DocAve® 6 High Availability

User Guide

Service Pack 7 Cumulative Update 1

Issued October2016

Table of Contents

What's New in This Guide 6

About DocAve High Availability 7

Submitting Documentation Feedback to AvePoint 8

Before You Begin 9

- Configuration 9
- AvePoint's Testing Policy and Environment Support 9
  - Supported Software Environments 9
  - Supported Hardware 9
  - Supported Backup and Recovery 10
  - Notable Environment Exceptions 10
- Agents 10
- Required Permissions 10
  - Common Permissions Required for All of the Four Sync Methods 10
  - Required Permissions for SQL Mirroring Method 13
  - Required Permissions for AlwaysOn Availability Group Method 15
  - Required Permissions for Log Shipping Method 18
  - Required Permissions for Platform Backup Log Shipping 21

Getting Started 24

- Launching High Availability 24
- User Interface Overview 25
- Navigating DocAve 26

About Throttle Control 28

- Configuring a Throttle Control Profile 28
- Managing Throttle Control Profiles 28

Configuring a Master Key Password for SQL Instance 30

About Custom Actions 31

- Configuring a Script Profile 31
- Managing Script Profiles 32
- Configuring a Command Profile 32

Managing Command Profiles 33
Configuring Cache Setting 34
About Connector Cache Setting 34
Configuring a Connector Cache Setting Profile 34
Managing Connector Cache Setting Profiles 34
About Log Shipping Cache Setting 35
Configuring a Log Shipping Cache Profile 35
Managing Log Shipping Cache Profiles 36
Configuring a High Availability Group 37
Using Wizard Mode to Create a Single Farm Mode High Availability Group 37
Using Wizard Mode to Create a Standby Farm Mode High Availability Group 45
Using Form Mode to Create a High Availability Group 54
Adding SharePoint Components to a Group 55
SQL Mirroring, Log Shipping, and Platform Backup Log Shipping 55
AlwaysOn Availability Group 57
Adding Custom Databases to a Group 59
Managing Groups 61
About the Dashboard 63
Distinguishing Failover-Ready Databases on Dashboard 65
High Availability for Specific Farm Components or Scenarios 67
Protecting Connector BLOB Data with High Availability 67
Creating a Standby Farm Mode High Availability Group Using Wizard Mode 67
Performing a Synchronization Job and Viewing the Standby Environment 70
Performing a Failover Job and Viewing the Standby Environment 73
Performing a Fallback Job 74
Protecting Storage Manager BLOB data with High Availability 75
Creating a Standby Farm Mode High Availability Group Using Wizard Mode 75
Performing a Synchronization Job and Viewing the Standby Environment 78
Performing a Failover Job and Viewing the Standby Environment 80
Performing a Fallback Job 81
Building up a Read-Only or Read-Write Standby Farm for the Maintenance of the Production Farm 82
Creating and Using Standby Farm Mode Group for Maintenance Purpose 82

Enabling Read-Only View for Standby Web Applications and Standby Service Applications after Synchronization 92
Creating Standby Farm Mode High Availability Groups 92
Performing Synchronization Jobs 93
Performing Failover Jobs 96
Performing a Fallback Job 98
Improving Performance of Log Shipping Fallback 99
Performing a Log Shipping Fallback 100
Forcing a Full Synchronization on a Database with a Broken LSN by Customizing Configuration File 104
Forcing a Full Synchronization or Full Fallback on Connector Data 105
Designating VSS Hardware Provider for Creating Snapshot 106
Defining Warm Standby Behavior of Service Applications 107
Designating Timeout Session for SQL Statement Execution in Log Shipping Incremental Data Transfer 108
Appendix A: SharePoint Components Supported for High Availability 109
SharePoint 2010 Components supported for High Availability 109
SharePoint 2013 Components Supported for High Availability 112
SharePoint 2016 Components Supported for High Availability 115
Appendix B: Checking the Pre-Scan Rules 118
SQL Mirroring 118
AlwaysOn Availability Group 119
Log Shipping 121
Platform Backup Log Shipping 122
Common Rules from Health Analyzer 123
Appendix C: Accessing Hot Key Mode 125
Common Access 125
Home Interface 126
Group Manager Interface 126
Appendix D: Protecting Data in Data ONTAP Storage System with High Availability SnapMirror Sync Method 128
Where to Install Agents for High Availability 128
Specific System Requirements and Limitations 128
Required Permissions for SnapMirror 129

Building a High Availability Group 131
Creating a Single Farm Mode High Availability Group Using Wizard Mode 132
Creating a Standby Farm Mode High Availability Group Using Wizard Mode 137
Managing Groups 141
Actions in the Dashboard 143
Distinguishing Failover-Ready Databases on Dashboard 144
SnapMirror Pre-Scan Rules 144
Customizing the Timeout Session for Getting Update-SnapMirror or Resync-SnapMirror 145
Customizing Where to Execute Filer SDK Toolkit 145
Connecting LUN Devices with FCP 146
Notices and Copyright Information 147

What's New in This Guide

The following updates have been made to this guide for High Availability in DocAve 6 Service Pack 7, Cumulative Update 1:

- Added instructions for adding multiple standby farms into High Availability group when you create a Standby farm mode High Availability group using AlwaysOn Availability Group sync method. For details, refer to Using Wizard Mode to Create a Standby Farm Mode High Availability Group.
- Updated the Maintenance Failover.

About DocAve High Availability

DocAve High Availability is a one-switch disaster recovery and maintenance 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 or later, High Availability supports the SharePoint 2010, SharePoint 2013, and SharePoint 2016 environments.

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.

AvePoint's Testing Policy and Environment Support

Supported Software Environments

AvePoint is committed to testing against all major versions and service packs of SharePoint as well as the latest versions of Windows Server and SQL Server, as Microsoft announces support and compatibility.

***Note:** AvePoint does not recommend or support installing DocAve on client operating systems.

Supported Hardware

AvePoint is committed to maintaining a hardware agnostic platform to ensure that DocAve operates on common Windows file sharing and virtualization platforms. To ensure that DocAve is hardware agnostic, AvePoint tests hardware that is intended to support SharePoint and DocAve infrastructure, storage targets, and hardware-based backup and recovery solutions, as supported by AvePoint's partnerships. AvePoint directly integrates with the following platforms: any Net Share, IBM Storwize Family, FTP, Amazon S3, AT&T Synaptic, Dropbox, Box, Caringo Storage, Del DX Storage, EMC Centra, HDS Hitachi Content Platform, Rackspace Cloud Files, TSM, Network File System, and Windows Azure Storage.

All other hardware platforms that support UNC addressable storage devices are supported.

***Note:** Due to changes in the IBM Tivoli Storage Manager API, DocAve 6 Service Pack 6 and later versions require that TSM Client version 7.1.2 is installed on the Control Service and Media Service servers.

***Note:** Most of the hardware partnerships referenced in this guide are intended to make use of advanced functionality (such as snapshot mirroring, BLOB snapshots, indexing, long-term storage, WORM storage, etc.), and are not indications that any changes to the product are required for basic support. In most cases, hardware can be supported with no change to the product.

Supported Backup and Recovery

DocAve supports BLOB backup storage according to the list of hardware platforms above. BLOB snapshot functionally, however, is currently only supported on OEM versions and NetApp hardware.

DocAve supports SQL content and Application database backups via the SharePoint Volume Shadow Copy Service (VSS) on all Windows and SQL server platforms listed above. DocAve also supports snapshot-based SharePoint VSS on all hardware listed above where the hardware partner has certified support with Microsoft.

DocAve supports application and configuration server backups for all the supported software environments listed above. DocAve 6 SP5 supports VM backup via Hyper-V/VMWare for the following operating systems: Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, and Microsoft Hyper–V Server 2012 R2.

Notable Environment Exceptions

The following are notable exceptions to the supported DocAve environments. The following represent environment level support information, not feature level support. Feature level support, specific to each feature, is described throughout this guide where applicable.

- High Availability only supports storing Connector BLOB data in Net Share or NFS devices because of a limitation with other devices' infrastructure-level methods.

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.

Common Permissions Required for All of the Four Sync Methods

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.
    - 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 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
    - Server Role of **public** to SQL Server
    - **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 High Availability jobs on
    - Server Role of **dbcreator** and **securityadmin** to SQL Server

      ***Note**: The Agent account configured on SQL Server must also have the **sysadmin** server role on the standby SQL Server for the following reasons:

        - If you want to perform the High Availability of Standby farm mode for Business Data Connectivity Service, Managed Metadata Service, or Search Service Application, this permission is required so that the Agent account configured on the SharePoint server that is included in the Agent group can be granted the **db_owner** role to the standby databases of those service applications.
        - If you want to perform the High Availability of Standby farm mode for a Web application, this permission is required so that the application pool user can be granted the **db_owner** role to the standby database.

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*.

VSS Writer account configured on the SQL Server:

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

SharePoint 2013/SharePoint 2016 application pool account configured on the SQL Server:

For SharePoint 2013 and SharePoint 2016, the standby application pool account must exist in the standby SQL Server and have the **db_owner** role for the production database. You can also grant the application pool account the server role of **sysadmin** in the standby SQL Server.

Service application pool account configured on the SharePoint Server:

If the High Availability group includes the PowerPoint Service Application, the service application pool account configured on the SharePoint server must have the **Write** permission to the *C:\ProgramData\Microsoft\SharePoint* directory in the SharePoint server for storing the temporary file of the Conversion job.

Agent account configured on the SharePoint Server to start the SP2010StorageOptimizationService.exe process or SP2013StorageOptimizationService.exe process

If you are about to synchronize the content database with BLOB data and related stub database together to the standby farm with read-only view enabled, make sure the Agent account to start the **SP2010StorageOptimizationService.exe** process or **SP2013StorageOptimizationService.exe** process on the SharePoint server has sufficient permissions in the following scenarios before performing the synchronization job.

- If the Agent account in the standby farm is a different user in the same domain as the Agent account in the production 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 the Agent account in the standby farm is in a different domain as the Agent account of the production farm, it is recommended making the domain in the production farm trusted by the domain in the standby farm and granting the Agent account in the standby farm the **db_owner** role in the production stub database. Otherwise, the Agent account in the standby farm must have the **sysadmin** role to the standby SQL instance, in order to make sure the standby stub files are readable.

Required Permissions for SQL Mirroring Method

Note that * indicates a permission specifically required for SQL Mirroring.

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 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
    - **Control Server** to the destination SQL instance
    - Server Role of **public** to SQL Server
    - Server role of **securityadmin** to the destination SQL Server. Note that this permission is required for provisioning Managed Metadata Service in the standby farm.
    - Server role of **dbcreator** or the **Alter Any Database** permission or **View Any Definition** permission to the 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** and **securityadmin** to SQL Server
    - *Permissions of **Create Endpoint** and **Alter Login** to SQL Server

- Server Role of **sysadmin** to the destination SQL Server.

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.** High Availability uses the Agent Temporary Buffer location to store the SQLite database file, which is used for Connector physical device mapping.
- **Read** and **Write** permissions to the directory of *…\AvePoint\DocAve6\Agent\Jobs*.
- ***Read** and **Write** permissions to the sparse file location

  ***Note**: If the spare file location is in File Share, the SQL Server Service account must be a member of the local **Administrators** or **Backup Operators**.

VSS Writer account configured on the SQL Server:

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

SharePoint 2013/SharePoint 2016 application pool account configured on the SQL Server:

For SharePoint 2013 and SharePoint 2016, you must ensure that the standby application pool account exists in the standby SQL Server and have the **db_owner** role to the production database, or you can grant the application pool account the server role of **sysadmin** in the standby SQL Server.

Service application pool account configured on the SharePoint Server:

If the High Availability group includes the PowerPoint Service Application, the service application pool account configured on the SharePoint server must have the **Write** permission to the *C:\ProgramData\Microsoft\SharePoint* directory in the SharePoint server for storing the temporary file of the Conversion job.

Agent account configured on the SharePoint Server to start the SP2010StorageOptimizationService.exe process or SP2013StorageOptimizationService.exe process

If you are about to synchronize the content database with BLOB data and related stub database together to the standby farm with read-only view enabled, make sure the Agent account to start the **SP2010StorageOptimizationService.exe** process or **SP2013StorageOptimizationService.exe** process on

the SharePoint server has sufficient permissions in the following scenarios before performing the synchronization job.

- If the Agent account in the standby farm is a different user in the same domain as the Agent account in the production 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 the Agent account in the standby farm is in a different domain as the Agent account of the production farm, it is recommended making the domain in the production farm trusted by the domain in the standby farm and granting the Agent account in the standby farm the **db_owner** role in the production stub database. Otherwise, the Agent account in the standby farm must have the **sysadmin** role to the standby SQL instance, in order to make sure the standby stub files are readable.

Required Permissions for AlwaysOn Availability Group Method

Note that * indicates a permission specifically required 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 permission of **View server state** to SQL Server
    - Database 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
    - Server Role of **public** to SQL Server
    - **Control Server** to the destination SQL instance
    - Server role of **securityadmin** to the destination SQL Server

***Note**: This permission is only required for provisioning Managed Metadata Service in the standby farm.

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
    - * Database Permission of **View Server State** to the SQL Server
    - * Database Permission of **Alter Availability Group** to the SQL Server

      ***Note**: The Agent account configured on SQL Server must also have the **sysadmin** server role on the standby SQL Server for the following reasons:

        - If you want to perform the High Availability of Standby farm mode for Business Data Connectivity Service, Managed Metadata Service, or Search Service Application, this permission is required so that the Agent account configured on the SharePoint server that is included in the Agent group can be granted the **db_owner** role to the standby databases of those service applications.
        - If you want to perform the High Availability of Standby farm mode for a Web application, this permission is required so that the application pool user can be granted the **db_owner** role to the standby database.

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** configured in **Control Panel** > **Agent Monitor** > **Configure,** and **Read** and **Write** permissions to the directory of *…\AvePoint\DocAve6\Agent\Jobs*.

High Availability uses the Agent Temporary Buffer location to store the SQLite database file, which is used for Connector physical device mapping.

VSS Writer account configured on the SQL Server:

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

SharePoint 2013/SharePoint 2016 application pool account configured on the SQL Server:

For SharePoint 2013 and SharePoint 2016, you must ensure that the standby application pool account exists in the standby SQL Server and have the **db_owner** role to the production database, or you can grant the application pool account the server role of **sysadmin** in the standby SQL Server.

Service application pool account configured on the SharePoint Server:

If the High Availability group includes the PowerPoint Service Application, the service application pool account configured on the SharePoint server must have the **Write** permission to the *C:\ProgramData\Microsoft\SharePoint* directory in the SharePoint server for storing the temporary file of the Conversion job.

Agent account configured on the SharePoint Server to start the SP2010StorageOptimizationService.exe process or SP2013StorageOptimizationService.exe process

If you are about to synchronize the content database with BLOB data and related stub database together to the standby farm with read-only view enabled, make sure the Agent account to start the **SP2010StorageOptimizationService.exe** process or **SP2013StorageOptimizationService.exe** process on the SharePoint server has sufficient permissions in the following scenarios before performing the synchronization job.

- If the Agent account in the standby farm is a different user in the same domain as the Agent account in the production 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 the Agent account in the standby farm is in a different domain as the Agent account of the production farm, it is recommended making the domain in the production farm trusted by the domain in the standby farm and granting the Agent account in the standby farm the **db_owner** role in the production stub database. Otherwise, the Agent account in the standby farm must have the **sysadmin** role to the standby SQL instance, in order to make sure the standby stub files are readable.

Required Permissions for Log Shipping Method

Note that * indicates a permission specifically required for Log Shipping method.

***Note**: If you are going to use the Log Shipping method to synchronize databases in an AlwaysOn Availability group, the required permissions for AlwaysOn Availability Group method must be met as well.

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.

      ***Note**: This permission is only required for provisioning Managed Metadata Service in the standby farm.

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
- * **Control Server** to the destination SQL instance

  ***Note**: The Agent account configured on SQL Server must also have the **sysadmin** server role on the standby SQL Server for the following reasons:

  - If you want to perform the High Availability of Standby farm mode for Business Data Connectivity Service, Managed Metadata Service, or Search Service Application, this permission is required so that the Agent account configured on the SharePoint server that is included in the Agent group can be granted the **db_owner** role to the standby databases of those service applications.
  - If you want to perform the High Availability of Standby farm mode for a Web application, this permission is required so that the application pool user can be granted the **db_owner** role to the standby database.

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 **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 **Read** and **Writer** permissions to the database file location (including the path in file share).

SharePoint 2013/SharePoint 2016 application pool account configured on the SQL Server:

For SharePoint 2013 and SharePoint 2016, you must ensure that the standby application pool account exists in the standby SQL Server and have the **db_owner** role to the production database, or you can grant the application pool account the server role of **sysadmin** in the standby SQL Server.

Service application pool account configured on the SharePoint Server:

If the High Availability group includes the PowerPoint Service Application, the service application pool account configured on the SharePoint server must have the **Write** permission to the *C:\ProgramData\Microsoft\SharePoint* directory in the SharePoint server for storing the temporary file of the Conversion job.

Agent account configured on the SharePoint Server to start the SP2010StorageOptimizationService.exe process or SP2013StorageOptimizationService.exe process

If you are about to synchronize the content database with BLOB data and related stub database together to the standby farm with read-only view enabled, make sure the Agent account to start the **SP2010StorageOptimizationService.exe** process or **SP2013StorageOptimizationService.exe** process on the SharePoint server has sufficient permissions in the following scenarios before performing the synchronization job.

- If the Agent account in the standby farm is a different user in the same domain as the Agent account in the production 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 the Agent account in the standby farm is in a different domain as the Agent account of the production farm, it is recommended making the domain in the production farm trusted by the domain in the standby farm and granting the Agent account in the standby farm the **db_owner** role in the production stub database. Otherwise, the Agent account in the standby farm must have the **sysadmin** role to the standby SQL instance, in order to make sure the standby stub files are readable.

Required Permissions for Platform Backup Log Shipping

Note that * indicates a permission specifically required for Platform Backup 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
    - **Control Server** to the destination SQL Server
    - Server role of **securityadmin** to the destination SQL Server

      ***Note**: This permission is only required for provisioning Managed Metadata Service in the standby farm.

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

  ***Note**: The Agent account configured on SQL Server must also have the **sysadmin** server role on the standby SQL Server for the following reasons:

  - If you want to perform the High Availability of Standby farm mode for Business Data Connectivity Service, Managed Metadata Service, or Search Service Application, this permission is required so that the Agent account configured on the SharePoint server that is included in the Agent group can be granted the **db_owner** role to the standby databases of those service applications.
  - If you want to perform the High Availability of Standby farm mode for a Web application, this permission is required so that the application pool user can be granted the **db_owner** role to the standby database.

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 **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 **Read** and **Writer** permissions to the database file location (including the path in file share).

SharePoint 2013/SharePoint 2016 application pool account configured on the SQL Server:

For SharePoint 2013 and SharePoint 2016, the standby application pool account must exist in the standby SQL Server and have the **db_owner** role for the production database. You can also grant the application pool account the server role of **sysadmin** in the standby SQL Server.

Service application pool account configured on the SharePoint Server:

If the High Availability group includes the PowerPoint Service Application, the service application pool account configured on the SharePoint server must have the **Write** permission to the *C:\ProgramData\Microsoft\SharePoint* directory in the SharePoint server for storing the temporary file of the Conversion job.

Agent account configured on the SharePoint Server to start the SP2010StorageOptimizationService.exe process, SP2013StorageOptimizationService.exe, or SP2016StorageOptimizationService.exe process

If you are about to synchronize the content database with BLOB data and related stub database together to the standby farm with read-only view enabled, make sure the Agent account to start the **SP2010StorageOptimizationService.exe** process, **SP2013StorageOptimizationService.exe** process, or **SP2016StorageOptimizationService.exe** process on the SharePoint server has sufficient permissions in the following scenarios before performing the synchronization job.

- If the Agent account in the standby farm is a different user in the same domain as the Agent account in the production 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 the Agent account in the standby farm is in a different domain as the Agent account of the production farm, it is recommended making the domain in the production farm trusted by the domain in the standby farm and granting the Agent account in the standby farm the **db_owner** role in the production stub database. Otherwise, the Agent account in the standby farm must have the **sysadmin** role to the standby SQL instance, in order to make sure the standby stub files are readable.

Getting Started

Refer to the sections below for important information on getting started with High Availability. Throttle Control, SQL Instance Settings, and Custom Actions are all optional settings in the Group Manager interface, but it is strongly recommended that you prepare these settings before creating a High Availability group.

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.

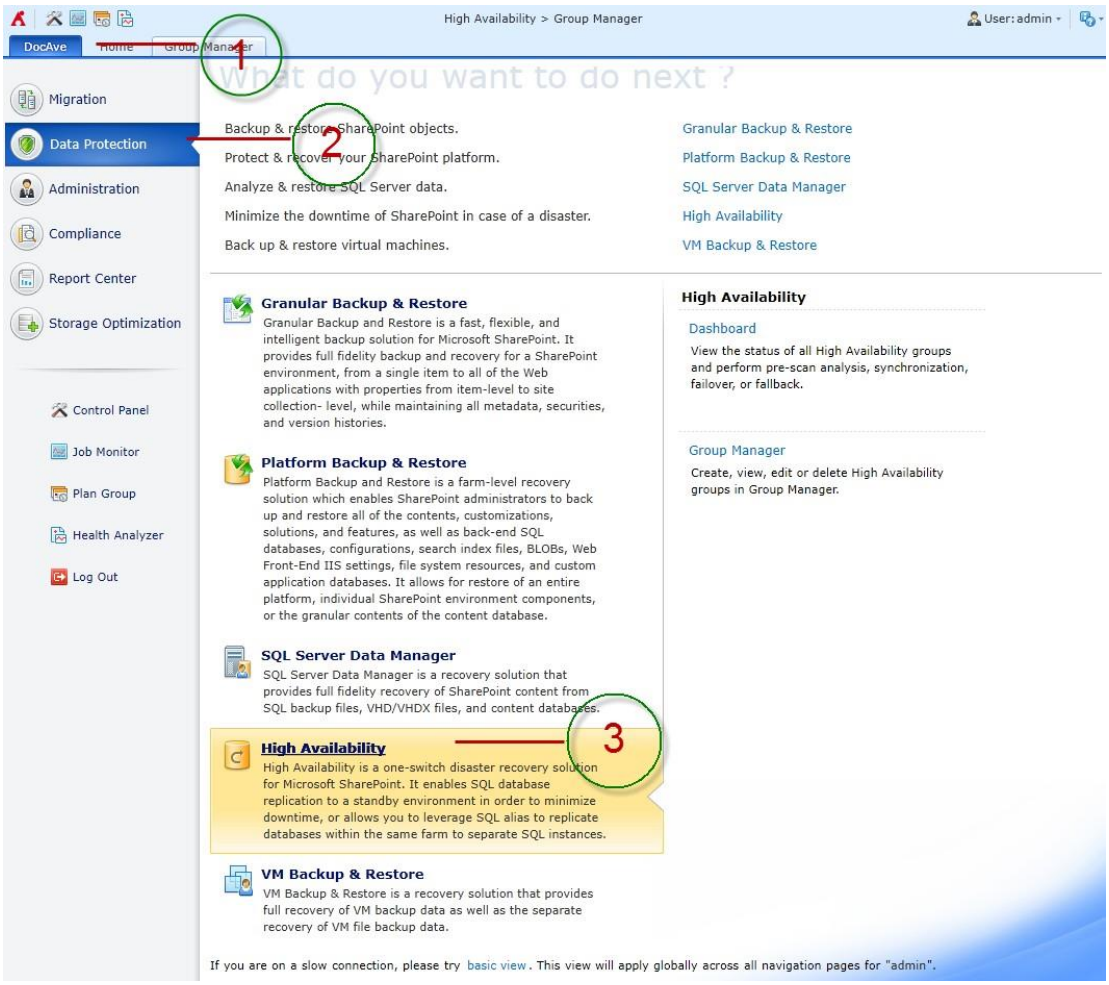

Figure 1: DocAve launch window.

User Interface Overview

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

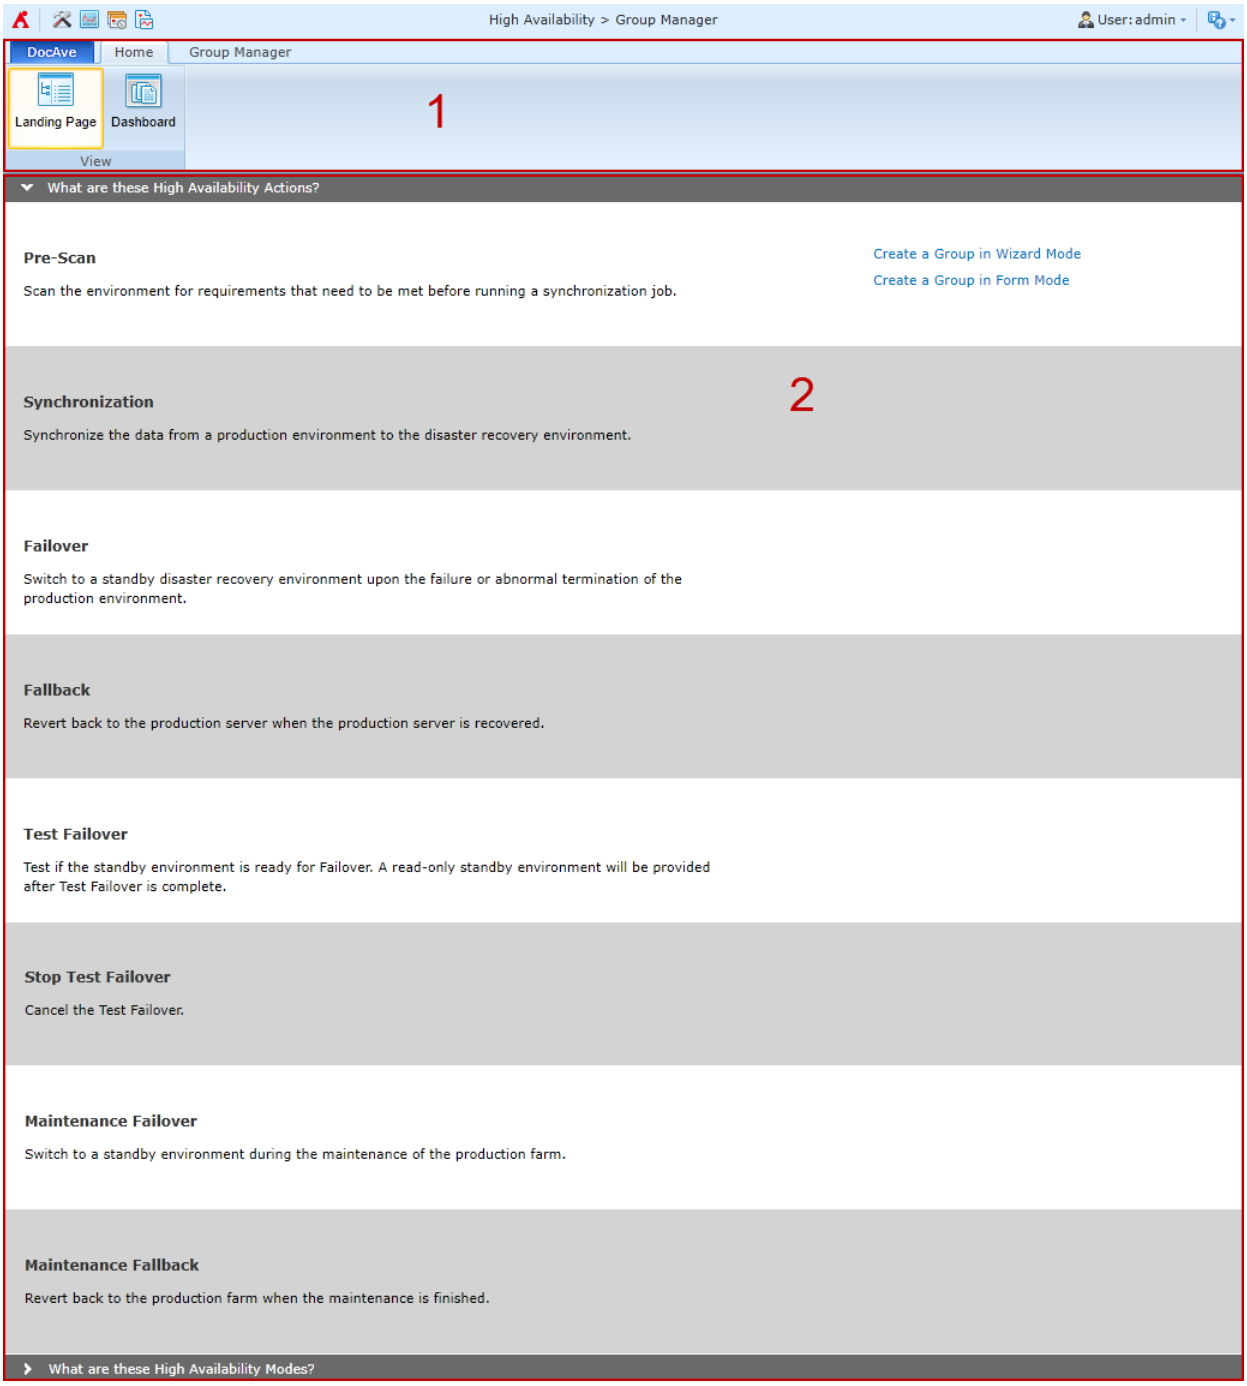

Figure 2: High Availability user interface.

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.

Navigating DocAve

DocAve mimics the look and feel of many Windows products, making for an intuitive and familiar working environment. While there are many windows, pop-up displays, and messages within DocAve products, they share similar features and are navigated in the same ways.

Below is a sample window in DocAve. It features a familiar, dynamic ribbon, and a searchable, content list view.

