 - Database Role of **public** to SQL Server
 - Database permission of **View server state** to SQL Server
 - Permission of **View Any Definition** to SQL Server
 - Server role of **dbcreator** or the **Alter Any Database** permission or **View Any Definition** permission to the SQL Server
 - Permission of **Control Server** to the destination SQL instance
 - Server role of **securityadmin** to the destination SQL Server

Agent account configured on the SQL Server

1. Local System Permissions
 - Member of the **Administrators** group
2. SQL Server Permissions
 - Database Role of **db_owner** for SQL Server master database
 - Database Role of **db_owner** for all of the databases you want to perform the High Availability jobs on
 - Database Role of **dbcreator** and **securityadmin** to SQL Server
 - Permission of **View Server State** to the SQL Server
 - Permission of **Alter Availability Group** to the SQL Server
 - Server Role of **sysadmin** in the SQL instance

SQL Server Service account configured on the SQL Server

The SQL Server Service account configured on the SQL Server must have the **Read** and **Write** permissions to the **Temporary Buffer** configured in **Control Panel > Agent Monitor > Configure**, and the **Read** and **Write** permissions to the directory of ... \AvePoint\DocAve6\Agent\Jobs.

VSS Writer account configured on the SQL Server

The VSS Writer account configured on the SQL Server must have the **Read** and **Writer** permissions to the database file location (including the path in file share).

Web application pool account configured on the SQL Server

If the standby application pool account does not exist in the standby SQL Server and does not have the **db_owner** role to the production database selected for High Availability, the standby application will not be readable, that is, the **Warm Standby** feature is not applicable. To ensure that the **Warm Standby** feature works, please grant Web application pool account the server role of **sysadmin** in the standby SQL Server.

Required Permissions for Log Shipping Method

Agent account configured on SharePoint servers that are included in the Agent group.

1. Local System Permissions
 - Member of the **Administrator** group
2. SharePoint Permissions
 - Member of **Farm Administrators** group
 - Full Control permission to the User Profile Service Application
3. SQL Permissions
 - Database Role of **db_owner** for SharePoint configuration database, and Central Administration content database
 - Database Role of **db_owner** for all of the databases that you want to perform High Availability jobs on
 - Server Role of **public** to SQL Server
 - Database permission of **View server state** to SQL Server
 - Database Role of **db_owner** for the master database or the **View Any Definition** permission to the SQL Server
 - Server role of **dbcreator** or the **Alter Any Database** permission or **View Any Definition** permission to the SQL Server
 - Permission of **Control Server** to the destination SQL Server
 - Server role of **securityadmin** to the destination SQL Server

Agent account configured on the SQL Server

1. Local System Permissions
 - Member of the **Administrators** group

2. SQL Server Permissions

- Database Role of **db_owner** for SQL Server master database
- Database Role of **db_owner** for all of the databases you want to perform the High Availability jobs on
- Server Role of **dbcreator**, **processadmin**, **securityadmin** to SQL Server
- The **Control Server** permission in the destination SQL instance
- Server Role of **sysadmin** in the SQL instance

SQL Server Service account configured on the SQL Server

The SQL Server Service account configured on the SQL Server must have **Read** and **Write** permissions to the **Temporary Buffer**, which is configured in **Control Panel > Agent Monitor > Configure**, and the **Read** and **Write** permissions to the directory of ...*\AvePoint\DocAve6\Agent\Jobs*.

VSS Writer account configured on the SQL Server

The VSS Writer account configured on the SQL Server must have the **Read** and **Writer** permissions to the database file location (including the path in file share).

Web application pool account configured on the SQL Server

If the standby application pool account does not exist in the standby SQL Server and does not have the **db_owner** role to the production database selected for High Availability, the standby application will not be readable, that is, the Warm Standby feature is not applicable. To ensure that the Warm Standby feature works, please grant Web application pool account the server role of **sysadmin** in the standby SQL Server.

Required Permissions for Platform Backup Log Shipping

Agent account configured on SharePoint servers that are included in the Agent group.

1. Local System Permissions
 - Member of the **Administrator** group
2. SharePoint Permissions
 - Member of **Farm Administrators** group
 - Full Control permission to the User Profile Service Application
3. SQL Permissions
 - Database Role of **db_owner** for SharePoint configuration database, and Central Administration content database
 - Database Role of **db_owner** for all of the databases that you want to perform High Availability jobs on

- Server Role of **public** to SQL Server
- Database permission of **View server state** to SQL Server
- Database Role of **db_owner** for the master database or the **View Any Definition** permission to the SQL Server
- Server role of **dbcreator** or the **Alter Any Database** permission or **View Any Definition** permission to the SQL Server
- Permission of **Control Server** to the destination SQL Server
- Server role of **securityadmin** to the destination SQL Server

Agent account configured on the SQL Server

1. Local System Permissions
 - Member of the **Administrators** group
2. SQL Server Permissions
 - Database Role of **db_owner** for SQL Server master database
 - Database Role of **db_owner** for all of the databases you want to perform the High Availability jobs on
 - Server Role of **dbcreator**, **processadmin**, **securityadmin** to SQL Server
 - The **Control Server** permission in the destination SQL instance
 - Server Role of **sysadmin** in the SQL instance

SQL Server Service account configured on the SQL Server

The SQL Server Service account configured on the SQL Server must have **Read** and **Write** permissions to the **Temporary Buffer**, which is configured in **Control Panel > Agent Monitor > Configure**, and the **Read** and **Write** permissions to the directory of ... \AvePoint\DocAve6\Agent\Jobs.

VSS Writer account configured on the SQL Server

The VSS Writer account configured on the SQL Server must have the **Read** and **Writer** permissions to the database file location (including the path in file share).

Web application pool account configured on the SQL Server

If the standby application pool account does not exist in the standby SQL Server and does not have the **db_owner** role to the production database selected for High Availability, the standby application will not be readable, that is, the Warm Standby feature is not applicable. To ensure that the Warm Standby feature works, please grant Web application pool account the server role of **sysadmin** in the standby SQL Server.

Getting Started

Refer to the sections below for important information on getting started with High Availability.

Launching High Availability

To launch High Availability and access its functionality, follow the instructions below.

1. Log in to DocAve. If you are already in the software, click the **DocAve** tab. The **DocAve** tab displays all modules on the left side of the window.
2. Click **Data Protection** to view the Data Protection modules.
3. Click **High Availability** to launch this module.

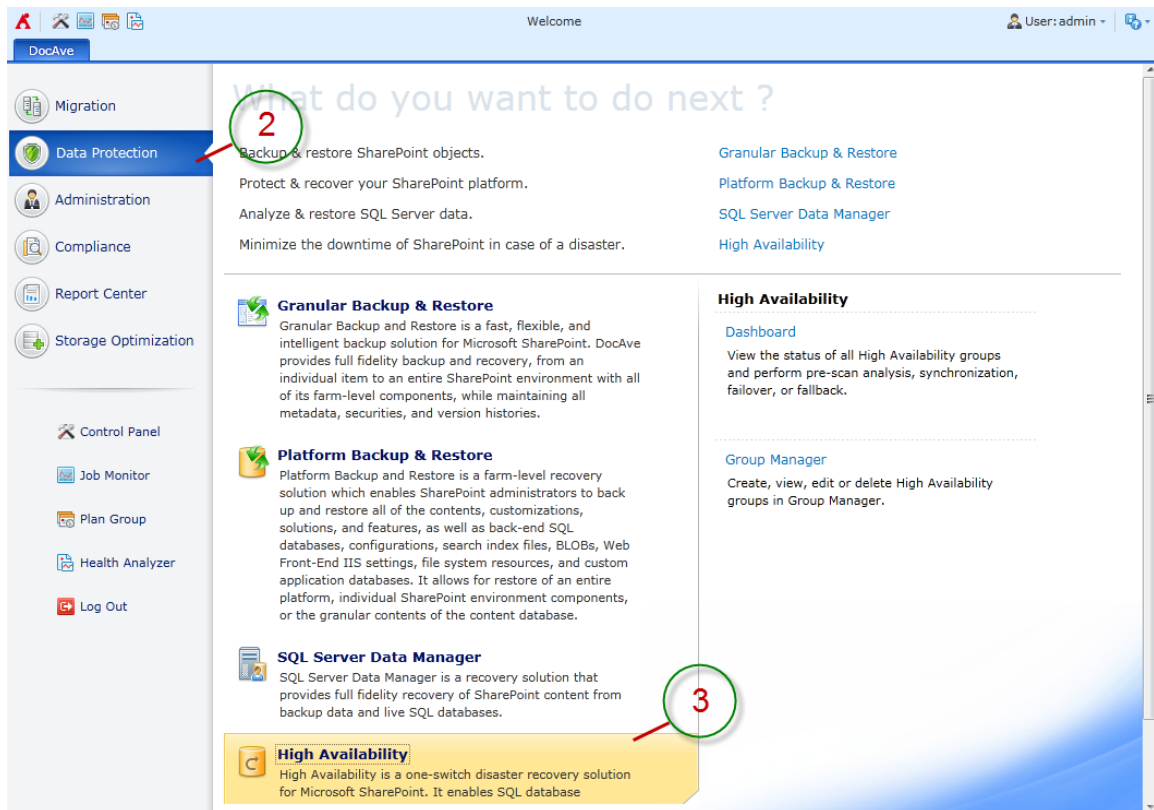


Figure 1: DocAve launch window.

