



DocAve® 6.0.1 Storage Manager

User Guide

Revision B
Issued April 2012

Table of Contents

Table of Contents.....	2
About DocAve Storage Manager	4
Complementary Products	4
Submitting Documentation Feedback to AvePoint	4
Before You Begin.....	5
Configuration	5
Agents	5
Licensing and Permissions	5
Local System Permissions	6
Getting Started.....	7
Launching Storage Manager	7
User Interface Overview	8
Selecting Farms and Nodes.....	9
Basic Steps to Configuring and Using Storage Manager.....	10
Configuring the BLOB Provider	11
EBS and RBS Providers: Differences and Recommendations.....	11
Enabling the BLOB Provider	12
Enabling the BLOB Provider Using the Tool.....	15
Configuring Logical Devices	16
Configuring Storage Manager General Settings	17
Configuring the Orphan BLOB Retention.....	17
Configuring the Processing Pool (Scheduled Storage Manager Only).....	18
Creating Storage Manager Rules	19
Creating Realtime Storage Manager Rules	19
Creating Scheduled Storage Manager Rules.....	21
Advanced Settings.....	23
Managing Created Rules.....	24
Inheriting and Stop Inheriting.....	24
Editing or Removing Storage Manager Rules	25

Editing Existing Storage Manager Rules.....	25
Removing Existing Rules from Nodes	26
Enabling and Disabling Rules	26
Disabling an Existing Rule.....	26
Enabling a Disabled Rule	27
Deleting Rules	27
Using Rule Viewer	28
Rule Based Viewing.....	28
Deleting Rules	28
Scope Based Viewing	29
Removing Rules.....	29
Converting Stubs to Content.....	30
How to Determine if the Data is Stub or Real Content.....	30
Customizing the Number of Storage Manager Job Threads.....	31
Checking the Job Status	31
Collecting Logs for Troubleshooting	32
DocAve Storage Manager Use Cases	33
DocAve Storage Optimization Suite in Conjunction with Tiered Storage.....	33
DocAve Storage Optimization and Data Protection Suites in Compliance Environments.....	36
Appendix A: Stub Database Inheritance	38
Appendix B: Coexistence of EBS and RBS	40
Index.....	43
Notices and Copyright Information	44

About DocAve Storage Manager

As SharePoint 2010 becomes the central repository for enterprise content within organizations, increasing user adoption and integration from data sources can cause unstructured data, called Binary Large Objects (BLOBs), to consume up to 95 percent of SharePoint SQL Server space. If ignored, BLOBs can lead to database performance and user experience degradation.

With DocAve Storage Manager, organizations can mitigate the negative consequences of exponential data growth by combining multiple real-time and scheduled business rules to externalize BLOB content based on file size, type, or other document properties, allowing SharePoint administrators to:

- Reduce SharePoint total cost of ownership by utilizing hierarchical storage management (HSM) systems for BLOB storage
- Improve user experience by optimizing SQL Server resources while maintaining seamless user access and interaction to externalized content
- Easily comply with information governance policies for content lifecycle management within or across multiple SharePoint farms with highly-customizable business rules
- Ensure all BLOBs work with SharePoint management functions and most third-party applications by leveraging Microsoft's fully supported External BLOB Storage (EBS) or Remote BLOB Store (RBS) APIs

Complementary Products

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

- DocAve Content Manager for SharePoint for restructuring or moving SharePoint content
- DocAve Replicator for SharePoint for copying SharePoint content within the same SharePoint farm or from one SharePoint farm to another. DocAve Replicator performs live, event-driven, scheduled, or offline replication. Synchronization and management of all content, configurations, and securities is performed with full fidelity.
- DocAve Report Center for SharePoint to examine pain points in the SharePoint infrastructure and report on SharePoint user behavior and changes
- DocAve Data Protection for setting backup and recovery points prior to adjusting SharePoint governance policies in this product

Submitting Documentation Feedback to AvePoint

AvePoint encourages customers to provide feedback regarding our product documentation. Click the following URL to access the **Submit Your Feedback** form on our Web site:

<http://www.avepoint.com/resources/documentation-feedback/?flush=1>

Before You Begin

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

Configuration

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

Agents

DocAve Agents are responsible for running DocAve jobs and interacting with the SharePoint object model. Agents enable DocAve Manager to communicate with the respective servers, allowing for Storage Manager to function properly.

Agent installation location depends upon the type of BLOB provider being used.

- RBS environment:
 - DocAve Agents must be installed on all of the front-end Web servers.
 - DocAve Agent must be installed on the server that installs the Office Web App service (Office Web App service includes Word Viewing Service Application, PowerPoint Service Application and Excel Calculation Services).
 - DocAve Agent must be installed on the server that installs the SQL Reporting Service.
- EBS environment:
 - DocAve Agents must be installed on the SharePoint Central Administration server and all of the front-end Web servers.
 - DocAve Agent must be installed on the server that installs the Office Web App service (Office Web App service includes Word Viewing Service Application, PowerPoint Service Application and Excel Calculation Services).
 - DocAve Agent must be installed on the server that installs the SQL Reporting Service.

***Note:** The use of system resources on a server increases when the installed agent is performing actions. This may affect server performance. However, if the agent installed on a server is not being used, the use of system resources is very low and, therefore, the effect on server performance is negligible.

For instructions on installing the DocAve Platform (DocAve Manager and DocAve Agents), see the [DocAve 6 Installation Guide](#).

Licensing and Permissions

The following permissions are required for the Storage Manager agent account; they ensure proper functionality of Storage Manager.

1. Local System Permissions: These permissions are automatically configured by DocAve during installation. Refer to [Local System Permissions](#) for a list of the permissions automatically configured upon installation. If there are no strict limitations within your organization on the permissions that can be applied, you can simply add the **DocAve Agent Account** to the local **Administrators** group to apply all of the required permissions.
2. SharePoint Permissions
 - User is a member of the Farm Administrators group. Since Administrator works across farms and on all SharePoint settings and configurations, this account is needed in order to provide the best and most complete quality of service.
 - Full control to all zones of all Web applications via User Policy for Web Applications
3. SQL Permissions
 - Database Role of db_owner for all the databases related with SharePoint, including content databases, stub databases, SharePoint configuration database and Central Admin database.
 - Database Role of dbcreator to SQL Server since DocAve must create a stub database before performing any Storage Manager job.

Local System Permissions

1. User is a member of the following local groups:
 - IIS_WPG (for IIS 6.0) or IIS_IUSRS (for IIS 7.0)
 - Performance Monitor Users
 - DocAve Users(The group is created by DocAve automatically and it has the following permissions)
 - Full control to the Registry of HKEY_LOCAL_MACHINE\SOFTWARE\AvePoint\DocAve6
 - Full control to the Registry of HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\EventLog6
 - Full Control to the Communication Certificate
 - Permission of *Log on as a batch job* (it can be found within Control Panel > Administrative Tools > Local Security Policy > Security Settings > Local Policies > User Rights Assignment)
 - Full Control Permission of DocAve Agent installation directory
2. Full Control to GAC in order to install Provider dll into GAC.
3. Full Control to Microsoft SQL Remote Blob Storage Folder to reconfigure maintainer configuration file.