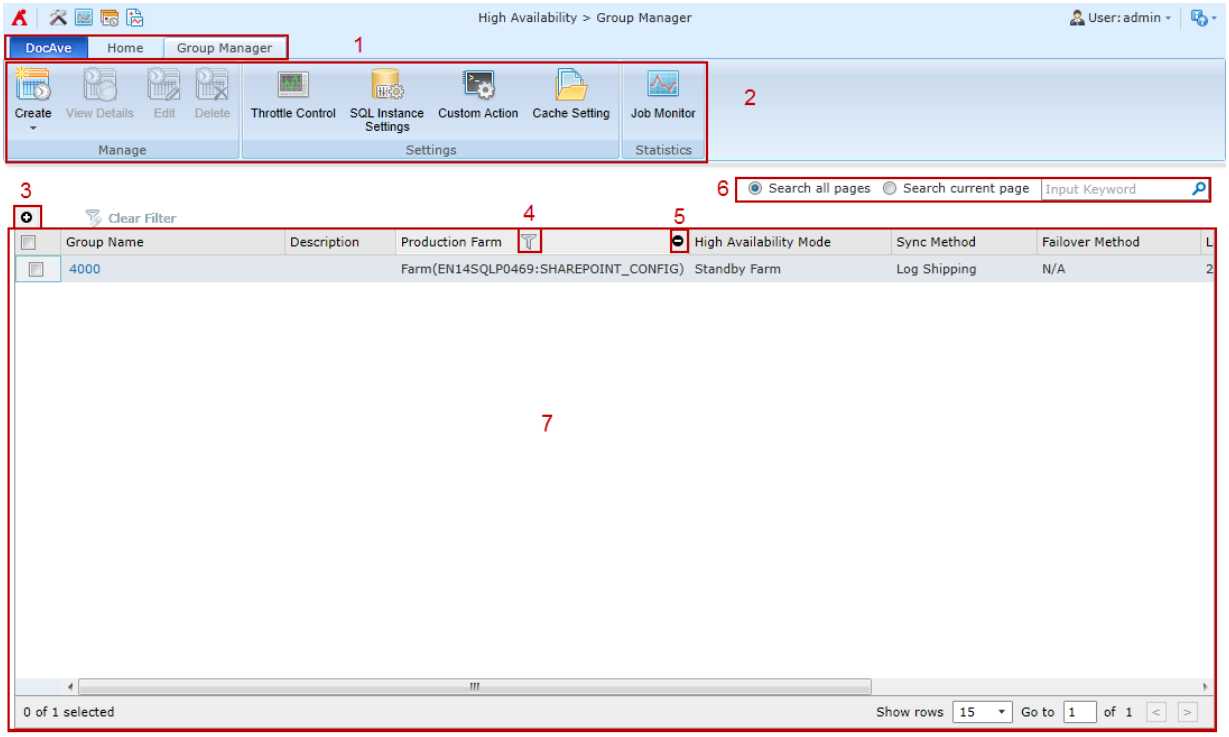

Figure 3: Navigating DocAve.

1. Ribbon Tabs— Allows users to navigate to the DocAve Welcome page and within the active module.
2. Ribbon Panes – Allows users to access the functionality of the active DocAve module.
3. Manage columns) – Allows users to manage which columns are displayed in the list. Click the manage columns) button, and then select the checkbox next to the column name in the drop-down list.
4. Filter the column) – Allows users to filter the information in the List View. Click the filter the column) button next to the column and then select the checkbox next to the column name.
5. Hide the column ( ) – Allows users to hide the selected column.
6. **Search** – Allows users to search the List View pane for a keyword or phrase. You can select **Search all pages** or **Search current page** to define the search scope.

***Note:** The search function is not case sensitive.

7. Management Pane – Displays the actionable content of the DocAve module.

About Throttle Control

Throttle Control ensures that jobs do not cause a noticeable degradation in network performance when running High Availability jobs. Throttle Control sets boundaries and limitations on jobs, which can help manage band-width for your networks.

***Note**: The Platform Backup Log Shipping sync method does not support throttle control settings.

Configuring a Throttle Control Profile

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

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.

1. **Profile Name** – Enter a profile name and the description for the Throttle Control profile.
2. **Run During** – Configure the job data transfer rate during working hours and non-working hours, respectively. Enter a positive integer into the text box and select **KB/Second** or **MB/Second** from the drop-down menu.
3. **Define Work Schedule** – Define the working schedule of the Throttle Control profile by **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.

Managing Throttle Control Profiles

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 way. You can 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, refer to Configuring a 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, refer to Configuring a 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. If the Throttle Control profile is currently being used by a High Availability group, the Throttle Control profile cannot be deleted.

Configuring a Master Key Password for SQL Instance

Using the **SQL Instance Settings** feature, you can configure a master key password for your SQL Server instance. If you do not configure the SQL Instance Settings, a default master key password will be automatically generated when creating endpoint with certificate for building mirroring relationship in the SQL Mirroring High Availability job, or performing synchronization for databases encrypted by Transparent Data Encryption (TDE).

***Note**: The SQL Instance Settings configured for a SQL Instance can only be modified but deleted after being configured.

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 SQL instance row, click **Configure** to configure the 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 master key password you enter must meet the Windows password policy requirements of the server that is running SQL Server.
5. Optionally, you can configure the location for storing the data files and log files of the SQL instance separately. By default, the default database locations for data and logs of this SQL instance is loaded.

   ***Note**: To configure the location for the SQL Cluster, make sure to enter the available shared disk 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.

About Custom Actions

The Custom Action feature, in conjunction with the command profile and script profile, executes the predefined scripts to execute the user-customized actions before or after a Synchronization, Failover, or Fallback.

Configuring a Script Profile

Script Profile allows users to select the script files in BAT, PS1, or EXE format from the Manager or Agent server and execute the scripts in the customized operation timing.

To configure a script profile, complete the steps below:

1. In the **Group Manager** tab, click **Custom Action** on ribbon. In the **Custom Action** interface, click the **Script Profile** tab, and then click **Create** to create a script profile.
2. In the pop-up window, enter the profile name and an optional description to the **Script Profile Name** field.
3. **Operation Timing** – Select when to execute this script from the drop-down list.
    - **Before Synchronization** – The script will be executed before starting the Synchronization.
    - **After Synchronization** – The script will be executed after the synchronization is finished.
    - **Before Failover** – The script will be executed before starting the Failover.
    - **After Failover** – The script will be executed after the Failover is finished.
    - **Before Fallback** – The script will be executed before starting the Fallback.
    - **After Fallback** – The script will be executed after the Fallback is finished.
4. **Script Type** –Select **Manager Script** to select and execute a script on the server where the Manager resides. Select **Agent Script** to specify an Agent and select a script file to execute on the selected Agent. If you select Agent Script, you are required to select an Agent from the drop-down list.
5. **Script File** – For **Manager Script**, all of the script files (.bat, .ps1, .exe) stored in the *…\AvePoint\DocAve6\Manager\Control\Config\HighAvailability\Scripts* directory on the Manager server will be loaded. For **Agent Script**, all of the script files (.bat, .ps1, .exe) stored in the *…\AvePoint\DocAve6\Agent\data\HighAvailability\Scripts* directory on the selected Agent server will be loaded. Select the script file from the drop-down list.
6. Click **Save** to save the configured script profile. Click **Cancel** to return to the script profile managing interface.

Managing Script Profiles

You can manage created script profiles via the **Custom Action** window. In the **Group Manager** tab, click **Custom Action** in the **Settings** group. The **Custom Action** interface appears. All of the previously-created **Command Profiles** are displayed in the main display pane. Click **Script Profile** tab on the ribbon to display all of the created script profiles.

In this interface, you can change the number of script profiles displayed per page. To change the number of the script profiles displayed per page, select the desired number from the **Show rows** drop-down menu in the lower right-hand corner. To sort the script profiles, click the column heading such as **Profile Name**, **Description**, **Script File**, **Script Type**, **Operation Timing**, and **Agent**. Perform the following actions from the ribbon to manage the script profiles:

- **Create** – Click the **Create** button on the ribbon to create a script profile.
- **View Details** – Select a previously created profile from the table, and then click **View Details** on the ribbon. You can view all of the settings of the selected profile, and click **Edit** to change the settings if desired.
- **Edit** – Select a script profile and click **Edit** on the ribbon to change the configurations for the selected script profile. You can change the settings and then click **OK** to save the modification, or click **Cancel** to exit without saving any changes.
- **Delete** – Select the script profile 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 script profile, or click **Cancel** to return without deleting the selected script profile. If the script profile is currently being added to a command profile, this script profile cannot be deleted.

Configuring a Command Profile

Command Profile allows users to select one or more script profiles and configure how to run the selected script profiles. You can include a predefined command profile to perform the customized features before or after the High Availability Synchronization, Failover, or Fallback. To configure a command profile, complete the steps below:

1. In the **Group Manager** tab, click **Custom Action** on ribbon. In the **Custom Action** interface, click the **Command Profile** tab, and then click **Create** to create a command profile.
2. In the pop-up window, enter a profile name and an optional description to the **Command Profile Name** field.
3. **Should High Availability wait for the script to complete?** – The feature only works on the Agent Scripts. Select **Yes** to let High Availability wait for the Agent scripts to complete. DocAve process and the Agent script process will be executed according to the operation timing you have selected, and DocAve process will continue after the user-defined Agent script is complete. Select **No** to execute the Agent script at the operation timing you selected in the script profile. High Availability will add a message to the job report that indicates the script process status. If

the status of the script in the job report is **Skipped**, which means the script execution is not completed, High Availability will not interfere the operation of the script until it finishes.

4. **Session Time Out** – If you selected **Yes** in step 3 above, designate a session time-out value. The script process will end automatically if it is not finished in the session.
5. **Processing Commands** –To execute the included scripts in sequential, select **Sequential**. To execute the included scripts in parallel, select **Parallel**.

   ***Note**: If you select **Sequential** as the way to run the scripts, you can define the order of the scripts being executed in the **Order** column of the **List of Scripts to Run** table.

6. **Add Script Profile** –Select one or more script profiles in the table on the left. Click the Add to Group ( ) button to add it or them to the table on the right. Click the delete ( ) button to remove the selected script profiles from the table.
7. To create a new script profile for this command profile, click the **Create a new script profile** link. Configure the script profile settings, and then save the profile.
8. Click **Save** to save the configuration to the command profile. Click **Cancel** to return to the command profile managing table.

Managing Command Profiles

For the created command profiles, you can manage them on the **Custom Action** window. In the **Group Manager** tab, click **Custom Action** in the **Settings** group. The **Custom Action** window appears. All of the previously-created command profiles are displayed in the main display pane. If you are already on the **Custom Action** interface, click **Command Profile** on the ribbon to display all of the previously created command profiles. In this interface, you can change the number of command profiles displayed per page. To change the number of the command profiles displayed per page, select the desired number from the **Show rows** drop-down menu in the lower right-hand corner. To sort the command profiles, click the column heading such as **Profile Name**, **Description**, and **Type**. Perform the following actions from the ribbon to manage the command profiles:

- **Create** – Click the **Create** button on the ribbon to create a command profile.
- **View Details** – Select a previously created profile from the table, and then click **View Details** on the ribbon. You can view all of the settings of the selected profile, and click **Edit** to change the settings if desired.
- **Edit** – Select a command profile and click **Edit** on the ribbon to change the configurations for the selected profile. You can change the settings and then click **OK** to save the modification, or click **Cancel** to exit without saving any changes.
- **Delete** – Select the command profile 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 command profile, or click **Cancel** to return without deleting the selected command profile. If the command profile is currently being used in a High Availability group, this command profile cannot be deleted.

Configuring Cache Setting

Cache Setting allows you to customize the location to store the necessary cache files used for mapping Connector devices and performing Log Shipping Fallback.

About Connector Cache Setting

Connector cache profile allows you to configure a local path or a UNC path to store the SQLite database file that is used for Connector device mapping.

Configuring a Connector Cache Setting Profile

To configure a Connector cache profile, complete the steps below:

1. In the **Group Manager** tab, click **Cache Setting** on the ribbon. In the **Cache Setting** interface, click the **Connector Cache Setting** tab, and then click **Create** to create a Connector cache profile.
2. In the pop-up window, enter the profile name and an optional description to the **Profile Name** field.
3. **Temporary Buffer** – Configure a local path or a UNC path to store the cache data used for Connector device mapping. Select the **Local Path** option or the **UNC Path** option.

   ***Note**: For SQL Server Cluster environment, you must use UNC path.

   - **Local Path** – Select an Agent from the drop-down list, and then enter a local path on that Agent to store the cache data.
   - **UNC Path** – Enter the UNC path. Select an existing managed account profile from the drop-down list or click the **New Managed Account Profile** link to create a new managed account profile for accessing the UNC path. For details on creating a new managed account profile, refer to Control Panel Reference Guide.
4. Click **Save** to save the configurations to this profile, or click **Cancel**.

Managing Connector Cache Setting Profiles

You can manage created Connector cache profiles via the **Cache Setting** window. In the **Group Manager** tab, click **Cache Setting** in the **Settings** group. The **Cache Setting** interface appears. All of the previously-created Connector cache profiles are displayed in the main display pane. If you are already on the **Cache Setting** interface, view all of the previously created Connector cache profiles under the **Connector Cache Setting** tab. In this interface, you can change the number of profiles displayed per page. To change the number of the profiles displayed per page, select the desired number from the **Show rows** drop-down menu in the lower right-hand corner. To sort the profiles, click the column heading such as **Profile Name**, **Description**, **Temporary Buffer**, and **Agent**. Perform the following actions from the ribbon to manage the cache profiles:

- **Create** – Click the **Create** button on the ribbon to create a Connector cache profile.

- **View Details** – Select a previously created profile from the table, and then click **View Details** on the ribbon. You can view all of the settings of the selected profile, and click **Edit** to change the settings if desired.
- **Edit** – Select a profile and click **Edit** on the ribbon to change the configurations for the selected profile. You can change the settings and then click **OK** to save the modification, or click **Cancel** to exit without saving any changes. Note that if you changed the location when editing a Connector cache profile, the cache data stored in the previous location will be transferred to the newly configured location while saving the modifications to the profile.
- **Delete** – Select the profile 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 profile, or click **Cancel** to return without deleting the selected profile. If the profile is currently being used in a High Availability group, this profile cannot be deleted.

About Log Shipping Cache Setting

Log Shipping cache profile allows you to designate an Agent to store the temporary log backup files for Log Shipping Fallback.

Configuring a Log Shipping Cache Profile

To configure a Log Shipping cache profile, complete the steps below:

1. In the **Group Manager** tab, click **Cache Setting** on the ribbon. In the **Cache Setting** interface, click the **Log Shipping Cache Setting** tab, and then click **Create** to create a Log Shipping cache profile.
2. In the pop-up window, enter the profile name and an optional description to the **Profile Name** field.
3. **Agent** – Select an Agent from the drop-down list to store the Log Shipping cache data. The Agents that already have cache setting configured will not be displayed in the drop-down list.
4. **Temporary Buffer** – Configure a local path or a UNC path to store the cache data used for the Log Shipping.
    - **Local Path** option – Enter a local path on that Agent to store the cache data.
    - **UNC Path** option – Enter the UNC path. Select an existing managed account profile from the drop-down list or click the **New Managed Account Profile** link to create a new managed account profile for accessing the UNC path. For details on creating a new managed account profile, refer to Control Panel Reference Guide.
5. Click **Save** to save the configurations to this profile, or click **Cancel**.

Managing Log Shipping Cache Profiles

You can manage Log Shipping cache profiles via the **Cache Setting** window. In the **Group Manager** tab, click **Cache Setting** in the **Settings** group. The **Cache Setting** interface appears. All of the Connector cache profiles are displayed in the main display pane. Click the **Log Shipping Cache Setting** tab on the ribbon to display all of the Log Shipping cache profiles.

In this interface, you can change the number of profiles displayed per page. To change the number of the profiles displayed per page, select the desired number from the **Show rows** drop-down menu in the lower right-hand corner. To sort the profiles, click the column heading such as **Profile Name**, **Description**, **Temporary Buffer**, and **Agent**. To manage the profiles use the following actions from the ribbon:

- **Create** – Click the **Create** button on the ribbon to create a profile.
- **View Details** – Select a profile from the table, and then click **View Details** on the ribbon. You can view all of the settings of the selected profile, and click **Edit** to change the settings if desired.
- **Edit** – Select a profile and click **Edit** on the ribbon to change the configurations for the selected cache profile. You can change the settings and then click **OK** to save the modification, or click **Cancel** to exit without saving any changes. Note that if you edited a Log Shipping cache profile, the cache data stored in the previous location will be transferred to the newly configured location while saving the modifications to the profile.
- **Delete** – Select the profile that you want to delete and click **Delete** on the ribbon. A confirmation window will pop up and ask if you are sure that you want to proceed with the deletion. If the profile currently being used in a High Availability group is deleted, the Fallback job will fail.

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.

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

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. AvePoint recommends adding the source Configuration database and Central Administration database to the High Availability Group.

The 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. Click **Next**.
2. The **General Settings** interface appears, where you can select the **High Availability Mode**, **Sync Method**, and configure the setting for Synchronization, Failover, and Fallback.
    a. **High Availability Mode** – Choose the High Availability Mode to use according to your environment. In this condition, select **Single farm** option.
    b. **Sync Method** – Select the **SQL Mirroring**, **AlwaysOn Availability Group**, **Log Shipping,** or **Platform Backup Log Shipping** from the sync method drop-down list. See below for more information on the four Sync Methods.
        - **Log Shipping** performs replication at the log level. The 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. High Availability improved the data transfer logic of Log Shipping Failover and Fallback in DocAve 6 Service Pack 5. For details, refer to Improving Performance of Log Shipping Fallback.

          ***Note**: The Log Shipping sync method supports protecting the databases in AlwaysOn Availability Group for High Availability.

          ***Note**: If the selected 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 Windows Authentication (default) or DocAve certificate. If you use SQL Mirroring as the sync method, you must select a sync mode, **Asynchronous** or **Synchronous** for performing the High Availability Synchronization. For more detailed information, refer to Database Mirroring Operating Modes.

  - **Asynchronous** − In **Asynchronous** mode, a high-performance mode, the production environment can commit transactions without an immediate response from the standby environment.
  - **Synchronous** − In **Synchronous** mode, a high-safety mode, the production environment will wait for an immediate response from the standby environment.

  ***Note**: If the endpoint is created by certificate when setting up the SQL Mirroring relationship, a login with **DocAve** as the prefix of the login name will be created in the SQL Server.

  ***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 databases 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 or several sets of corresponding secondary databases. Optionally, secondary databases can be made available for read-only access and/or some backup operations.

***Note**: When using the AlwaysOn Availability Group as the sync method to create a High Availability Single farm mode group, you must add all of the databases in the production farm to the same AlwaysOn Availability Group and create a SQL Alias to point to the listener name after performing Failover.

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

c. **Would you like to include Storage Optimization data?** – Choose whether or not to include the Storage Manager BLOB data or Connector BLOB data in this group for High Availability configuration. If you select the **Connector BLOB Data** option, you are required to choose the way in which Connector BLOB data is synchronized. If you want to synchronize the Connector BLOB data using DocAve High Availability, select the **Use High Availability to synchronize Connector BLOB data** option. If you prefer to use your own way for Connector BLOB data synchronization, deselect this option. In this case, after each Synchronization, you must go to the Connector cache location designated in the selected Connector cache profile, find the SQLite database file with mapping information, and place the Connector BLOB data in the proper location according to the recorded mapping information.

***Note**: Only when the **Storage Manager BLOB Data** or **Connector BLOB Data** option is selected, can the associated stub databases of the content databases be displayed in the tree.

***Note**: High Availability jobs do not support automatically transferring Storage Manager BLOB data. Therefore, you must manually copy the Storage Manager BLOB data to the mapped destination physical devices after each High Availability job.
After the High Availability Failover job completes, the newly added Storage Manager BLOB data will at first be written to the device designated in Storage Manager rules; if this write fails, the BLOB data will be written to the mapped physical device that is designated in the High Availability group. Make sure that the BLOB data in the mapped device will be transferred to the production physical device after Fallback.

***Note**: High Availability only supports the Connector BLOB data stored in a Net Share or NFS device.

***Note**: Since the Synchronization job of **Platform Backup Log Shipping** method synchronizes the Connector BLOB data in the production environment directly, instead of using the Platform Backup data, you must minimize the job interval between the

Platform Backup job and the High Availability Synchronization job to make sure the data consistency.

***Note**: If the Agent account does not have Full Control permission to the standby Web application where the Connector libraries reside, the High Availability Fallback job may fail to load the Connector data from the standby farm and will not fallback the Connector data.

d. **Production Farm** – Select the SharePoint farm from the drop-down list.

e. **Production Agent Group** – Select a production Agent group from the drop-down list for load balancing and job performance.

f. **Failover Method** – Select the failover method:

- If you selected **Log Shipping** or **Platform Backup Log Shipping** as the **Sync Method**, only the **SQL Alias** method is available.
- If you selected **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 selected the **SharePoint Failover Server** method, DocAve will use the failover server configured in the SharePoint environment for the mirrored databases.

  ***Note:** The **SharePoint Failover Server** is available only if you selected **SQL Mirroring** as the **Sync Method**.

- If you selected 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.

g. **Failover Options** – Select your desired failover option:

- If you selected **SQL Mirroring** as the **Sync Method** and **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 **Log Shipping** as the **Sync Method** and **SQL Alias** as the **Failover Method**, you can select the **Keep database in read-only mode to perform a one-way Failover** option to perform a one-way Failover with read-only standby environment. If you want to perform a Fallback after the one-way Failover, you must manually set the status of the standby databases to normal.
- If you selected **Platform Backup Log Shipping** as the **Sync Method** and **SQL Alias** as the **Failover Method**, you can select the **Keep database in read-only mode** option to keep the secondary database in read-only mode.
- If you selected **Log Shipping** as the **Sync Method** and **SQL Alias** as the **Failover Method**, you can select the **Perform an incremental synchronization job before Failover** option to run an incremental synchronization job before the Failover to minimize the data loss, however, the performance of Failover job will be affected.

Selecting the **Perform an incremental synchronization job before Failover** option will enable the **Snapshot Retention Policy** settings. Refer to the instructions below to configure Snapshot Retention Policy settings:

- Enter the maximum number of snapshots you wish to retain into the text box. By default, the number is **32**. It is not recommended that this value is larger than **64**.
- Click **Advanced Settings** to further customize the threshold policy. Define which policy DocAve High Availability will apply while the amount of snapshots on your device exceeds the entered value.
  - **Skip the incremental synchronization job and do nothing to the oldest snapshot and job** – The current group's new incremental synchronization job will be skipped, and all of the snapshots are kept.
  - **Delete the earliest snapshot of this group and start the new job** – The current group's oldest snapshots are deleted directly from the storage, and a new job for the current group starts.

***Note**: If you are using Microsoft SQL Cluster and you select the **Perform an incremental synchronization job before Failover** option for the High Availability group, AvePoint recommends using the same UNC path for **Temporary Buffer** of each Agent installed on the nodes of SQL Cluster, which is configured through **Control Panel** > **Agent Monitor** > **Configure**. Otherwise, the Log Shipping Fallback job may not be able to get the incremental synchronization data from the temporary location.

h. **Log Shipping Fallback Options** – This field is only available for the **Log Shipping** method. The Log Shipping Fallback operation will rename the original production database, and you can choose to keep or delete the renamed original production database after Fallback job is finished.

   ***Note**: If you have enough free disk space, AvePoint recommends selecting the **Keep the renamed original production database after the Fallback job has completed** option in case that the data resynchronized from standby farm has errors contained.

i. **Connector Data Fallback Options** – This field is only available when you select to include the Connector BLOB data of the selected components into the group. You can choose whether or not to fallback the Connector data from the standby farm to production while performing the Fallback job.

j. 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 A: SharePoint Components Supported for High Availability

4. **Connector Device Mapping** – This field is only available when you have selected to synchronize the Connector BLOB data. You must configure the mapping settings to map the production physical device that contains Connector BLOB data to the standby physical device while adding the components to the group.

    - All of the configured physical device mappings are displayed in the table of right pane.
    - Select a Connector cache profile from the drop-down list, or create a new one to store the SQLite database file used for the Connector device mapping. For details on creating a Connector cache profile, refer to Configuring a Connector Cache Setting Profile.
    - The Agent you select in the **Add to Group** step used to execute the synchronization of the Connector BLOB data is displayed as the default Agent.
    - The standby physical device selected in the **Add to Group** step is displayed as the default standby physical device.

    You can change them as desired in this field. The mapping information will be recorded and updated in a SQLite database file under the cache location specified in the selected Connector cache profile.

5. **SQL Alias Settings** – If you set **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 edit the port for the SQL Alias.

    b. In the pop-up window, the **Dynamically determine port** option is selected by default. With this option selected, you can use the SQL Browser to determine the port dynamically. Also, you can deselect the **Dynamically determine port** option and enter an integer into the **Port** text box to specify a port for the SQL Alias.

    c. Click **Save** to save the configuration. Click **Cancel** to go back without saving any changes.

6. **SharePoint Server Settings** – Select the SharePoint servers from the table, and the SQL Alias 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.

7. **Throttle Control** – Select a throttle control profile from the drop-down list to define the data transfer rate during 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.

8. **Custom Action** – Select a command profile to run the scripts you selected. Select a previously created command profile from the drop-down list, and click the **View** link to view the detailed settings of the selected command profile. Also, you can create a new command profile by selecting **New Command Profile** from the drop-down list. For details on custom actions, refer to About Custom Action.

9. **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.
    - **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** – Synchronize all of the selected data from the production environment to the standby environment. Designate 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** – Synchronize data changes 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. Select the time range of the job occurrence in the **Range of Recurrence** field. Designate 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, but synchronize nothing.

            - **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 define 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.

10. **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.
11. Select **Next**. The **Pre-Scan** interface appears.

12. 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.
13. 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.
14. 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 to ensure the availability by using a standby farm to synchronize the SharePoint components between one SharePoint instance and your disaster recovery (DR) instance. In order to accomplish this, a secondary farm must be configured with a DocAve High Availability Agent installed.

Standby farm mode requires you to select the production farm and the standby farm respectively. After a 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**. Select **Next**.
2. The **General Settings** interface appears, where you will select the High Availability Mode, Sync Method, and configure the general settings for Synchronization, Failover, and Fallback.
    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. **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 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. High Availability improved the data transfer logic of Log Shipping Failover and Fallback in DocAve 6 Service Pack 5. For details, refer to Improving Performance of Log Shipping Fallback

          ***Note**: If the selected 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 and standby servers by Windows Authentication (default) or by DocAve certificate. If you use SQL Mirroring as the sync method, you must select a sync mode, **Asynchronous** or **Synchronous** for performing the High Availability Synchronization. For more detailed information, refer to Database Mirroring Operating Modes.

- **Asynchronous** – In **Asynchronous** mode, a high-performance mode, the production environment can commit transactions without an immediate response from the standby environment.
- **Synchronous** – In **Synchronous** mode, a high-safety mode, the production environment will wait for an immediate response from the standby environment.

***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.
  You can add more than one standby farm for the production farm if you create a Standby farm mode High Availability group using AlwaysOn Availability Group sync method.

  ***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 selected components to the destination. This method requires users to have Platform Backup data of the selected components prepared.

c. **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. Consider the following issues when configuring this option:

- If you are about to synchronize the content database with BLOB data and related stub database together to the standby farm with read-only view enabled, make sure the Agent account to start the **SP2010StorageOptimizationService.exe** process or **SP2013StorageOptimizationService.exe** process has sufficient permissions in the following scenarios before performing the synchronization job.
  - If the Agent account in the standby farm is a different user in the same domain as the Agent account in the production 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 the Agent account in the standby farm is in a different domain as the Agent account of the production farm, it is recommended making the domain in the production farm trusted by the domain in the standby farm and granting the Agent account in the standby farm the **db_owner** role in the production stub database. Otherwise, the Agent account in the standby farm must have the **sysadmin** role to the standby SQL instance, 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**. For details on selecting a standby farm and a secondary replica in AlwaysOn Availability Group for read-only view, refer to Click Next to go to the **Advanced Settings** step.
- **Secondary Replica for Read-Only View**.

  ***Note**: After the synchronization for SharePoint 2010 environment, the standby site collection will appear as read-write, but it is actually read-only. Although the **Database Read-Only** property appears as "**No**", the database is actually read-only in the AlwaysOn Availability Group Secondary Replica.
- If you design this group for maintenance purpose, AvePoint strongly recommends selecting the **Enable read-only view after synchronization** option, in case that any service application database in Mirroring session or AlwaysOn Availability Group may fail the upgrade and maintenance.
- By default, the service applications in the group with the **Enable read-only view after synchronization** option selected will not be created in the standby farm

for implementing warm standby after synchronization. If you want to make the service application accessible after synchronization in the standby farm, refer to Defining Warm Standby Behavior of Service Applications for details to change the warm standby behavior of service applications.

***Note**: The service application can only be synchronized to standby environment once, since the service application databases synchronized to the standby environment in this case is normal. If you want to perform the Synchronization for the service applications in this group again, you must manually delete the service application from the standby farm.

d. **Would you like to include Storage Optimization data?** – Choose whether or not to include the Storage Manager BLOB data or Connector BLOB data of the selected components in this group for High Availability configuration, and choose whether or not to automatically enable the EBS/RBS setting during Failover for the standby content database.

   ***Note**: Only when the **Storage Manager BLOB Data** or **Connector BLOB Data** option is selected, can the associated stub databases of the content databases be displayed in the tree.

   - **Storage Manager BLOB Data** – Select the **Storage Manager BLOB Data** option to include the Storage Manager BLOB data into the synchronization. The **Automatically enable EBS/RBS during Failover** option appears and it is selected by default.

     ***Note**: High Availability jobs do not support automatically transferring Storage Manager BLOB data. Therefore, you must manually copy the Storage Manager BLOB data to the mapped destination physical devices after each High Availability job.
     After the High Availability Failover job completes, the newly added Storage Manager BLOB data will at first be written to the device designated in Storage Manager rules; if this write fails, the BLOB data will be written to the mapped physical device that is designated in the High Availability group. Make sure that the BLOB data in the mapped device will be transferred to the production physical device after Fallback.

   - **Connector BLOB Data** – Select the **Connector BLOB Data** option to include the Connector BLOB Data into the synchronization. You must select whether or not to use High Availability to synchronize Connector BLOB data, which will allow DocAve High Availability to perform the synchronization of BLOB data. If you prefer to use your own way to synchronize Connector BLOB data, deselect this option. In this case, after each Synchronization, you must go to the Connector cache location designated in the selected Connector cache profile, find the SQLite database file with mapping information updated, and place the Connector BLOB data in the proper location according to the recorded mapping information.

***Note**: High Availability only supports the Connector BLOB data stored in a Net Share or NFS device.

***Note**: Since the Synchronization job of **Platform Backup Log Shipping** method synchronizes the Connector BLOB data in the production environment directly, instead of using the Platform Backup data, you must minimize the job interval between the Platform Backup job and the High Availability Synchronization job to make sure the data consistency.

***Note**: If the Agent account does not have Full Control permission to the standby Web application containing the Connector libraries, the High Availability Fallback job may fail to load the Connector data from the standby farm and cannot fallback the Connector data from the standby farm to the production.

- **Automatically enable EBS/RBS during Failover** – Choose whether to automatically enable EBS/RBS in the standby environment during Failover. With this option selected, if the production farm has EBS enabled, the EBS settings will be enabled in the standby farm during Failover. If the production content database has RBS enabled, the RBS settings will be enabled for the standby content database during Failover; however, the standby content database with RBS enabled will not be kept in read-only mode.

e. **Production Farm** – Select the production farm that you want to perform the High Availability from the drop-down list.

f. **Production Agent Group** – Select a production Agent group from the drop-down list for load balancing and job performance.

g. **Standby Farm** – Select the Standby farm from the drop-down list as you wish.

h. **Standby Agent Group** – Select a standby Agent group from the drop-down list for load balancing and job performance improvement.

   i. ***Note**: If you select **AlwaysOn Availability Group** as the sync method, the **Standby Farm and Standby Agent Group** field will appear, instead of the **Standby Farm** field and **Standby Agent Group** field. You can add more than one standby farm for the High Availability of your business. Follow the steps below to add multiple standby farms: Click **Add a Standby Farm** beneath the table on the right side. A row will be added to the table.

   ii. Select a standby farm from the drop-down list under the **Standby Farm** column, and then select an Agent group for this standby farm to execute the High Availability jobs. You can click **View** to view the details of the selected Agent group, or click the delete) button to remove this record.

   iii. Repeat the steps above to add multiple standby farms.

i. **Failover Options** – Select the desired failover options according to the **Sync Method** you selected.