User Interface Overview

After clicking **High Availability**, the High Availability user interface launches with the **Home** tab active.

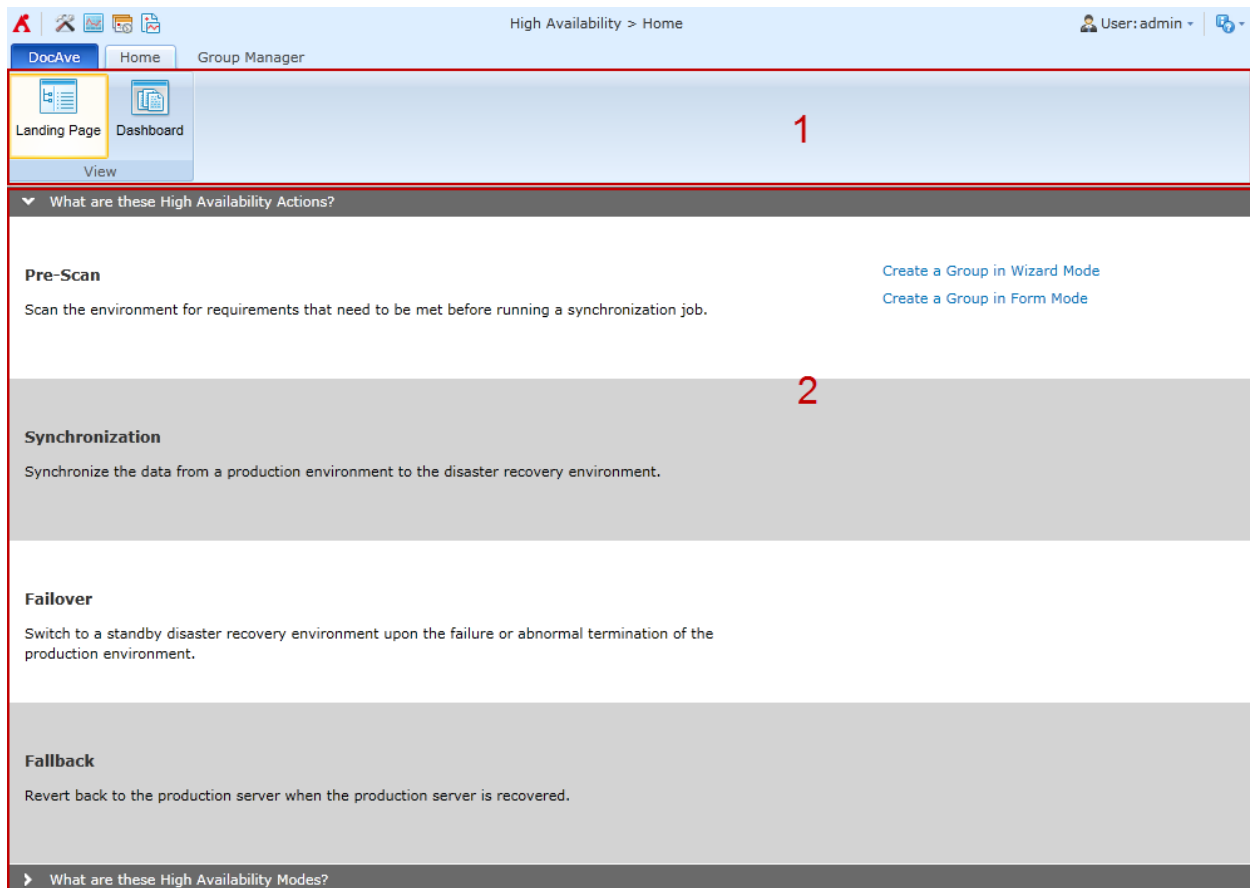


Figure 2: High Availability user interface.

1. The **ribbon** shows the available actions for High Availability.
2. The **workspace** shows explanations and the content that is used during the configuration of actions performed in DocAve High Availability.

About Throttle Control

For low-bandwidth networks, throttle control is available to set boundaries and limitations on jobs. Throttle Control ensures that jobs do not cause a noticeable degradation in network performance when running these jobs.

To access the Throttle Control interface, select the **Group Manager** tab. On the **Group Manager** tab, select **Throttle Control** on the ribbon.

The **Throttle Control** interface displays all of the Throttle Control profiles that you have previously created.

In this interface, you can change the number of Throttle Control profiles displayed per page and the order in which they are displayed. To change the number of Throttle Control profiles displayed per page, select the desired number from the **Show rows** drop-down menu in the lower right-hand corner.

Customize how these Throttle Control profiles are displayed in a number of different ways:

- **Search** – Filters the previously created Throttle Control profiles displayed by the keyword you designate. The keywords must be contained in a column value. At the top of the Throttle Control profiles viewing pane, enter the keyword for the Throttle Control profiles you wish to display. You can select **Search all pages** or **Search current page** to define the search scope. **Search all pages** means that the Throttle Control profiles in all of the pages whose names and descriptions contain the keywords will be displayed; while **Search current page** means that only the Throttle Control profiles in the current page whose names and descriptions contain the keywords will be displayed. Note that the search function is not case sensitive.
- **Manage columns** (⊕) – Manage which columns are displayed in the list so that only the information you want to see is displayed. Click the manage columns button (⊕), and then check the checkbox next to the column name to have that column shown in the list.
- **Hide the column** (⊖) – Click the hide the column button (⊖) in the column title to hide the column.

Perform the following actions in the **Throttle Control** interface:

- **Create** – Click **Create** on the ribbon to create a new Throttle Control profile. For details on creating a new Throttle Control profiles, see [Configuring Throttle Control Profile](#).
- **Edit** – Click **Edit** on the ribbon to change the configurations for the selected Throttle Control profile. For details on editing configurations for Throttle Control profile, see [Configuring Throttle Control Profile](#).
- **View Details** – Select a Throttle Control profile, and then click **View Details** on the ribbon to view the detailed information about the settings on the selected profile. You can click **Edit** on the **View Details** interface to edit the profile settings.
- **Delete** – Click **Delete** on the ribbon. A confirmation window will pop up and ask whether you are sure that you want to proceed with the deletion. Click **OK** to delete the selected

Throttle Control profile, or click **Cancel** to return to the **Throttle Control** interface without deleting the selected Throttle Control profiles.

Configuring Throttle Control Profile

To create a new Throttle Control profile, select **Create** on the ribbon. To modify a previously configured Throttle Control profile, select the Throttle Control profile, and then select **Edit** on the ribbon.

In the **Create Throttle Control Profile** interface or **Edit Throttle Control Profile** interface, configure the following settings:

1. **Profile Name** – Enter a profile name and the description for the Throttle Control profile you are about to create or edit.
2. **Run During** – Specify the job running speed during working hours and non-working hours, respectively. Enter a positive integer into the text box and select the **KB/Second** or **MB/Second** from the drop-down menu to specify the desired unit of network bandwidth.
3. **Define Work Schedule** – Define the working schedule of the Throttle Control profile you are creating/editing. Specify the **Working hours** and **Working days**.
4. Select **Save** to save the configurations and return to the Throttle Control interface, or select **Cancel** to return to the Throttle Control interface without saving any changes.

Configuring SQL Instance Settings

Optionally, you can configure SQL Instance Settings to specify a master key password for the SQL Server instance. If you do not configure the SQL Instance Settings, a default master key password will be automatically created when performing the High Availability actions related to SQL Mirroring method, or performing the synchronization for the databases encrypted by Transparent Data Encryption (TDE).

To configure the SQL instance settings, complete the following steps:

1. To access the **SQL Instance Settings** interface, select the **Group Manager** tab, and then select the **SQL Instance Settings** button on the ribbon. The **SQL Instance Settings** interface appears.
2. Select an Agent from the **SQL Agent** drop-down list. All of the selected Agent's SQL instances are displayed in the table.

Alternatively, you can use the search function on the left-top of the SQL instance table to search the SQL instance.

3. Under the **Edit** column across the specified SQL instance row, click **Configure** to configure the specified SQL instance settings.
4. In the **Edit** page, enter the master key password for this SQL instance into the **Master Key Password** text box. The specified master key password must meet the Windows password policy requirements of the computer that is running the instance of SQL Server.
5. Optionally, you can configure the location for the data and logs in the SQL instance separately. The default value is the default database locations for data and logs of this SQL instance. To configure the location for the SQL Cluster, make sure to enter the available shared path of your SQL Cluster environment.
6. Select **Save** to save the configurations you have made and go back to the **SQL Instance Settings** interface. Select **Cancel** to go back without saving any changes.

Configuring a High Availability Group

The High Availability Group feature provides an integrated wizard/form interface for creating and editing High Availability mode and settings for High Availability functions such as Pre-Scan, Synchronization, Failover, and Fallback.