Getting Started

SharePoint and the DocAve platform modules have common functionality. While some of this shared functionality is covered in this guide, the primary focus of this document is the functionality that is specific to the DocAve module.

For information on the shared functionality not covered in this document, refer to SharePoint Help.

Refer to the sections below for important information on getting started with Storage Manager.

Launching Storage Manager

To launch Storage Manager 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 **Storage Optimization** to view the Storage Optimization modules.
3. Click Realtime Storage Manager or Scheduled Storage Manager to launch the appropriate module.

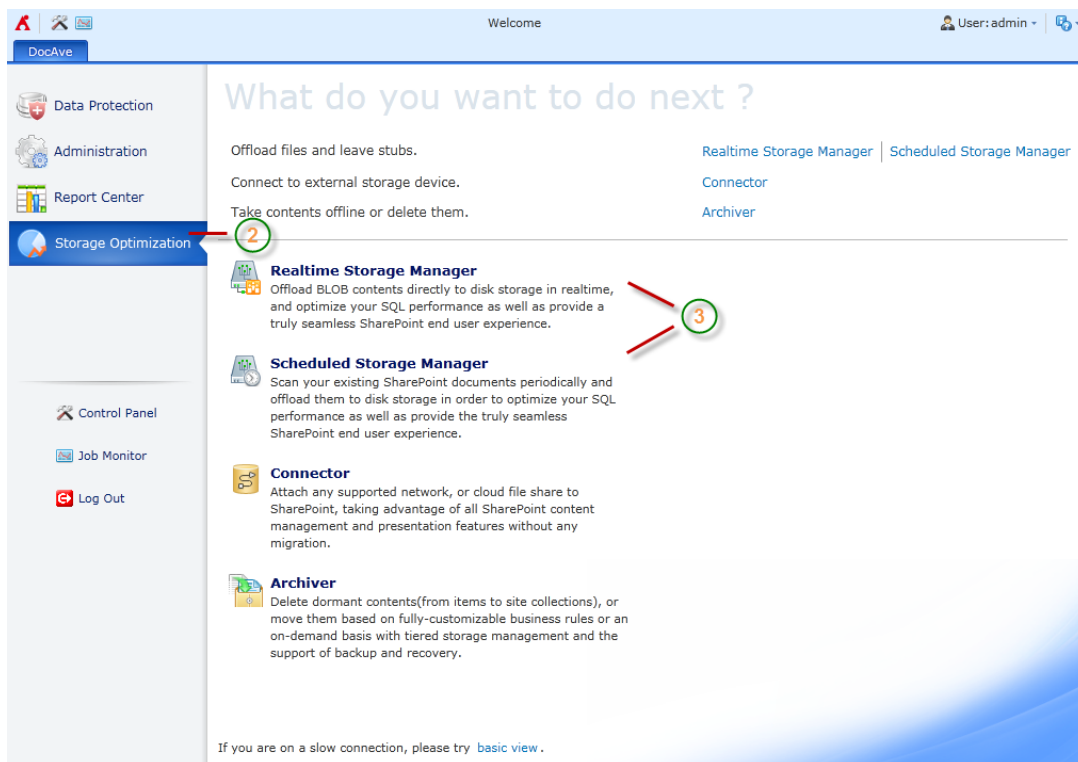


Figure 1: DocAve module launch window.

User Interface Overview

After selecting **Realtime** or **Scheduled Storage Manager**, the Storage Optimization suite user interface launches with the appropriate **Storage Manager** tab active. This tab displays your farm environment and allows for quick access to a list of Storage Manager features.

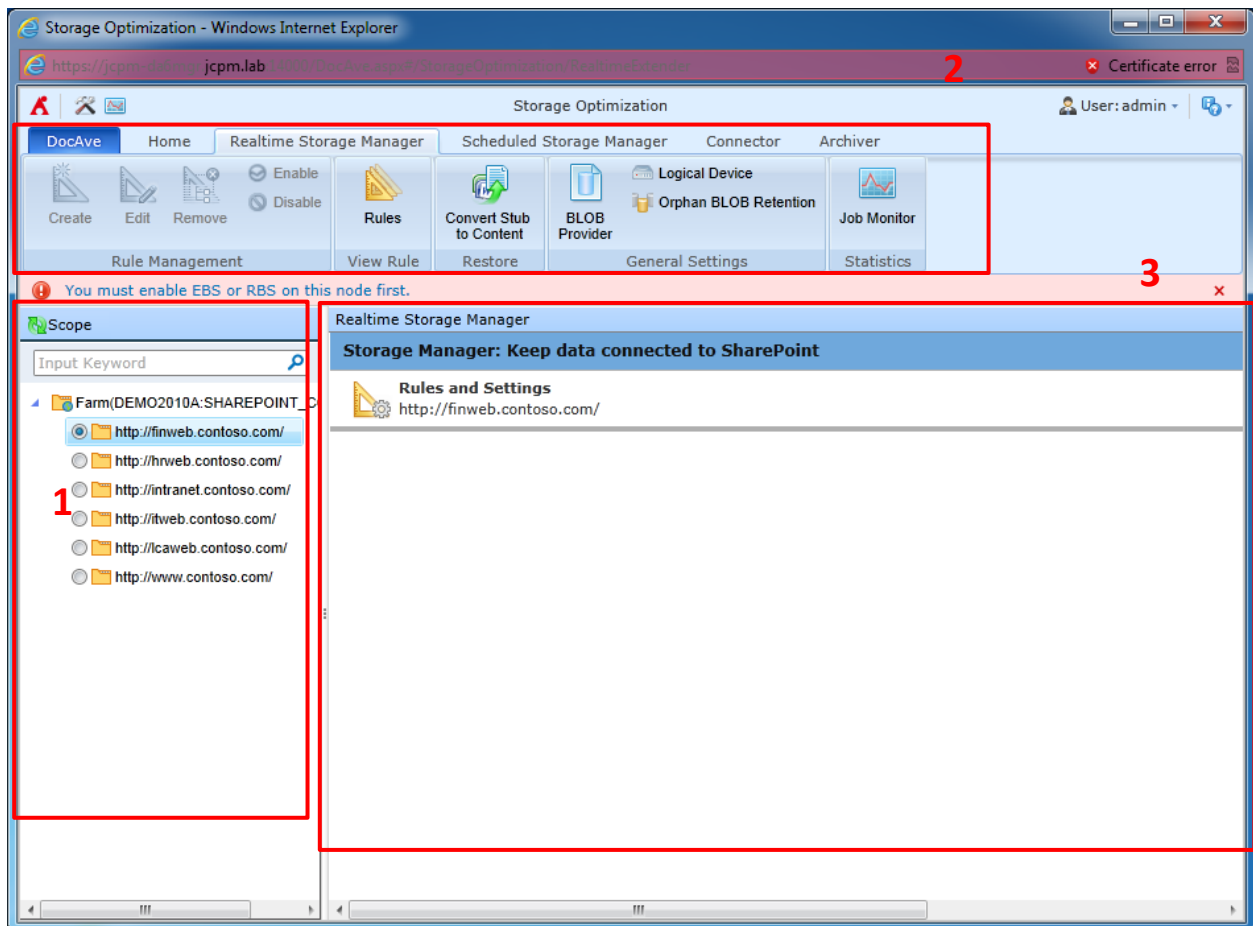


Figure 2: Realtime Storage Manager user interface.