- **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 Platform Backup Log Shipping method.
- **Keep database in read-only mode to perform a one-way Failover** – DocAve will perform a one-way Log Shipping Failover and keep the standby databases in read-only mode. If you want to perform a Fallback after the Failover, you must manually set the status of the standby databases to normal. This option is specific for the Log Shipping method.

  ***Note:** This option does not support the service application 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.

- If you select the **Log Shipping** as the **Sync Method**, you can select the **Perform an incremental synchronization job before Failover** option to minimize the data loss, but the performance of Failover job will be affected.

  Selecting the **Perform an incremental synchronization job before Failover** option will enable the **Snapshot Retention Policy** settings. Refer to the instructions below to configure the Snapshot Retention Policy settings:

  - Enter the maximum number of snapshots you wish to retain into the text box. By default, the number is **32**. It is not recommended that this value be larger than **64**.
  - Click **Advanced Settings** to further customize the threshold policy. Define which policy DocAve High Availability will apply if the amount of snapshots on your device exceeds the entered value.
    - **Skip the incremental synchronization job and do nothing to the oldest snapshot and job** – The current group's new incremental synchronization job will be skipped, and all of the snapshots are kept.

- **Delete the earliest snapshot of this group and start the new job** – The current group's oldest snapshots are deleted directly from the storage, and a new job for the current group starts.

***Note**: If you are using Microsoft SQL Cluster and you select the **Perform an incremental synchronization job before Failover** option for the High Availability group, AvePoint recommends configuring the same UNC path for **Temporary Buffer** of each Agent installed on the nodes of SQL Cluster, which is configured through **Control Panel** > **Agent Monitor** > **Configure**. Otherwise, the Log Shipping Fallback job may not be able to get the incremental synchronization data from the temporary location.

j. **Log Shipping Fallback Options** – This field is only available for the **Log Shipping** method. The Log Shipping Fallback operation will rename the original production database, and you can choose to keep or delete the renamed original production database after Fallback job is finished.

***Note**: If you have enough free disk space, AvePoint recommends selecting the **Keep the renamed original production database after the Fallback job has completed** option in case that the data resynchronized from standby farm has errors contained.

k. **Connector Data Fallback Options** – This field is only available when you select to include the Connector BLOB data of the selected components into the group. You can choose whether or not to fallback the Connector data from the standby farm to the production while performing the Fallback job.

l. 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 A: SharePoint 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. Click **Next** to go to the **Advanced Settings** step.

5. **Secondary Replica for Read-Only View** – This field is only available when you have selected the **Enable read-only view after synchronization** option for having read-only standby farms after Synchronization job.

   a. Click **Add a Standby Farm** to add a record for the standby farm that you want to make it read-only after the synchronization. The AlwaysOn Availability Groups where the SharePoint databases of the standby farm are included will be listed.

b. Select a secondary replica of the AlwaysOn Availability Group. Synchronization job will set the **Readable secondary** option to **Yes** in SQL Server to make the selected secondary replicas read-only.

***Note**: The selected secondary replica will be used as the primary replica after Failover by default, and it cannot be changed.

You can click the Delete) button to remove the record if you do not want that standby farm to be read-only after Synchronization.

6. **Connector Device Mapping** – This field is only available when you have selected to synchronize the Connector BLOB data. You must configure the mapping settings to map the production physical device that contains Connector BLOB data to the standby physical device while adding the components to the group.

   - All of the configured physical device mappings are displayed in the table of right pane.
   - Select a Connector cache profile from the drop-down list, or create a new one to store the SQLite database file used for the Connector device mapping. For details on creating a Connector cache profile, refer to Configuring a Connector Cache Setting Profile.
   - The Agent that you selected in the **Add to Group** step, which is used to execute the synchronization of the Connector BLOB data, is displayed as the default Agent.
   - The standby physical device selected in the **Add to Group** step is displayed as the default standby physical device.

   You can change the device mapping settings as desired in this field. The mapping information will be recorded and updated in a SQLite database file under the cache location specified in the selected Connector cache profile.

7. **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.
8. **Custom Action** – Select a command profile to run the scripts as you desired. Select a previously created command profile from the drop-down list, and you can click the **View** link to view the detailed settings of the selected command profile. Also, you can create a new command profile by selecting **New Command Profile** from the drop-down list. For details on custom actions, refer to About Custom Action.
9. **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. Designate 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 between 1 and 65535 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. Select the time range of the job occurrence in the **Range of Recurrence** field. Designate a start time earlier that 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 stand by 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 define 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.

10. **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.
11. Select **Next**. The **Pre-Scan** interface appears.
12. 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.
13. 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.
14. 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 users who are familiar with creating 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 on the configurable options.

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 SharePoint components to a High Availability Group using the SQL Mirroring, Log Shipping, or Platform Backup Log Shipping Sync Methods, 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**: The databases with the status of Mirroring, Read-Only, Restoring, or Standby Read-Only will be grayed out on the tree and unable to be selected.

   ***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, refer to Adding Custom Databases to a Group.
3. If you are selecting databases, you must enter and test the **Destination SQL instance**. Enter the destination SQL instance name into the text box. Click **Test**. After a successful test, 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 a 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 synchronization, the **Database Snapshot Location** field will appear with the default locations provided. You can configure the locations to store the database snapshot. 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.

   ***Note**: Database snapshot is the key to the read-only view of the standby environment. The database snapshots created by DocAve High Availability are named

"**DatabaseName_JobID_HASnapshot**" The database snapshot is updated according to the High Availability Synchronization job schedule in order to keep the snapshot data up-to-date. The database snapshot in the standby farm is similar to the production database, but it includes the stub related information of the production database. The stubs in the standby farm can be accessed too, but in read-only format. AvePoint recommends not storing database snapshots in the same volume as the database location.

5. If you selected to synchronize the Storage Manager BLOB data, you must configure the settings in the **Standby Physical Device** field. The device types of the physical devices configured in DocAve Manager Control Panel that contain BLOB data are all listed in the **Device Type** column of the table in the right pane. For each device type listed in the table, select a standby physical device with the same device type from the Standby Physical Device drop-down list. 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**: If you add the **Include New** node under the Web application into the High Availability group, the physical device that is created after saving the group whose device type does not have standby physical device configured in the group will not be included in the Synchronization. To include the Storage Manager BLOB data that is stored in the physical device whose device type has no standby physical device configured in the current group, you must edit the group and edit the settings for the **Include New** node. In the window for editing settings of **Include New**, the newly discovered device type will be displayed under **Device Type** column. Select a standby physical device in the same device type for mapping production physical devices of this device type and save the configuration. The settings will be applied to all of the content databases under the same Web application, that is, the selected standby physical device will be used by default if any content database under the same Web application has configured to use physical device in this device type to store Storage Manager BLOB data.

6. If you selected to synchronize the Connector BLOB data, you must configure the settings in the **Standby Agent and Physical Device for Connector BLOB Data** field.

   a. Select an Agent to execute the Connector BLOB synchronization and process the newly added Connector data.

      ***Note**: For **Single farm** High Availability mode, AvePoint recommends using the Agent with SharePoint Web front-end server installed in the production farm to execute jobs or selecting an Agent installed on the server that is not in this production farm. For **Standby farm** High Availability mode, AvePoint recommends using an Agent that is not affected by the Failover job and not included in either production farm or standby farm.

   b. Select a standby physical device from the drop-down list for mapping the Connector BLOB data in the production farm. You can click **View** to go to the **Storage Configuration** interface 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**: The selected Agent and standby physical device will be used by default to map the newly added Connector data of the production environment.

Click **Save** to save the configurations and go back to the **Add to Group** step. Click **Cancel** to go back to the **Add to Group** step without saving any configurations.

7. The selected nodes displays in the table. 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 Storage Manager BLOB data for the selected components and configured the standby physical device, the **Standby Physical Device** field will appear. This field will display the selected 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 or service applications in the selected nodes for the Standby farm mode High Availability Group, AvePoint recommends that you view and configure the settings of the standby Web applications or service 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 SQL Server databases 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. Select an existing AlwaysOn Availability Group from the **AlwaysOn Availability Group** drop-down list.

4. From the table to the right of the **Secondary Replicas for Synchronization** field, select one or more secondary availability replicas as the synchronization destinations. You can click **Configure** under the **File Location** column to edit the location for storing the database files of the corresponding replica.
5. Choose a replica from the **Primary Replica for Failover** drop-down list as the primary replica to host the primary database after the failover. This field is only available for **Single farm** mode High Availability Group.
6. Enter the availability group listener name into the **Availability Group Listener** text box for the connectivity of the availability group.
    - In **Single farm** mode configurations, you can enter the SQL Alias or IP address of the listener into the text box.
    - In **Standby farm** mode configurations, you will not configure the Availability Group Listener and this field will not appear on the interface.
7. If you selected to synchronize the Storage Manager BLOB data, you must configure the settings in the **Standby Physical Device** field. The device types of the production physical devices that contain BLOB data are all listed in the **Device Type** column of the table in the right pane. For each device type listed in the table, select a standby physical device with the same device type from the **Standby Physical Device** drop-down list. You can click **View** to go to the **Storage Configuration** interface 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**: If you add the **Include New** node under the Web application into the High Availability group, the physical device that is created after saving the group whose device type does not have the standby physical device configured in the group will not be included in the Synchronization. To include the Storage Manager BLOB data that is stored in the physical device whose device type has no standby physical device configured in the current group, you must edit the group and edit the settings for the **Include New** node. In the window for editing settings of **Include New**, the newly discovered device type will be displayed under the **Device Type** column. Select a standby physical device in the same device type for this device type and save the configuration. The settings will be applied to all of the content databases under the same Web application, that is, the selected standby physical device will be used by default if any content database under the same Web application has been configured to use the physical device in this device type to store Storage Manager BLOB data.

8. If you selected to synchronize the Connector BLOB data, you must configure the settings in the **Standby Agent and Physical Device for Connector BLOB Data** field.
    a. Select an Agent to execute the Connector BLOB synchronization and store the SQLite database file.

        ***Note**: For Single farm High Availability mode, AvePoint recommends using the Agent with SharePoint Web front-end server installed in the production farm to execute jobs

or selecting an Agent installed on the server that is not in this production farm. For Standby farm High Availability mode, AvePoint recommends using an Agent that is not affected by the Failover job and not included in either production farm or standby farm.

b. Select a standby physical device from the drop-down list for mapping the Connector BLOB data in the production farm. 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**: The selected Agent and standby physical device will be used by default to map the newly added Connector physical device of the production environment.

9. 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.
10. The selected nodes displays in the table. 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 Storage Manager BLOB data for the selected components and configure the standby physical device, the **Standby Physical Device** field will appear. This field will display the selected 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 corresponding node from the High Availability Group.

***Note**: If there are Web applications or service applications in the selected nodes for the Standby farm mode High Availability Group, AvePoint recommends that you view and configure the settings of the standby Web applications or service applications first before you go to the next step. In addition, if you have added multiple standby farms into this group, check and configure the Web application settings and service application settings in each standby farm.

Adding Custom Databases to a Group

To implement High Availability functions in the SQL Server databases of third-party applications, but not part of your SharePoint component hierarchy, you can use the **Custom Database** node to add the database to the farm component tree in order to synchronize the database to the standby environment during the High Availability Synchronization job.

***Note**: The databases in the status of Mirroring, Read-Only, Restoring, or Standby Read-Only are not supported being added to the group and will not be displayed in the database tree.

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** (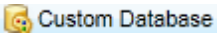

) button to display the SQL servers that have DocAve Agent installed. Click the Agent name to load all of the SQL instances.
3. Click the SQL instance name to display all of the databases that reside in the SQL instance.
4. Select the desired databases to add to the farm component tree.
5. Click **Save** 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

Use Group Manager to manage High Availability groups. After launching the DocAve High Availability interface, click **Group Manager** next to the **Home** tab. In the **Group Manger** 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**. 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. If the group includes the service applications with the **Enable read-only view after synchronization** option selected, and a Synchronization job has been run using this group, the settings for the service applications in this group cannot be changed; if this group does not have the **Enable read-only view after synchronization** option selected, the settings of the service applications will be not allowed to change after Failover. You can refer to the message under the **Last Action Status** column in **Dashboard** to tell whether the group has been run. If the group is currently being used to run a High Availability job, this group cannot be edited.
- **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. If the group is currently being used to run a High Availability job, this group cannot be deleted.
- **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, refer to About Throttle Control.
- **SQL Instance Settings** – Click **SQL Instance Settings** to configure the SQL instance settings. For more information on SQL instance settings, refer to Configuring a Master Key Password for SQL Instance.

- **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, refer to DocAve 6 Job Monitor Reference Guide.

About the Dashboard

The Dashboard interface provides a detailed view of the production farms, standby databases, High Availability Groups, and High Availability jobs. You can perform High Availability jobs from the ribbon, such as Pre-Scan, Synchronization, Failover, Fallback, Maintenance Failover, and Maintenance 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 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: **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 determine if 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 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.

- **Pre-Scan** – AvePoint recommends performing a Pre-Scan job before synchronization. Select the checkbox beside the desired nodes, and then you can select **Pre-Scan** on the ribbon to check if 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 B: 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.

  ***Note**: You can force High Availability to perform Full Synchronization on Connector data in each Synchronization job via modifying the configuration file. For details, refer to Forcing a Full Synchronization or Full Fallback on Connector Data.

- **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.

***Note**: After you perform a Failover job, the synchronization schedule of the group will be disabled, and the scheduled Synchronization jobs will be **Skipped** after the Failover job is performed. Performing a Fallback job will re-enable the Synchronization schedule.

- **Fallback** – Revert back to the production server when the production server is recovered. Select the checkbox beside 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 a fallback to the production server after you perform the High Availability Failover job.

  ***Note**: You can force High Availability to perform a Full Fallback on Connector data in the Fallback job via modifying configuration file. For details, refer to Forcing a Full Synchronization or Full Fallback on Connector Data.

- **Test Failover** – Test if the standby farm is ready for Failover. You can get a read-only standby farm after the Test Failover.

  ***Note**: Test Failover does not support service applications.

- **Stop Test Failover** – Roll back to the original status of the production farm and standby farm before performing the Test Failover.

  ***Note:** The Test Failover job and Stop Test Failover job will disable the Synchronization schedule settings of the groups, skipping the Synchronization jobs. You can perform a Pre-Scan job or a Synchronization job to enable the group's synchronization schedule settings, after the Stop Test Failover is completed.

For the purposes of upgrading and maintaining production and standby farms, High Availability provides **Maintenance Failover** and **Maintenance Fallback** for the Standby farm High Availability mode. To perform Maintenance Failover or Maintenance Fallback, all of the Standby farm mode High Availability groups created for the same production farm must be selected.

- **Maintenance Failover** – Switches to a standby environment during the maintenance of the production farm. You can choose to have a **Read-only** standby environment or a **Read-Write** standby environment. By selecting the **Read-write** option, Maintenance Failover will bring the production database offline and enable users to read and write in the standby environment while the production farm is upgrading or is under

maintenance. If it is your first time to perform Maintenance Failover, click **Maintenance Tutorials** on the ribbon of the Maintenance Failover window to learn more.

***Note**: To ensure the successful update of the production SharePoint farm:

If the High Availability groups use SQL Mirroring, the Maintenance Failover job will set the production service application databases to normal, and Maintenance Fall back will skip these databases.

If the High Availability groups use the AlwaysOn Availability Group sync method, the Maintenance Failover job for a **Read-Only** standby environment will remove the database from the new primary replica, and Maintenance Fallback will skip these databases.

If the High Availability groups use the AlwaysOn Availability Group sync method, the Maintenance Failover job for a **Read-Write** standby environment will remove database from every replica, and Maintenance Fallback will skip these databases.

- **Maintenance Fallback** – Reverts back to the production farm when maintenance of the production farm is finished. Before the upgrade and maintenance of a standby farm, a Maintenance Fallback job must be performed after the maintenance of the production farm. The Maintenance Fallback window displays the Maintenance Failover mode you selected for a read-only or read-write standby farm, and the change in that field is not allowed. If you have performed a Maintenance Failover to produce a read-write standby environment, Maintenance Fallback will detach the standby content database and attach the content database in the production environment allowing the new or updated data to synchronize to the production farm.

***Note**: Maintenance Failover and Maintenance Fallback will disable the Synchronization schedule settings of the groups, skipping the Synchronization jobs. You can perform a Synchronization job to enable group's synchronization schedule settings, after the upgrade and maintenance of the production farm and standby farm.

Distinguishing Failover-Ready Databases on Dashboard

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

Sync Method	Warm Standby Database Status	Cold Standby Database Status	Last Action Status	Comment
	Restoring	Restoring	Synchronization Failed	Cannot failover
	Does Not Exist	Does Not Exist		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	Standby/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
Platform Backup Log Shipping	Standby/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

High Availability for Specific Farm Components or Scenarios

Refer to the section below for details on using High Availability for carrying out specific farm components:

- Protecting Connector BLOB Data with High Availability
- Protecting Storage Manager BLOB data with High Availability
- Building up a Read-Only or Read-Write Standby Farm for the Maintenance of the Production Farm
- Enabling Read-Only View for Standby Web Applications and Standby Service Applications after Synchronization
- Improving Performance of Log Shipping Fallback

Protecting Connector BLOB Data with High Availability

This section introduces how to use High Availability to protect the Connector BLOB data in the production farm.

Creating a Standby Farm Mode High Availability Group Using Wizard Mode

Complete the steps to create a High Availability group including Connector BLOB data:

1. In the **Group Manager** tab, click **Create** on the ribbon and the select **Wizard Mode** from the drop-down list.
2. Enter a group name and an optional description for the group you are creating. Click **Next**.
3. In the **General Settings** page, configure the following settings:
    - **High Availability Mode** – Select **Standby farm** option.
    - **Sync Method** – Select the **Log Shipping** method from the **Sync Method** drop-down list.
    - **Would you like to make the standby environment read-only?** – Select the **Enable read-only view after synchronization**.
    - **Would you like to include Storage Optimization data?** – Select the Connector BLOB Data option and choose whether or not to allow High Availability to synchronize the Connector BLOB data.
        - **Automatically enable EBS/RBS during Failover** – Choose whether or not to automatically enable EBS/RBS in the standby environment during Failover.
    - **Failover Options** – Select the **Keep database in read-only mode to perform a one-way Failover** option to perform a one-way Log Shipping Failover and keep the standby databases in read-only mode.

- **Log Shipping Fallback Options** – The **Delete the renamed original production database after the Fallback job has completed** option is selected by default.
- **Connector Data Fallback Options** – Select the **Fallback Connector data** option.

For more details on each setting in this step, refer to Using Wizard Mode to Create a Standby Farm Mode High Availability Group.

4. Click **Next**. The **Add to Group** page appears.
5. Expand the farm components tree, and select the Web application with Connector BLOB data. The associated stub database will be automatically selected. Click the Add to Group (>>) button to add the selected Web application and stub database to the right pane.

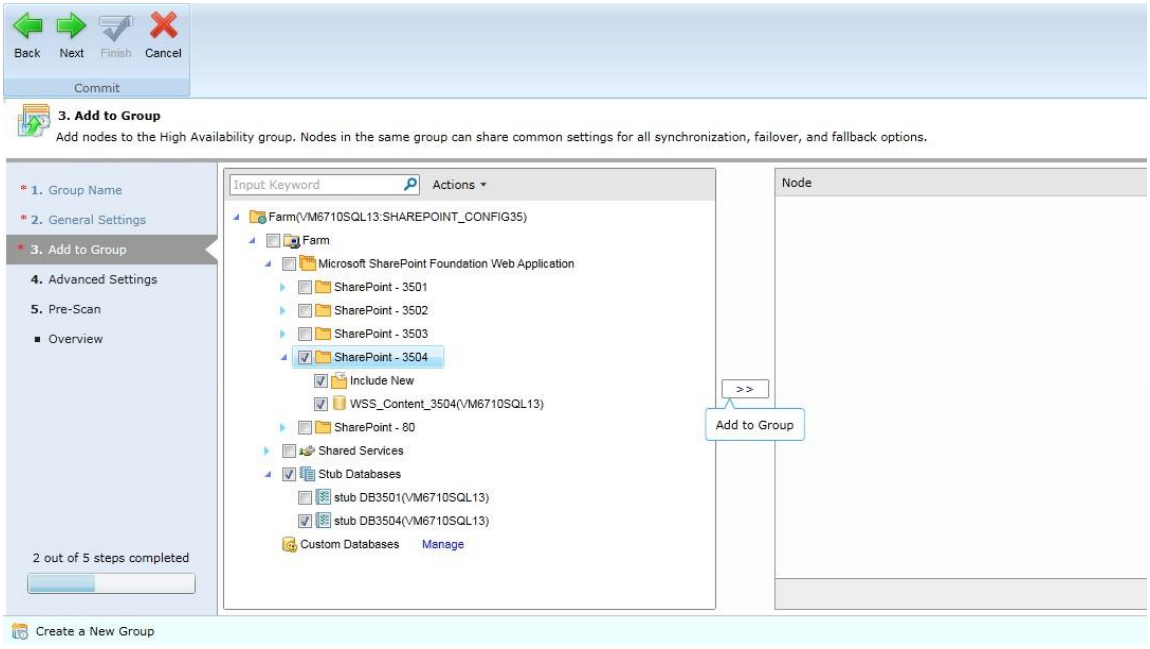

Figure 4: Adding the Web application with Connector BLOB data and stub database to the High Availability group.

6. An interface appears for configuring the standby SQL Instance, standby database location, and standby Agent and physical device for Connector BLOB data.

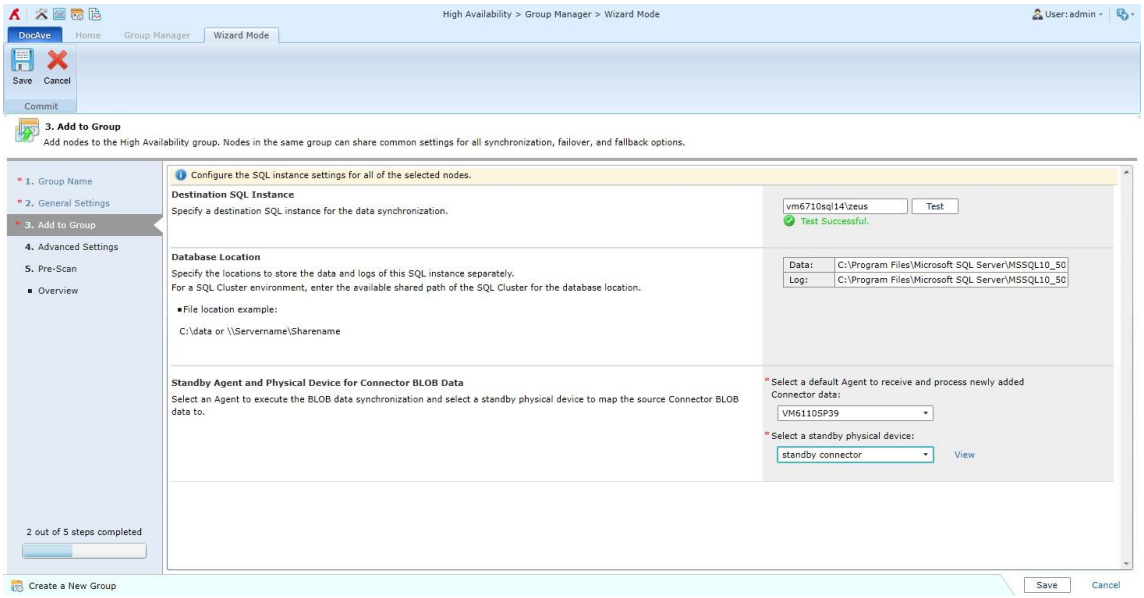

Figure 5: Configuring the settings for the components added to the High Availability group.

a. Enter the standby SQL Instance name into the text box in the Destination SQL Instance field, and then click **Test**.

b. If test is successful, the default database location of the SQL Instance will be displayed in the **Database Location** field. You can change the database file location in the **Database Location** field.

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

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

c. In the **Standby Agent and Physical Device for Connector BLOB Data** field, specify an Agent to execute the Connector BLOB synchronization. AvePoint recommends selecting an Agent that is not included in both of the production and standby farms.

d. Select a standby physical device from the drop-down list for mapping the Connector BLOB data of the selected Web application. DocAve High Availability only supports the Connector BLOB data stored in Net Share or NFS storage. 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**: The selected Agent and standby physical device will be used by default to map the newly added Connector physical device of the production environment.

e. Click **Save** to save the configurations, and go back to the **Add to Group** page.

7. Click **Next** in the **Add to Group** page. The **Advanced Settings** page appears.
8. In the **Connector Device Mapping** field, a table shows all of the Connector BLOB data mapping settings. The production servers with Connector BLOB data stored are displayed under the **Source Server with Connector BLOB** column. By default, the selected Agent is used to execute the Connector BLOB synchronization for all of the listed production servers, and the selected standby physical device is used for mapping all of the Connector BLOB data stored in the production server. You can change the mapping settings in the table as desired. Beneath the table, select a Connector cache profile from the drop-down list, or create a new one to designate a desired location for storing the SQLite database file used for Connector device mapping. For details on creating a Connector cache profile, refer to Configuring a Connector Cache Setting Profile.
9. Configure the Throttle Control, Custom Action, Schedule, and Notification settings. For details, refer to Using Wizard Mode to Create a Standby Farm Mode High Availability Group.
10. Click **Next**. The **Pre-Scan** interface appears.
11. Click **Next** to skip scanning the environment for the requirements that need to be met before running a synchronization job using Log Shipping method. It is recommended to skip this step, save the group configurations, and go to the **Dashboard** to run the Pre-Scan job. For details on Log Shipping pre-scan rules, refer to Log Shipping.

    ***Note**: AvePoint does not recommend running Pre-Scan job within the group creation wizard, since the job will take a while to complete.

12. In the **Overview** interface, you can select **Edit** in the field that you want to modify and you are brought to the corresponding step to make your changes.
13. Click **Finish** to exit the wizard and save the group. The group will be displayed in the **Group Manager** interface. A message bar will appear informing you that the group has been successfully created. You can click the **Dashboard** link on the message bar or navigate to **Home** > **Dashboard** to access the **Dashboard** interface to perform a Synchronization job using the created group.

Performing a Synchronization Job and Viewing the Standby Environment

To perform a Synchronization job using the group you created in the section above, complete the steps below:

1. In the **Dashboard** interface, select the group and then click **Synchronization** on the ribbon.

2. After the Synchronization job is finished, log into the standby SQL Instance to view the status of the standby databases. The standby databases are in **Standby/Read-Only** status.

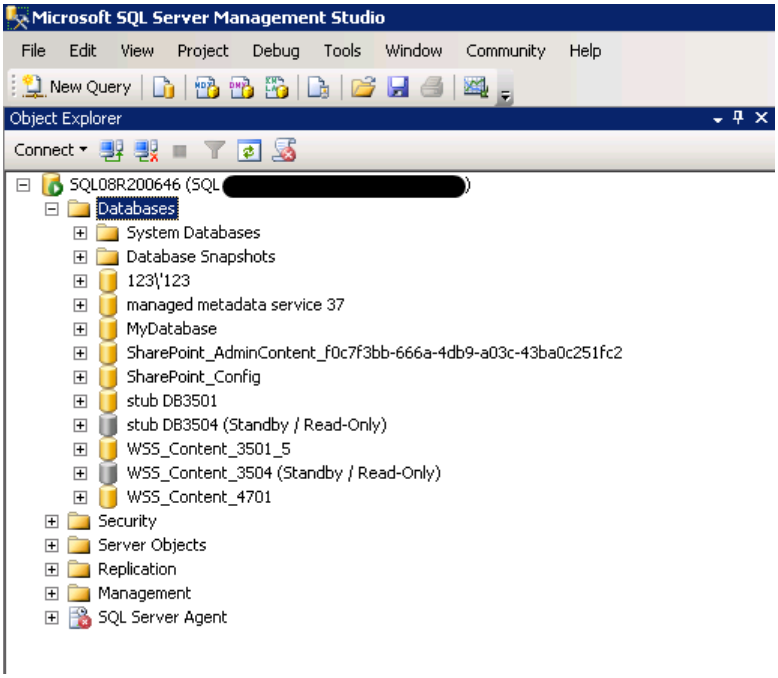

Figure 6: The status of the standby databases.

If the **Use High Availability to Synchronize Connector BLOB data** option is not selected, High Availability did not synchronize Connector BLOB data in the production server to the standby physical device, but stored the mapping settings in the SQLite database. Complete the settings below to manually copy and paste the Connector BLOB data to the destination physical device according to the mapping settings:

1. Log into the Agent selected in the Connector cache profile, and then go to the path where to store the SQLite database file.
2. Find the folder with the corresponding plan ID in the *temporary buffer location\HighAvailability\Connector device mapping files* directory.
3. Open the **HAConnectorMapping_PlanID.db** file.

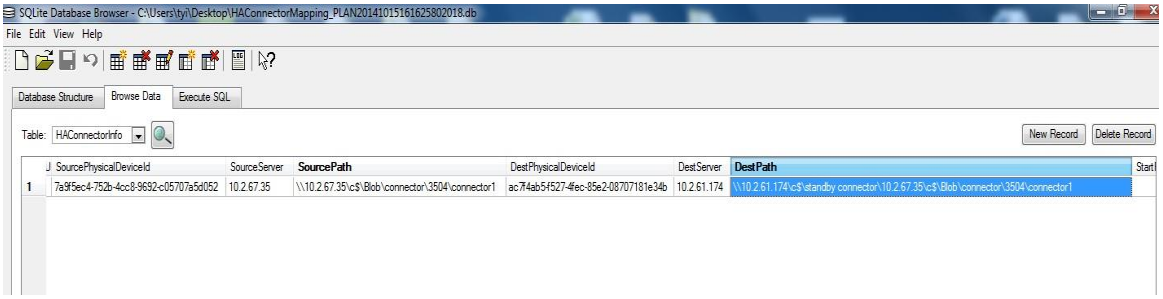

Figure 7: Viewing the data in the SQLite database.

4. Find the mapping settings information and manually copy and paste the Connector BLOB data from the production server to the standby physical device path displayed in the **DestPath** column. Besides, you can use the third-party tool, for example, Distributed File System (DFS). For details on DFS, refer to the Distributed File System article on Microsoft website.

5. After manually placing the Connector BLOB data into the standby physical device, access the site collection in the standby farm. The Connector data is read-only.