For more information on creating a High Availability Group, refer to the following instructions.

Using Wizard Mode to Create a Single Farm Mode High Availability Group

Use Single farm mode when you want to ensure the availability of a single SharePoint farm. You are recommended to add the source Configuration database, and Central Administration database to the High Availability Group for Single farm mode.

Wizard Mode provides you with step-by-step guidance on how to configure a new group. To configure a Single Farm mode High Availability group using Wizard Mode, complete the following steps:

***Note:** A red * in the user interface indicates a mandatory field or step.

1. *Enter a **Group Name** and optional **Description**, if desired. Select **Next**.
2. The **General Settings** interface appears, where you will specify the **High Availability Mode**, **Sync Method**, **Failover Method**, and **Failover Option**.
 - a. **High Availability Mode** – Choose which High Availability Mode you are about to use according to your environment. In this condition, select **Single farm** option.
 - b. **Production Farm** – Select the SharePoint farm that you want to perform the High Availability from the drop-down list.
 - c. **Production Agent Group** – Select a production agent group from the drop-down list for load balancing and job performance.
 - d. **Sync Method** – Select the **SQL Mirroring**, **AlwaysOn Availability Group**, **Log Shipping**, or **Platform Backup Log Shipping** from the sync method drop-down list. Refer to the instructions below for more information on the four Sync Methods.
 - **Log Shipping** performs a replication at the log level. Log Shipping method sends transaction log backups from a primary database on a primary server instance to the secondary databases on separate secondary server instances.

***Note:** If the node has been included in a Platform Backup plan with the **Use DocAve Platform Backup and Restore as the only backup method for SharePoint database** option selected, it cannot be added to the High Availability Group with Log Shipping method.

- **SQL Mirroring** performs a replication at the transaction level. SQL Mirroring is implemented on a per-database basis and works only with databases that use the full recovery model. The SQL Mirroring method creates endpoints in both the production server and the standby server by DocAve certificate.

***Note:** If you use SQL Mirroring as the High Availability Sync Method to synchronize a large number of databases, DocAve High Availability may only synchronize some of the databases to the standby server. If this occurs, it is due to limitations of your SQL Server hardware. The factors that affect the overall performance and limit the number of databases that can be mirrored on the servers include:

- Amount of RAM available on the principal and mirror servers.
- Processing power of the principal and mirror servers.
- Bandwidth available for the I/O subsystem on the principal and mirror servers
- Network bandwidth between the two servers.
- Volume of transactions on the principal database that generate transaction log records (that is, transactions that change the database in any way).

For more details on the issues above, see <http://support.microsoft.com/kb/2001270>.

- **AlwaysOn Availability Group**, introduced in SQL Server 2012 as an enterprise-level alternative to SQL Mirroring, maximizes the availability of a set of user database for an enterprise. An availability group supports a failover environment for a discrete set of user databases, known as availability databases, which failover together. An availability group supports a set of read-write primary databases and one to four sets of corresponding secondary databases. Optionally, secondary databases can be made available for read-only access and/or some backup operations.

***Note:** DocAve recommends not applying the SQL Mirroring method and the AlwaysOn Availability Group method on the same SQL instance.

- **Platform Backup Log Shipping** performs a replication based on the data of the Platform Backup jobs. This High Availability sync method collaborates with the Platform Backup function to replicate the backup data of the specified components to the destination. This method requires users to have Platform Backup data of the specified components prepared.

e. **Failover Method** – Select the failover method according to the sync method you specified.

- If you select the **Log Shipping** or **Platform Backup Log Shipping** as the **Sync Method**, only the **SQL Alias** method is available.

- If you select the **SQL Mirroring** as the **Sync Method**, you can select the **SQL Alias** method or **SharePoint Failover Server** method from the drop-down list. If you select the **SharePoint Failover Server** method, DocAve will use the failover server configured in SharePoint environment for the mirrored databases.
***Note:** The **SharePoint Failover Server** is available only if you select the **SQL Mirroring** as the **Sync Method**.
 - If you select the **AlwaysOn Availability Group** as the **Sync Method**, the **AlwaysOn Availability Group** is available. An availability group fails over at the level of an availability replica.
- f. **Failover Options** – Select your desired failover option according to the **Sync Method** and the **Failover Method** that you specified.
- If you select **SQL Mirroring** as the **Sync Method**, and the **SQL Alias** as the **Failover Method**, the **Keep SQL mirroring relationship** option is selected by default to keep the SQL mirroring session in the event of a disaster.
 - If you select the **Log Shipping** or **Platform Backup Log Shipping** as the **Sync Method**, and the **SQL Alias** as the **Failover Method**, you can select the **Keep database in read-only mode** option to keep the read-only status to the secondary database. This option is specific for the Log Shipping and Platform Backup Log Shipping method.
***Note:** The **Failover Options** are not available for the other combinations of the Sync Method and Failover Method
- g. **Would you like to synchronize the stub database?** – Specify whether or not to synchronize the related stub database together with the selected farm components to the standby environment. To synchronize the stub database to the standby environment, select the **Sync related stub database for the selected components** option.
- h. Select **Next**. The **Add to Group** interface appears.
3. ***In the Add to Group** interface, select from the farm tree the SharePoint components to add to this High Availability Group. Refer to the [Adding SharePoint Components to a Group](#) for more information. For detailed information on which SharePoint components are supported in Single farm mode High Availability, see [Appendix B – SharePoint 2010 Components Supported for High Availability](#).
4. **SQL Alias Settings** – If you set the **SQL Alias** as the **Failover Method**, you are required to configure the **SQL Alias Settings**. To configure the SQL Alias settings, complete the following steps:
- a. Click **Edit** under the **Port** column to specify a port number for the corresponding SQL Server instance.

- b. In the pop-up window, enter an integer into the **Port** text box to specify a port for the SQL Alias, or you can select the **Dynamically determine port** checkbox to use the SQL Browser to determine the port dynamically.
 - c. Click **Save** to save the configuration. Click **Cancel** to go back without saving any changes.
5. **SharePoint Server Settings** – Select the SharePoint servers from the table, and the SQL Alias specified above will be created on your selected SharePoint Servers separately when running a failover job.

***Note:** DocAve recommends selecting all of the SharePoint servers listed in the table to make sure that the SQL Alias information can be added to all of the SharePoint servers in your farm.
6. **Throttle Control** – Select a throttle control profile from the drop-down list to define the job running speed during the working hours and non-working hours. You can click **View** to go to the **Throttle Control** interface to view the selected throttle control profile. For detailed information on throttle control profile, see [About Throttle Control](#).
7. **Schedule Selection** – Define whether or not to configure a custom schedule to run the synchronization jobs.
 - **No schedule** – If you select this option, you are required to manually run the synchronization jobs based on this group.
 - **Configure the schedule myself** – If you select this option, you are required to configure the custom schedule settings for the synchronization jobs. Click **Add Schedule** to set up a schedule. The **Add Schedule** interface appears. Select a type of the recurring synchronization job. The two options are **Initialize/Reinitialize** and **Synchronize**.
 - **Initialize/Reinitialize** – It will replicate the entire data from the production server to the standby server. Specify the start time for this type of synchronization job in the **Range of Recurrence** field. Click **Save**, and then this schedule will display in the table. Click **Cancel** to go back to the **Advanced Settings** page.
 - **Synchronize** – It will make the data changes synchronous between production database and standby database based on the data stored in the production databases at the scheduled time. Enter an integer into the interval text box in the **Schedule Settings** field, and select **Minutes, Hours, Days, Weeks, or Months**. The data changes will be synchronized based on this schedule. Specify the time range of the job occurrence in the **Range of Recurrence** field. Specify a start time that cannot be earlier than the current time, and then designate when to stop synchronize the data in this group.

***Note:** If you select the **SQL Mirroring** or **AlwaysOn Availability Group** as the **Sync Method** in the **Single farm** mode, this schedule option will only take effect when you select the **Include New** checkbox in the **Add to Group** data tree and there are newly added databases in the selected nodes. Otherwise, the scheduled jobs will run in vain.

- **No end date** – Select this option to synchronize the data changes from the production database to the standby database until you stop it manually.
 - **End after _ occurrences** – Select this option to have this type of synchronization job stop after the defined occurrence that you configured in the text box.
 - **End by** – Select this option to specify the exact date and time for the data changes stopped being synchronized. If desired, you can change the time zone by clicking the current time zone link to select the desired one.
8. **Notification** – You can select a previously created notification profile from the drop-down list, or you can click the New Notification Profile link to create a new notification profile. For detailed instructions, refer to the [DocAve 6 Control Panel Reference Guide](#). Click **View** to view the information of the notification profile.
 9. Select **Next**. The **Pre-Scan** interface appears.
 10. Select **Scan** to scan the environment for the requirements that need to be met before running a synchronization job. Or you can skip this step, save the group configurations, and then run the Pre-Scan job for this group in Dashboard.
 11. Select **Next**. The **Overview** interface appears. You can Select **Edit** in the field that you want to modify, and then you are brought to the wizard to make your changes.
 12. Select **Finish** after you have confirmed your configurations. Select **Cancel** to go back to the Group Manager interface without saving any changes.