1. The **SharePoint tree (Scope panel)** displays all content within your farm(s). Use this panel to select the content that you wish to perform actions on. Selecting content often reveals new tabs and functionality on the **ribbon**.
2. The **ribbon** shows the available actions and wizards for the selected nodes. This content is dynamic; it will often change depending on what is selected on the SharePoint tree and in the **workspace** area.
3. The **workspace** shows all form-based content that is used during the configuration of actions performed in DocAve products.

Selecting Farms and Nodes

To select farms and nodes:

1. From the **Scope** panel on the left, double-click the farm that contains the relevant SharePoint content.
2. Select the relevant content from which you want to perform further operations by clicking the radio button(s) to the left of the content.
3. After selecting content, you will be able to perform the procedures described throughout this guide.

Basic Steps to Configuring and Using Storage Manager

The following are the basic steps required in order to properly use Storage Manager. Click the link to jump to the corresponding section.

- [Configure and enable the BLOB Provider \(EBS or RBS\).](#)
- [Configure a logical device.](#)
- [Define the Orphan Blob Retention policy and configure the Processing Pool.](#)
- [Create Storage Manager rules.](#)

Configuring the BLOB Provider

In order to use DocAve Storage Manager, the BLOB Provider must be enabled.

A binary large object (BLOB) is a collection of binary data stored as a single entity in a database management system. By default, any file or attachment that is uploaded into SharePoint is stored as a BLOB in the content database. By configuring the BLOB provider, you decide which BLOBs are kept in the database and which are stored externally.

The BLOB Provider feature intercepts SharePoint database traffic and redirects all of the BLOB traffic to the external BLOB storage; what remains in SharePoint is a stub of the data. The two types of providers are External BLOB Store (EBS) provider and Remote BLOB Store (RBS) provider; both are described below.

***Note:** If there is insufficient space on the BLOB storage device, then when uploading a document that meets the criteria for Realtime Storage Manager externalization, the upload fails.

EBS and RBS Providers: Differences and Recommendations

EBS is a COM interface provided by Microsoft SharePoint Server 2007 and 2010, while RBS is a set of standardized Advanced Programming Interfaces (APIs) that are incorporated as an add-on feature pack for Microsoft SQL Server.

***Note:** RBS can be run on the local server running Microsoft SQL Server 2008 R2, SQL Server 2008, SQL Server 2008 R2 Express, or SQL Server 2008 R2 Developer edition. To run RBS on a remote server, you must be running SQL Server 2008 R2 Enterprise edition. RBS is not supported for Microsoft SQL Server 2005.

As mentioned, EBS was an earlier attempt by Microsoft in SharePoint 2007 (available as a hot fix to MOSS 2007 SP1 and later) to help customers externalize their BLOBs. However, EBS was hard to implement and had some limitations. Microsoft introduced EBS as an immediate aid; it was designed in a way that is an evolutionary approach, where administrators can move to RBS later. EBS will continue to be supported for SharePoint 2010, but it is on the deprecation list, which means its support will end in a future release of SharePoint. Microsoft recommends using RBS in SharePoint 2010 not only because EBS is on the deprecation list, but also because RBS is more powerful and maintainable. The table below illustrates some of the advantages of RBS over EBS.

***Note:** If the user does not want to use a third-party vendor, RBS is the only option.

Feature	RBS	EBS
BLOB store scope	RBS can be enabled at content database level and Web application level. Each content	EBS can be enabled only at the farm level.

	database can have its own BLOB store. RBS is more flexible.	
Number of providers	Multiple RBS providers can be in the same SharePoint farm.	Only one EBS provider per SharePoint farm.
Interface	Managed. RBS is a purely .NET-based solution. From a technology perspective, RBS fits in to .NET quite nicely.	Unmanaged. EBS relies on a legacy COM interface.
Migrating BLOBs from SQL Server stores to BLOB stores and vice versa	Windows PowerShell	Custom
SharePoint interface	SharePoint 2010 ships with many Windows PowerShell command lets that can be used to manage RBS installation and configuration.	None

For more information related to BLOB externalization best practices, refer to the following AvePoint whitepaper:

[http://www.avepoint.com/assets/pdf/sharepoint_whitepapers/Optimize SharePoint Storage with BLOB Externalization.pdf](http://www.avepoint.com/assets/pdf/sharepoint_whitepapers/Optimize_SharePoint_Storage_with_BLOB_Externalization.pdf).

Enabling the BLOB Provider

To enable the BLOB Provider on your server(s), follow the instructions below.

***Note:** When only the RBS provider is enabled, the Realtime Storage Manager rule is only available at the Web application level. When only the EBS provider is enabled, the Realtime Storage Manager rule is available at the Web application level and site collection level. When both the RBS provider and the EBS provider are enabled, RBS is used; refer to [Appendix B](#) for more information.

To access the BLOB Provider, click either the **Realtime Storage Manager** or **Scheduled Storage Manager** tab > **BLOB Provider**. The **BLOB Provider** page appears in the workspace.

***Note:** You can also use the **AgentToolSP2010StorageEBS.exe** and **AgentToolSP2010StorageRBS.exe** to enable the corresponding BLOB provider. For more information, refer to [Enabling the BLOB Provider Using the Tool](#).

1. **Install BLOB Provider Binaries** on the specified farm(s). This page displays the information for all of the servers in the specified farm. By default, the BLOB Provider is installed with the Agent

installation. If no BLOB Provider is installed, **Not Installed** is displayed in the **BLOB Provider Binaries** column. If necessary, click **Install** to install the corresponding BLOB Provider. Click **Next** when finished. The **Configure Stub Database** page appears.

2. **Configure Stub Database** – Specify where to store all of the Connector stubs' information. Click the farm name to expand the tree and to select the farm level, Web application level, and/or the content database level.