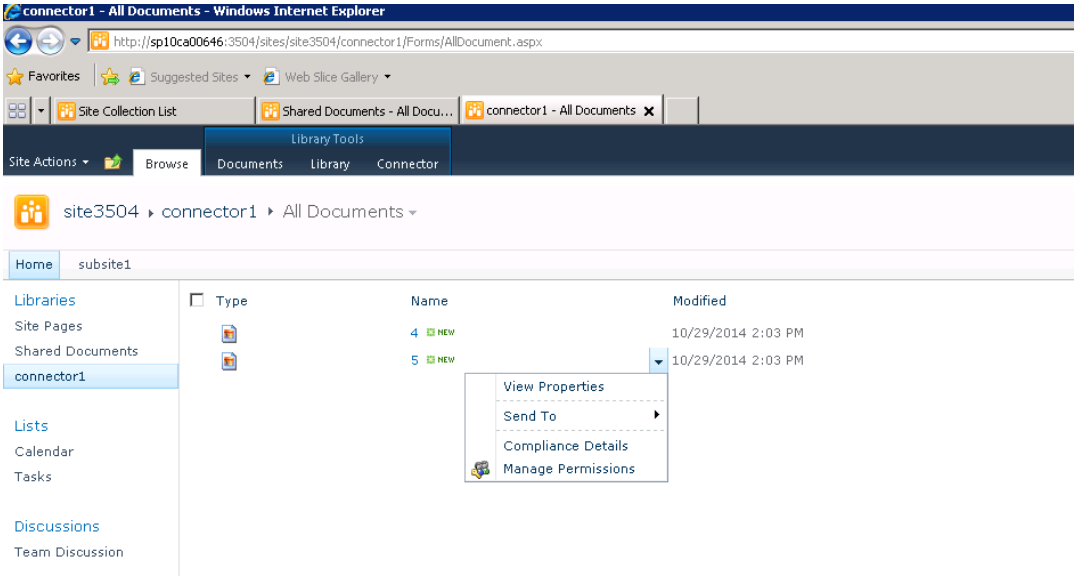

Figure 8: The Connector data in the standby site collection is read-only.

Performing a Failover Job and Viewing the Standby Environment

Complete the steps below to perform a Failover job:

1. In the **Dashboard** interface, select the group that has successfully performed the Synchronization job and then select **Failover** on the ribbon.
2. The **Failover** window appears. In this case, do not select the **Keep database in read-only mode to perform a one-way Failover** option. Click **OK** to perform the Failover job. If you select the **Keep database in read-only mode to perform a one-way Failover** option, the status of the standby databases will still be **Standby/Read-Only**, and the standby Web application will be read-only. If this option is not selected, the standby databases will be normal and the standby Web application will be read-write, as shown below.

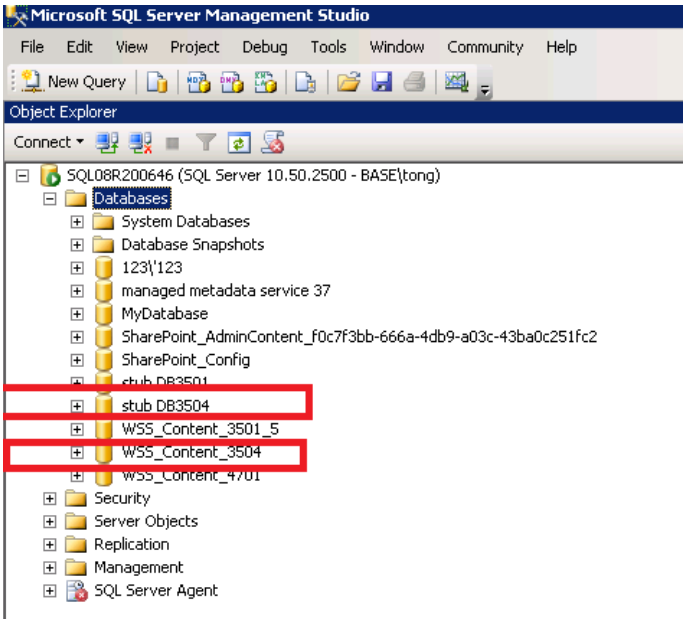

Figure 9: Viewing the status of the standby databases.

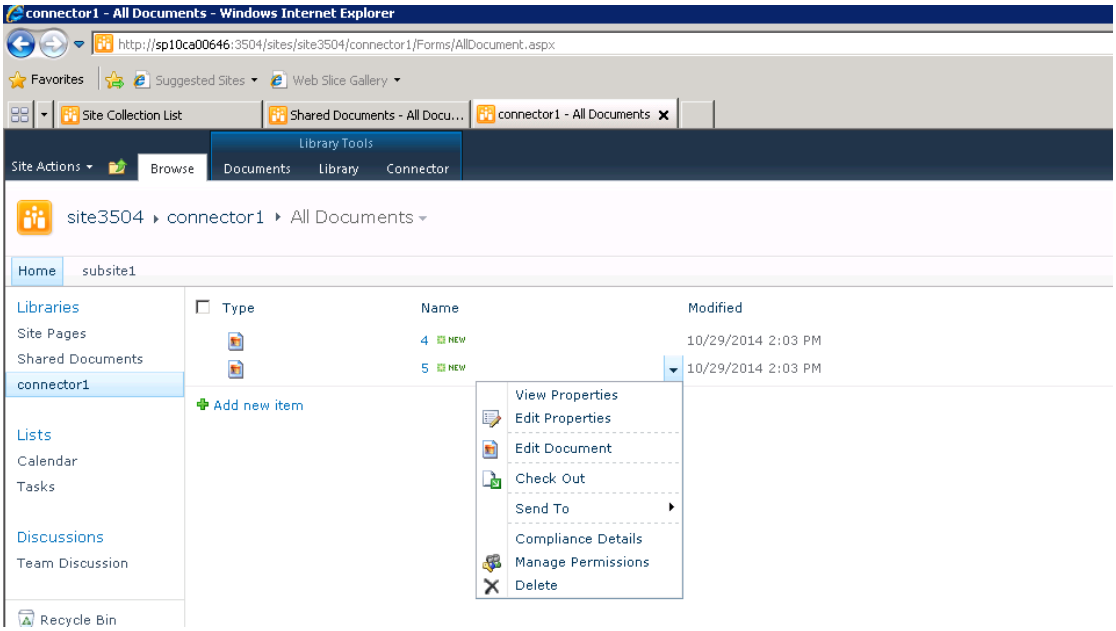

Figure 10: The Connector data in the standby site collection is read-write.

Performing a Fallback Job

In the **Dashboard** interface, select the group that has successfully performed the Failover job. Click **Fallback** on the ribbon to perform a Fallback job. Additionally, you can choose whether or not to fallback Connector data. After the Fallback job, both of the production and standby farm are read-write.

Protecting Storage Manager BLOB data with High Availability

This section introduces how to use High Availability to protect the Storage Manager BLOB data in the production farm.

Creating a Standby Farm Mode High Availability Group Using Wizard Mode

Complete the following steps to create a High Availability group that includes Storage Manager BLOB data:

1. In the **Group Manager** tab, click **Create** on the ribbon and the select **Wizard Mode** from the drop-down list.
2. Enter a group name and an optional description for the group you are creating. Click **Next**.
3. In the **General Settings** page, configure the following settings:
    - **High Availability Mode** – Select **Standby farm** option.
    - **Sync Method** – Select the **Log Shipping** method from the **Sync Method** drop-down list.
    - **Would you like to make the standby environment read-only?** – Select the **Enable read-only view after synchronization** option to make the standby environment readable after the synchronization. The Storage Manager stubs in the standby farm will be readable after synchronization.
    - **Would you like to include Storage Optimization data?** – Select the **Storage Manager BLOB Data** option.
        - **Automatically enable EBS/RBS during Failover** – Choose whether or not to automatically enable EBS/RBS in the standby environment during Failover. It is recommended to select this option to enable the EBS or RBS settings during Failover.
    - **Failover Options** – Select the **Keep database in read-only mode to perform a one-way Failover** option to perform a one-way Log Shipping Failover and keep the standby databases in read-only mode.

    For more details on each setting in this step, refer to Using Wizard Mode to Create a Standby Farm Mode High Availability Group.
4. Click **Next**. The **Add to Group** page appears.

5. Expand the farm components tree, and select the Web application with Storage Manager BLOB data and the associated stub database will be automatically selected. Click the Add to Group (>>) button to add the selected Web application and stub database to the right pane.

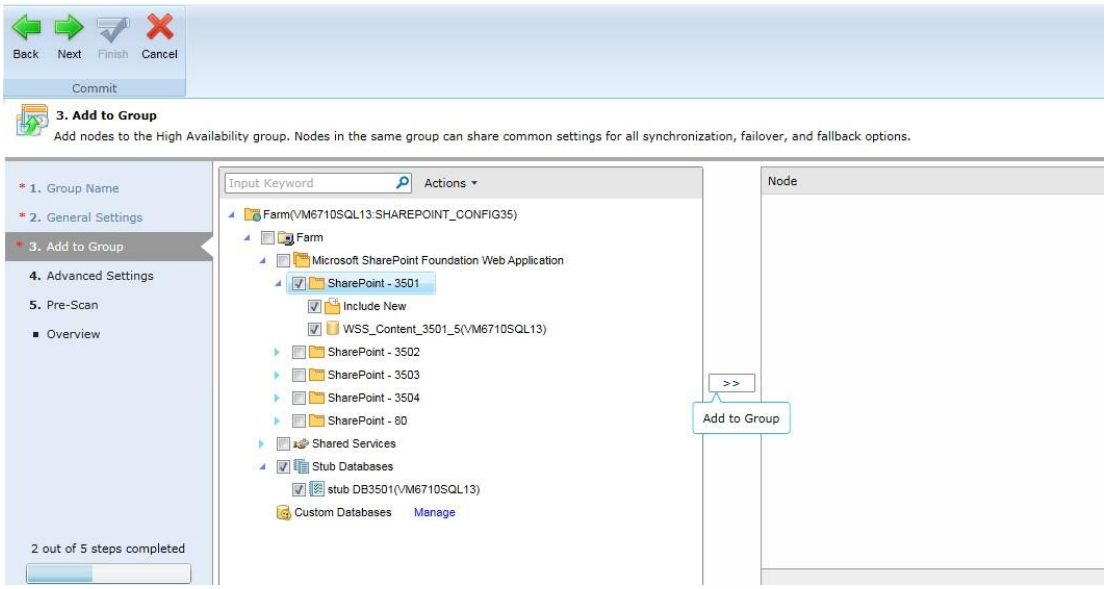

Figure 11: Adding the Web application with Storage Manager BLOB data and stub database to the High Availability group.

6. An interface appears for configuring the standby SQL Instance, standby database location, and standby physical device.

   a. Enter the standby SQL Instance name into the text box in the **Destination SQLInstance** field, and then click **Test**.

   b. If test is successful, the default database location of the SQL Instance will be displayed in the **Database Location** field. You can change the database location, if desired.

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

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

   c. In the **Standby Physical Device** field, the storage types of the physical devices created in the Control Panel are displayed under the **Device Type** column. To map the Storage Manager BLOB data that is stored in the production device, select a standby 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**: The selected standby physical device will, by default, be used to map the newly added Storage Manager physical device of the production environment according to the device type.

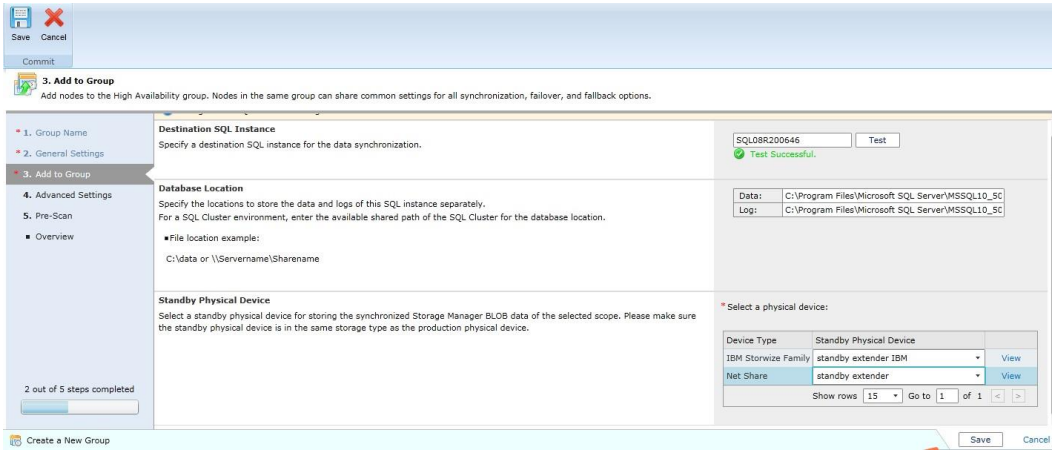

Figure 12: Configuring settings for standby databases and standby physical device of Storage Manager BLOB data.

d. Click **Save** to save the configurations, and go back to the **Add to Group** page.

7. You can click **Edit** next to the database in the right pane to configure the settings for the standby database.
8. In the **Physical Device Mapping** field, a table shows the SharePoint components containing the Storage Manager BLOB data, the production physical device storing the Storage Manager BLOB data, and the mapped standby physical device to store the Storage Manager BLOB data in the standby farm. You can select another standby physical device from the drop-down list for mapping the Storage Manager BLOB data in the corresponding production physical device. Click **Save** to save the configurations for the standby database.
9. Click **Next** in the **Add to Group** page. The **Advanced Settings** page appears.
10. Configure the Throttle Control, Custom Action, Schedule, and Notification settings. For details, refer to Using Wizard Mode to Create a Standby Farm Mode High Availability Group.
11. Click **Next**. The **Pre-Scan** interface appears.
12. Click **Next** to skip scanning the environment for the requirements that need to be met before running a synchronization job using Log Shipping method. It is recommended to skip this step, save the group configurations, and go to the **Dashboard** to run the Pre-Scan job. For details on Log Shipping pre-scan rules, refer to Log Shipping.

    ***Note**: AvePoint does not recommend running Pre-Scan job within the group creation wizard, since the job will take a while to complete.

13. In the **Overview** interface, you can select **Edit** in the field that you want to modify.
14. Click **Finish** to exit the wizard, and save the group. The group will be displayed in the **Group Manager** interface. A message bar will appear informing you that the group is successfully created. You can click the **Dashboard** link on the message bar, or navigate to **Home** > **Dashboard** to access the **Dashboard** interface to perform a Synchronization job using the created group

Performing a Synchronization Job and Viewing the Standby Environment

To perform a Synchronization job using the group you created in the section above, complete the steps below:

1. In the **Dashboard** interface, select the group and then click **Synchronization** on the ribbon.
2. After the Synchronization job is finished, log into the standby SQL Instance to view the status of the standby databases. The standby databases are in **Standby/Read-Only** status.

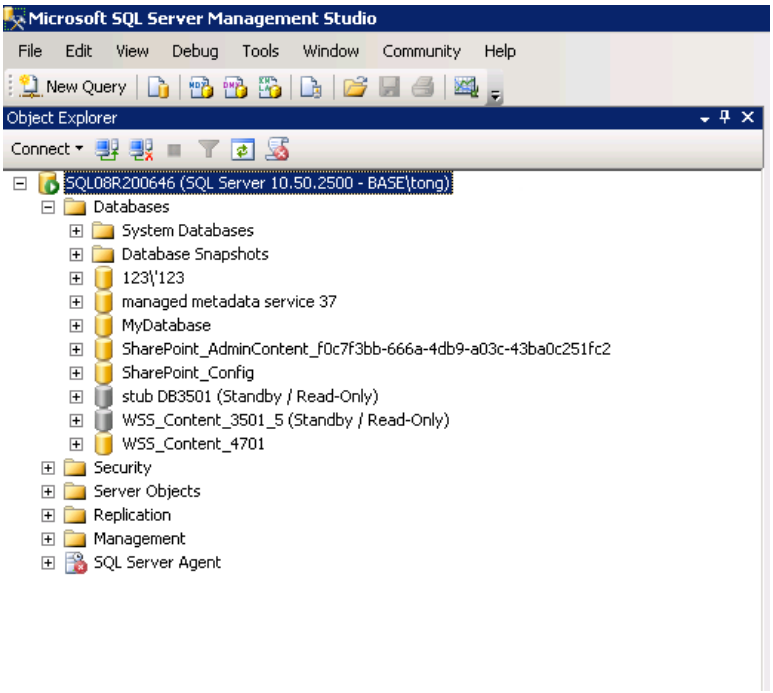

Figure 13: The status of the standby databases.

3. Since High Availability does not support synchronizing the Storage Manager BLOB data to the standby physical device, you must manually copy and paste the Storage Manager BLOB data from the production physical device to the destination physical device according to the mapping settings configured in the group.

4. After manually placing the Storage Manager BLOB data into the standby physical device, access the site collection in the standby farm. The Storage Manager data is read-only.

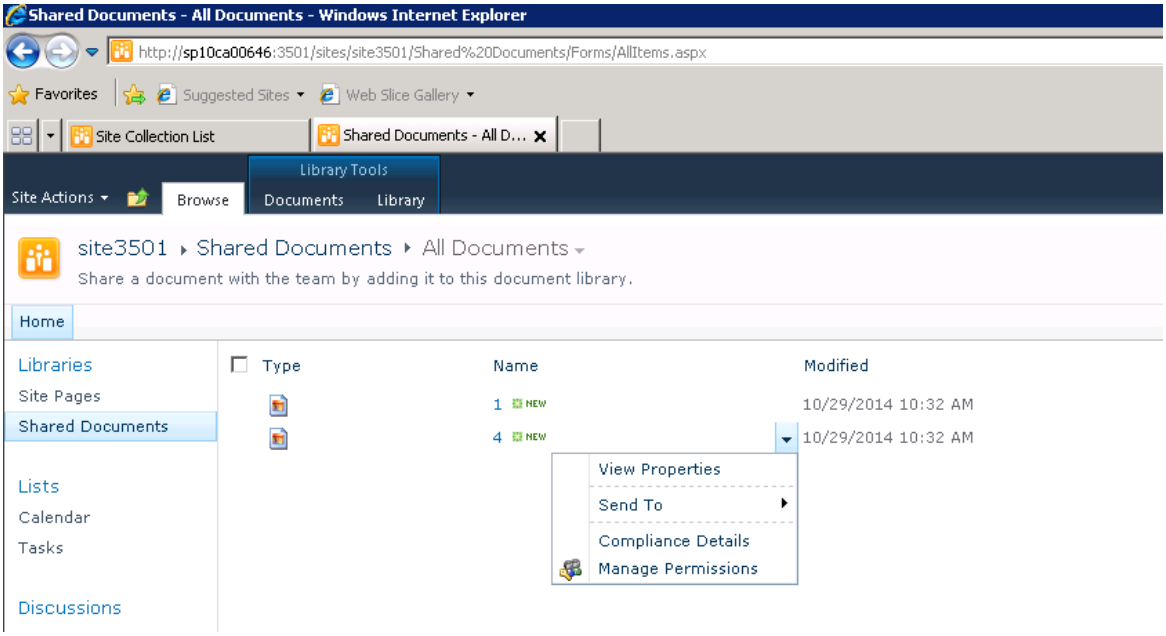

Figure 14: The Storage Manager data in the standby site collection is read-only.

Performing a Failover Job and Viewing the Standby Environment

Complete the steps below to perform a Failover job:

1. In the **Dashboard** interface, select the group that has successfully performed the Synchronization job and then select **Failover** on the ribbon.
2. The **Failover** window appears. If you select the **Keep database in read-only mode to perform a one-way Failover** option, the status of the standby databases will still be **Standby/Read-Only**, and the standby Web application will be read-only. If this option is not selected, the standby databases will be normal and the standby Web application will be read-write. In this case, do not select the **Keep database in read-only mode to perform a one-way Failover** option. Click **OK** to perform the Failover job.

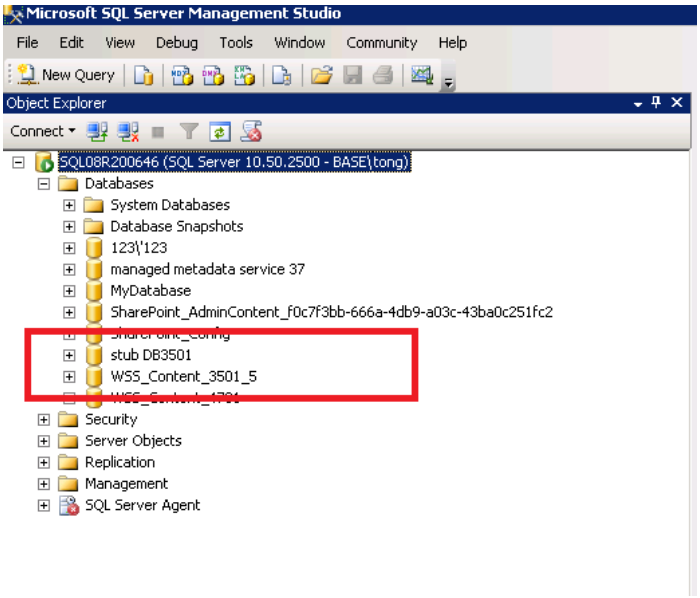

Figure 15: Viewing the status of the standby databases.

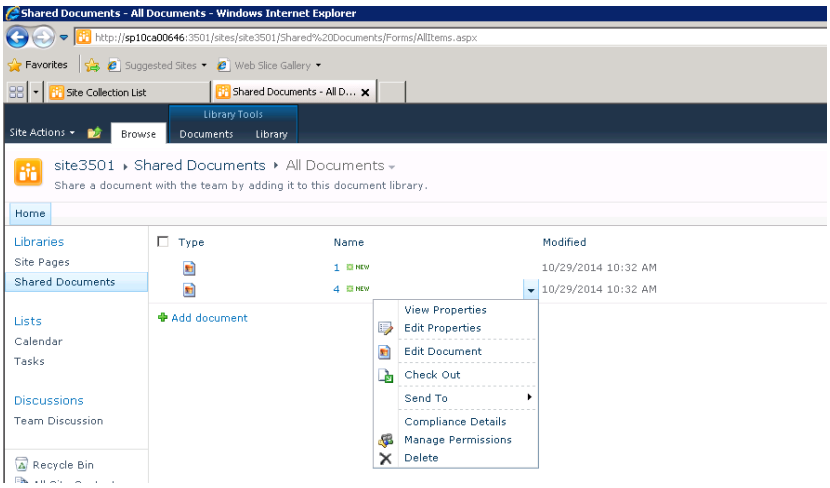

Figure 16: The Storage Manager data in the standby site collection is read-write.

Performing a Fallback Job

In the **Dashboard** interface, select the group that has successfully performed the Failover job. Click **Fallback** on the ribbon to perform a Fallback job. After the Fallback job, both of the production and standby farm are read-write.

***Note**: Since DocAve High Availability does not support synchronizing Storage Manager BLOB data, you must make the changes of the Storage Manager BLOB data in the standby farm successfully synchronized to the production farm after the Fallback job using other methods.

Building up a Read-Only or Read-Write Standby Farm for the Maintenance of the Production Farm

Sometimes, you may need to perform upgrade or maintenance on your SharePoint production environment, but you will not be able to access the production farm during the upgrade or maintenance. High Availability Failover gives you the ability to access nearly the same data as the production farm from a standby farm; however, after the Failover job completes, the production farm is not accessible, which means you cannot perform upgrade or maintenance on you production farm. When you want to perform an upgrade or maintenance to your production farm, but have a read-only or read-write standby farm, use High Availability Maintenance Failover and Maintenance Fallback.

Refer to the section below to walk through the process of using High Availability for maintenance purpose in an example scenario.

Creating and Using Standby Farm Mode Group for Maintenance Purpose

Refer to the section below to walk through the whole process of High Availability Failover and Fallback.

1. Create a **Standby farm** mode High Availability group. In this case, select **Log Shipping** as the **Sync Method**, and deselect the **Enable read-only view after synchronization** option. For details on each group setting, refer to Using Wizard Mode to Create a Standby Farm Mode High Availability Group.

   ***Note**: If you want to have a read-only standby farm with accessible service applications after the Maintenance Failover, select the **Enable read-only view after synchronization** option for the group including the service applications and then minimize the time interval between the first Synchronization job and Maintenance Failover job to avoid data loss.

2. After the group is successfully created, navigate to **Home** > **Dashboard**. Select the group and click **Synchronization** on the ribbon in the **Dashboard** to perform Synchronization job.
3. You can go to Job Monitor to view the job status and job information of the Synchronization job.
4. After the Synchronization job is finished, log into the SharePoint Central Administration of the standby SharePoint farm. You will find the standby Web applications are created. Log into the standby SQL Instance to view the status of the standby databases. Since the **Enable read-only view after synchronization** option is deselected in this case, the status of the standby Web application content databases and service application databases is **Restoring**. The service applications are not created in the standby farm after synchronization, that is, the service applications cannot be accessed in standby farm.

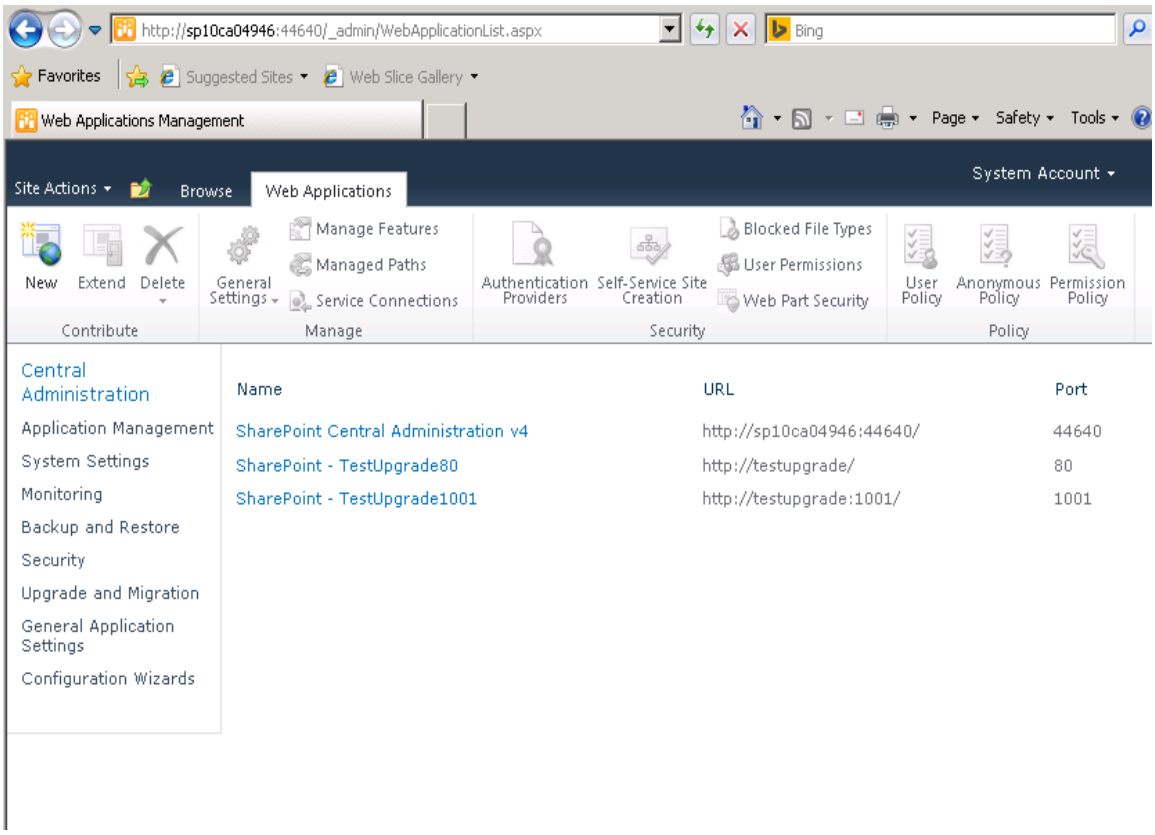

Figure 17: The standby Web applications are created in the standby farm.

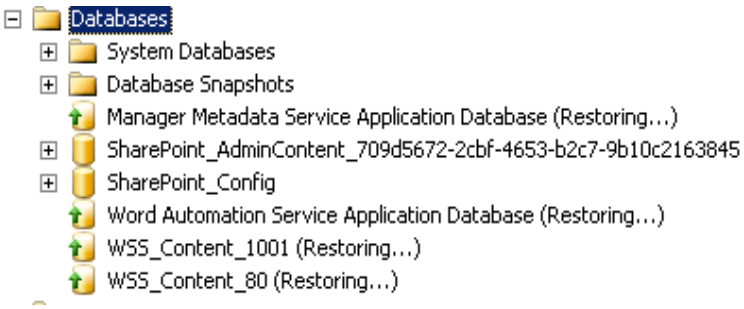

Figure 18: The status of the standby databases.

5. In the **Dashboard**, select all of the Standby farm mode groups of the same production farm. The **Maintenance Failover** button is available for selection. Click **Maintenance Failover** on the ribbon, and the **Maintenance Failover** window appears.

6. Click **Maintenance Tutorials** to read the brief introduction on the whole process of performing a High Availability for upgrade or maintenance purpose. Click **Close** to exit the window.

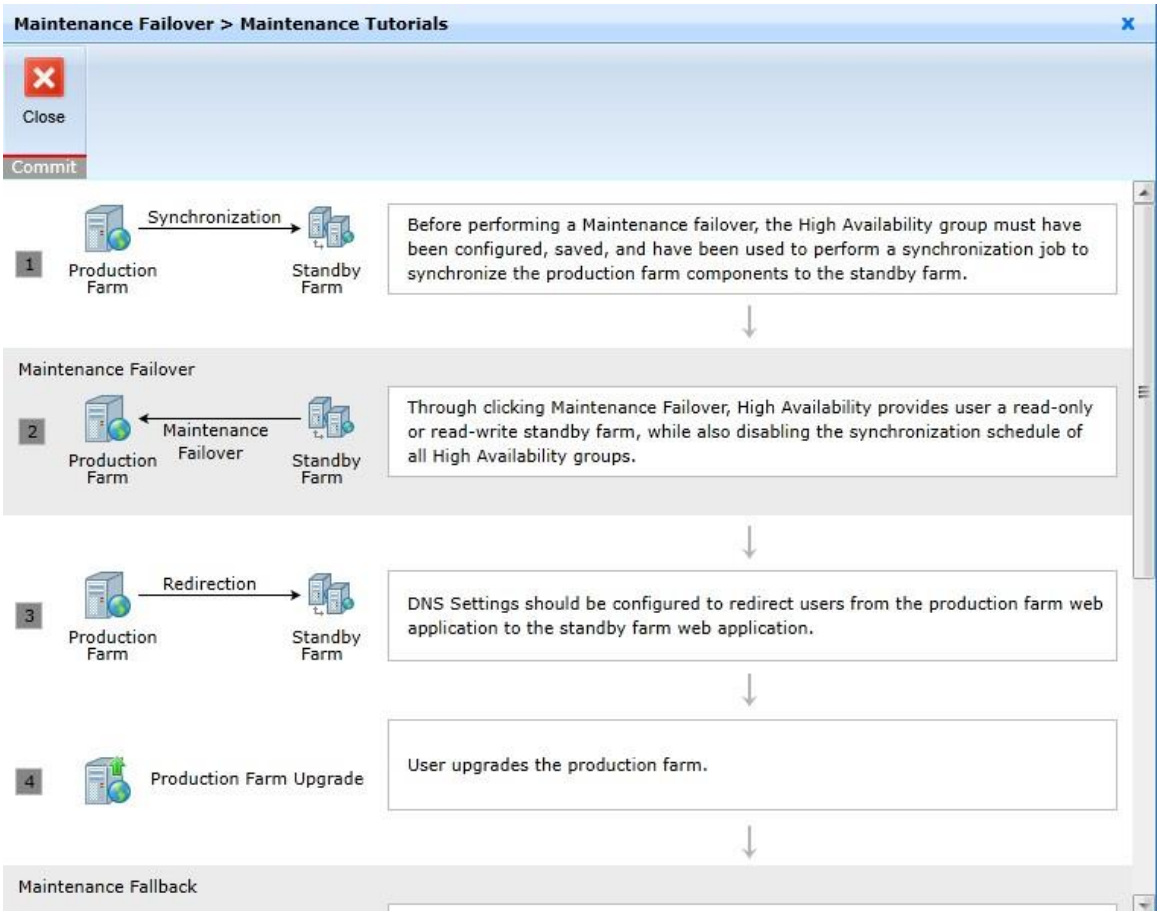

Figure 19: Reading the Maintenance Tutorials.