Using Wizard Mode to Create a Standby Farm Mode High Availability Group

Use Standby farm mode when you want to ensure the availability by using a standby farm to synchronize the SharePoint components between one SharePoint instance and your disaster recovery (DR) one. In order to accomplish this, a secondary farm must be configured with a DocAve High Availability Agent installed.

Standby farm mode requires you to specify the production farm and the standby farm respectively. Once after the failover, you need to use the port of the standby Web application/service application to access the database.

Wizard Mode provides you with step-by-step guidance on how to configure a new group. Follow the instructions below to configure a Standby farm mode High Availability group using Wizard Mode.

***Note:** A red * in the user interface indicates a mandatory field or step.

1. *Enter a **Group Name** and optional **Description**, if desired. Select **Next**.

2. The General Settings interface appears, where you will specify the High Availability Mode, Sync Method, and Failover Options.
 - a. **High Availability Mode** – Choose which High Availability Mode you are about to use according to your environment. In this condition, select **Standby farm** option.
 - b. **Production Farm** – Select the production farm that you want to perform the High Availability from the drop-down list.
 - c. **Production Agent Group** – Select a production Agent group from the drop-down list for load balancing and job performance.
 - d. **Standby Farm** – Select the Standby farm from the drop-down list as you wish.
 - e. **Standby Agent Group** – Select a standby Agent group from the drop-down list for load balancing and job performance improvement.
 - f. **Sync Method** – Select the **SQL Mirroring**, **AlwaysOn Availability Group**, **Log Shipping**, or **Platform Backup Log Shipping** from the **Sync Method** drop-down list. Refer to the instructions below for more information on the four Sync Methods.
 - **Log Shipping** performs a replication at the log level. Log Shipping method sends transaction log backups from a primary database on a primary server instance to one or more secondary databases on separate secondary server instances.
***Note:** if the node has been included in a Platform Backup plan, it cannot be added to the High Availability Group with Log Shipping method.
 - **SQL Mirroring** performs a replication at the transaction level. SQL Mirroring is implemented on a per-database basis and works only with databases that use the full recovery model. The SQL Mirroring method creates endpoints in both the production and standby servers by DocAve certificate.
***Note:** If you use SQL Mirroring as the High Availability Sync Method to synchronize a large number of databases, DocAve High Availability may only synchronize some of the databases to the standby server. If this occurs, it is due to limitations of your SQL Server hardware. The factors that affect the overall performance and limit the number of databases that can be mirrored on the servers include:
 - Amount of RAM available on the principal and mirror servers.
 - Processing power of the principal and mirror servers.
 - Bandwidth available for the I/O subsystem on the principal and mirror servers
 - Network bandwidth between the two servers.

- Volume of transactions on the principal database that generate transaction log records (that is, transactions that change the database in any way).

For more details on the issues above, see <http://support.microsoft.com/kb/2001270>.

- **AlwaysOn Availability Group**, introduced in SQL Server 2012 as an enterprise-level alternative to SQL Mirroring, maximizes the availability of a set of user database for an enterprise. An availability group supports a failover environment for a discrete set of user databases, known as availability databases, which failover together. An availability group supports a set of read-write primary databases and one to four sets of corresponding secondary databases. Optionally, secondary databases can be made available for read-only access and/or some backup operations.

***Note:** DocAve recommends not applying the SQL Mirroring method and the AlwaysOn Availability Group method on the same SQL instance.
 - **Platform Backup Log Shipping** performs a replication based on the data of the Platform Backup jobs. This High Availability sync method collaborates with the Platform Backup function to replicate the backup data of the specified components to the destination. This method requires users to have Platform Backup data of the specified components prepared.
- g. **Would you like to make the standby environment read-only?** – If you want to make the standby environment readable, select the **Enable read-only view after synchronization** option. Pay attention to the following issues:
- If you are about to synchronize the content database with BLOB data and related stub database together to the standby farm, the Agent account in the standby farm must have the **db_owner** role in the production stub database in order to make sure the standby stub files are readable.
 - If you have selected the SQL Mirroring sync method and your system is currently under I/O-heavy load, read-only view for SQL Mirroring method may take longer than expected and have an impact on the transaction throughput for your system.
 - If you have selected the AlwaysOn Availability Group as the sync method, DocAve will automatically set the **Readable secondary** option in the SQL Server for the secondary replicas of the AlwaysOn Availability Group to **Yes**.

***Note:** After the synchronization for SharePoint 2010 environment, the standby site collection will appear as writeable, but it is actually read-only. The **Database Read-Only** property appears as “**No**”. However, the database is actually read-only in the AlwaysOn Availability Group Secondary Replica.

- h. **Failover Options** – Select the desired failover options according to the **Sync Method** you specified.
- **Bring production database offline** – DocAve will bring the primary databases offline and bring the secondary database online once after the failover.
***Note:** This option does not support service application databases.
 - **Keep SQL mirroring relationship** – DocAve will keep the SQL mirroring session in the event of a disaster. This option is specific for the SQL Mirroring method.
 - **Keep database in read-only mode** – DocAve will keep the database in read-only status to the secondary database. This option is specific for the Log Shipping and Platform Backup Log Shipping method.
***Note:** This option does not support the service application databases and stub databases.
***Note:** If you have selected the **Log Shipping** method and the **Keep database in read-only mode** option for the High Availability group and you want to keep the changes in both of the production farm and standby farm after the Failover, make sure you remove the schedules for the replication group and disable the connection in Distributed File System (DFS) Management.
- i. **Would you like to synchronize the stub database** – Specify whether or not to synchronize the related stub database together with the selected farm components to the standby environment. To synchronize the stub database to the standby environment, select the **Sync related stub database for the selected components** option.
- j. Click **Next**. The **Add to Group** interface appears.
3. ***In the Add to Group** interface, select from the farm tree the SharePoint components to add to this High Availability Group. Refer to [Adding SharePoint Components to a Group](#) for more information. For detailed information on which SharePoint components are supported in Standby farm mode High Availability, see [Appendix B – SharePoint 2010 Components Supported for High Availability](#).
- *Note:** DocAve High Availability cannot synchronize the Partition Mode setting of the service application to the standby environment. If your added service application is in Partition Mode, The service application will be synchronized to the standby environment without the Partition Mode setting.
4. **Throttle Control** – Select a throttle control profile from the drop-down list to define the job running speed during the working hours and non-working hours. You can click **View** to go to the **Throttle Control** interface to view the selected throttle control profile. For detailed information on throttle control profile, refer to [About Throttle Control](#).
5. **Schedule Selection** – Define whether or not to configure a custom schedule to run the synchronization jobs.

***Note:** For the Log Shipping and Platform Backup Log Shipping sync methods, it is recommended to configure the synchronization schedules to keep the databases in sync.

- **No schedule** – If you select this option, you are required to manually run the synchronization jobs based on this group.
 - **Configure the schedule myself** – If you select this option, you are required to configure the custom schedule settings for the synchronization jobs. Click **Add Schedule** to set up a schedule. The **Add Schedule** interface appears. Select a type of the recurring synchronization job. The two options are **Initialize/Reinitialize** and **Synchronize**.
 - **Initialize/Reinitialize** – It will replicate the entire data from the production server to the standby server. Specify the start time for this type of synchronization job in the **Range of Recurrence** field. Click **OK**, and then this schedule will display in the table. Click **Cancel** to go back to the **Advanced Settings** page.
 - **Synchronize** – It will make the data changes synchronous between production database and standby database based on the data stored in the production databases at the scheduled time. Enter an integer into the interval text box in the **Schedule Settings** field, and select by **Hours, Days, Weeks, or Months**. The data changes will be synchronized based on this schedule. Specify the time range of the job occurrence in the **Range of Recurrence** field. Specify a start time earlier than the current time, and then designate when to stop synchronize the data in this group.
 - **No end date** – Select this option to synchronize the data changes from the production database to the standby database until you stop it manually.
 - **End after _ occurrences** – Select this option to have this type of synchronization job stop after the defined occurrence that you configured in the text box.
 - **End by** – Select this option to specify the exact date and time for the data changes stopped being synchronized. If desired, you can change the time zone by clicking the current time zone link to select the desired one.
6. **Notification** – You can select a previously created notification profile from the drop-down list, or you can click the New Notification Profile link to create a new notification profile. For detailed instructions, refer to the [DocAve 6 Control Panel Reference Guide](#). Click **View** to view the information of the notification profile.
 7. Select **Next**. The **Pre-Scan** interface appears.
 8. Select **Scan** to scan the environment for the requirements that need to be met before running a synchronization job. Or you can skip this step, save the group configurations, and go to the **Dashboard** to run the Pre-Scan job.
 9. Select **Next**. The **Overview** interface appears. You can select **Edit** in the field that you want to modify, and then you are brought to the wizard to make your changes.

10. Select **Finish** after you have confirmed your configurations. Select **Cancel** to go back to the Group Manager interface without saving any changes.