For ease of use, it is recommended to configure one stub database for the entire farm. In some cases (records management, for example, where there may be millions of files), it may be necessary to configure different stub databases down to the Web application or content database level. For more information on this process, see [Appendix A: Stub Database Inheritance](#).

***Note:** You must ensure that all stub databases are properly backed-up in case the database is corrupted.

3. After selecting which nodes you want to link to a stub database, click **Configure** in the **Manage** group on the ribbon. You can also click the **Configure** button beside each tree node to configure the stub database for the specified node.

- **Configure Stub Database** – Specify the **Database Server** (where the stub database resides) and **Database Name** for the stub database. See [Appendix A: Stub Database Inheritance](#) for additional information on stub database inheritance.

***Note:** Once the stub database of a content database is configured and saved, the configuration cannot be changed.

By default, the database server used by the farm is loaded automatically. Considering the management and maintenance of the stub databases, it is recommended to use the default database server. You must specify other database servers if any one of the following conditions is applicable:

- The account does not have enough permission to connect to the default SQL Server.
- The account does not have enough permission to create the stub database on the default SQL Server.
- The default SQL Server is located on a machine in another location and the network status is "poor".
- **Authentication** – Select the authentication method used to access the database.
 - **Windows authentication** (recommended; the default option) – Use this method when you want the user identity to be confirmed by Windows.
 - **SQL authentication** – SQL Server confirms the user identity according to the specified account and password. The specified account must be added to the **sysadmin** role in SQL Server.
- Click **OK** to proceed to the **Enable BLOB Provider** page.

4. **Enable BLOB Provider** – Shows the **BLOB Provider Status** of the farms that have DocAve agents installed. Enable or disable the BLOB Provider on the specified farm in this step. Click **Configure** in the **Action** column; the following options appear.

- **Enable RBS for Farm (farm name)** – Displays the component of the farm and the schedule of enabling RBS.
 - **Tree Structure** – Click the farm name to expand the tree; you can view the RBS status and enable the RBS on the specified content database by selecting the check-box in the **Enable** column, if necessary. The RBS can be enabled at the Web application level and content database level. If you enable RBS at the Web application level, RBS for both the existing and newly-added content databases under the specified Web application is enabled.
 - **Include New Content Databases** – Selecting the **Enable** checkbox at this level enables RBS for all the newly-added content databases; RBS will be automatically enabled by the scheduled jobs. In order to use this function, a schedule must be configured in the **Schedule** field. Note that in order to enable this selection, you must have a stub database configured at the Web application level, and you must select **Include new content databases** when [configuring the BLOB Provider](#).
 - **Schedule** (available when **Enable RBS for Farm** is selected) – This schedule checks for newly-added content databases; the purpose is to enable RBS. Choose whether to enable RBS based on a schedule. Storage Manager runs a search on your farm for new content databases according to your specified schedule. If any new content databases are found, Storage Manager enables the RBS for them.
 - **No schedule** – Enables the RBS immediately.
 - **Configure the schedule myself** – Configure a schedule and enable RBS according to the schedule. Select a **Start time** and **Interval** value.

Note the following:

- If the parent node does not have a stub database, the stub database of the grandparent node will be applied to the newly-added content database.
 - The stub database rule will be applied to the newly-added content database immediately after it is added to the corresponding Web application. It is not related to the schedule specified here.
- **Enable EBS for Farm (farm name)** – Shows the EBS status of the farm and specify whether to enable the EBS. Select **Enable** to enable EBS and select **Disable** to disable it. **If your environment is SharePoint 2010, it is recommended that you enable RBS for the farm.**

***Note:** If both EBS and RBS are enabled for a SharePoint object, RBS will be used first. For more information, refer to [Appendix B](#).

5. **Overview** – Shows all the information of the farms, you can click **Edit** to edit the configuration. Click **Finish** to enable/disable the EBS Provider immediately or save the configuration of the RBS provider without enabling/disabling it. Click **Finish and Run Now** to save the configuration and then enable/disable the corresponding BLOB Provider immediately.

***Note:** For enabling RBS provider without using a schedule, you must click **Finish and Run Now** to enable it immediately. If you click **Finish**, only the configuration of the RBS provider is saved, but the RBS BLOB Provider will not be enabled.

Enabling the BLOB Provider Using the Tool

This section describes the steps needed to enable EBS or RBS using the .exe tool. Refer to [EBS and RBS Providers: Differences and Recommendations](#) for more information on these providers.

Enabling EBS Using the Tool

1. Access the installation path of DocAve Agent. The default path is ...*\AvePoint\DocAve6\Agent\bin*.
2. Locate AgentToolSP2010StorageEBS.exe, right-click on it, and select Run as administrator.
3. The following buttons are shown in the tool.
 - In the Check EBS Status field:
 - **Check EBS** – Checks whether EBS is enabled on the farm.
 - **Enable EBS** – Enables EBS on the farm. You can choose to restart the IIS now or later in the pop-up.
 - **Disable EBS** – Disables EBS on the farm. You can choose to restart the IIS now or later in the pop-up.
 - In the Install Blob Com field:
 - **Install** – Installs BLOB COM on the farm. You can choose to restart the IIS now or later in the pop-up.
 - **Uninstall** – Uninstalls BLOB COM from the farm. You can choose to restart the IIS now or later in the pop-up.
 - In the Check whether the Blob Com has been installed correctly field:
 - **Check** – Checks the status of the items listed in the left field.

Enabling RBS Using the Tool

1. Access the installation path of DocAve Agent. The default path is ... \AvePoint\DocAve6\Agent\bin.
2. Locate AgentToolSP2010StorageRBS.exe, right-click on it, and select Run as administrator.
3. The following buttons are shown in the tool.
 - In the **Remote Blob Storage Installation Status** field:
 - **Check** – Checks the installation status of RBS in this farm.
 - **Install** – Installs RBS on the farm.
 - **Uninstall** – Uninstalls RBS from the farm.
 - After verifying the RBS installation status using the three options above, you can perform the following actions in the **Remote Blob Storage Enable Status** field:
 - **Browse** – Generates/Refreshes the tree structure of the farm. The tree is detailed to the content database level.
 - **Check** – Select some SharePoint components on the tree and click this button to check whether RBS is enabled for the selected components.
 - **Enable** – Enables RBS on the specified components.
 - **Disable** – Disables RBS on the specified components.

Configuring Logical Devices

Because Storage Manager stores the uploaded BLOB content to an external device (leaving only a stub of the data), you must first configure one or more physical device(s) and then set up a logical device. Storage Manager can write to any Net Share, EMC Centera, or Cloud Storage device, depending on the external location where the BLOB content is stored.

In addition, DocAve has the ability to treat multiple storage devices (physical devices) as a single logical unit (logical device) when saving backup data. This is especially useful for very large Storage Manager scopes, as many small drives can be combined. A logical device must be defined before creating an archive Storage Manager plan.