7. In the **Maintenance Failover Mode** field, choose to make the standby environment **Read-only** or **Read-write**. If you select Read-write option, the production databases will be brought offline after the Maintenance Failover, and you can choose whether or not to perform an incremental synchronization before Failover. For details, refer to Perform an incremental synchronization before Failover. Click **OK** to perform the Maintenance Failover job, and then click **OK** on the confirmation window.

   ***Note**: The Synchronization job started after the Maintenance Failover job will be skipped. You can perform a Synchronization job later to enable the Synchronization schedule after the maintenance of the production and standby farms.

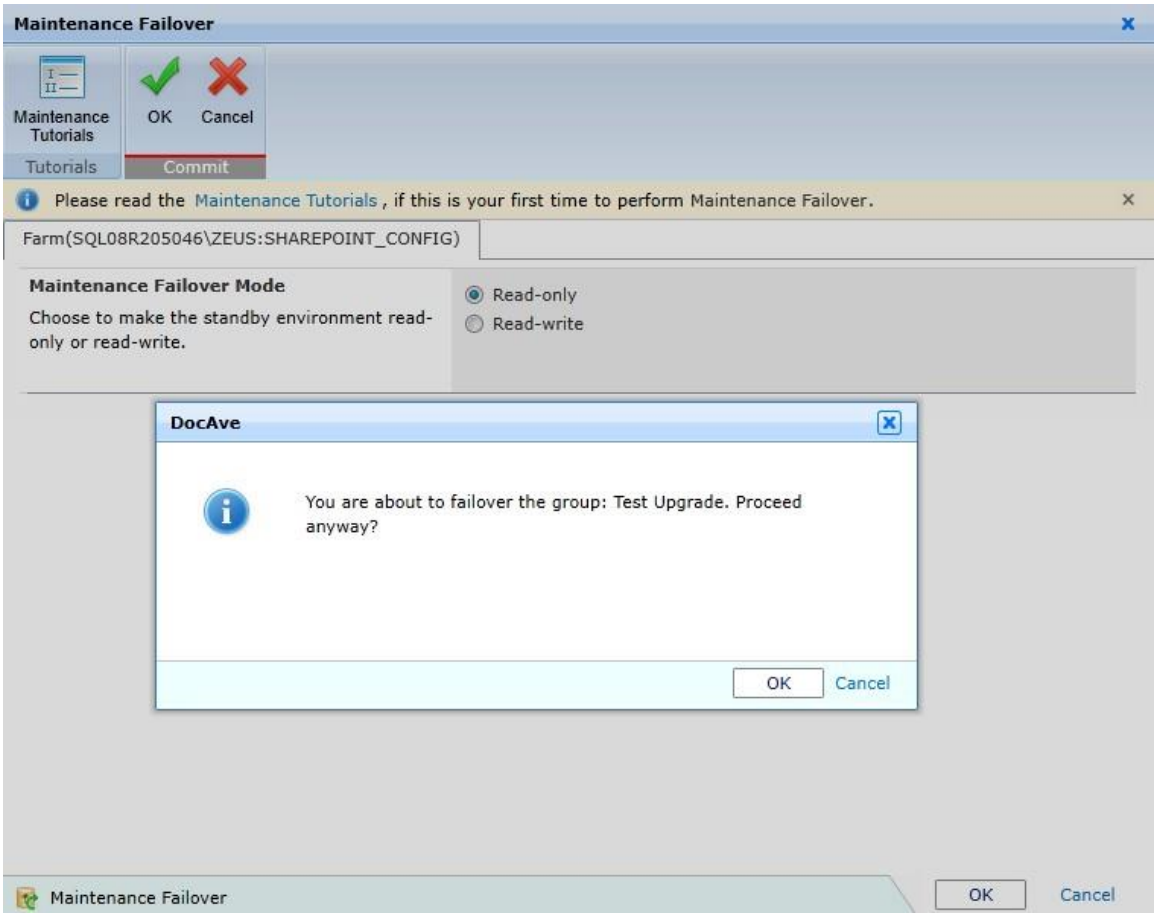

Figure 20: Performing Maintenance Failover job to have a read-only standby farm.

8. After the Maintenance Failover job is finished, log into the standby SQL Instance to view the status of the Web application content databases and service application databases.

    - If you selected the **Read-only** option in the **Maintenance Failover Mode** field, the status of the standby Web application content databases is **Standby/Read-Only**; the status of the standby service application databases is normal and the service applications are created in the standby farm after Maintenance Failover.

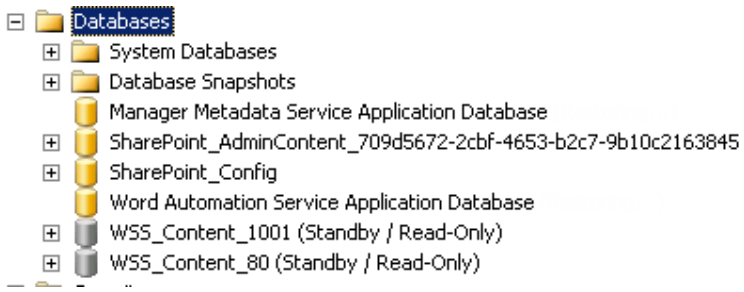

Figure 21: Viewing the status of the standby databases.

Access the site collection: *http://TestUpgrade:1001/sites/TestSite*. You will find that you are still accessing the production site collection that is read-write.

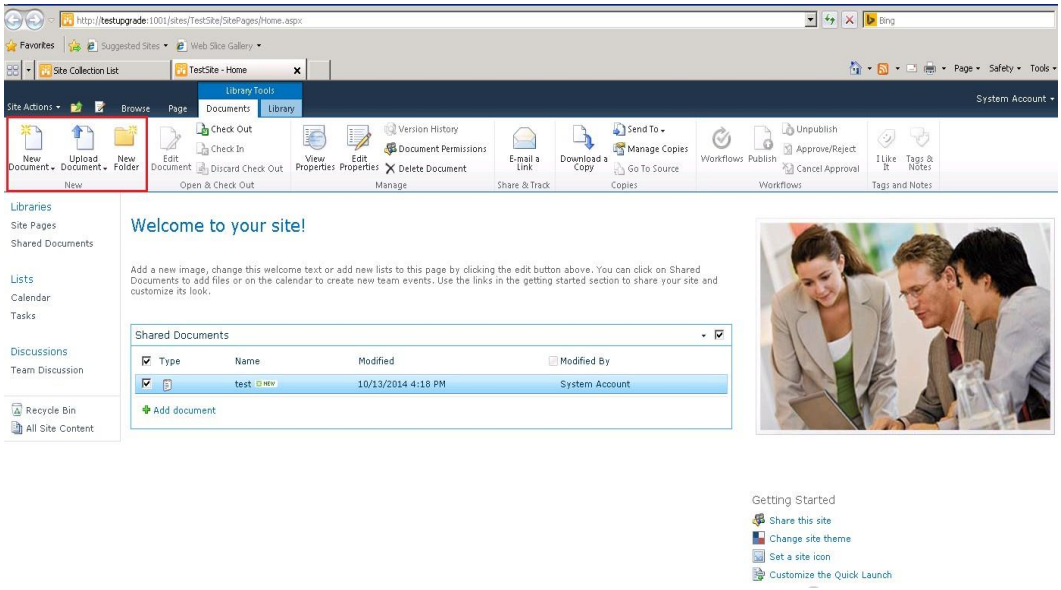

Figure 22: The production site collection is read-write.

- If you selected the **Read-write** option in the **Maintenance Failover Mode** field, the status of standby Web application content databases and the service application databases is **normal**.

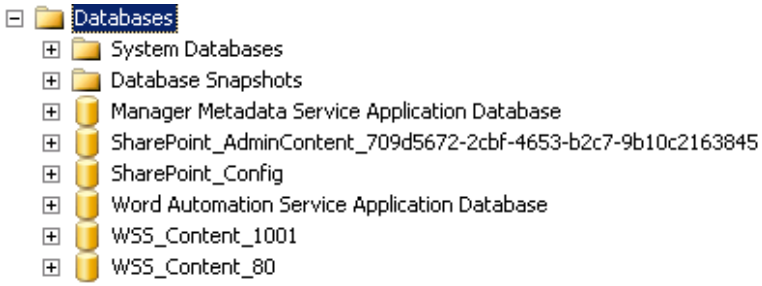

Figure 23: Viewing the status of the standby databases.

Access the site collection: *http://TestUpgrade:1001/sites/TestSite*. You will find that you are not able to access the site collection, since the production Web application content databases are brought offline, detached.

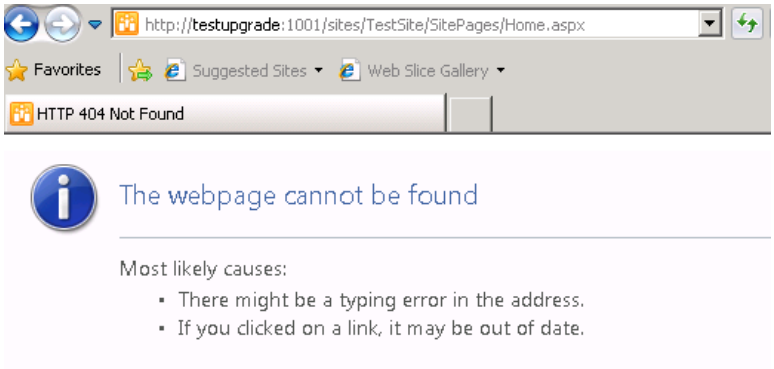

Figure 24: The production site collection is unavailable.

9. Change the target host of the **TestUpgrade** alias, in the DNS, to the standby host.

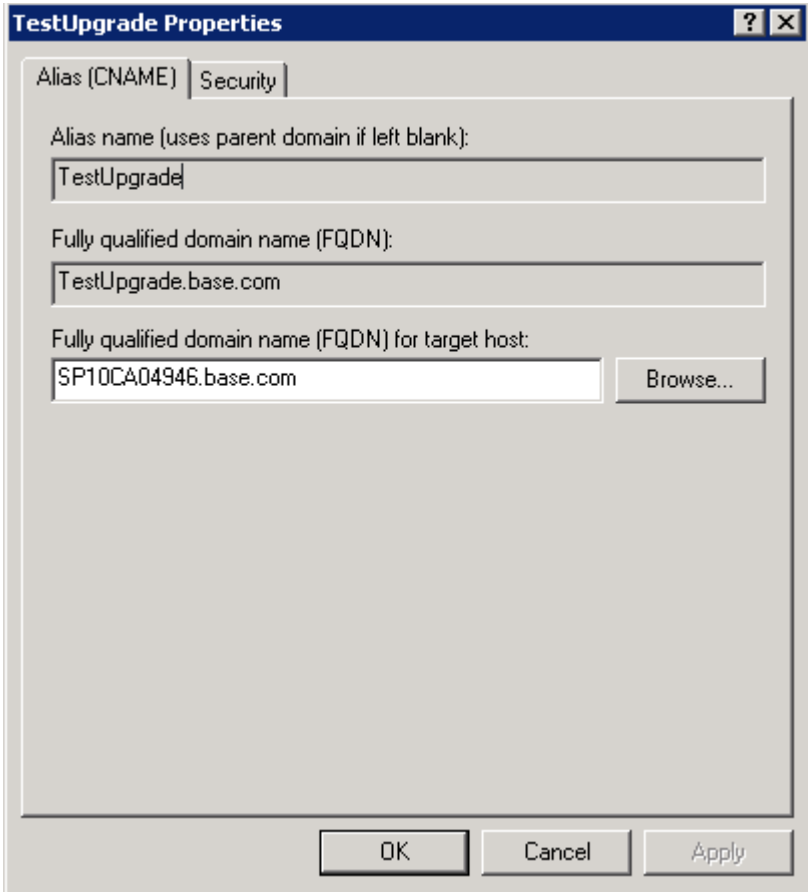

Figure 25: Changing the target host of this Alias to the standby host.

10. Access the site collection: *http://TestUpgrade:1001/sites/TestSite* again.

    - If the **Maintenance Failover Mode** is set to **Read-only**, you will find that you are accessing the read-only standby site collection.

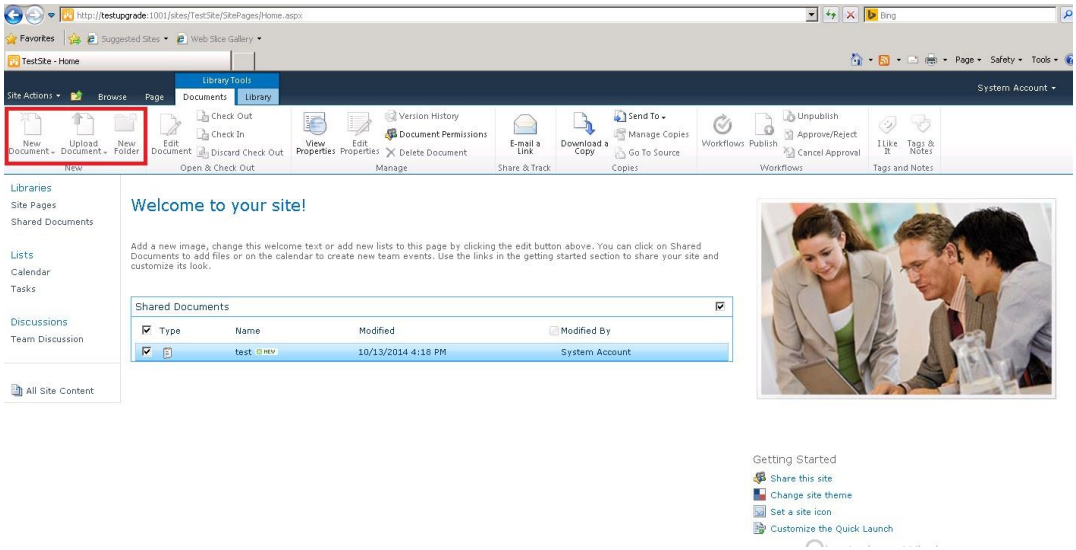

Figure 26: The standby site collection is read-only.

- If the **Maintenance Failover Mode** is set to **Read-write**, you will find that you are accessing the read-write standby site collection.

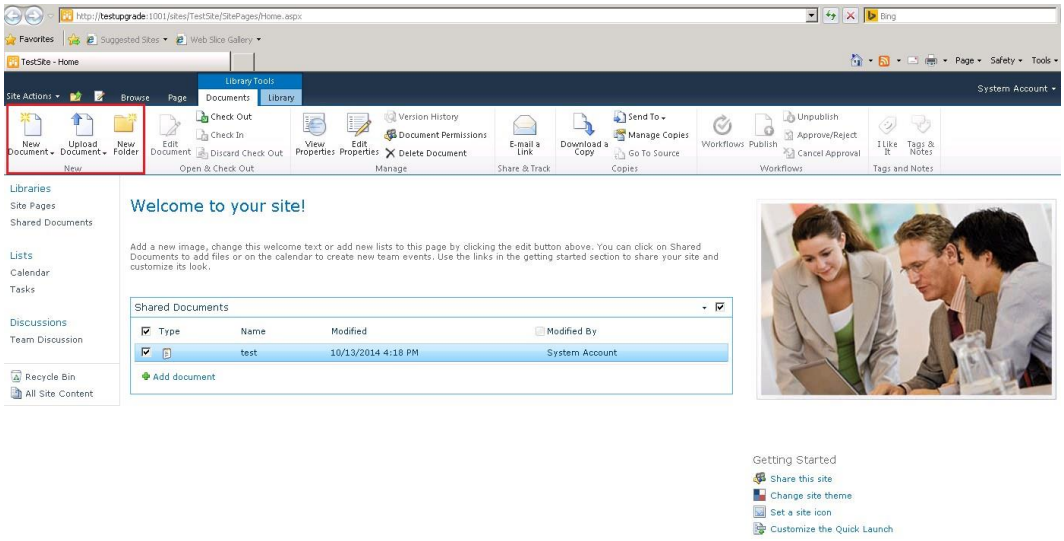

Figure 27: The standby site collection is read-write.

Add a new document into the standby site collection.

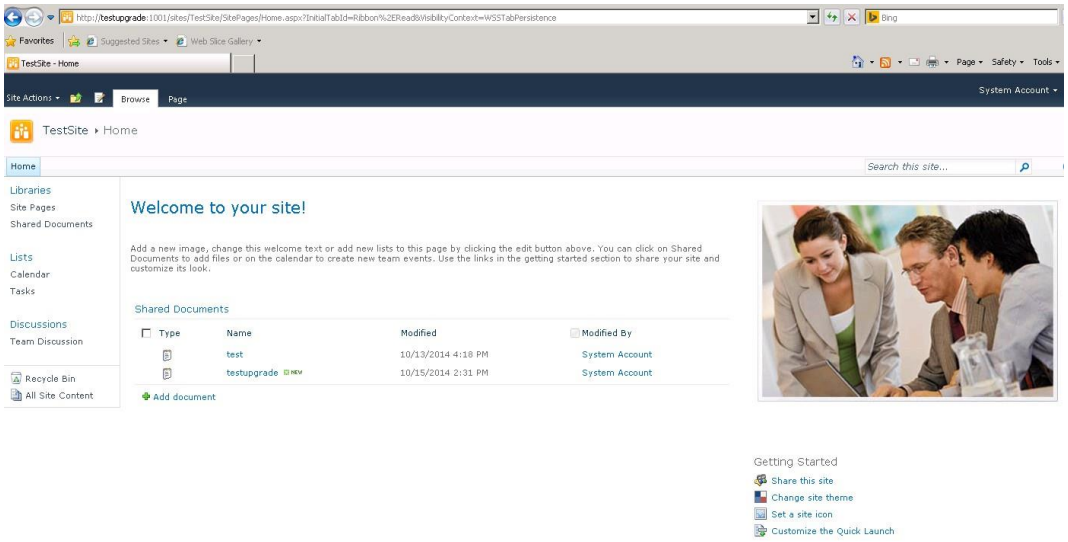

Figure 28: Adding a new file into the standby site collection.

***Note**: If you are still accessing the production site collection, log into the Agent where SharePoint server resides and run **cmd.exe** as administrator to execute **IPConfig /flushdns**.

11. Now, you can perform the upgrade or maintenance of your SharePoint production farm.
12. After the upgrade or maintenance of the production farm is finished, log into DocAve Manager, and then navigate to **Data Protection** > **High Availability** > **Home** > **Dashboard**. Select all of the nodes in the groups for the same production farm, and then click **Maintenance Fallback**. The **Maintenance Fallback** window appears.
13. In the **Maintenance Fallback** window, the Maintenance Failover mode you selected to have a **Read-only** or **Read-write** standby farm after Maintenance Failover is selected and this field is disabled for modification. Click **OK** to perform the Maintenance Fallback job.
14. After the Maintenance Fallback job is finished, log into the standby SQL Instance to view the status of the standby databases.

    - If the **Maintenance Failover Mode** is set to **Read-only**, the status of the standby databases is **Restoring**.

      Access the site collection: *http://testupgrade:1001/sites/TestSite.* You will find that the site collection is unavailable, as you are accessing the standby site collection and the standby content database is in **Restoring** status.

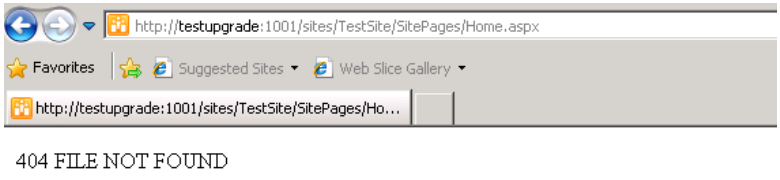

Figure 29: Cannot access the site collection.

- If the **Maintenance Failover Mode** is set to **Read-write**, the status of the standby databases is **normal**.

  Access the site collection: *http://testupgrade:1001/sites/TestSite.* You will find that the site collection is unavailable, as you are accessing the standby site collection and the standby content database is detached.

Figure 30: Cannot access the site collection.

15. Change the target host of the **TestUpgrade** alias, in the DNS, to the production host.

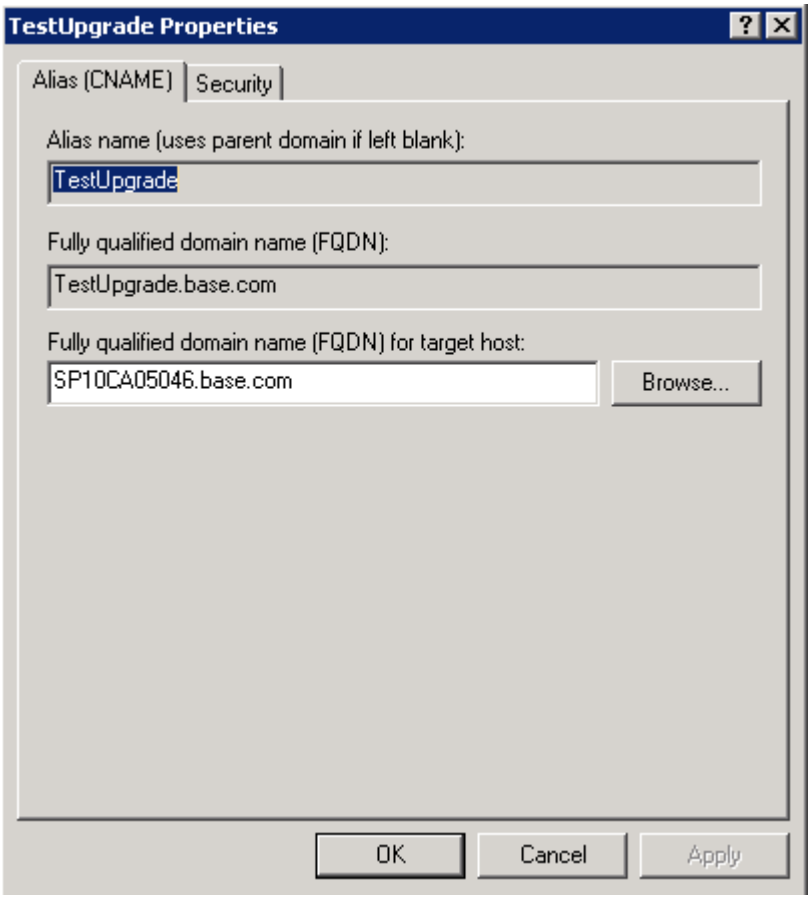

Figure 31: Changing the target host of the Alias to production host.

16. Access the site collection: *http://TestUpgrade:1001/sites/TestSite* again.

    - If the **Maintenance Failover Mode** is set to **Read-only**, you will find that you are accessing the production site collection normally.

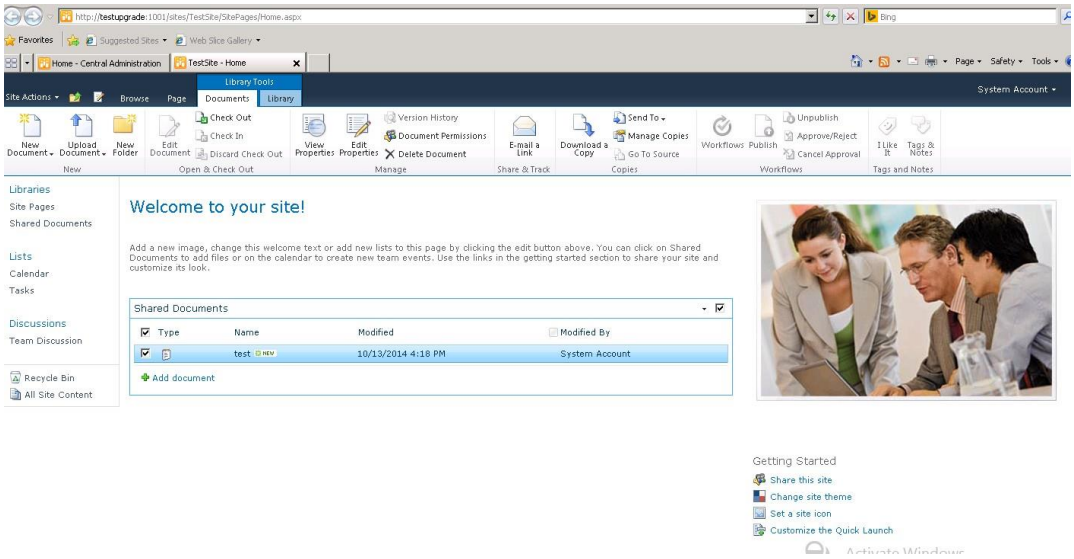

Figure 32: The production site collection is working normally.

- If the **Maintenance Failover Mode** is set to **Read-write**, you will find that you are accessing the production site collection normally, and the changes made in the standby site collection during the upgrade or maintenance of the production farm is synchronized to the production farm.

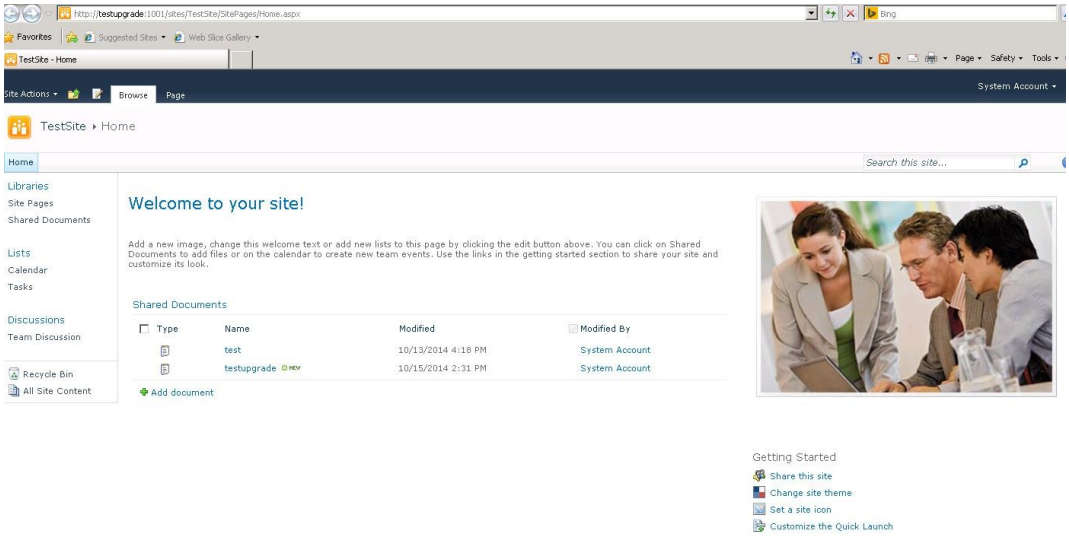

Figure 33: The production site collection is working normally, and the changes made in standby site collection are successfully synchronized to the production farm.

***Note**: If you still cannot access the standby site collection, log into the Agent where the SharePoint server resides and run **cmd.exe** as administrator to execute **IPConfig /flushdns**.

17. Now, the standby farm is ready for upgrade or maintenance.
18. After the upgrade or maintenance of the standby farm is finished, perform a Synchronization job using the group to enable the synchronization schedule settings configured in the group.

Enabling Read-Only View for Standby Web Applications and Standby Service Applications after Synchronization

Selecting the **Enable read-only view after synchronization** option will allow users to access the standby environment in a read-only view after a synchronization job; however, this option makes the standby service applications read-write after they are created in the standby environment by a Synchronization job. After being created in the standby environment, the services applications cannot be synchronized again. Therefore, the data changes of the production service application made after the first synchronization job cannot be synchronized to the standby farm, which will cause the data loss of the service applications after Failover. To solve this issue and in the same time enable the read-only view of the standby Web applications after synchronization, refer to the instructions below.

***Note**: Although the instructions below can help you have a read-only view for the Web applications after the synchronization and minimize the data loss related to the service applications between the first synchronization and the Failover, you will not be able to access the associated service applications in the standby farm after the synchronization.

Creating Standby Farm Mode High Availability Groups

In this section, you are about to create two High Availability groups: one for Web applications and another for service applications. Refer to the instructions below:

- Creating a standby farm mode group for Web applications – Select the **SQL Mirroring** as the sync method with **Asynchronous** mode, and select the **Enable read-only view after synchronization** option. For the details on all of the group settings, refer to Using Wizard Mode to Create a Standby Farm Mode High Availability Group.
- Creating a standby farm mode group for service applications – Select the **SQL Mirroring** as the sync method with **Asynchronous** mode, and deselect the **Enable read-only view after synchronization** option. For the details on all of the group settings, refer to Using Wizard Mode to Create a Standby Farm Mode High Availability Group.

After the two groups are successfully created, the groups are displayed in the **Plan Manager** interface.

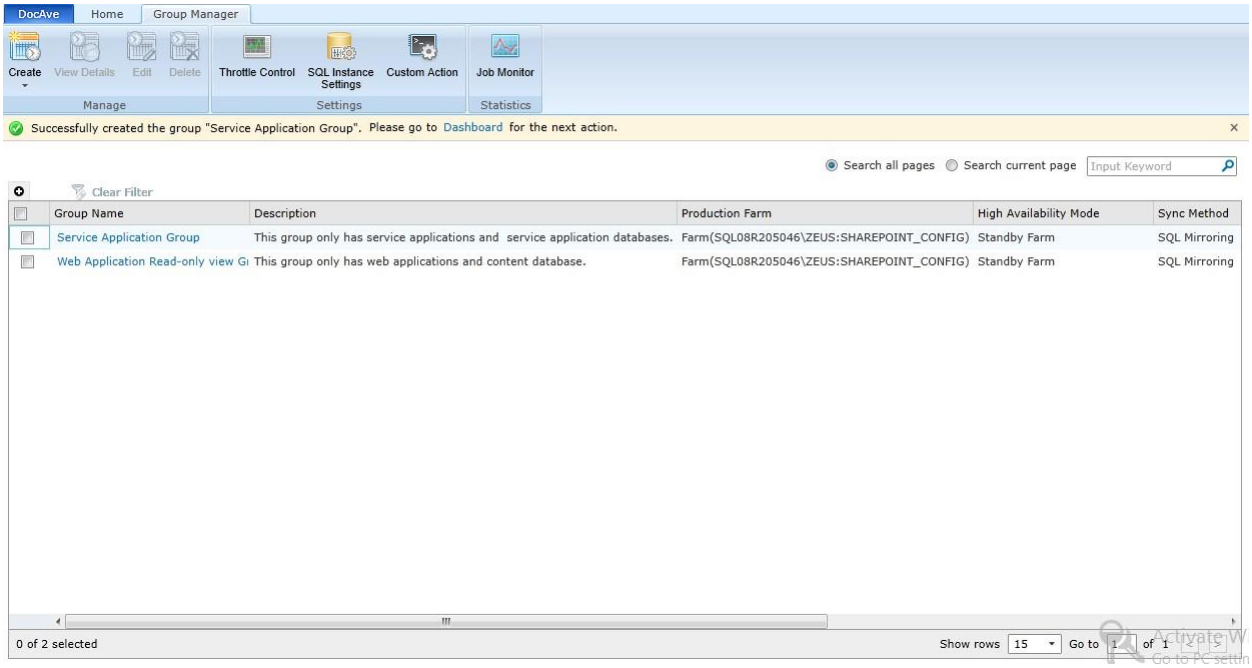

Figure 34: Viewing the groups in Plan Manager interface.