Using Form Mode to Create a High Availability Group

Form Mode is recommended for the users who are familiar with creating the High Availability Groups. To create a High Availability Group in Form Mode, click **Group Manager > Create > Form Mode**. Refer to the [Using Wizard Mode to Create a Single Farm Mode High Availability Group](#), or [Using Wizard Mode to Create a Standby Farm Mode High Availability Group](#) for detailed information according to your desired High Availability Mode.

***Note:** If this is your first time to create a High Availability Group or you think you would benefit from the descriptions of each group component, AvePoint recommends you use the Wizard Mode.

Adding SharePoint Components to a Group

Refer to the section below for more information on how to add the SharePoint components to a High Availability group according to different sync methods.

SQL Mirroring, Log Shipping, and Platform Backup Log Shipping

To add the SharePoint components to High Availability Group using the Sync Method of SQL Mirroring, Log Shipping, or Platform Backup Log Shipping, complete the following steps:

1. Click the farm node to load the tree, or click the **Actions** button to choose **Show selected nodes** or **Show all nodes** from the drop-down list. You can also enter the keyword of the desired nodes into the text box on top of the farm tree. Click the magnifying glass (🔍) to search the nodes according to the keyword you designate.
2. Select the desired nodes and then select the **Add to Group** button.

***Note:** Optionally, you can add any desired databases in the SQL Servers to the farm component tree and then add them to the group to achieve the High Availability function. For more instructions on adding custom databases to group, refer to [Adding Custom Databases to a Group](#).

3. If there are databases being selected, you must specify and test the **Destination SQL instance**. Enter the destination SQL instance name into the text box. Click **Test**. After successfully tested the entered SQL instance, the default location for data files and log files will be displayed in the **Database Location** field. You can customize the database location as necessary.

***Note:** If you want to use the shared path or file share path as the database location, change the logon user of the SQL VSS Writer Service to domain account. In addition, make sure the logon user (domain account) has the following two permissions:

- Read and Write permissions to the shared path/file share path.
- Permissions required by the Agent accounts configured on SQL Server for the corresponding sync method.

***Note:** If you have selected the SQL Mirroring method and selected the option to make the standby environment read-only after the synchronization, you are required to configure the **Database Snapshot Location**. Make sure the version of the entered SQL Server/Instance support the database snapshot feature. For details, refer to the Microsoft TechNet article on [Features Supported by the Editions of SQL Server 2012 – High Availability](#).

4. Customize the path for storing the standby database and database snapshot as necessary.




***Note:** Database snapshot is the key to the read-only view of the standby environment. The database snapshots created by DocAve High Availability is attached with the “**_HASnapshot**” as the postfix. The database snapshot will be generated based on the High Availability Synchronization job schedule to keep the data most updated. The database snapshot in the

standby farm is almost the same with the production database, including stub related information. The stubs in the standby farm can be accessed too, but only readable. AvePoint recommends not storing the database snapshot in the same volume of the database location.

5. If you have selected to synchronize the related stub database, you are required to specify a standby physical device for each device type for synchronizing the BLOB data from the production physical devices. Select a physical device from the drop-down list for each device type. You can click **View** to view the details of your selected physical device. Also, you can select **New Physical Device** from the drop-down list to create a new physical device in the Control Panel. For details on creating a physical device, refer to [DocAve 6 Control Panel Reference Guide](#).

***Note:** Make sure the selected standby physical device is in the same storage type as the physical devices in the production environment in case that the SharePoint contents may not be able to access after the Failover.


6. Click **Save** to save the configurations and go back to the **Add to Group** interface. Click **Cancel** to go back to the **Add to Group** interface without saving any configurations.
7. The selected nodes displays in the table below. You can perform the following actions:

- Click the settings () button besides the node name to edit the destination settings.
***Note:** If you have selected to synchronize the stub database for the selected components and configured the standby physical device for the destination SQL instance, the **Physical Device Mapping** field will appear. This field will display the specified standby physical device for each production physical device. You can choose to change the standby physical device to the physical device of your existing BLOB storage.
- Click the delete () button on the **Delete** column header to remove all of the added nodes from the group or click the delete () button next to the corresponding node to remove the node from the High Availability Group.

***Note:** If there are Web applications in the selected nodes for the Standby farm mode High Availability Group, DocAve High Availability recommends you to view and configure the settings of the standby Web applications first before you go to the next step.

AlwaysOn Availability Group

To add the SharePoint components to High Availability Group using the Sync Method of AlwaysOn Availability Group, complete the following steps:


1. Click the farm node to load the tree, or click the **Actions** button to choose **Show selected nodes** or **Show all nodes** from the drop-down list. As well, you can enter the keyword of the desired nodes into the text box on top of the farm tree. Click the magnifying glass () button to search the nodes according to the keyword you designate.



2. Select the desired nodes and then select the **Add to Group** button. You are brought to an interface to configure the AlwaysOn Availability Group settings for the SQL instance of the selected nodes.

***Note:** Optionally, you can add any desired databases in the SQL Servers to the farm component tree and then add them to the group. For more instructions on adding custom databases to group, refer to [Adding Custom Databases to a Group](#).

3. All of the corresponding databases of the selected nodes in this SQL instance displays in the **Databases** field.
4. Select an existing AlwaysOn Availability Group from the **AlwaysOn Availability Group** drop-down list to add the selected database to.
5. Select one or more secondary availability replicas as the synchronization destinations from the table in the right of the **Secondary Replicas for Synchronize** field.
6. Choose a replica from the **Primary Replica for Failover** drop-down list as the primary replica to host the primary database after the failover.
7. Enter the availability group listener name into the **Availability Group Listener** text box for the connectivity of the specified availability group.
 - In the **Single farm** mode configurations, you can enter the SQL Alias or IP address of the listener into the text box.
 - In the **Standby farm** mode configurations, you will not configure the Availability Group Listener and this field will not appear on the interface.

If you have selected the stub database to add to the group, you are required to specify a **Standby Physical Device** to synchronize the BLOB data in the production physical devices.

8. Select **Save** to save the configurations and go back to the **Add to Group** interface. Select **Cancel** to go back to the **Add to Group** interface without saving any changes.
9. The selected nodes displays in the table below. You can perform the following actions:
 - Click the settings () button besides the node name to edit the destination settings.


***Note:** If you have selected to synchronize the stub database for the selected components and configure the standby physical device for the destination SQL instance, the **Physical Device Mapping** field will appear. This field will display the specified standby physical device for each production physical device. You can choose to edit the standby physical device.
 - Click the delete () button on the **Delete** column header to remove all of the added nodes from the group or click the delete () button next to the corresponding node to remove the corresponding node from the High Availability Group.

***Note:** If there are Web applications in the selected nodes for the Standby farm mode High Availability Group, DocAve High Availability recommends that you view and configure the settings of the standby Web applications first before you go to the next step.

Adding Custom Databases to a Group

To implement the High Availability capability to the databases that are in your SQL Server but not your SharePoint components, you can manage the **Custom Database** node to add the database to the farm component tree. Synchronize the database to the standby environment by adding it to the High Availability group. Therefore, if the custom database is crushed, the standby database will work after failover.

To add custom databases to the farm component tree, complete the following steps:

1. In the expanded farm component tree, click the **Manage** button beside the **Custom Database** node. The **Custom Database** window appears.
2. Click the **Custom Database** ( Custom Database) button to display the SQL servers that have DocAve Agent installed. Click the Agent name to load all of the SQL instances.
3. Click the SQL instance name to display all of the databases that reside in the SQL instance.
4. Select the desired databases to add to the farm component tree.
5. Click **Save** to save your tree selection and exit.
6. In the farm component tree, the databases you selected are displayed under the **Custom Database** node.

Managing Groups

For the created High Availability Groups, you can use Group Manager to manage them. After launching the DocAve High Availability interface, click **Group Manager** next to the **Home** tab. In the **Group Manager** interface, all of the previously created groups are displayed in the main display pane.

In this interface, you can change the number of the groups displayed per page. To change the number of the groups displayed per page, select the desired number from the **Show rows** drop-down menu in the lower right-hand corner. To sort the groups, click the column heading such as **Group Name**, **Description**, **Production Farm**, **High Availability Mode**, **Sync Method**, **Failover Method**, and **Last Modified Time**.

Customize how these groups are displayed in a number of different ways:

- **Search** – Filter the groups to be displayed by the keyword entered in the **Input Keyword** text box (the keyword is not case sensitive). The keyword must be contained in a column value. You can select **Search all pages** or **Search current page** to define the search scope. **Search all pages** means that the groups in all of the pages whose column contains the keywords will be displayed; while **Search current page** means that only the groups in the current page whose column contains the keywords will be displayed.
- **Filter the column** (🔍) – Filter which groups in the list is displayed. Unlike Search, you can filter whichever groups using designated criteria, rather than search based on a keyword. Hover over the **Production Farm**, **High Availability Mode**, **Sync Method**, or **Failover Method** column heading, and then select the filter the column button (🔍). Select the checkboxes to have the groups in that category shown in the list.
- **Manage columns** (⊕) – Manage which columns are displayed in the list so that only the information you want to see is displayed. Click the manage columns button (⊕), and then check the checkbox next to the column name to have that column shown in the list.
- **Hide the column** (⊖) – Click the hide the column button (⊖) in the column title to hide the column.