***Note:** Instead of using a Storage Policy, DocAve Storage Manager uses the logical device to ensure the stubs can be accessed properly from SharePoint, regardless of the DocAve Manager service status.

For instructions on defining physical devices and setting up logical devices, refer to the [DocAve 6 Control Panel Reference Guide](#).

Configuring Storage Manager General Settings

Before creating Realtime or Scheduled Storage Manager rules, define the orphan BLOB retention policy (optional but recommended) and configure the processing pool (mandatory for Scheduled Storage Manager).

Configuring the Orphan BLOB Retention

If you decide to remove the Storage Manager stubs from SharePoint, the BLOB content remains in the external storage by default. Configure the **Orphan BLOB Retention** function to have DocAve remove orphan BLOB content in external storage after the corresponding stubs are removed from SharePoint.

Though this setting is optional, it is recommended that you use this function to better optimize the external storage. It is important to note that if there is insufficient space on the BLOB storage device, then when uploading a document that meets the criteria for Realtime Storage Manager externalization, the upload fails.

Complete the steps below to configure orphan BLOB retention rules.

1. Click the Realtime Storage Manager tab or Scheduled Storage Manager tab.
2. In the General Settings group, click Orphan BLOB Retention.
3. Click the farm name to expand the tree and select the Web application(s) where you want to configure the BLOB retention rules. If you select the entire farm node, the BLOB retention rule you configure will be applied to all the existing Web applications in the specified farm.

***Note:** The BLOB retention rule configured for the farm node is not inherited by a Web application that is created after the rule is applied. In this case, configure the BLOB retention rule for the newly-created Web application separately.

4. Click **Configure**. The **Enable Orphan BLOB Retention** option appears. Selecting **No** disables the **Orphan BLOB Retention** function for the selected scope. Selecting **Yes** displays the following options.
 - **Schedule** – Configure the schedule for the Orphan BLOB Retention.
 - **Start Time** – The scheduled BLOB Retention job starts on the time specified.
 - **Interval** – Set an interval for recurring rules based on a schedule using the option of **Day(s)**, **Week(s)** or **Month(s)**.
 - **Advanced Settings** – Specify a delay time.
5. The orphan BLOB content is retained in the external device for the time period specified here. During this time period, if the deleted Storage Manager stubs are restored back to SharePoint, then the restored stubs can still be accessed normally.
6. After the specified time elapses, the content is deleted. Use this option to prevent the deletion of orphan BLOB content when Storage Manager stubs are deleted accidentally. The deletion can be delayed using the option of **Day(s)**, **Week(s)** or **Month(s)**.

***Note:** It is not recommended that you use the SQL Server RBS Maintainer tool to delay the deletion time manually, because the orphan BLOB content might be deleted unexpectedly. If this occurs, the BLOB content will become orphan stubs with no related BLOB content after the backed-up stubs are restored back to SharePoint.

7. Click **OK** to save the configuration and click **Cancel** to cancel the operation.
8. After saving the BLOB retention rule, the corresponding **Orphan BLOB Retention** column will display **Yes** beside the selected nodes. Click **Yes** to view the current BLOB retention rule on the respective SharePoint component.

Configuring the Processing Pool (Scheduled Storage Manager Only)

The Processing Pool feature allows you to control the maximum number of Storage Manager jobs that can be run at the same time. Normally, a Storage Manager job is fairly resource-intensive, so running multiple Storage Manager jobs simultaneously may affect the performance of the server. To avoid this condition, use the Processing Pool feature.

Storage Manager jobs that are added into the Processing Pool become threads. The number of threads you allow in the processing pool is the maximum number of Storage Manager jobs that can be run simultaneously. The remaining jobs are placed in a queue.


Each SharePoint farm has a default processing pool named **Default_Farm(FarmName)**. The number of threads set in the default processing pool is **5**.

To create a new Processing Pool:

1. Click Scheduled Storage Manger tab > Processing Pool in the General Settings group.
2. Click **Create** in **Manage** group on the **Processing Pool** tab. The **Processing Pool** page appears.
3. Enter a **Processing Pool Name** and an optional **Description** used to distinguish the new processing pool.
4. Select the **Farm** from the drop-down list. The processing pool can only be used for the Scheduled Storage Manager rules configured for this farm.
5. Select an **Agent Group** from the drop-down list. The agents in the specified agent group are used by this Processing Pool to perform the Scheduled Storage Manager jobs. For more information about creating agent groups, refer to the [DocAve 6 Control Panel Reference Guide](#).
6. Specify the **Number of Threads** that will be used for the Scheduled Storage Manager jobs. For example, if you enter **8** in this field, then 8 Scheduled Storage Manager jobs can be run at the same time.

Creating Storage Manager Rules

After completing the pre-configuration steps, you are ready to create Storage Manager rules for your selected scope.

***Note:** If the newly-created SharePoint objects are not displayed on the tree, right-click the corresponding node and click  to refresh the tree.

Once a Storage Manager job is run and the BLOB content in SharePoint is stored to the external device, what remains in SharePoint is a stub, not the actual data. The data can be viewed and modified as though it is being accessed from a SharePoint SQL Server, when in reality the data is stored on an external device.

Refer to the sections below for information on creating Storage Manager rules for [Realtime Storage Manager](#) and [Scheduled Storage Manager](#).

Creating Realtime Storage Manager Rules

Realtime Storage Manager rules take certain BLOB content being uploaded into SharePoint SQL Server and redirect that content to the specified external device. This is done in “real time”; that is, the BLOB content is never actually uploaded into SQL. As a result, document size is the only criteria by which to specify which content gets moved to the external device.

***Note:** When only the RBS provider is enabled, the Realtime Storage Manager rule is only available at the Web application level. When only the EBS provider is enabled, the Realtime Storage Manager rule is available at the Web application level and site collection level. If both the RBS provider and the EBS provider are enabled, refer to [Appendix B](#) for more information.

***Note:** If there is insufficient space on the BLOB storage device, then when uploading a document that meets the criteria for Realtime Storage Manager externalization, the upload fails.

To create Realtime Storage Manager rules:

1. Click the **Realtime Storage Manager** tab and [select the content](#) from the Scope panel.
2. Click Create in the Rule Management group.

***Note:** When multiple rules are configured at the same level, the uploaded BLOB content is moved to the external device as long as it triggered any one of the Storage Manager rules. The rules will be automatically ordered according to the criteria; rules with bigger criteria will be checked first.

3. **Rule Name** – Select a method for creating the Storage Manager rule. In this case, we select **Create a new rule**.