Performing Synchronization Jobs

To perform the Synchronization job using the two groups you created in the section above, complete the steps below:

1. In the **Dashboard** interface, select the two groups and all of the nodes included in the two groups. Click **Synchronization** on the ribbon and click **OK** on the confirmation window.

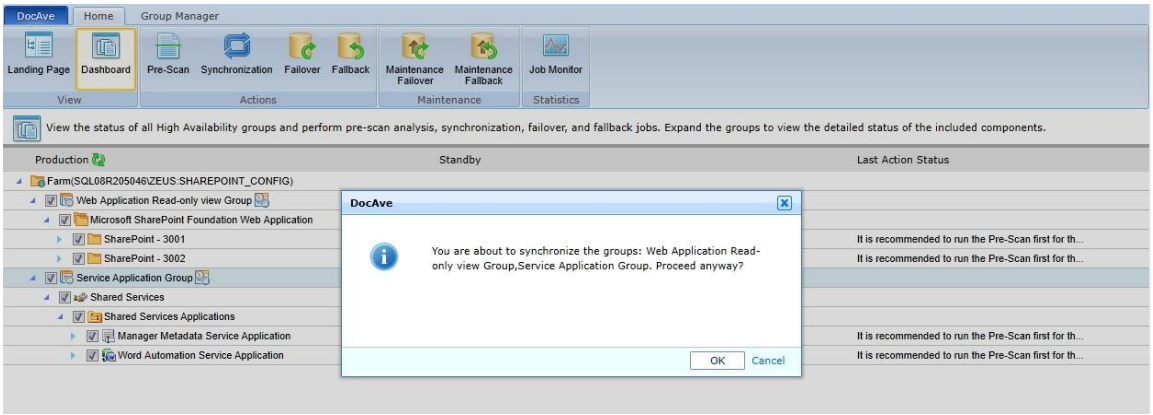

Figure 35: Performing Synchronization jobs.

2. Navigate to Job Monitor to view the job status and job information.

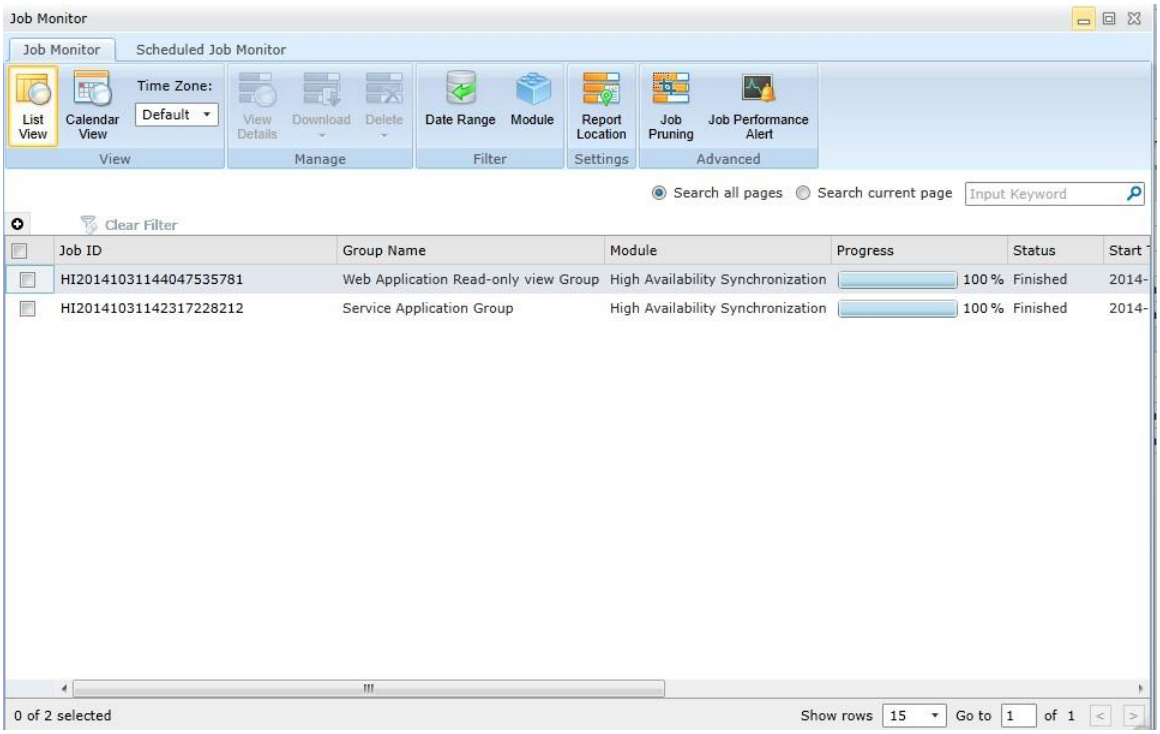

Figure 36: Viewing job status and job information in Job Monitor.

3. After the Synchronization jobs are finished, log into the production SQL instance to view the status of the production databases. The status of the production content databases and service application databases are **Principal, Synchronized**.

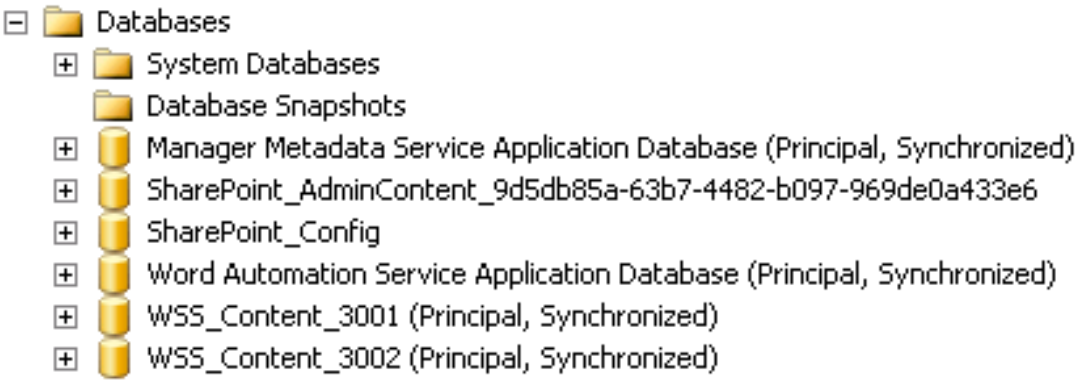

Figure 37: The status of the production databases.

4. Log into the standby SQL Instance to view the status of the standby databases. The standby content databases and service application databases are in **Mirror, Synchronized/Restoring…** status.

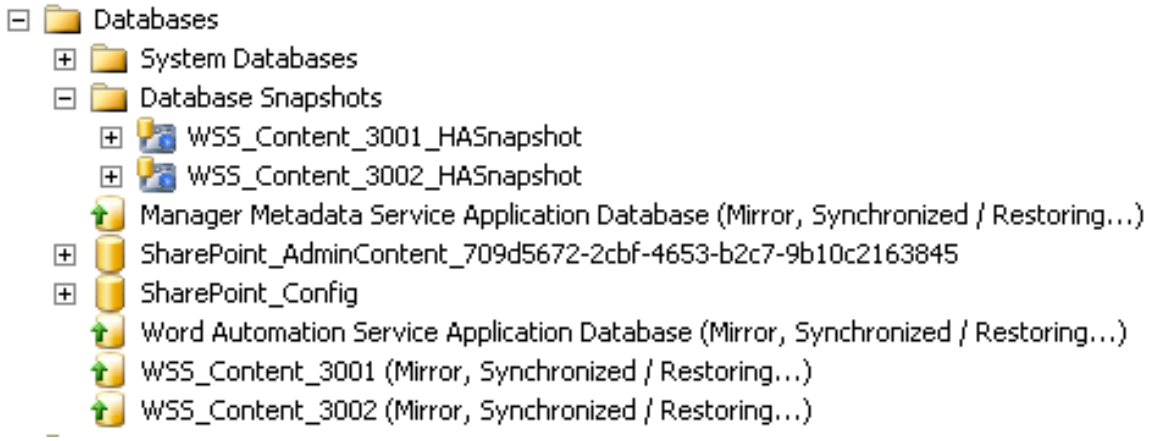

Figure 38: The status of the standby databases.

5. Log into SharePoint Central Administration of the standby farm. You will find that the standby Web applications were successfully created in the standby farm, since the High Availability group created for Web applications has the **Enable read-only view after synchronization** option selected; however, the service applications were not created in the standby farm, since the High Availability group created for service applications had the **Enable read-only view after synchronization** option deselected.

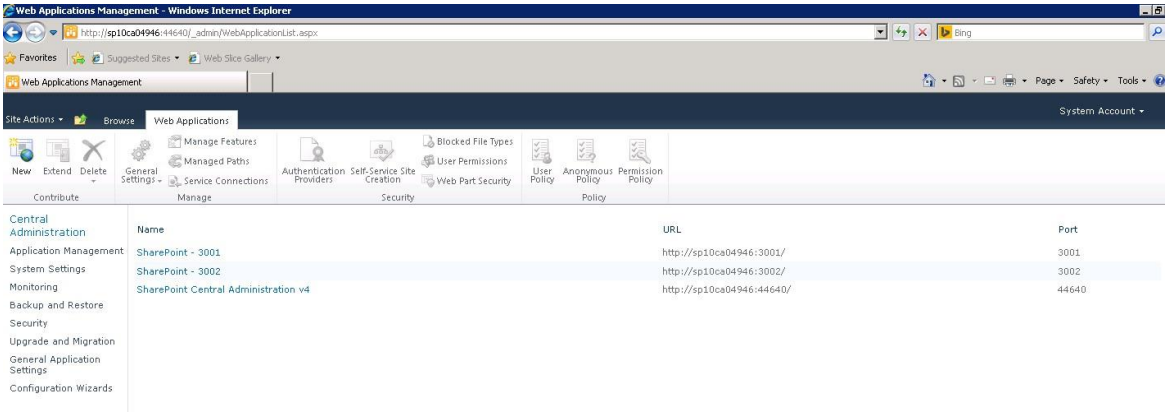

Figure 39: Viewing the standby Web applications.

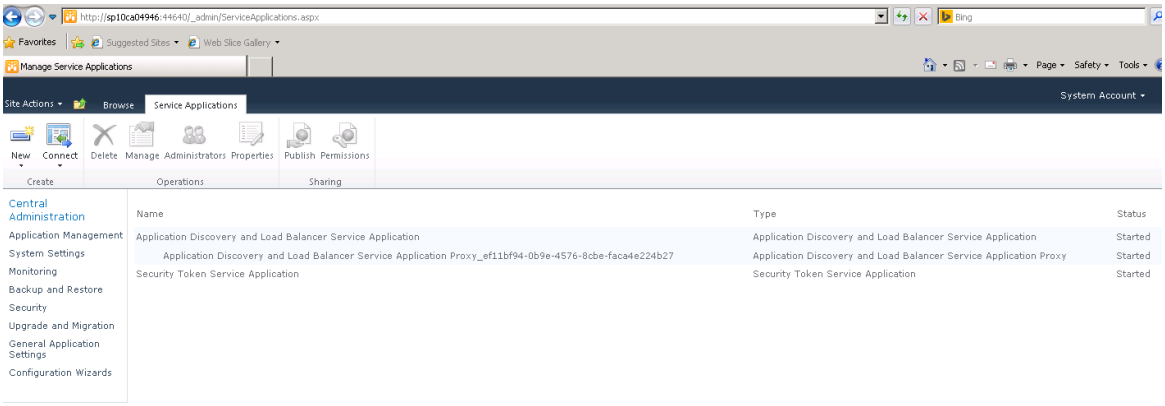

Figure 40: Standby service applications have not been created yet.

Performing Failover Jobs

Complete the steps below to perform Failover jobs:

1. In the **Dashboard** interface, select the all of the nodes in the two groups that have successfully performed the Synchronization jobs and then select **Failover** on the ribbon.
2. The **Failover** window appears. Select the **Keep SQL Mirroring relationship** option. Click **OK** to run the Failover jobs.
3. Go to Job Monitor to view the job status and job information.

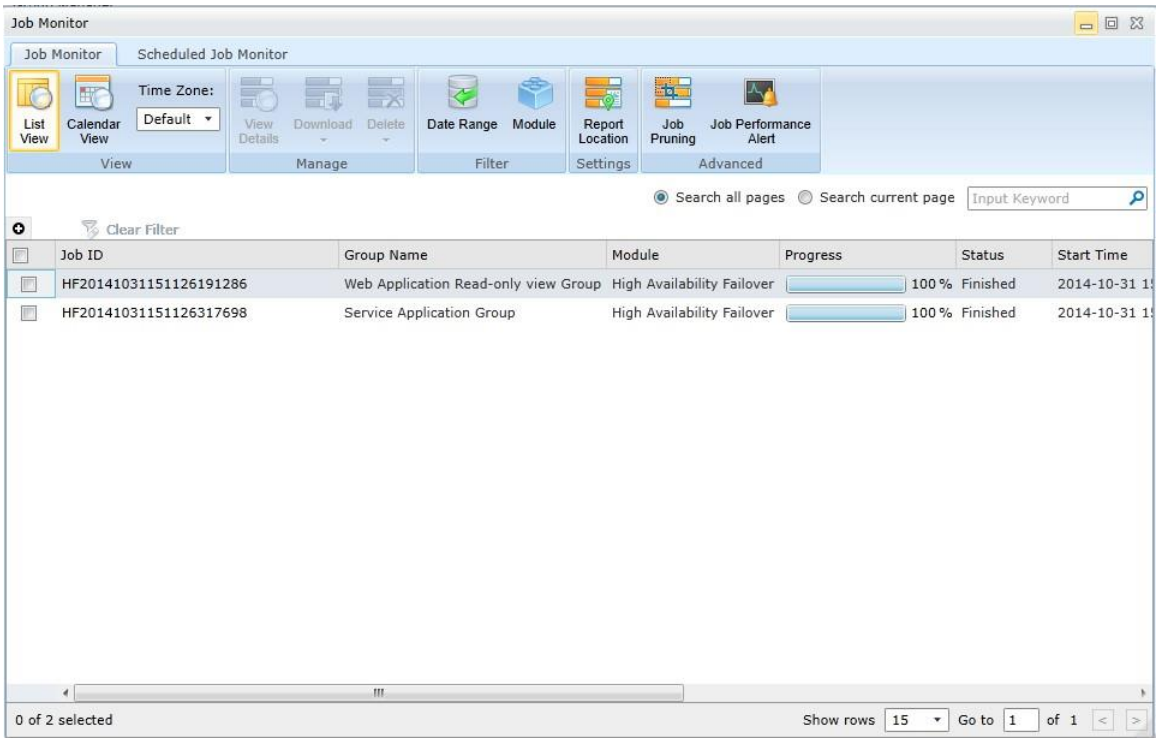

Figure 41: Viewing the job status and job information in Job Monitor.

4. After the Failover jobs are finished, log into the production SQL instance to view the status of the production databases. The status of the production content databases and service application databases are **Mirror, Synchronized/Restoring…**.

Figure 42: The status of the production databases.

5. Log into the standby SQL Instance to view the status of the standby databases. The standby content databases and service application databases are in **Principal, Synchronized** status.

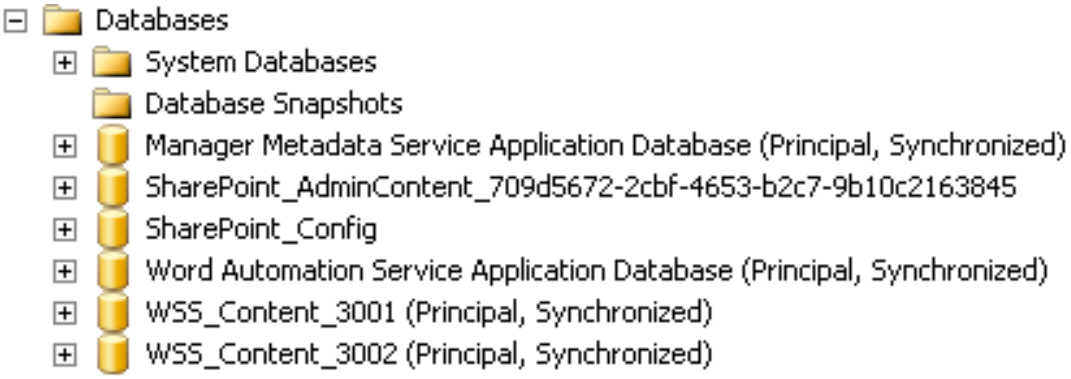

Figure 43: The status of the standby databases.

6. Log into SharePoint Central Administration of the standby farm. You will find that the standby service applications are successfully created in the standby farm.

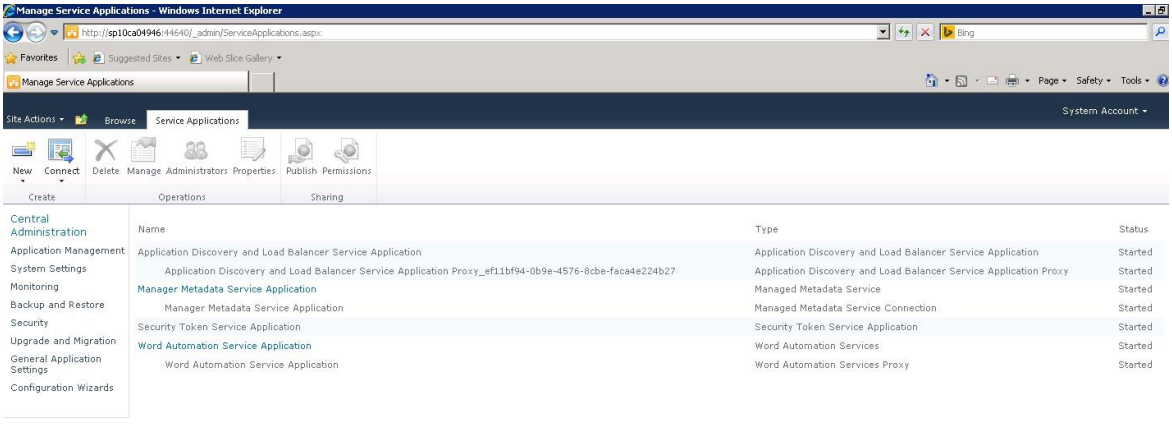

Figure 44: Viewing the standby service applications.

Performing a Fallback Job

In the **Dashboard** interface, select all of the nodes in the two groups that has successfully performed the Failover job. Click **Fallback** on the ribbon to perform Fallback jobs. After the Fallback jobs are finished, the production farm is in use again.

Improving Performance of Log Shipping Fallback

Compared with the other three Sync Methods, Log Shipping is more popularly used because it is less dependent on SQL Server versions and has more simple configurations and operations. Nevertheless, the Log Shipping method has its own limits on the data synchronization.

To minimize the data loss after Log Shipping Fallback, High Availability improved the performance of Log Shipping in DocAve 6 Service Pack 5.

One of the two pieces of logic may be applied for Log Shipping data transfer as described in the following two scenarios:

- **Scenario One** – If the production farm is available before Failover, refer to the information below:
    - **Overview** – With the **Perform an incremental synchronization job before Failover** option selected, a full backup snapshot will be created for the production databases in the production farm and a backup for transaction log file will be created before the Failover. The snapshot data and the transaction log backup file will be synchronized to the standby farm during Failover. The Log Shipping Fallback job will first back up the transaction logs for the data changes in standby farm and then use this transaction log backup file to synchronize the data changes back to the production farm.

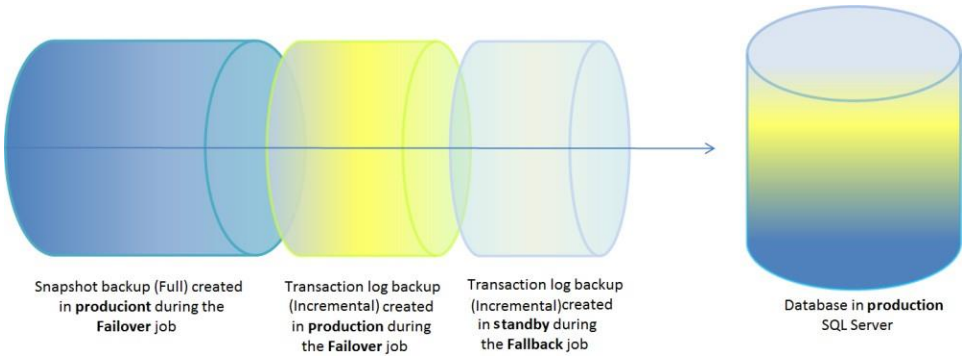

Figure 45: Synchronizing log backup file from standby farm to production farm in Log Shipping Fallback.

    - **Advantage –** The advantage of using this procedure is that it enhances performance and leverages the transaction log backup file transfer to make the data consistent between the standby farm and production farm after Fallback job, which mainly relies on the VSS snapshot technology and the consistency of the SQL Server database Log Sequence Number (LSN).
- **Scenario Two** – If the production farm is not available before Failover or the **Perform an incremental synchronization job before Failover** option is deselected, refer to the information below:
    - **Overview** – Log Shipping Failover will bring the standby farm online, with no data transferred, even if the **Perform an incremental synchronization job**

before Failover option is selected when the production farm is already unavailable. Log Shipping Fallback job will create a full backup snapshot in the standby farm and transfer the snapshot data to the production farm, which is exactly the same kind of synchronization used to move data from standby farm to the production farm.

Figure 46: Synchronizing full backup snapshot data from standby farm to production farm in Log Shipping Fallback.

- **Advantage** –This procedure is easy to perform. Additionally, the data will retain reliability if the data is completely transferred to the production farm. In addition, compared with the procedure above, this procedure saves disk space.

Performing a Log Shipping Fallback

Complete the steps below to walk through the two procedures of Log Shipping Fallback:

1. Create a High Availability group using **Log Shipping** sync method to include the content database, service database, stub database, and custom databases. For details about the settings for creating a group, refer to Configuring a High Availability Group.
2. Navigate to **Home** > **Dashboard**, and select the group you created in the last step to perform a Synchronization job. After the Synchronization job is finished, the status of the standby databases will be:
    - **Restoring**, if the **Enable read-only view after synchronization** option is not selected.
    - **Standby/ReadOnly**, if the **Enable read-only view after synchronization** option is selected.
3. Select the group in the **Dashboard**, after the Synchronization job is successfully finished. Select **Failover** on the ribbon. The **Failover** window appears.
    - For Scenario One in Improving Performance of Log Shipping Fallback, follow the instructions below:
        iv. Select the **Perform an incremental synchronization job before Failover** option, and then click **OK** to perform the Log Shipping Failover job.

v. After the Failover job is finished, log into the production SQL Server where the selected production databases reside. You will find a snapshot is created in the production SQL Server, and then log into the Agent server where the Log Shipping cache settings is configured. In the configured cache location, a folder named as **HighAvailability** is created storing the transaction log backup file.

The full path of the transaction log backup file is *…\CacheSetting\HighAvailability\Database log backup files\AgentName\PlanID\SnapshotID_SQLInstanceName_DatabaseName*.

***Note**: To find the cache location, navigate to **High Availability** > **Group Manager** > **Cache Setting**> **Log Shipping Cache Profile**, select the corresponding Log Shipping cache profile, and then click **View Details**.

- For Scenario Two in Improving Performance of Log Shipping Fallback, follow the instructions below:

i. If your production farm is already unavailable before performing Failover, the Log Shipping Failover job will bring the standby farm online without performing an incremental synchronization job before Failover regardless of whether you selected the **Perform an incremental synchronization job before Failover** option or not. If your production farm is available before performing the Failover, deselect the **Perform an incremental synchronization job before Failover** option, and then click **OK** to perform the Log Shipping Failover job.

ii. After the Failover job is finished, log into the production SQL Server where the selected production databases reside. You will find no snapshot was created in the production SQL Server and the **HighAvailability** folder was not created under the Agent Temporary Buffer location.

4. After the Failover job is successfully finished, access the production SharePoint farm. You will see the data in the production Web applications and the production content databases.

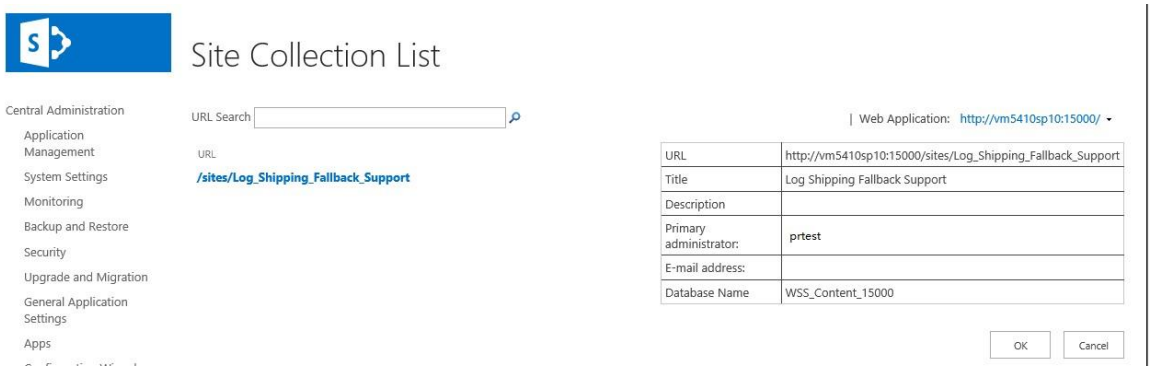

Figure 47: Data in the production Web application.

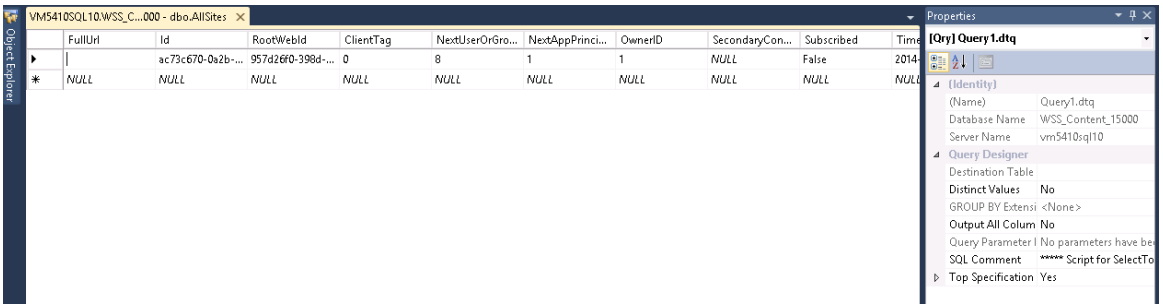

Figure 48: Data in dbo.AllWebs table of the production content database.

5. Go back to the DocAve Manager interface. Navigate to **Data Protection** > **High Availability**> **Home** > **Dashboard**. Select the group, and click **Fallback**. The **Fallback** window appears.
6. Choose whether or not to keep the renamed original production database after Fallback job is finished.

   ***Note**: If you have enough free disk space, AvePoint recommends selecting the **Keep the renamed original production database after the Fallback job has completed** option in case that the data resynchronized from standby farm contains errors. If this option is selected, a database named **OrigianlDatabaseName_HA_PlanID** will be created in the production SQL Server after Fallback.

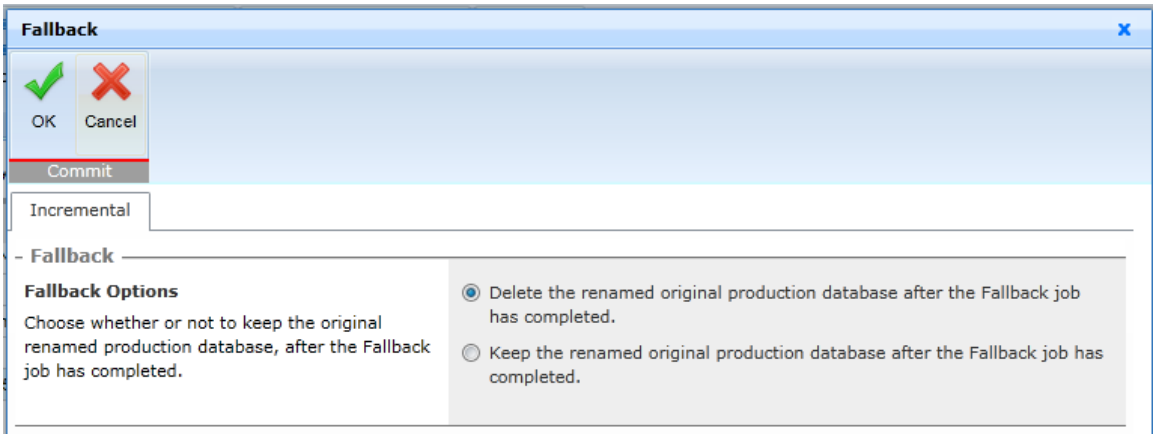

Figure 49: The Fallback window.

7. As to Scenario One, the snapshot is deleted from the production SQL Server and the transaction log backup file is deleted from the Agent temporary buffer location after the Fallback job.

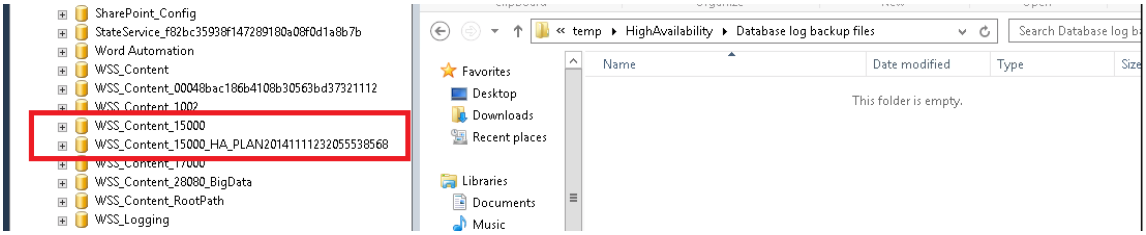

Figure 50: The renamed database is created in the production SQL Server and the transaction log backup file is deleted from the Agent temporary buffer location.