Perform the following actions in the **Group Manager** interface:

- **Create** – Click the **Create** button on the ribbon and then choose to create a group in the wizard mode or form mode.
- **View Details** – Click the group name or select a group and click **View Details**. The group settings are displayed on the **View Details** interface. When you want to modify the group settings, click **Edit** on the ribbon.
- **Edit** – Select a group and click **Edit** on the ribbon to change the configurations for the selected group. You can modify the High Availability Group settings except the **High Availability Mode**, **Production Farm**, **Standby Farm**, **Sync Method**, and **Failover Method**. If you have run the synchronization job, failover job, or fallback job on the group, the related settings of the Web application and database are not editable. You can refer to the message under the **Last Action Status** column in **Dashboard** to tell whether the group has been run.

- **Delete** – Select the groups that you want to delete and click **Delete** on the ribbon. A confirmation window will pop up and ask whether you are sure that you want to proceed with the deletion. Click **OK** to delete the selected groups, or click **Cancel** to return to the **Group Manager** interface without deleting the selected groups.
- **Throttle Control** – Click **Throttle Control** to go to the **Throttle Control** interface. You can manage the previously created throttle control profiles, or create a new throttle control profile. For more information on Throttle Control, see [About Throttle Control](#).
- **SQL Instance Settings** – Click **SQL Instance Settings** to configure the SQL instance settings. For more information on SQL instance settings, see [Configuring SQL Instance Settings](#).
- **Job Monitor** – Click **Job Monitor** to go to the **Job Monitor** interface. You can view the status of the jobs. This is useful for monitoring jobs or troubleshooting for errors. For more information on Job Monitor, see [DocAve 6 Job Monitor Reference Guide](#).

About the Dashboard

The Dashboard interface provides a detailed view for the production farms, standby databases, the High Availability Groups, as well the High Availability actions. You can perform the available High Availability actions from the ribbon, such as Pre-Scan, Synchronization, Failover, and Fallback.

To access the **Dashboard** in the High Availability **Home** interface, click **Dashboard** on the ribbon. You can view all of the previously created High Availability Groups grouped by different production farms. Click the view details () button besides the group name to go to the **View Details** interface. You can view detailed information of the group settings. Click **Edit** to modify the group settings, or click **Close** to return to the **Dashboard** interface.

In the **Dashboard** interface, click the group name to show all of the SharePoint components added in the group. There are four columns displayed in the **Dashboard** pane, respectively **Production**, **Standby**, **Last Action Status**, and **Last Action Time**. The **Last Action Status** shows the status of the last operation performed on the database. To judge whether the database is failover-ready through its **Last Action Status**, refer to [Distinguishing Failover-Ready Databases on Dashboard](#). The **Last Action Time** shows the finish time of the last action performed on the database.

You can perform the following actions in the **Dashboard** interface:

- **Pre-Scan** – It is recommended to perform the Pre-Scan job first before synchronization to scan the environment according to the predefined pre-scan rules. Select the checkbox ahead of the desired nodes, and then you can select **Pre-Scan** on the ribbon to check whether the group can perform a synchronization job. A pop-up window appears to ask for confirmation. Click **OK** to run a Pre-Scan job for the selected nodes based on the predefined pre-scan rules. Click **Cancel** to leave this pop-up window without starting the Pre-Scan job. For more information on the Pre-Scan rules, refer to [Appendix D – Checking the Pre-Scan Rules](#).
- **Synchronization** – Synchronize the data from the production environment to the disaster recovery environment. Select the checkbox ahead of the desired nodes, and then you can select **Synchronization** on the ribbon to make the data synchronous between the production server and standby server.
- **Failover** – Switch to a standby disaster recovery environment upon the failure or abnormal termination of the production environment. Select the checkbox ahead of the desired nodes, and then you can select **Failover** on the ribbon to enable the standby server when the production server is down. If you select the groups using AlwaysOn Availability Group method to run the Failover, DocAve High Availability will failover the included AlwaysOn Availability Group in the SQL Server.

***Note:** AvePoint recommends refreshing the groups in the Dashboard before you perform the High Availability Failover.
- **Fallback** – Revert back to the production server when the production server is recovered. Select the checkbox ahead of the desired nodes, and then you can select

Fallback to synchronize the data from the standby server to the production server. If you select the groups using AlwaysOn Availability Group method to run the Fallback, DocAve High Availability will fallback the included AlwaysOn Availability Group in the SQL Server.

***Note:** You can only perform the fallback to switch back to the production server after you perform the High Availability Failover job.

***Note:** You can select multiple groups together or select some components from the group to perform the High Availability pre-scan, synchronization, failover, or fallback. It is strongly recommended that you add the components in the same farm to the same group, in case the multiple processes lead to the failure of jobs.

Distinguishing Failover-Ready Databases on Dashboard

In High Availability Dashboard, all of the groups and components added to the groups will be displayed with its last action status and time. Through Dashboard, you can tell when and which High Availability operation is performed on the database, and whether the database is ready to perform Failover. For details on how to distinguish the failover-ready databases, refer to the table below.

***Note:** There might be abnormal standby database status which is depending on various factors of High Availability jobs hard to be listed in the table below.

Sync Method	Warm Standby Database Status	Cold Standby Database Status	Last Action Status	Comment
SQL Mirroring	Mirroring	Mirroring	Synchronization Succeeded	Failover ready
			Synchronization Skipped	Failover ready
	Does Not Exist	Does Not Exist	Synchronization Failed	Cannot failover
				Cannot failover
AlwaysOn Availability Group	Synchronized	Synchronized	Synchronization Succeeded	The status of Synchronized or Synchronizing depends on the Availability Mode. For more information, refer to Availability Modes (AlwaysOn Availability Groups) .
	Synchronizing	Synchronizing		
Log Shipping	Read-Only	Restoring	Synchronization Succeeded	Failover ready

Sync Method	Warm Standby Database Status	Cold Standby Database Status	Last Action Status	Comment
	Restoring	Restoring	Synchronization Failed	Does not recommend performing failover. It might cause data loss.
	Does Not Exist	Does Not Exist		Cannot failover
Platform Backup Log Shipping	Read-Only	Restoring	Synchronization Succeeded	Failover ready
	Restoring	Restoring	Synchronization Failed	Does not recommend performing failover. It might cause data loss.
	Does Not Exist	Does Not Exist		Cannot failover

Configuration File Settings

High Availability provides a configuration file to help users achieve advanced features beyond the DocAve interface. The configuration file resides in the ...*AvePoint\DocAve6\Agent\data\SP2010\HighAvailability\SP2010HAConfiguration.xml* directory.

***Note:** The **SP2010HAConfiguration.xml** file is used for both of SharePoint 2010 environment and SharePoint 2013 environment.

Refer to the section below for detailed instruction on the features that can be achieved by using High Availability configuration file.

Forcing a Full Synchronization on a Database with a Broken LSN

If you use the Log Shipping or Platform Backup Log Shipping as the sync method, you must carefully prevent the LSN (Log Sequence Number) of the database from being broken. To solve this agonizing issue, DocAve High Availability provides you a method to perform a full synchronization (initialization) on the database whose LSN is broken to ensure the data consistency between the production environment and standby environment.

To perform a full synchronization job on the database with a broken LSN (Log Sequence Number), complete the steps below:

1. Navigate to the ...*AvePoint\DocAve6\Agent\data\SP2010\HighAvailability\SP2010HAConfiguration.xml* directory on the SQL Server where the database whose LSN is broken resides.
2. Open the **SP2010HAConfiguration.xml** file with Notepad.
3. Find the **LSNBrokenForceFull** attribute, and then change its value to **True**.
4. Save the configuration file, and close it.

The next synchronization job will perform the full synchronization on the database whose LSN is broken.

Appendix A – Accessing Hot Key Mode

In order to work faster and improve your productivity, DocAve supports hot key mode for you to perform corresponding actions quickly by only using your keyboard. To access hot key mode from the **DocAve** interface, press the key combination of **Ctrl +Alt + Z** (simultaneously) on your keyboard. And then press **1** to direct to the products pane. Press **D** to access the **Data Protection** product suite, and then press **H** to access the **High Availability** interface.

Common Access

Each time you want to access hot key mode from anywhere in the **High Availability** interface, press the **Ctrl +Alt + Z** key combination to go to the **Common Access** page that displays the hot keys to access the desired functions.

The following table provides a list of hot keys for common access. From anywhere in the **High Availability** interface, press the **Ctrl +Alt + Z** key combination to access to the **Common Access** page, and then press the specified hot key to access to the corresponding product. For example, continue pressing **H** to jump to the **Home** tab.

Functionality Name and Hot Key	
Home	H
Group Manager	G
DocAve Home Page	1
DocAve Online Community	2
Control Panel	3
Job Monitor	4
Plan Group	5
Health Analyzer	6
Account Information	9
Help and About	0

Home Interface

The following table provides a list of hot keys for the functionalities on the **Home** Interface.