- **Reuse an existing rule** – Uses all the settings of the Storage Manager rule that you select from the drop-down list. This option expands the scope of the original rule. Use this option if you want to use the same Storage Manager rule in a different location.
- **Copy from an existing rule** – Copies the existing settings from the Storage Manager rule selected from the drop-down list. Update the rule if necessary and save this rule with a new name. This option can help you create a new rule quickly.
- **Create a new rule** – Creates a new Storage Manager rule. You must enter a name for the new rule.

***Note:** It is recommended that you create a common rule which can be reused directly or used after some modifications. This can help you make the most use of a created rule.

4. **Database Selection** (shown only when RBS is enabled) – Choose the scope of the rule by selecting the content database.

***Note:** The option(s) in this field will be greyed out if the corresponding database does not have a stub database or does not have RBS enabled.

- **Include new content databases** – Applies the rule on any newly-added content databases. Select this option if you want to enable the same rule on the newly-added content databases automatically, without having to manually configure them. Note that in order to enable this selection, you must have a stub database configured at the Web application level, and you must select **Include new content databases** when [configuring the BLOB Provider](#).
 - **Select all** – Selects all of the existing content databases.
5. **Criteria** – Specify the size trigger on the BLOB content to be uploaded to the external device using the option of **KB** or **MB**. If the size of the file is greater than or equal to the threshold specified here, the BLOB content is uploaded to the external device.
 6. **Storage** – Specify a logical device in the drop-down list to which the data will be stored. If desired, click **New Logical Device** to create a new logical device. For the information about creating the storage, refer to [Configuring Logical Devices](#).
 7. Select an optional **Data Compression** setting. Select the **Compression** checkbox to enable data compression, and then choose a compression level using the slider. A low compression level results in a faster compression rate but a larger data set, while a high compression level results in a slower compression rate but a smaller, better quality data set. Note that small data sets occupy more system resources and cause slower job times. The compression is performed on the SharePoint Server (**SharePoint Agent**).
 8. Select the **Encryption** checkbox (optional) to enable data encryption. Encryption protects sensitive materials; however, note that encrypting data causes slower archive times. The encryption is performed on the SharePoint Server (**SharePoint Agent**).
 9. After saving the Realtime Storage Manager rule, it is automatically enabled. View it in the **Rules and Settings** field on the right side of the workspace.

Creating Scheduled Storage Manager Rules

Use Scheduled Storage Manager to set more refined criteria as to what BLOB content gets uploaded to the external device, and when. The Scheduled Storage Manager criteria selections are more diverse. In addition, as opposed to Realtime Storage Manager rules, the scope of Scheduled Storage Manager rules is not restricted by the BLOB provider being used.

1. Click the **Scheduled Storage Manager** tab and [select the content](#) from the Scope panel.
2. Click **Create** in the **Rule Management** group.
3. **Rule Name** – Select a method for creating the Storage Manager rule. In this case, we select to create a new rule.
 - **Reuse an existing rule** – Uses all the settings of the Storage Manager rule that you select from the drop-down list. This option expands the scope of the original rule. Use this option if you want to use the same Storage Manager rule in a different location.
 - **Copy from an existing rule** – Copies the existing settings from the Storage Manager rule selected from the drop-down list. Update the rule if necessary and save this rule with a new name. This option can help you create a new rule quickly.
 - **Create a new rule** – Creates a new Storage Manager rule. You must enter a name for the new rule.

***Note:** It is recommended that you create a common rule that can be reused directly or used after some modifications. This can help you make the most use of a created rule.

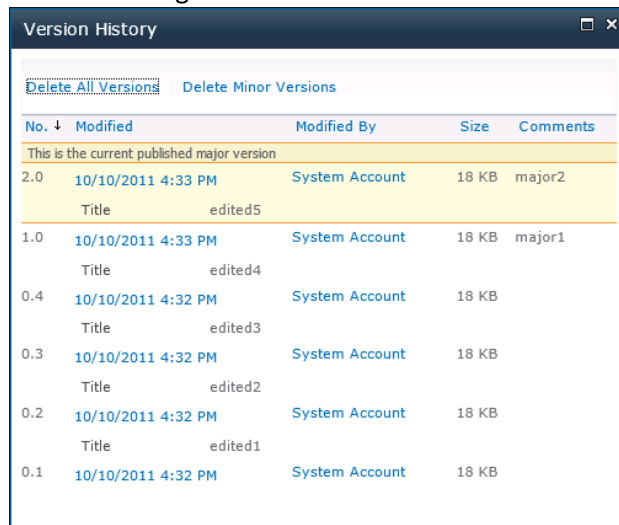
4. **Criteria** – Allows you to select specific objects or data to be uploaded to the external device within the following SharePoint levels: document level, document version level, and attachment level. Each level has a unique set of rules that apply to enhance configurations. The following table details the support status of the criteria that users can configure for each level:

For detailed examples of each filter rule, refer to the [DocAve 6 Control Panel Reference Guide](#).

Rule	Level		
	Document	Document Version	Attachment
Name	√	×	√
Title	×	√	×
Size	√	√	√
Modified Time	√	√	×
Created Time	√	×	√
Modified By	√	√	×
Created By	√	×	√
Content Type	√	×	×
Column(Text)	√	×	√
Column(Number)	√	×	√
Column(Yes/No)	√	×	√
Column(Date and Time)	√	×	√

Keep the Latest Version

The specified number of latest versions are kept in SharePoint; the other versions are stored to the external storage.



The screenshot shows a 'Version History' window with a table of document versions. The table has columns for 'No.', 'Modified', 'Modified By', 'Size', and 'Comments'. The current published major version is 2.0. Below it are minor versions 1.0, 0.4, 0.3, 0.2, and 0.1. The table is as follows:

No. ↓	Modified	Modified By	Size	Comments
This is the current published major version				
2.0	10/10/2011 4:33 PM	System Account	18 KB	major2
	Title	edited5		
1.0	10/10/2011 4:33 PM	System Account	18 KB	major1
	Title	edited4		
0.4	10/10/2011 4:32 PM	System Account	18 KB	
	Title	edited3		
0.3	10/10/2011 4:32 PM	System Account	18 KB	
	Title	edited2		
0.2	10/10/2011 4:32 PM	System Account	18 KB	
	Title	edited1		
0.1	10/10/2011 4:32 PM	System Account	18 KB	

Keep latest 3 major and minor versions – Versions 0.1 and 0.2 are stored to the external storage.