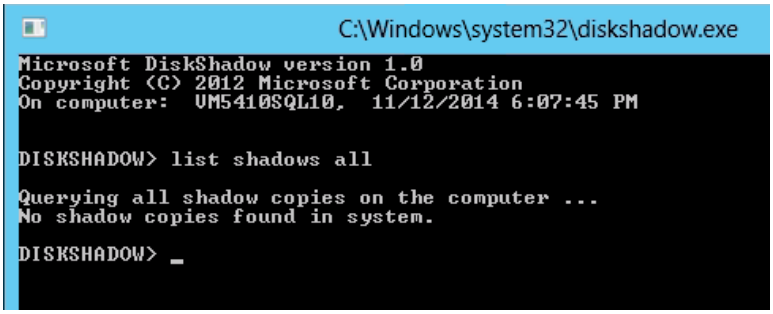

Figure 51: The snapshot is deleted from the production SQL Server.

8. Access to the production SharePoint farm, you will find that the site collection added to the standby farm after Failover is successfully synchronized to the production farm.

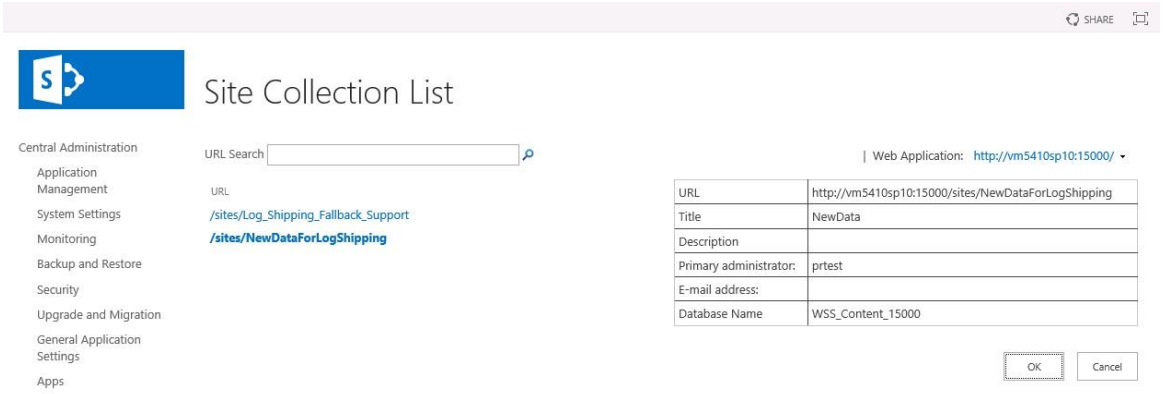

Figure 52: Viewing the production SharePoint farm.

Forcing a Full Synchronization on a Database with a Broken LSN by Customizing Configuration File

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

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

If you use **Log Shipping** or **Platform Backup Log Shipping** as the sync method, you must be sure not to break the database LSN (Log Sequence Number). DocAve High Availability provides you with 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**HighAvailability* directory on each Agent where the SharePoint Server resides in the production farm.
2. Open the **AgentCommonHAConfiguration.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.

Forcing a Full Synchronization or Full Fallback on Connector Data

If you include the Connector data into the High Availability group, you can configure the **AgentCommonHAConfiguration.xml** file to force High Availability to perform Full Synchronization or Full Fallback on Connector data. Complete the steps below:

1. Navigate to the …*AvePoint**DocAve6**Agent**data**HighAvailability* directory on the Agent server.

   ***Note**: To force a Full Synchronization, edit the configuration file on each Agent of the production Agent group. To force a Full Fallback, edit the configuration file on each Agent of the standby Agent group.

2. Open the **AgentCommonHAConfiguration.xml** file with Notepad.
3. Find the **SyncConnectorConfig** node, and then change the value of the **ForceFull** attribute to **True**.
4. Save the configuration file, and close it.

Designating VSS Hardware Provider for Creating Snapshot

To select the VSS Hardware Provider that will be used to created snapshots in High Availability Synchronization or Fallback jobs, modify the **AgentCommonHAConfiguration.xml** configuration file.

1. Navigate to the *…\AvePoint\DocAve6\Agent\data\HighAvailability* directory on each Agent server that is included in the Production Agent Group and Standby Agent Group selected in this High Availability group.
2. Open the **AgentCommonHAConfiguration.xml** file with Notepad.
3. Find the **<VSSProviderConfig>** node, and locate its child node: **VolumeConfig**.
    - **ProviderOption** – Enter either **0**, **1**, or **2** for this option. The default value is **1**, which indicates that the system provider will be used to create a snapshot. The value **0** indicates that the system allocated provider will be used. The value **2** indicates that a user designated provider will be used. If you set the value of this attribute to **2**, you must add the **ProviderID** attribute to enter the GUID of the provider you want to use.

      ***Note**: If you choose to designate a VSS hardware provider to create a snapshot, High Availability will recreate the snapshot using the system provider when the preferred provider failed to create the snapshot. If the volume does not support the system provider, the VSS writer will randomly select a provider. If the hardware provider you choose to create snapshot is not a user-designated one, the snapshot creation will not be executed again after the snapshot creation failed. If you have more than one volume for creating snapshots, all of the snapshots will be recreated once High Availability detects one snapshot creation has failed.

    - **ProviderID** – Enter the GUID of the provider you want to use. The value of this attribute is case sensitive and this attribute is only valid when the **ProviderOption** attribute is set to **2**.
4. Save the configuration file and close it.

Defining Warm Standby Behavior of Service Applications

By default, if the High Availability group including service applications has selected the **Enable read-only view after synchronization** option, the service applications will not be created in the standby farm after the Synchronization job for implementing warm standby. If you want High Availability to create the service applications in the standby farm and make them accessible after Synchronization for warm standby, complete the steps below:

1. Navigate to the *…\AvePoint\DocAve6\Agent\data\HighAvailability* directory on each Agent installed with SharePoint Server.
2. Open the **AgentCommonHAConfiguration.xml** file with Notepad.
3. Find the **<ServiceConfig SupportWarmStandby="false"/>** node.
4. Set the value of **SupportWarmStandby** attribute to **true**.
5. Additionally, if you want to publish the service applications that are shared across farms to the standby farm for warm standby, set the value of the **NeedCross** attribute to **true**. If the value is **false**, a new service application will be created in the standby farm even if the service application in the production farm is shared across farms.

   ***Note:** The service applications shared across farms can be **Business Data Connectivity Service**, **App Management Service Application** (only available SharePoint 2013 and SharePoint 2016), **Managed Metadata Service**, **User Profile Service Application**, **Search Service Application**, and **Secure Store Service**. If the production farm and standby farm are in different domains, the **User Profile Service Application** requires the trust to be enabled between both of the domains; the **Business Data Connectivity Service** and **Secure Store Service Application** administration features requires the domain of the consuming farm must be trusted by the domain of the publishing farm.

   ***Note**: It is not recommended to modify the service applications in the standby farm, since the changes will be reflected in both of the consuming farm and publishing farm.

6. Save the configuration file and close it.

Designating Timeout Session for SQL Statement Execution in Log Shipping Incremental Data Transfer

Log Shipping sync method leverages the SQL statement to execute the incremental data transfer. By default, the timeout session of executing the SQL statement is 3600 seconds, which can be modified through configuration file. In addition, High Availability will automatically estimate the size of the incremental data to transfer and come up with a timeout session value. The timeout session evaluated according to the data size to transfer will be compared with the value of the **CommandTimeOut** attribute in the **AgentCommonHAConfiguration.xml** file. The longer time will be applied as the real timeout session for SQL statement execution.

Refer to the table below for the timeout sessions in different ranges of data size.

Data Size	Timeout Session
< 50 GB	1 Hour
50 GB <= Data Size < 100 GB	2 Hour
100 GB <= Data Size < 200 GB	3 Hour
200 GB <= Data Size < 500 GB	10 Hour
>= 500 GB	24 Hour

In addition, if you do not want to set a timeout session for executing the SQL statement in Log Shipping incremental data transfer, set the value of the **CommandTimeOut** attribute to **0** in the **AgentCommonHAConfiguration.xml** file.

To customize the timeout session for SQL statement execution, complete the steps below:

1. Navigate to the …*AvePoint**DocAve6**Agent**data**HighAvailability* directory on the SQL Server Agent.
2. Open the **AgentCommonHAConfiguration.xml** file with Notepad.
3. Find the **<TSQLConfig CommandTimeOut="3600"/>** node. You can change the value of the **CommandTimeOut** attribute as desired. The time unit is **Second**. Setting the value to **0** means there is no time limit on the SQL statement execution for Log Shipping incremental data transfer.
4. Save the changes and close this file.

Appendix A: SharePoint Components Supported for High Availability

Refer to the following three sections on the SharePoint 2010, SharePoint 2013, and SharePoint 2016 components supported for High Availability.

***Note**: High Availability does not support SharePoint built-in workflow, workflows defined by SharePoint designers, or the workflows defined by Nintex workflow.

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**	**Read-Only Standby**
Configuration Database		Supported	Not Supported	Not Supported
Central Administration Database		Supported	Not Supported	Not Supported
Web Application	Content Database	Supported	Supported	Supported
Service Application	**Databases of Service Application**			
Access Service	This Service Application does not have a database	Not Supported	Supported	Supported
Application Registry Service	Application registry service database	Supported	Not Supported	Not Supported
Business Data Connectivity Service	Bdc_Service_DB(Business Data Connectivity Database)	Supported	Supported	Supported
Excel Services Application	This Service Application does not have a database	Not Supported	Supported	Supported
Managed Metadata Service	Managed Metadata Service_Database(Service Application Database)	Supported	Supported	Supported
PerformancePoint Service Application	PerformancePoint Service Application_Database(Microsoft.PerformancePoint.Scorecards.BIMonitoringServiceDatabase)	Supported	Supported	Supported
Search Service Application	Search_Service_Application_DB(Administration Database)	Supported	Supported	Supported

SharePoint Components		Single Farm Mode	Standby Farm Mode	Read-Only Standby
	Search_Service_Application_CrawlStoreDB(Crawl Database) Search_Service_Application_PropertyStoreDB(Property Database)			
Secure Store Service	Secure_Store_Service_DB(Secure Store Service Database)	Supported	Supported	Supported
State Service	StateService_Database(State Service Database Settings)	Supported	Not Supported	Not Supported
Visio Graphics Service	This Service Application does not have a database	Not Supported	Supported	Supported
Usage and Health data collection	WSS_Logging	Supported	Not 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. If the standby farm is in a different domain that is not trusted by the production domain, it is not 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. If the standby farm is in a different domain that is not trusted by the production domain, it is not supported.
Web Analytics Service Application	WebAnalyticsServiceApplication_StagingDB(Web Analytics Staging Database) WebAnalyticsServiceApplication_ReportingDB(Web	Supported	Supported	Supported

SharePoint Components		Single Farm Mode	Standby Farm Mode	Read-Only Standby
	Analytics Warehouse Database)			
Word Automation Service	WordAutomationServices_Database(Microsoft.Office.Word.Server.Service.QueueDatabase)	Supported	Supported	Supported
Subscription Service	Subscription Settings Database	Supported	Supported	Supported
Session State Service	SessionStateService_Database	Supported	Not Supported	Not Supported
SharePoint Foundation Search	WSS_Search_Database	Supported	Not Supported	Not Supported

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	Read-Only Standby
Configuration Database		Supported	Not Supported	Not Supported
Central Administration Database		Supported	Not Supported	Not Supported
Web Application	Content Database	Supported	Supported	Supported
Service Application	**Databases of Service Application**			
App Management Service Application	AppMng_Service_DB (App Management Database)	Supported	Supported ***Note**: It is not supported, if the standby farm is in a different domain that is not trusted by the production domain.	Supported ***Note**: It is not supported, if the standby farm is in a different domain that is not trusted by the production domain.
Access Service Application	This Service Application does not have a database	Not Supported	Supported	Supported
Access Service Application 2010	This Service Application does not have a database	Not Supported	Supported	Supported
Business Data Connectivity Service	Bdc_Service_DB(Business Data Connectivity Database)	Supported	Supported	Supported
Excel Services Application	This Service Application does not have a database	Not Supported	Supported	Supported
Managed Metadata Service	Managed Metadata Service_Database(Service Application Database)	Supported	Supported	Supported
Machine Translation Service Application	Machine Translation Service Database	Supported	Not Supported	Not Supported
PerformancePoint Service Application	PerformancePoint Service Application_Database(Microsoft.PerformancePoint.Scorecards.BIMonitoringServiceDatabase)	Supported	Supported	Supported

SharePoint Components		Single Farm Mode	Standby Farm Mode	Read-Only Standby
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
State Service	StateService_Database(State Service Database Settings)	Supported	Not Supported	Not Supported
Secure Store Service	Secure_Store_Service_DB(Secure Store Service Database)	Supported	Supported	Supported
Session State Service	SessionStateService_Database	Supported	Not Supported	Not Supported
Security Token Service Application	This Service Application does not have a database	Not Supported	Not Supported	Not Supported
Visio Graphics Service	This Service Application does not have a database	Not Supported	Supported	Supported
Usage and Health data collection	WSS_Logging	Supported	Not 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**: It is not supported, if the standby farm is in a different domain that is not trusted by the production domain.	Supported ***Note**: It is not supported, if the standby farm is in a different domain that is not trusted by the production domain.
Word Automation Service	WordAutomationServices_Database(Microsoft.Office.Word.Server.Service.QueueDatabase)	Supported	Supported	Supported

SharePoint Components		Single Farm Mode	Standby Farm Mode	Read-Only Standby
Workflow Service Application	This Service Application does not have a database	Not Supported	Not Supported	Not Supported
Work Management Service Application	This Service Application does not have a database	Not Supported	Not Supported	Not Supported
PowerPoint Service Application	This Service Application does not have a database	Not Supported	Supported	Supported
Subscription Service	Subscription Settings Database	Supported	Supported	Supported

***Note**: The SharePoint 2013 service applications supported by DocAve High Availability whose databases are added to the AlwaysOn Availability group may experience data loss during Failover, if the AlwaysOn Availability group uses asynchronous-commit for disaster recovery. For details, refer to the Supported high availability and disaster recovery options for SharePoint databases (SharePoint 2013) article on Microsoft website.

SharePoint 2016 Components Supported for High Availability

Refer to the table below for more information on the SharePoint 2016 components supported by the Single farm mode High Availability and the Standby farm mode High Availability:

SharePoint Components		Single Farm Mode	Standby Farm Mode	Read-Only Standby
Configuration Database		Supported	Not Supported	Not Supported
Central Administration Database		Supported	Not Supported	Not Supported
Web Application	Content Database	Supported	Supported	Supported
Service Application	**Databases of Service Application**			
App Management Service Application	AppMng_Service_DB (App Management Database)	Supported	Supported ***Note**: It is not supported, if the standby farm is in a different domain that is not trusted by the production domain.	Supported ***Note**: It is not supported, if the standby farm is in a different domain that is not trusted by the production domain.
Access Service Application	This Service Application does not have a database	Not Supported	Supported	Supported
Access Service Application 2010	This Service Application does not have a database	Not Supported	Supported	Supported
Business Data Connectivity Service	Bdc_Service_DB(Business Data Connectivity Database)	Supported	Supported	Supported
Managed Metadata Service	Managed Metadata Service_Database(Service Application Database)	Supported	Supported	Supported
Machine Translation Service Application	Machine Translation Service Database	Supported	Not Supported	Not Supported
Project Service	This Service Application does not have a database	Not Supported	Supported	Supported
PerformancePoint Service Application	PerformancePoint Service Application_Database(Microsoft.PerformancePoint.Scorecards.BIMonitoringServiceDatabase)	Supported	Supported	Supported

SharePoint Components		Single Farm Mode	Standby Farm Mode	Read-Only Standby
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
State Service	StateService_Database(State Service Database Settings)	Supported	Not Supported	Not Supported
Secure Store Service	Secure_Store_Service_DB(Secure Store Service Database)	Supported	Supported	Supported
Session State Service	SessionStateService_Database	Supported	Not Supported	Not Supported
Security Token Service Application	This Service Application does not have a database	Not Supported	Not Supported	Not Supported
Visio Graphics Service	This Service Application does not have a database	Not Supported	Supported	Supported
Usage and Health data collection	WSS_Logging	Supported	Not 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**: It is not supported, if the standby farm is in a different domain that is not trusted by the production domain.	Supported ***Note**: It is not supported, if the standby farm is in a different domain that is not trusted by the production domain.
Word Automation Service	WordAutomationServices_Database(Microsoft.Office.Word.Server.Service.QueueDatabase)	Supported	Supported	Supported

SharePoint Components		Single Farm Mode	Standby Farm Mode	Read-Only Standby
Workflow Service Application	This Service Application does not have a database	Not Supported	Not Supported	Not Supported
Work Management Service Application	This Service Application does not have a database	Not Supported	Not Supported	Not Supported
PowerPoint Service Application	This Service Application does not have a database	Not Supported	Supported	Supported
Subscription Service	Subscription Settings Database	Supported	Supported	Supported

***Note**: The SharePoint 2016 service applications supported by DocAve High Availability whose databases are added to the AlwaysOn Availability group may experience data loss during Failover, if the AlwaysOn Availability group uses asynchronous-commit for disaster recovery. For details, refer to the Supported high availability and disaster recovery options for SharePoint databases (SharePoint 2013) article on Microsoft website.

Appendix B: 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 SQL Server version for SQL Mirroring – Check if the versions of the SQL Servers on the production farm and standby farm are the same, and if the version is SQL 2005 SP1 or later.
- Check the free disk space in the standby server – Check if there is enough free disk space in the standby server.
- Check the status of the production database for SQL Mirroring – Check if the status of the production database is normal. If the production database status is principle, and the standby database status is mirroring, the database will be skipped.
- Check the status of the standby database for SQL Mirroring – Check if the standby database exists.
- Check the recovery model of the selected nodes – Check if the recovery mode of the selected nodes is **Full**.
- Check the connection between the production server and the standby server – Check if the production server can successfully connect to the standby server using the specific port.
- Check the connection between the standby server and the production server – Check if the standby server can successfully connect to the production server using the specific port.
- Check the Agent account permission of **db_owner** role to database – Check if the Agent account has the **db_owner** database role to the production databases.
- Check the production endpoint for SQL Mirroring – Check if there is no endpoint in the production SQL Instance, or there is the endpoint created by Windows Authentication, or named as **DocAveSQLMIRROR**.
- Check the standby endpoint for SQL Mirroring – Check if there is no endpoint in the standby SQL Instance, or there is the endpoint created by Windows Authentication, named as **DocAveSQLMIRROR**.
- Check if the RBS Filestream Provider is supported by this sync method – Check if the RBS Filestream Provider feature is enabled for the databases. SQL Mirroring sync method does not support RBS Filestream Provider.

- Check Data Service Port – Check if the Agent in production SQL Server can connect the Agent in standby SQL Server using the port of data transfer.
- Check for illegal database names in the product SQL Instance – Check if there are any databases in the production SQL Instance with blank space before or after the database name.
- Check for illegal database names in the standby SQL Instance – Check if there are any databases in the standby SQL Instance with blank space before or after the database name.
- Check SQL Server VSS Writer logged-in user permission – Check if the SQL Server VSS Writer logged-in user has sufficient permission.
- Check File Share Shadow Copy provider – Check if the File Share Shadow Copy provider exists when the database file location is in the UNC path.
- Check free space for Connector BLOB data – Check if there is enough free space in the standby physical device for storing the Connector BLOB data.

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 if there are conflicting production site collections using the same URL in the content databases that are added into the same standby Web application.
- Check for a site conflict in standby farm – Check if 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 if 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 SQL Server version for AlwaysOn Availability Group – Check if 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 if there is enough free disk space in the standby server.
- Check the status of the production database for AlwaysOn Availability Group – Check if 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 for AlwaysOn Availability Group – Check if the standby database exits. 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 if the recovery mode of the selected nodes is **Full**.
- Check the production SQL instance role in the AlwaysOn Availability Group – Check if the production SQL instance role in the AlwaysOn Availability Group is **Primary**.
- Check the Agent account permission of **db_owner** role to database – Check if the Agent account has the **db_owner** database role to the production databases.
- Check if the RBS Filestream Provider is supported by this sync method – Check if the RBS Filestream Provider feature is enabled for the databases.
- Check Data Service Port – Check if the Agent in production SQL Server can connect the Agent in standby SQL Server using the port of data transfer.
- Check for illegal database names in the product SQL Instance – Check if there are any databases in the production SQL Instance with blank space before or after the database name.
- Check for illegal database names in the standby SQL Instance – Check if there are any databases in the standby SQL Instance with blank space before or after the database name.
- Check SQL Server VSS Writer logged-in user permission – Check if the SQL Server VSS Writer logged-in user has sufficient permission.
- Check File Share Shadow Copy provider – Check if the File Share Shadow Copy provider exists when the database file location is in the UNC path.
- Check free space for Connector BLOB data – Check if there is enough free space in the standby physical device for storing the Connector BLOB data.

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 if there are conflicting production site collections using the same URL in the content databases that are added into the same standby Web application.
- Check for a site conflict in standby farm – Check if 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 if 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 SQL Server version for Log Shipping – Check if 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 if there is enough free disk space in the standby server.
- Check the status of the production database for Log Shipping – Check if the status of the production database is normal.
- Check the status of the standby database for Log Shipping – Check if the standby database exits or the status of the standby database is **Read-only** or **Restoring**.
- Check the recovery model of the selected nodes – Check if the recovery mode of the selected nodes is **Full** or **Bulk-logged**.
- Check the Log Sequence Number (LSN) – Check if the LSN of the production database and standby database matches or not.
- Check the Agent account permission of **db_owner** role to database – Check if the Agent account has the **db_owner** database role to the production databases.
- Check if the RBS Filestream Provider is supported by this sync method – Check if the RBS Filestream Provider feature is enabled for the databases.
- Check Data Service Port – Check if the Agent in production SQL Server can connect the Agent in standby SQL Server using the port of data transfer.
- Check for illegal database names in the product SQL Instance – Check if there are any databases in the production SQL Instance with blank space before or after the database name.
- Check for illegal database names in the standby SQL Instance – Check if there are any databases in the standby SQL Instance with blank space before or after the database name.
- Check SQL Server VSS Writer logged-in user permission – Check if the SQL Server VSS Writer logged-in user has sufficient permission.
- Check File Share Shadow Copy provider – Check if the File Share Shadow Copy provider exists when the database file location is in the UNC path.
- Check free space for Connector BLOB data – Check if there is enough free space in the standby physical device for storing the Connector BLOB data.

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 if there are conflicting production site collections using the same URL in the content databases that are added into the same standby Web application.
- Check for a site conflict in standby farm – Check if 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 if 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 SQL Server version – Check if the major version of the SQL Servers on the production farm and standby farm is the same.
- Check free disk space on the standby server for Platform Backup Log Shipping – Check if there is enough free disk space in the standby server.
- Check the status of the production database for Log Shipping – Check if the status of the production database is normal.
- Check the status of the standby database – Check if the standby database exits or the status of the standby database is **Read-only** or **Restoring**.
- Check the recovery model of the selected nodes – Check if the recovery mode of the selected nodes is **Full** or **Bulk-logged**.
- Check the Log Sequence Number (LSN) for Platform Backup Log Shipping – Check if the LSN of the backed up database and standby database matches or not.
- Check the Platform Backup history – Check if the Platform Backup job record for the selected farm components exists.
- Check the Agent account permission of **db_owner** role to database – Check if the Agent account has the **db_owner** database role to the production databases.
- Check for available Media service – Check if there is available Media service.
- Check if the RBS Filestream Provider is supported by this sync method – Check if the RBS Filestream Provider feature is enabled for the databases. SQL Mirroring sync method does not support RBS Filestream Provider.
- Check for illegal database names in the product SQL Instance – Check if there are any databases in the production SQL Instance with blank space before or after the database name.
- Check for illegal database names in the standby SQL Instance – Check if there are any databases in the standby SQL Instance with blank space before or after the database name.

- Check SQL Server VSS Writer logged-in user permission – Check if the SQL Server VSS Writer logged-in user has sufficient permission.
- Check File Share Shadow Copy provider – Check if the File Share Shadow Copy provider exists when the database file location is in UNC path.
- Check free space for Connector BLOB data – Check if there is enough free space in the standby physical device for storing the Connector BLOB data.

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 if there are conflicting production site collections using the same URL in the content databases that are added into the same standby Web application.
- Check for a site conflict in standby farm – Check if 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 if 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 Heath 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 if 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 if 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 is a Member of IIS User Group − Check if 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 if 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 if 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 if 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 if 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 if 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 if 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 if 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 if the Agent version matches the Manager version. The Agent version must be the same as the Manager version for communication.

Appendix C: 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 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
		Fallback	B
		Test Failover	TF
		Stop Test Failover	SF
		Maintenance Failover	MF
		Maintenance Fallback	MB
		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	CG	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					

Functionality Name and Hot Key										
SQL Instance Settings	SI	Close		X						
		Configure	Save	S						
			Cancel	C						
Custom Action	CA	Command Profile	C	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					
		Script Profile	S	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					
Cache Setting	CS	Connector Cache Setting	C	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					
		Log Shipping Cache Setting	L	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					
Job Monitor	J									

Appendix D: Protecting Data in Data ONTAP Storage System with High Availability SnapMirror Sync Method

If you are using a Data ONTAP storage system and your data resides in volumes with NetApp SnapMirror relationship set up, it is recommended that you use High Availability SnapMirror sync method to protect your data in LUNs or CIFS Shares.

To use the High Availability SnapMirror sync method in SnapManager 8.2 for SharePoint, you must apply a DocAve 6 Service Pack 5 Cumulative Update 1 or later High Availability license with SnapMirror enabled. For details on how to apply a license, refer to the *License Manager* section in *SnapManager for SharePoint Control Panel User's Guide*.

Note: The SQL Mirroring, Log Shipping, and AlwaysOn Availability Group sync methods are also available to use in SnapManager 8.2 for SharePoint High Availability, but only after applying a DocAve High Availability license.

Refer to the following sections for details on the permissions, configurations, and operations required to use the High Availability SnapMirror sync method.

Where to Install Agents for High Availability

Refer to the *Appendix A: Where to Install DocAve Agents* section in the DocAve 6 Installation Guide for details on where to install SnapManager for SharePoint Agents for use with High Availability.

Specific System Requirements and Limitations

Refer to the section below for information about system requirements and limitations for using High Availability in a Data ONTAP storage system.

Since Windows Server 2008 cannot create snapshots for CIFS Share devices, read-only view is not supported for a standby environment after a synchronization job when the production databases in the CIFS Share are located in SQL Server running Windows Server 2008.

To enable read-only view for a standby environment after synchronizing with a database in a CIFS share, the Agent account used to install the Agent on the SQL Server must be a SQL Server Service account.

To perform a High Availability job on a SharePoint 2010 farm, .NET Framework 4.5 must be installed on SQL Server.

Does not support High Availability for production farms with SnapDrive 7.1 installed. In this case, the clone creation operation will fail.

Required Permissions for SnapMirror

Refer to the section below for the permissions required to use SnapMirror sync 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
    - Full Control permission to the Web 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.

      ***Note**: This permission is only required for provisioning Managed Metadata Service in the standby farm.

      ***Note:** SnapManager for SharePoint requires the use of the db_owner role for content databases. If RBS is enabled, the Web Application Services account must have the SP_DATA_ACCESS role and the db_owner role in order to work with RBS content stored in content databases. For more information on SharePoint database roles see http://technet.microsoft.com/en-us/library/ee748631%28v=office.15%29.aspx.

Agent account configured on 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
    - **Control Server** to the destination SQL instance
    - Server Role of **sysadmin** in the SQL instance

      ***Note**: The Agent account to execute the High Availability job must have the **db_owner** database role to the standby databases, otherwise, the Agent account configured on the SQL Server must be granted with the server role of **sysadmin** to the destination SQL Server.
      ***Note**: **Read** and **Write** permissions to the **Temporary Buffer**, which is configured in **Control Panel** > **Agent Monitor** > **Configure**. High Availability uses the Agent Temporary Buffer location to store the SQLite database file for Connector.

Agent account configured to access the storage system:

The Agent account configured to access the storage system must be:

- A member of the local **Administrators** group, if the storage system is Data ONTAP 7.X or 7 mode of Data ONTAP 8.X.
- A member of **Ontapi admin** group, if the storage system is Cluster mode of Data ONTAP 8.X.

SQL Server Service account configured on SQL Server:

The SQL Server Service account configured on the SQL Server must have **Read** and **Write** permissions to the following paths:

- CIFS share path where database files reside
- The directory of *…\NetApp\SMSP8\Agent\Jobs* or *…\IBM\SMSP8\Agent\Jobs*.

VSS Writer account configured on SQL Server:

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

- A member of the local **Administrators** group.
- Server role of **sysadmin** to the SQL Server

- **Read** and **Write** permissions to the database file location (including the path in file share).

SharePoint 2013 application pool account configured on SQL Server:

For SharePoint 2013, you must ensure that the standby application pool account exists in the standby SQL Server and that it has the **db_owner** role to the production database, or you can grant the application pool account the server role of **sysadmin** in the standby SQL Server.

Service application pool account configured on a SharePoint Server:

If the High Availability group includes the PowerPoint Service Application, the service application pool account configured on the SharePoint server must have **Write** permission to the *C:\ProgramData\Microsoft\SharePoint* directory in the SharePoint server. This is to store the temporary file generated during a conversion job.

Agent account configured on the SharePoint Server to start the SP2010StorageOptimizationService.exe process or SP2013StorageOptimizationService.exe process

If you are about to synchronize a content database with BLOB data and its related stub database to a standby farm with read-only view enabled, make sure that the Agent account on the SharePoint server has sufficient permissions in the following scenarios:

- If the Agent account in the standby farm is different from the Agent account in the production farm but is in the same domain, the Agent account in the standby farm must have the **db_owner** role in the production stub database. This is in order to make sure the standby stub files are readable.
- If the Agent account in the standby farm is different from the Agent account in the production farm and is in a different domain, it is recommended that you make the production farm domain trusted by the standby farm domain, and that you grant the Agent account in the standby farm the **db_owner** role in the production stub database. If the recommendation above is not feasible, the Agent account in the standby farm must have the **sysadmin** role to the standby SQL instance for the standby stub files to be readable.

Building a High Availability Group

High Availability group is designed to share common settings for the nodes selected from the SharePoint farm component tree. Each High Availability group contains at least one database node, and each node can only exist in one group.

There are two modes to build a High Availability group: Form Mode and Wizard Mode. Use the Wizard if this is your first time building a High Availability group; use Form Mode if you are familiar with building the group.

Follow the instructions below to create a Single farm or Standby farm High Availability group using Wizard Mode.