Functionality Name and Hot Key					
Landing Page	L				
Dashboard	D	Pre-Scan	P		
		Synchronization	S		
		Failover	F	OK	O
				Cancel	C
		Fallback	B		
Job Monitor	J				

Group Manager Interface

The following table provides a list of the hot keys for the functionalities on the **Group Manager** interface.

Functionality Name and Hot Key							
Create	C	Wizard Mode	W	Back	B		
				Next	N		
				Finish	F		
				Cancel	C		
		Form Mode	F	Save	S		
				Cancel	C		
View Details	V	Edit	E	Save	S		
				Cancel	C		
		Close	X				
Edit	E	Save	S				
		Cancel	C				
Delete	D						
Throttle Control	T	Create	C	Save	S		
				Cancel	C		
		View Details	V	Edit	E	Save	S
						Cancel	C
				Close	X		
		Edit	E	Save	S		
				Cancel	C		
		Delete	D				
		Close	X				
		SQL Instance Settings	S	Close	X		
Configure	Save			S			
	Cancel			C			
Job Monitor	J						

Appendix B – SharePoint 2010 Components Supported for High Availability

Refer to the table below for more information on the SharePoint 2010 components supported by the Single farm mode High Availability and the Standby farm mode High Availability:

SharePoint Components		Single Farm Mode	Standby Farm Mode
Configuration Database		Supported	Not Supported
Central Administration Database		Supported	Not Supported
Web Application	Content Database	Supported	Supported
Service Application	Databases of Service Application		
Access Service	This Service Application does not have a database	Not Supported	Supported
Application Registry Service	Application registry service database	Supported	Not Supported
Business Data Connectivity Service	Bdc_Service_DB(Business Data Connectivity Database)	Supported	Supported
Excel Services Application	This Service Application does not have a database	Not Supported	Supported
Managed Metadata Service	Managed Metadata Service_Database(Service Application Database)	Supported	Supported
PerformancePoint Service Application	PerformancePoint Service Application_Database(Microsoft.PerformancePoint.Scorecards.BIMonitoringServiceDatabase)	Supported	Supported
Search Service Application	Search_Service_Application_DB(Administration Database) Search_Service_Application_CrawlStoreDB(Crawl Database) Search_Service_Application_PropertyStoreDB(Property Database)	Supported	Supported * Supported for all sync methods except Always on AG. * Note: After Failover or Fallback, please run a full crawl to make sure the index files are up to date and available for query in the standby farm.
Secure Store Service	Secure_Store_Service_DB(Secure Store Service Database)	Not Supported	Not Supported

SharePoint Components		Single Farm Mode	Standby Farm Mode
State Service	StateService_Database(State Service Database Settings)	Supported	Not Supported
Visio Graphics Service	This Service Application does not have a database	Not Supported	Not Supported
Usage and Health data collection	WSS_Logging	Supported	Not Supported
User Profile Service Application	User Profile Service Application_ProfileDB (Microsoft.Office.Server.Administration.ProfileDatabase) User Profile Service Application_SyncDB(Microsoft.Office.Server.Administration.SynchronizationDatabase) User Profile Service Application_SocialDB (Microsoft.Office.Server.Administration.SocialDatabase)	Supported	Supported *Note: The User Profile Service Application_SyncDB will not be displayed in the farm tree, and the production synchronization database will not be synchronized to the standby farm.
Web Analytics Service Application	WebAnalyticsServiceApplication_StagingDB(Web Analytics Staging Database) WebAnalyticsServiceApplication_ReportingDB(Web Analytics Warehouse Database)	Supported	Supported
Word Automation Service	WordAutomationServices_Database (Microsoft.Office.Word.Server.Service.QueueDatabase)	Supported	Supported
Subscription Service	Subscription Settings Database	Supported	Supported
Session State Service	SessionStateService_Database	Supported	Not Supported
SharePoint Foundation Search	WSS_Search_Database	Supported	Not Supported

Appendix C – SharePoint 2013 Components Supported for High Availability

Refer to the table below for more information on the SharePoint 2013 components supported by the Single farm mode High Availability and the Standby farm mode High Availability:

SharePoint Components		Single Farm Mode	Standby Farm Mode
Configuration Database		Supported	Not Supported
Central Administration Database		Supported	Not Supported
Web Application	Content Database	Supported	Supported
Service Application	Databases of Service Application		
App Management Service Application	AppMng_Service_DB (App Management Database)	Supported	Supported
Access Service Application	This Service Application does not have a database	Not Supported	Supported
Access Service Application 2010	This Service Application does not have a database	Not Supported	Supported
Business Data Connectivity Service	Bdc_Service_DB(Business Data Connectivity Database)	Supported	Supported
Excel Services Application	This Service Application does not have a database	Not Supported	Supported
Managed Metadata Service	Managed Metadata Service_Database(Service Application Database)	Supported	Supported
Machine Translation Service Application	Machine Translation Service Database	Supported	Not Supported
PerformancePoint Service Application	PerformancePoint Service Application_Database(Microsoft.PerformancePoint.Scorecards.BIMonitoringServiceDatabase)	Supported	Supported
Search Service Application	Search_Service_Application_DB(Administration Database) Search_Service_Application_CrawlStoreDB(Crawl Database) Search_Service_Application_AnalyticsReportingStoreDB(Analytics Reporting Database) Search_Service_Application_LinksStoreDB(Link Database)	Supported	Supported * Supported for all sync methods except Always on AG. *Note: After Failover or Fallback, please run a full crawl to make sure the index files are up to date and available for query

SharePoint Components		Single Farm Mode	Standby Farm Mode
			in the standby farm.
State Service	StateService_Database(State Service Database Settings)	Supported	Not Supported
Secure Store Service	Secure_Store_Service_DB(Secure Store Service Database)	Supported	Supported
Session State Service	SessionStateService_Database	Supported	Not Supported
Security Token Service Application	This Service Application does not have a database	Not Supported	Not Supported
Visio Graphics Service	This Service Application does not have a database	Not Supported	Supported
Usage and Health data collection	WSS_Logging	Supported	Not Supported
User Profile Service Application	User Profile Service Application_ProfileDB (Microsoft.Office.Server.Administration.ProfileDatabase) User Profile Service Application_SyncDB(Microsoft.Office.Server.Administration.SynchronizationDatabase) User Profile Service Application_SocialDB (Microsoft.Office.Server.Administration.SocialDatabase)	Supported	Supported
Word Automation Service	WordAutomationServices_Database (Microsoft.Office.Word.Server.Service.QueueDatabase)	Supported	Supported
Workflow Service Application	This Service Application does not have a database	Not Supported	Not Supported
Work Management Service Application	This Service Application does not have a database	Not Supported	Not Supported
PowerPoint Service Application	This Service Application does not have a database	Not Supported	Supported
Subscription Service	Subscription Settings Database	Supported	Supported

Appendix D – Checking the Pre-Scan Rules

This section provides you detailed information on the Pre-Scan rules defined for the four different Sync Methods: **SQL Mirroring**, **AlwaysOn Availability Group**, **Log Shipping**, and **Platform Backup Log Shipping**.

SQL Mirroring

Refer to the section below for detailed information on the Pre-Scan rules for SQL Mirroring method:

- Check the version of SQL Server – Check whether the versions of the SQL Servers on the production farm and standby farm are the same, and whether the version is SQL 2005 SP1 or later.
- Check the free disk space in the standby server – Check whether there is enough free disk space in the standby server.
- Check the status of the production database – Check whether the status of the production database is normal. If the production database status of the SQL Mirroring is principle, and the standby database status is mirroring, the database will be skipped.
- Check the status of the standby database – Check whether the standby database exists.
- Check the recovery model of the selected nodes – Check whether the recovery mode of the selected nodes is **Full**.
- Check the connection between the production server and the standby server – Check whether or not the production server can successfully ping to the standby server.
- Check the connection between the standby server and the production server – Check whether or not the standby server can successfully ping to the production server.
- Check the endpoint for SQL mirroring – Check whether there is no endpoint in the production and standby SQL instance, or there is the endpoint named as **DocAveSQLMIRROR**.
- Check the Agent account permission of **db_owner** role to database – Check whether the Agent account has the **db_owner** database role to the production databases.

The following Pre-Scan rules are specific for the Standby farm mode High Availability:

- Check the schema version of the database – If the schema version does not exist, this rule will be passed successfully. If the schema version exists, and it is equal to or lower than the standby SharePoint version, this rule will be passed successfully.
- Check for a site conflict in production farm – Check whether or not there are conflicting production site collections using the same URL in the content databases that are added into the High Availability group.
- Check for a site conflict in standby farm – Check whether or not there is a conflicting production site using the same URL as any in the standby farm.

- Check the conflicting port of Web application in standby farm – Check whether or not there is Web application in the standby farm using the same port as the production Web application included in the High Availability group.

AlwaysOn Availability Group

Refer to the section below for detailed information on the Pre-Scan rules for AlwaysOn Availability Group method.