***Note:** The current version is always kept when using this rule. The number specified in the rule does not contain the current version. In other words, the current version and the specified number of versions (in the rule) are kept.

x

v

x

- After configuring one rule, click **Add** to add it to the existing rules. Click the **X** following each rule to delete it.
 - When 2 or more rules are configured, change the logical relationship between the rules by clicking the logic icon following the rule. There are two logics: **And** and **Or**. The default logic is **And**.
 - **And** – The content that meets all the rules is uploaded to the external device.
 - **Or** – The content that meets any one of the rules is uploaded to the external device.
 - Under the **Basic Filter Conditions** field, view the rules' logical relationship. For example, if the logical relationship is ((1 And 2) Or 3) in the **Basic Filter Condition** field, the contents which meet both the filter rule 1 and filter rule 2, or meet the filter rule 3 is uploaded to the external device.
5. **Storage** – Specify a logical device in the drop-down list to store the data. If desired, click **New Logical Device** to create a new logical device. For more information about creating logical devices, refer to [Configuring Logical Devices](#).
 6. Select an optional **Data Compression** setting. Select the **Compression** checkbox to enable data compression, and then choose a compression level using the slider. A low compression level results in a faster compression rate but a larger data set, while a high compression level results in a slower compression rate but a smaller, better quality data set. Note that small data sets

occupy more system resources and cause slower job times. The compression is performed on the SharePoint Server (**SharePoint Agent**).

7. Select the **Encryption** checkbox (optional) to enable data encryption. Encryption protects sensitive materials; however, note that encrypting data causes slower archive times. The encryption is performed on the SharePoint Server (**SharePoint Agent**).

After saving the Scheduled Storage Manager rule, it is automatically enabled. View it in the **Rules and Settings** field on the right side of the workspace.

Advanced Settings

Once a Scheduled Storage Manager rule is created and saved, additional options can be configured from the **Rules and Settings** field in the workspace.

Configure the following options:

1. Specify the order of the rules using the **Order** drop-down list. When multiple Scheduled Storage Manager rules are configured for a particular SharePoint object, the content that meets any one of the applied rules is uploaded to an external device. You can specify which rule will be checked first by changing their order.
2. **Processing Pool** – Specify a processing pool to be used by the Scheduled Storage Manager rules. Click **New Processing Pool** to create a new pool. Refer to [Processing Pool](#) for more information.
3. **Notification** – Configure email notification settings. Using the drop-down boxes, select the type of report (**Summary** or **Detailed**), enter the recipient's e-mail address, and click **Add**. Repeat this procedure for any additional recipients.
4. **Schedule** – Select when to store the BLOB content to the external storage.
 - **No Schedule** – Run the plan immediately after finishing the wizard.
 - **Configure the schedule myself** – Configure a schedule yourself and run the Scheduled Storage Manager job according to the specified start time and interval.
5. Click **Apply** to apply the Scheduled Storage Manager job settings to the selected node. After the Scheduled Storage Manager job settings are applied to the selected node(s), the **Run Now** button becomes enabled. Or, click **Apply and Run Now** to save the Scheduled Storage Manager job settings and run the Scheduled Storage Manager job immediately (this button combines the functions of **Apply** and **Run Now**).
6. Click **Run Now** to run the Scheduled Storage Manager job immediately.

Managing Created Rules

Once created and applied, Storage Manager rules can be edited, removed, disabled (enabled), and deleted. In addition, inheritance can be applied or stopped. See the sections below for information on managing Storage Manager rules.

Inheriting and Stop Inheriting

After [creating a Storage Manager rule](#) on the selected node, the sub-nodes automatically inherit the rule that is applied to their parent node.

Stop Inheriting logically separates the rule in the lower level-node from the upper-level node; this is useful, for example, in ensuring that often-used data is not uploaded to an external device accidentally.

When configuring rules for the first time, you can configure rules directly at any level that the Realtime/Scheduled Storage Manager allows. After one rule has been configured for a particular level (Level A), you can still configure rules directly to levels that are higher than Level A.

However, if you want to configure rules at levels that are lower than Level A, you must first break the rule inheritance. To break this inheritance, click **Stop Inheriting** in the **Rule Management** group.

To manually apply a Storage Manager rule inheritance to a node:

1. Select the node that will inherit the parent node's Storage Manager rule.
2. Click **Inherit** in the **Rule Management** group to inherit the rule. After clicking **Inherit**, this button changes to **Stop Inheriting**.

Helpful Notes on Using Inherit

- By default, the Storage Manager rule configured for a higher level is inherited by the lower levels.
- You cannot create a new rule on a node if it inherits the Storage Manager rules of the higher level. Inheritance must first be broken.
- The inherited Storage Manager rule cannot be edited or removed; it can only be viewed.
- Once you have broken the rule inheritance on a node, if you choose to inherit the rules of the higher level again, all of the rules that are added after breaking the inheritance on this node will be removed from the following nodes:
 - The node that inherits the higher level rules again.
 - The nodes that inherit the rules from the node above.

***Note:** The newly-added rules in the specified level(s) are not deleted. They still exist in the Storage Manager Rule list for re-applying if necessary.

The Conditions of Stop Inheriting

- When breaking the rule's inheritance at a specified level, the inheritance is only broken at this level. The rule's inheritance of the lower levels is not broken.
- After the rule's inheritance is broken, you can apply new rules and edit the corresponding rule settings at the lower level.
- If you break the rule inheritance on a node and then disable the inherited rule, this rule can still be used by the upper level.
- If you break the rule inheritance on a node and then edit the inherited rule, the changes will affect all of the SharePoint objects which used this rule.
- If you run a Scheduled Storage Manager job at a higher level after breaking the rule's inheritance on a specified node, this node will not be included in the scope of the specified job, even if it still uses this rule.

Editing or Removing Storage Manager Rules

Existing Storage Manager rules can be edited or removed from nodes. See the information below.

Editing Existing Storage Manager Rules

Certain created Storage Manager rule settings can be modified using the **Edit** feature. The **Rule Name Description, Criteria, Storage Policy, Data Compression, and Data Encryption** settings can all be edited. The **Rule Name** is fixed and cannot be changed using the **Edit** feature.

To edit an existing Storage Manager rule:

1. [Select the node](#) in the **Scope** panel that contains the rule or rules you would like to edit.
2. Select a rule from the **Rules and Settings** workspace by selecting the checkbox to the left of the rule.
3. Click **Edit** in the **Rule Management** group.
4. If the current rule you are editing is associated with several SharePoint objects, you can view associated objects for this rule in the pop up window. If necessary, click **Continue** in the **Commit** group to continue the editing.
5. The **Storage Manager Rule** page appears; you can view the settings for this rule. See [Creating Storage Manager Rules](#) for more information.

Removing Existing Rules from Nodes

When a created rule is no longer useful for the node(s) in the farm tree, you can remove this rule from the node. Use the **Remove** feature to simplify rules management and lessen unnecessary created rules in the nodes.

To remove an existing Storage Manager rule:

1. [Select the node](#) in the **Scope** panel that contains the rule or rules you would like to remove.
2. Select one or more rules from the **Rules and Settings** workspace by selecting the checkbox(es) to the left of the rule.
3. Click **Remove** in the **Rule Management** group.
4. The rule or rules are removed from the selected node and sub-nodes that inherit this rule in the tree.