Creating a Single Farm Mode High Availability Group Using Wizard Mode

To create a single farm mode High Availability group using the wizard, complete the steps below:

1. In the **Group Manager** tab, click **Create** on the ribbon and the select **Wizard Mode** from the drop-down list.
2. Enter a group name and an optional description for the group you are creating. Click **Next**.
3. In the **General Settings** page, configure the following settings:

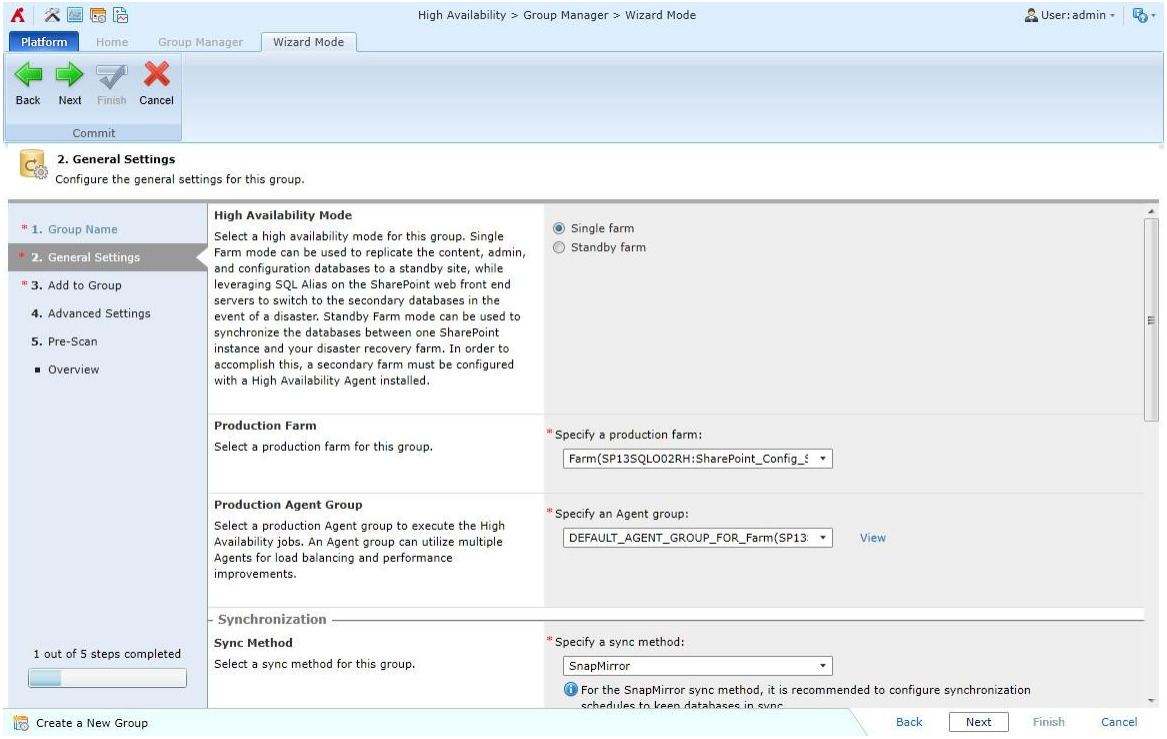

Figure 53: The General Settings interface (1).

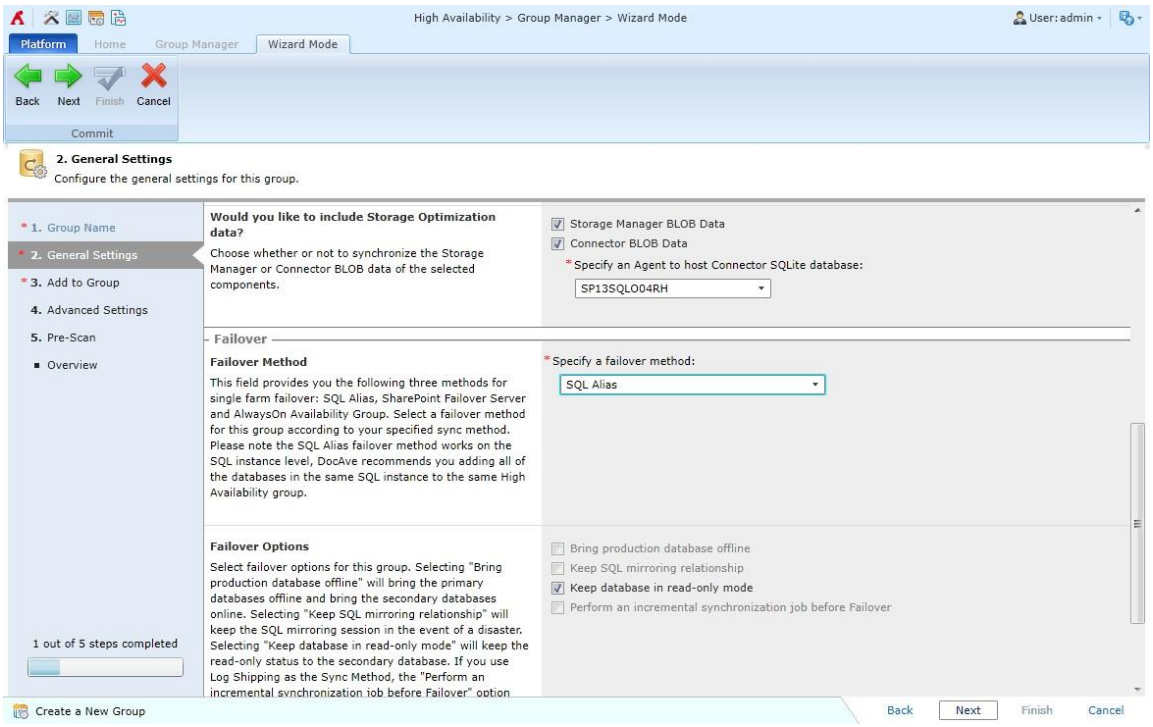

Figure 54: The General Settings interface (2).

a. **High Availability Mode** – Choose which High Availability Mode to use according to your environment. For this process, select **Single farm** option.

b. **Production Farm** – Select the production farm from the drop-down list that you want to keep continuously available and operational.

c. **Production Agent Group** – Select a production Agent group from the drop-down list to execute the job. An Agent group can utilize multiple Agents for load balancing and job performance improvements.

d. **Sync Method** – Select **SnapMirror** from the sync method drop-down list.

e. **Would you like to include Storage Optimization data?** – When one or both of these options are selected, a High Availability Synchronization job synchronizes Storage Manager BLOB data and Connector BLOB data to the SnapMirror secondary storage. If you select the **Connector BLOB Data** option, you are required to select an Agent to host the SQLite database that records the production device information. The SQLite database will be updated in each synchronization job, and will be used when performing Failover.

   ***Note**: Only when the **Storage Manager BLOB Data** or **Connector BLOB Data** option is selected can the associated stub databases be displayed in the tree.

f. **Failover Method** – Select the SQL Alias method for the Single farm High Availability group to synchronize the production database to a SQL Alias. It is recommended that

you select the same SQL Alias for all of the production databases in the same SQL instance.

g. **Failover Options** – Select your desired failover option. For a read-only standby environment, select the **Keep database in read-only mode** option.

4. Click **Next**. The **Add to Group** interface appears.
5. Select the SharePoint components to add to the High Availability group. Complete the steps below to add desired nodes into the group:

   a. Click the farm node to load the tree. You can also enter a keyword into the text box on top of the farm tree. Click the magnifying glass) to search the nodes according to the keyword.

   b. Select the desired nodes and then click **Add to Group** >>). You can click the **Actions** button to choose **Show selected nodes** or **Show all nodes** from the drop-down list to customize the farm tree display.

   Refer to the table below for the storage status indicated by the icons:

Type	Database
In LUN or on SMB 3 Share with SnapMirror enabled	
In LUN or on SMB 3 Share without configuring SnapMirror	
The database is not available	
Invalid. No database or index in LUN, or no database on SMB 3 Share	
No Agent is installed, the database cannot be added to group	

***Note**: The database statuses listed in the table above are only visible when the **Sync Method** of the High Availability group is **SnapMirror**. In the Data ONTAP cluster mode storage system, to successfully create a snapshot for the database in LUN, you must assign the aggregate to Vserver.

***Note**: Optionally, you can add any desired database in SQL Servers to the farm component tree and then add them to the group to achieve High Availability. For more instructions, refer to Adding Custom Databases to a Group.

   c. If you are selecting databases, you must enter the **Destination SQL instance** for the data synchronization. In the **Add to Group** window, enter the destination SQL instance name into the **Destination SQL Instance** text box. All of the volumes that store the selected components' data are listed below the **Destination SQL Instance** field. Complete the steps below to configure the standby volume settings for the production volume:

i. **Standby Volume** – All of the available SnapMirror Secondary volumes for this production volume are displayed in the drop-down list. Select a standby volume from the drop-down list.

ii. **Storage System Profile** – Select a storage system profile that is configured for connecting to a standby storage system where the selected standby volume resides.

iii. **Mount Point Option** – Choose to automatically assign a drive letter for mounting LUN to the standby storage system or designate a mount point path.

If you are including Storage Manager BLOB data or Connector BLOB data into this group, you must configure standby volume settings for the BLOB data in the production farm. Complete the steps below:

a. View the production BLOB information displayed in the **Production BLOB Information** field.

b. **Standby Volume** – All of the available SnapMirror Secondary volumes for this production volume are displayed in the drop-down list. Select a standby volume from the drop-down list.

c. **Storage System Profile** – Select a storage system profile that is configured for connecting the standby storage system where the selected standby volume resides.

6. Click **Test** to test the configurations. After successfully testing the configurations, click **Save** to return the **Add to Group** interface. You can click **Edit** next to the database in the right pane to configure the settings for the standby database.
7. Click **Next** to enter the **Advanced Settings** step.
8. **SQL Alias Settings** – If you set **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 enter a port number for the corresponding SQL Server instance.
    b. In the pop-up window, enter a whole number 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.
9. **SharePoint Server Settings** – Select the SharePoint servers from the table, and the SQL Alias will be created on your selected SharePoint Servers separately when running a failover job.

    ***Note**: It is recommended to select 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.
10. **Custom Action** – Select a command profile to run the scripts you selected. Select a previously created command profile from the drop-down list, and click the **View** link to view the detailed settings of the selected command profile. Also, you can create a new command profile by

selecting **New Command Profile** from the drop-down list. For details on custom actions, refer to About Custom Actions.

11. **Schedule Selection** – Define whether or not to configure a custom schedule to run the synchronization jobs.

    ***Note**: It is recommended that you configure a synchronization schedule to run a Synchronization job on the group using the SnapMirror sync method in order to keep databases in sync. The synchronization job schedule will be disabled after Failover is performed, and the schedule will then be enabled after any other type of job is run.

    - **No schedule** – If you select this option, you must manually run the synchronization job each time.
    - **Configure the schedule myself** –Click **Add Schedule** to configure custom schedule settings for the synchronization job. The **Add Schedule** interface appears. Select the type of recurring synchronization job. The two options are **Initialize/Reinitialize** and **Synchronize**.
        - **Initialize/Reinitialize** – Synchronizes all of the selected data from the production environment to the standby environment. Configure the start time in the **Range of Recurrence** field. Click **Save**; this schedule will display in the table. Click **Cancel** to go back to the **Advanced Settings** page without saving anything.
        - **Synchronize** – Synchronizes data changes between the production database and standby database. Enter a whole number 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. Select the time range of the job occurrence in the **Range of Recurrence** field. Designate a start time that cannot be 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 manually stop it.
            - **End after _ occurrences** – Select this option to have the synchronization job stop after the defined occurrence.
            - **End by** – Select this option to define the exact date and time to stop running the job. If desired, you can change the time zone by clicking the time zone link and changing it to the desired one.
12. **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 on creating a new notification profile, refer to the *SnapManager for SharePoint Control Panel Reference Guide*. Click **View** to view the information of the notification profile.
13. Click **Next**. The **Pre-Scan** interface appears.
14. Click **Scan** to scan the environment for the requirements that need to be met before running a synchronization job. You can skip this step, save the group configurations, and then run the Pre-

Scan job for this group in **Dashboard**, and it is recommended that you run the Pre-Scan job from the Dashboard after saving the configurations to the group. For details on the SnapMirror sync method pre-scan, refer to SnapMirror Pre-Scan Rules.

15. Click **Next**. The **Overview** interface appears. If desired, click **Edit** in the field that you want to modify.
16. Click **Finish** after you have confirmed your settings. Click **Cancel** to go back to the Group Manager interface without saving any changes.

Creating a Standby Farm Mode High Availability Group Using Wizard Mode

Complete the steps below to create a standby farm mode High Availability group using Wizard Mode:

1. In the **Group Manager** tab, click **Create** on the ribbon and the select **Wizard Mode** from the drop-down list.
2. Enter a group name and an optional description for the group you are creating. Click **Next**.
3. In the **General Settings** page, configure the following settings:
    - **High Availability Mode** – Choose the High Availability Mode. For this process, select **Standby farm** option.
    - **Production Farm** – Select the production farm from the drop-down list.
    - **Production Agent Group** – Select a production Agent group from the drop-down list to execute the job. An Agent group can utilize multiple Agents for load balancing and job performance improvements.
    - **Standby Farm** – Select the Standby farm from the drop-down list.
    - **Standby Agent Group** – Select a standby Agent group from the drop-down list to execute the job. An Agent group can utilize multiple Agents for load balancing and job performance improvements.
    - **Sync Method** – Select **SnapMirror** from the sync method drop-down list.
    - **Would you like to make the standby environment read-only?** – If you want to make the standby environment read-only, select the **Enable read-only view after synchronization** option. Consider the following issues when configuring this option:

      The service applications in the group with **Enable read-only view after synchronization** option selected can only be synchronized to standby environment once, since the service application databases synchronized to the standby environment in this case is normal. If you want to perform the Synchronization for the service applications in this group again, you must manually delete the service application from the standby farm.

      If you have performed a synchronization job on the group with the **Enable read-only view after synchronization** option selected, you cannot deselect this option to perform later Synchronization jobs. Otherwise, the Synchronization jobs will fail.

- **Would you like to include Storage Optimization data?** – When one or both of these options are selected, a High Availability Synchronization job synchronizes Storage Manager BLOB data and Connector BLOB data to the SnapMirror secondary storage. If you select the **Connector BLOB Data** option, you are required to select an Agent to host the SQLite database that records the production device information. The SQLite database will be updated in each synchronization job, and will be used when performing Failover.

  ***Note**: Only when the **Storage Manager BLOB Data** or **Connector BLOB Data** option is selected, can the associated stub databases of the content databases be displayed in the tree.

  - **Storage Manager BLOB Data** – Select the **Storage Manager BLOB Data** option to include the Storage Manager BLOB data into the synchronization. The **Automatically enable EBS/RBS during Failover** option appears and it is selected by default.
  - **Connector BLOB Data** – Select the **Connector BLOB Data** option to include the Connector BLOB Data into the synchronization. If you select the **Connector BLOB Data** option, you are required to select an Agent to host the SQLite database that records the production device information. The SQLite database will be updated in each synchronization job, and will be used when performing Failover.
  - **Automatically enable EBS/RBS during Failover** – Choose whether or not to automatically enable EBS/RBS in the standby environment during Failover. With this option selected, if the production farm has EBS enabled, the EBS settings will be enabled in the standby farm during Failover. If the production content database has RBS enabled, the RBS settings will be enabled for the standby content database during Failover; however, the standby content database with RBS enabled will not be kept in read-only mode.

- **Failover Options** – Select your desired failover option. For a read-only standby environment, select the **Keep database in read-only mode** option.

4. Click **Next**. The **Add to Group** page appears.
5. Expand the farm components tree, and select the nodes as desired. Click the Add to Group (>>) button to add the selected nodes to the right pane. For details on which SharePoint components are supported in Single farm mode or Standby mode High Availability, refer to Appendix A: SharePoint Components Supported for High Availability.

   ***Note**: SMSP 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.

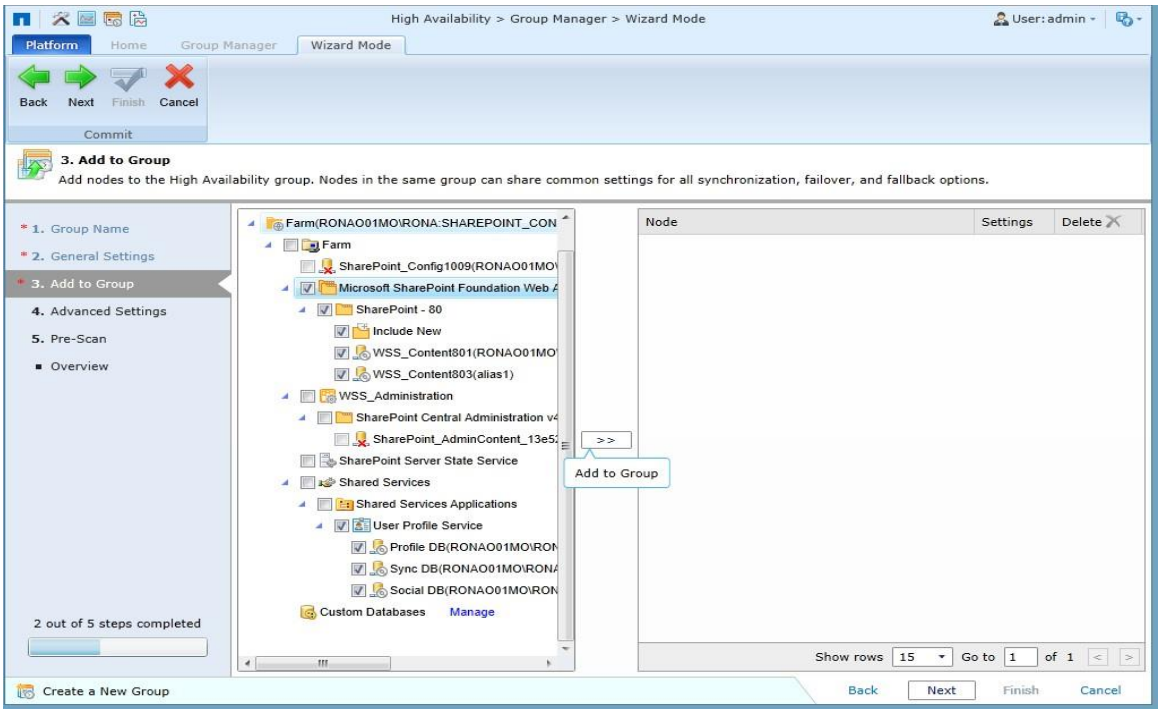

Figure 55: Adding the components to the group.

6. An interface appears for configuring the standby SQL instance, standby volume, mount point path.

    a. Enter the standby SQL instance name into the text box in the **Destination SQLInstance** field.

    b. View the production volume information displayed in a table. Complete the steps below to configure the standby volume settings for the production volume:

        i. **Standby Volume** – All of the available SnapMirror Secondary volumes of this production volume are displayed in the drop-down list. Select a standby volume from the drop-down list.

        ii. **Storage System Profile** – Select a storage system profile that is configured for connecting the standby storage system where the selected standby volume resides.

        iii. **Mount Point Option** – Choose to automatically assign a drive letter for mounting LUN to the standby storage system or designate a mount point path.

    If you are including the Storage Manager BLOB data or Connector BLOB data into this group, you must configure standby volume settings for the BLOB data in the production farm. Complete the steps below:

a. View the production BLOB information displayed in a table of **Production BLOB Information** field.

b. **Standby Volume** – All of the available SnapMirror Secondary volumes of this production volume are displayed in the drop-down list. Select a standby volume from the drop-down list.

c. **Storage System Profile** – Select a storage system profile that is configured for connecting the standby storage system where the selected standby volume resides.

7. Click **Test** to test the configurations in the window. After successfully tested the configurations, click **Save** to return the **Add to Group** interface. You can click **Edit** next to the database in the right pane to configure the settings for the standby database.
8. Click **Next** in the **Add to Group** page. The **Advanced Settings** page appears.
9. **Custom Action** – Select a command profile to run the scripts you selected. Select a previously created command profile from the drop-down list, and click the **View** link to view the detailed settings of the selected command profile. Also, you can create a new command profile by selecting **New Command Profile** from the drop-down list. For details on custom actions, refer to About Custom Actions.
10. **Schedule Selection** – Define whether or not to configure a custom schedule to run the synchronization jobs.

    ***Note**: It is recommended that you configure a synchronization schedule to run the Synchronization job on the group with SnapMirror sync method in order to keep database in sync. The synchronization job schedule will be disabled after Failover is performed and the schedule will then be enabled after any other type of job is run.

    - **No schedule** – If you select this option, you are required to manually run the synchronization jobs.
    - **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** – Synchronizes all of the selected data from the production environment to the standby environment. Designate 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** – Synchronizes data changes between production database and standby database. Enter a whole number 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. Select the time range of the job occurrence in the **Range of Recurrence** field. Designate a start time that cannot be 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 manually stop it.
- **End after _ occurrences** – Select this option to have synchronization job stop after the defined occurrence.
- **End by** – Select this option to define the exact date and time to stop running the job. If desired, you can change the time zone by clicking the current time zone link to select the desired one.

11. **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 *SnapManager for SharePoint Control Panel Reference Guide*. Click **View** to view the information of the notification profile.
12. Click **Next**. The **Pre-Scan** interface appears.
13. Click **Next** to skip scanning the environment for the requirements that need to be met before running a synchronization job. For details on the SnapMirror pre-scan rules, refer to SnapMirror Pre-Scan Rules.

    ***Note**: It is not recommended to run Pre-Scan job within the group creation wizard, since the job will take a while to complete.

14. In the **Overview** interface, you can click **Edit** in the field that you want to modify.
15. Click **Finish** to exit the wizard, and save the group. The group will be displayed in the **Group Manager** interface. A message bar will appear informing you that the group is successfully created. You can click the **Dashboard** link on the message bar, or navigate to **Home** > **Dashboard** to access the **Dashboard** interface to perform the Pre-Scan and Synchronization job using the created group.

Managing Groups

Use Group Manager to manage High Availability groups. After launching the High Availability interface, click **Group Manager** next to the **Home** tab. In the **Group Manger** 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**.

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. If the group includes the service applications with the **Enable read-only view after synchronization** option selected, and a Synchronization job has been run using this group, the settings for the service applications in this group cannot be changed; if this group does not have the **Enable read-only view after synchronization** option selected, the settings of the service applications will not be allowed to change after Failover. You can refer to the message under the **Last Action Status** column in **Dashboard** to tell whether the group has been run. If the group is currently being used to run a High Availability job, this group cannot be edited.
- **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. If the group is currently being used to run a High Availability job, this group cannot be deleted.
- **Storage System Profile** – Click **Storage System Profile** to go to the **Storage Configuration** interface of Control Panel to manage or create the storage system profile. User can use a predefined storage system profile to access the storage system. For details on managing or configuring storage system profile, refer to the *Storage Configuration* section in *SnapManager for SharePoint Control Panel Reference Guide*.
- **Custom Action** – Integrated with script profiles and command profiles, custom action allows you to run the predefined scripts to execute the customized actions before or after Synchronization, Failover, or Fallback. For details on managing or configuring script profiles and command profiles, refer to About Custom Actions.
- **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, refer to Job Monitor User Guide.

Actions in the Dashboard

The Dashboard interface provides a detailed view of the production farms, standby databases, High Availability Groups, and High Availability jobs. You can perform High Availability jobs from the ribbon, such as Pre-Scan, Synchronization, Failover, Fallback, Maintenance Failover, and Maintenance 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 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: **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 determine if 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 select multiple groups together or select some components from the group to perform a 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 multiple processes lead to the failure of jobs.

- **Pre-Scan** – It is recommended you perform a Pre-Scan job before synchronization. Select the checkbox beside the desired nodes, and then click **Pre-Scan** on the ribbon to check if the group can perform a synchronization job. A pop-up window appears, asking for confirmation. Click **OK** to run a Pre-Scan job for the selected nodes. Click **Cancel** to leave this pop-up window without starting the Pre-Scan job. For more information on the Pre-Scan rules, refer to SnapMirror Pre-Scan Rules.
- **Synchronization** – Synchronize the data from the production environment to the disaster recovery environment. Select the checkbox beside the desired nodes, and then click **Synchronization** on the ribbon to synchronize data between the production server and standby server. For the Synchronization job with SnapMirror sync method, the timeout duration for updating SnapMirror relationship is, by default, 30 minutes. If the result is not obtained during this time period, the Synchronization job will fail or finish with exceptions. You can customize the timeout session duration by modifying a configuration file. For details, refer to Customizing the Timeout Session for Getting the Result of Update-SnapMirror or Resync-SnapMirror.
- **Failover** – Switch to a standby disaster recovery environment upon the failure or abnormal termination of the production environment. Select the checkbox beside the desired nodes, and then click **Failover** on the ribbon to enable the standby server. In the **Failover** window, you can change the standby volume and reconfigure the storage system profile. Since you selected groups that use the SnapMirror sync method to run the Failover, High Availability will failover the entire production volume to the standby

volume, regardless of whether or not the databases in the production volume are included in this group.

***Note**: It is recommended that you refresh 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 beside the desired nodes, and then you can click **Fallback** to synchronize the data from the standby server to the production server. For the Fallback job with SnapMirror sync method, the timeout duration for resynchronizing SnapMirror relationship is, by default, 30 minutes. If the result is not obtained in this time period, the Fallback job will fail or finish with exceptions. You can customize the timeout session by modifying the configuration file. For details, refer to Customizing the Timeout Session for Getting the Result of Update-SnapMirror or Resync-SnapMirror.

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

For the purposes of upgrading and maintaining production and standby farms, High Availability provides **Maintenance Failover** and **Maintenance Fallback** for Standby farm High Availability mode. To perform Maintenance Failover or Maintenance Fallback, all of the Standby farm mode High Availability groups created for the same production farm must be selected. For detailed instructions on Maintenance Failover and Maintenance Fallback, refer to the Maintenance Failover, Maintenance Fallback, and Building up a Read-Only or Read-Write Standby Farm for the Maintenance of the Production Farm.

Distinguishing Failover-Ready Databases on Dashboard

In High Availability Dashboard, all of the groups and components will be displayed with their last action status and time. Through the Dashboard, you can see which High Availability operation is performed on the database and when, and if the database is ready to perform a Failover. For details on how to distinguish the failover-ready databases in the High Availability group using SnapMirror sync method, refer to the table below.

Warm Standby Database Status		Cold Standby Database Status	Last Action Status	Comment
Read-Only (content database, custom database, and stub database)	Normal (service application database)	Does Not Exist	Synchronization Succeeded	Failover ready
Does Not Exist		Does Not Exist	Synchronization Failed	Does not recommend performing failover. It might cause data loss.

SnapMirror Pre-Scan Rules

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

- Check if the database is in LUN or CIFS Share – Checks if the database included in the High Availability group is in a LUN or CIFS Share.
- Check the status of the standby database for SnapMirror – Checks if the status of the standby database is read-only or if the standby database does not exist.
- Check the status of the production database for SnapMirror – Checks if the status of the production database is normal or **Principal**, or if the production database is included in AlwaysOn Availability Group.
- Check if SnapDrive is installed – Checks if SnapDrive has already been installed with SQL Server and DocAve Agent.
- Check if the SnapMirror relationship exists – Checks if the SnapMirror relationship exists, and checks if the status is **SnapMirrored**.
- Check the sync mode of the SnapMirror relationship – Checks if the sync mode of SnapMirror relationship is **Async** mode.

Customizing the Timeout Session for Getting Update-SnapMirror or Resync-SnapMirror

Complete the steps below to customize the timeout session for updating SnapMirror or resynchronizing SnapMirror. Note that you will have to perform the procedure below on each Agent server in the Production Agent Group and each Agent server in the Standby Agent Group.

1. Locate the **NetApp** or **IBM** folder, and navigate to the *SMSP8\Agent\data\HighAvailability* directory.
2. Open the **AgentCommonHAConfiguration.xml** file with Notepad.
3. Find the **<SnapMirror UpdateTimeOutSeconds="1800" ResyncTimeOut="1800"/>** node.
    - **UpdateTimeOutSeconds** – By default, the value is **1800**, representing **1800 Seconds**, (**30 Minutes**). You can change the value of this attribute, but note that the timeout session cannot be shorter than 30 seconds.
    - **ResyncTimeOut** – By default, the value is **1800**, representing **1800 Seconds** (**30 Minutes**). You can change the value of this attribute, but note that the timeout session cannot be shorter than 30 seconds.
4. Save the configuration file, and close it.

Customizing Where to Execute Filer SDK Toolkit

By default, the Filer SDK Toolkit commands are executed on the Media server at random. You can configure the **AgentCommonHAConfiguration.xml** file to execute the SDK Toolkit commands on a Media server or on the Agent server that browses the farm component tree and executes the High Availability

jobs. Note that you will have to perform the procedure below on each Agent server. Complete the steps below to customize where to execute the Filer SDK Toolkit commands:

1. Locate the **NetApp** or **IBM** folder, and navigate to the *SMSP8\Agent\data\HighAvailability* directory.
2. Open the **AgentCommonHAConfiguration.xml** file with Notepad.
3. Find the **<ServiceHostConfig PreferredMediaServer="" UseMediaHost="true"/>** node.
    - **UseMediaHost** – Customize whether or not to run Filer SDK Toolkit commands on Media server.
        - **True** means to run the SDK Toolkit commands on Media server. By default, the Media server where to run the commands is picked randomly. You can enter a Media server name as the value of **PreferredMediaServer** attribute to run commands as preferred.
        - **False** means to run the SDK Toolkit commands on the Agent server that is used to browse the farm component tree and execute the High Availability jobs.
4. Save the configuration file, and close it.

Connecting LUN Devices with FCP

Complete the steps below to modify the **AgentCommonHAConfiguration.xml** file for using Fibre Channel Protocol to connect the LUN devices. Note that you will have to perform the procedure below on each SQL Server Agent.

1. Locate the **NetApp** or **IBM** folder, and navigate to the *SMSP8\Agent\data\HighAvailability* directory.
2. Open the **AgentCommonHAConfiguration.xml** file with Notepad.
3. Find the **<LUNConnectionConfig PreferredInitiatorName="" UseFibreChannel="false"/>** node.
    - **UseFibreChannel** – Choose whether or not to use FCP to connect LUN devices. The default value is **False**, which means not using FCP to connect the LUN devices. To use FCP to connect LUN devices, set the value of this attribute to **True**. You can enter an initiator ID as the value of **PreferredInitiatorName** attribute to use the defined initiator as preferred. If you do not assign a preferred initiator or the initiator ID is incorrect, the system initiator will be used.
4. Save the configuration file, and close it.

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 © 2013-2016 AvePoint, Inc. All rights reserved. All materials contained in this publication are protected by United States and international copyright laws 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. Notwithstanding the foregoing, to the extent any AvePoint material in this publication is reproduced or modified in any way (including derivative works and transformative works), by you or on your behalf, then such reproduced or modified materials shall be automatically assigned to AvePoint without any further act and you agree on behalf of yourself and your successors, assigns, heirs, beneficiaries, and executors, to promptly do all things and sign all documents to confirm the transfer of such reproduced or modified materials to AvePoint.

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, Office, Office 365, SharePoint, Windows PowerShell, SQL Server, Outlook, Windows Server, Active Directory, and Dynamics CRM 2013 are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

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 without 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