- Check the version of SQL Server – Check whether the major versions of the SQL Server in the production farm and standby farm are the same, and the version is equal to or later than SQL Server 2012.
- Check the free disk space in the standby server – Check whether there is enough free disk space in the standby server.
- Check the status of the production database – Check whether the production database exists in the AlwaysOn Availability Group, or the database status is normal (that is, there is no suffix of status attached behind the database name).
- Check the status of the standby database – Check whether the standby database exists. If the standby database is included in the availability group or does not exist in the standby SQL Server, this rule will pass.
- Check the recovery model of the selected nodes – Check whether the recovery mode of the selected nodes is **Full**.
- Check the status of the production SQL instance in the availability group – Check whether the status of the production SQL instance in the availability group is **Primary**.
- Check the Agent account permission of **db_owner** role to database – Check whether the Agent account has the **db_owner** database role to the production databases.

The following Pre-Scan rules are specific for the Standby farm mode High Availability:

- Check the schema version of the database – If the schema version does not exist, this rule will be passed successfully. If the schema version exists, and it is equal to or lower than the standby SharePoint version, this rule will be passed successfully.
- Check for a site conflict in production farm – Check whether or not there are conflicting production site collections using the same URL in the content databases that are added into the High Availability group.
- Check for a site conflict in standby farm – Check whether or not there is a conflicting production site using the same URL as any in the standby farm.
- Check the conflicting port of Web application in standby farm – Check whether or not there is Web application in the standby farm using the same port as the production Web application included in the High Availability group.

Log Shipping

Refer to the section below for detailed information on the Pre-Scan rules for Log Shipping method.

- Check the version of SQL Server – Check whether the major version of the SQL Servers on the production farm and standby farm is the same.
- Check the free disk space in the standby server – Check whether there is enough free disk space in the standby server.
- Check the status of the production database – Check whether the status of the production database is normal.
- Check the status of the standby database – Check whether the standby database exists or the status of the standby database is **Read-only** or **Restoring**.
- Check the recovery model of the selected nodes – Check whether the recovery mode of the selected nodes is **Full** or **Bulk-logged**.
- Check the Log Sequence Number (LSN) – Check whether the LSN of the production database and standby database matches or not.
- Check the Agent account permission of **db_owner** role to database – Check whether the Agent account has the **db_owner** database role to the production databases.

The following Pre-Scan rules are specific for the Standby farm mode High Availability:

- Check the schema version of the database – If the schema version does not exist, this rule will be passed successfully. If the schema version exists, and it is equal to or lower than the standby SharePoint version, this rule will be passed successfully.
- Check for a site conflict in production farm – Check whether or not there are conflicting production site collections using the same URL in the content databases that are added into the High Availability group.
- Check for a site conflict in standby farm – Check whether or not there is a conflicting production site using the same URL as any in the standby farm.
- Check the conflicting port of Web application in standby farm – Check whether or not there is Web application in the standby farm using the same port as the production Web application included in the High Availability group.

Platform Backup Log Shipping

Refer to the section below for detailed information on the Pre-Scan rules for Platform Backup Log Shipping method.

- Check the version of SQL Server – Check whether the major version of the SQL Servers on the production farm and standby farm is the same.
- Check the free disk space in the standby server – Check whether there is enough free disk space in the standby server.

- Check the status of the production database – Check whether the status of the production database is normal.
- Check the status of the standby database – Check whether the standby database exists or the status of the standby database is **Read-only** or **Restoring**.
- Check the recovery model of the selected nodes – Check whether the recovery mode of the selected nodes is **Full** or **Bulk-logged**.
- Check the Log Sequence Number (LSN) – Check whether the LSN of the backed up database and standby database matches or not.
- Check the Platform Backup history – Check whether the Platform Backup job record for the selected farm components exists.
- Check the Agent account permission of **db_owner** role to database – Check whether the Agent account has the **db_owner** database role to the production databases.

The following Pre-Scan rules are specific for the Standby farm mode High Availability:

- Check the schema version of the database – If the schema version does not exist, this rule will be passed successfully. If the schema version exists, and it is equal to or lower than the standby SharePoint version, this rule will pass successfully.
- Check for a site conflict in production farm – Check whether or not there are conflicting production site collections using the same URL in the content databases that are added into the High Availability group.
- Check for a site conflict in standby farm – Check whether or not there is a conflicting production site using the same URL as any in the standby farm.
- Check the conflicting port of Web application in standby farm – Check whether or not there is Web application in the standby farm using the same port as the production Web application included in the High Availability group.

Common Rules from Health Analyzer

High Availability Pre-Scan applies some common rules from Health Analyzer to help check whether or not the Agent account permission settings and server settings in your SharePoint environment are met with the prerequisites of running DocAve High Availability.

High Availability Pre-Scan applies the following common rules from Health Analyzer:

- Agent Account has Full Control to DocAve Certificate – Check whether or not the Agent account has the **Full Control** permission to the DocAve certificate. The Agent account requires this permission to communicate safely with the Control service.
- Agent Account has Full Control to Agent Registry – Check whether or not the Agent account has the **Full Control** permission to the Agent registry. The Agent account requires this permission to access the Agent registry created while installing Agent to get the installation directory.

- Agent Account has Full Control to the Eventlog Node in Registry – Check whether or not the Agent account has the **Full Control** permission to the **eventlog** node in **Registry**. The Agent account requires this permission to write the event log.
- Agent Account is a Member of IIS User Group – Check whether or not the Agent account is a member of IIS user group. The Agent account requires this permission to use WCF Port Sharing service.
- Agent Account has Full Control to Agent Installation Directory – Check whether or not the Agent account have the **Full Control** permission to the Agent installation directory. The Agent account requires this permission to record logs and job data in this directory.
- Agent Account is a Member of Local Administrators Group – Check whether or not the Agent account is a member of the local Administrators group. The Agent account requires this permission to do some operations. For example, add assembly to the Global Assembly Cache.
- Agent Account has Log on as a Batch Job User Right in Local Security Policy – Check whether or not the Agent account has the **Log on as a batch job** user right in **Local Security Policy**. The Agent account requires this user right to start a new process.
- Agent Account is a Member of Performance Monitor Users Group – Check whether or not the Agent account is a member of the **Performance Monitor Users** group. The Agent account requires this permission to use .NET Performance Counter API, which is required by .NET API. SharePoint API also uses it internally.
- Agent Account is a Member of Farm Administrators Group – Check whether or not the Agent account is a member of the Farm Administrators group. The Agent account requires this permission to operate SharePoint API and the farm level nodes.
- Agent Account has the **db_owner** Database Role in SharePoint Configuration Database – Check whether or not the Agent account has the **db_owner** database role in the SharePoint Configuration database. The Agent account requires this permission to obtain data from the SharePoint Configuration database.
- Agent Account has the **db_owner** Database Role in SharePoint Central Administration Content Database – Check whether or not the Agent account has the **db_owner** database role in the SharePoint Central Administration content database. Agent account requires this permission to access the Central Administration site.
- Agent Version Matches Manager Version – Check whether or not the Agent version matches the Manager version. The Agent version must be the same as the Manager version for communication.

Notices and Copyright Information

Notice

The materials contained in this publication are owned or provided by AvePoint, Inc. and are the property of AvePoint or its licensors, and are protected by copyright, trademark and other intellectual property laws. No trademark or copyright notice in this publication may be removed or altered in any way.

Copyright

Copyright © 2012-2014 AvePoint, Inc. All rights reserved. All materials contained in this publication are protected by United States copyright law and no part of this publication may be reproduced, modified, displayed, 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 AvePoint, 3 Second Street, Jersey City, NJ 07311, USA or, in the case of materials in this publication owned by third parties, without such third party's consent.

Trademarks

AvePoint[®], DocAve[®], the AvePoint logo, and the AvePoint Pyramid logo are registered trademarks of AvePoint, Inc. with the United States Patent and Trademark Office. These registered trademarks, along with all other trademarks of AvePoint used in this publication are the exclusive property of AvePoint and may not be used without prior written consent.

Microsoft, MS-DOS, Internet Explorer, Microsoft Office SharePoint Servers 2007/2010/2013, SharePoint Portal Server 2003, Windows SharePoint Services, Windows SQL Server, and Windows are either registered trademarks or trademarks of Microsoft Corporation.

Adobe Acrobat and Acrobat Reader are trademarks of Adobe Systems, Inc.

All other trademarks contained in this publication are the property of their respective owners and may not be used such party's consent.

Changes

The material in this publication is for information purposes only and is subject to change without notice. While reasonable efforts have been made in the preparation of this publication to ensure its accuracy, AvePoint makes no representation or warranty, expressed or implied, as to its completeness, accuracy, or suitability, and assumes no liability resulting from errors or omissions in this publication or from the use of the information contained herein. AvePoint reserves the right to make changes in the Graphical User Interface of the AvePoint software without reservation and without notification to its users.

AvePoint, Inc.
Harborside Financial Center, Plaza 10
3 Second Street, 9th Floor
Jersey City, New Jersey 07311
USA