***Note:** The removed rule or rules still exist in the Storage Manager Rule list for re-applying if necessary.

Enabling and Disabling Rules

Once rules are created and applied to a node, they can be enabled (to be included in a run) and disabled (to be omitted from a run) as desired. Disabling is useful, for example, when some of the rules do not need to be run in a particular job. As opposed to [deleting the rule](#), which removes it from the Storage Manager Rule list entirely, disabling the rule omits it from that particular run but saves it in the Storage Manager Rule list for future use.

See below for information on disabling and enabling rules.

Disabling an Existing Rule

After creating and applying the rule, the **Status** of rule in the **Rules and Settings** pane reads **Enabled** by default. In some situations, it may be desired to disable a rule for a particular job.

To disable an enabled Storage Manager rule:

1. [Select the node](#) in the **Scope** panel that contains the rule or rules you would like to disable.
2. Select a rule or rules from the **Rules and Settings** workspace by selecting the checkbox(es) to the left of the rule.
3. Click **Disable** from the **Rule Management** group to disable the selected rule(s). The rule status changes to **Disabled** in the **Status** column.

***Note:** If one Storage Manager rule is disabled at a specified level, the corresponding rule inherited by the lower levels will also be disabled.

Enabling a Disabled Rule

To enable an existing disabled Storage Manager rule:

1. [Select the node](#) in the **Scope** panel that contains the rule or rules you would like to enable.
2. Select a rule or rules from the **Rules and Settings** workspace by selecting the checkbox(es) to the left of the rule.
3. Click **Enable** from the **Rule Management** group to enable the selected rule(s). The rule status changes to **Enabled** in the **Status** column.

Deleting Rules

Delete a created rule when it no longer has a use for any node within the entire farm tree. Deleting a rule permanently removes the rule from the Storage Manager Rules list.

To delete a Storage Manager rule, use the Rule Viewer. For more information, refer to [Deleting Rules](#).

***Note:** Since the delete feature is an operation that cannot be undone, be sure that the rule you are deleting is absolutely no longer needed.

Using Rule Viewer


Use the Rule Viewer to search for and manage created Storage Manager rules. The basic operations described below are:

- Viewing and searching for rules using a Rule Based view
- Viewing and searching for rules using a Scope Based view
- Deleting rules that are useless to the whole farm tree
- Removing rules from the selected scope(s)



To launch the Rule Viewer, click **Rules** in the **View Rule** group. The Rule Viewer page appears. See the sections below for information on using the Rule Viewer.

Rule Based Viewing

The **Rule Based** table view initially displays all rules sorted alphabetically in the first (**Rule Name**) column and includes two other columns: **Scope** and **Module**.

Note that only the **Module** column can be deselected from the column list by clicking the  icon, deselecting **Module**, and clicking **OK**.

When in **Rule Based** view, you can use the following functions:

- **Search all pages** – Input a keyword and click  to search for the related rules from the first page to the last page.
- **Search current page** – Input a keyword and click  to search for the related rules in the page you select.

***Note:** Keywords are not case sensitive, and the search accepts partial word input. However, use of wildcards is not supported.
- **Show rows** –The number of the rules that are displayed on each page is **15** by default. You can reset the number from the drop-down list. Once the number of rules exceeds the number that can be included in a single page, they will be displayed in the next page. You can enter a specific **Number** to directly go to the corresponding page or click < or > icon to page forwards or page backwards.

Deleting Rules


To delete a Storage Manager rule in the **Rule Based** view, follow the steps below.

1. Click the Realtime Storage Manager tab or the Scheduled Storage Manager tab.
2. In the **View Rule** group, click **Rules**.
3. In the pop-up window, click **Rule Based** in the **View** group.



4. Select a rule or rules you want to delete in the table by selecting the checkbox(es) to the left of the rule.
5. Click **Delete** in the **Manage** group to delete the selected rule(s). The rule is permanently deleted from the Storage Manager Rules list.

Scope Based Viewing

The **Scope Based** table view initially displays all rule-related URLs sorted alphabetically in the first (**Scope**) column and includes the **Rule Name** column as well. The **Rule Name** column lists every rule applied to the scope in the same row.

Note that the **Rule Name** column can be deselected from the column list by clicking the  icon, deselecting **Rule Name**, and clicking **OK**.

When in **Scope Based** view, you can use the following functions:

- **Search all pages** – Input a keyword and click  to search for the related rules from the first page to the last page.
- **Search current page** – Input keyword and click  to search for the related rules in the page you select.
- **Show rows** – The number of the rules that are displayed on each page is **15** by default. You can reset the number from the drop-down list. Once the number of rules exceeds the number that can be included in a single page, they will be displayed in the next page. You can enter a specific **Number** to directly go to the corresponding page or click **<** or **>** icon to page forwards or page backwards.

Removing Rules

To remove a Storage Manager rule in the **Scope Based** view, follow the steps below.

1. Click the Realtime Storage Manager tab or Scheduled Storage Manager tab.
2. In the **View Rule** group, click **Rules**.
3. In the pop-up window, click **Scope Based** in **View** group.
4. Select a scope or scopes where you want to remove the rules in the table by selecting the checkbox(es) to the left of the scope.

Click **Remove** in the **Manage** group to remove all the rule(s) from the selected scope(s).

For more information, refer to [Removing Existing Rules from Nodes](#).

Converting Stubs to Content

After running a Storage Manager job, should you want to convert the stubs back to data on SharePoint's SQL Server, use the **Convert Stub to Content** function. Refer to [How to Determine the Data is Stub or Real Content](#) for information on identifying stubs and content.

1. Click Realtime or Scheduled Storage Manager tab > Convert Stub to Content in the Restore group. A pop-up window appears.
2. In the pop-up window, select the scope where you want to perform the stub restore. After expanding the tree, you can enter the criteria in the textbox and click the magnifier to search for the specified object. Press **Enter**; only the object that you specified in the textbox will be displayed on the expanded tree.
3. The tree can be expanded down to the item level. Click **Items** and all stubs of the items stored in the external storage are displayed in the **Stub Browser** area. The name of the stub, the type of the stub, and the size of the stub's real data are displayed.
4. Choose to convert the stubs immediately or on a specified time in the **Schedule** field; configure the following settings.
 - **Convert now** – Choose this option if you want to convert the stubs immediately.
 - **Configure the schedule myself** – Choose this option to convert the stubs based on a schedule. If this option is selected, the following option appears:
 - **Schedule Settings** – Specify the start time of the converting job.
5. Click **OK** to save the configuration and the stubs of the selected scope will be converted to the real data after the converting job.

Refer to [Customizing the Number of Storage Manager Job Threads](#) for information on setting the number of threads (jobs) that the **Convert Stubs to Content** feature runs simultaneously.

How to Determine if the Data is Stub or Real Content

There are no identifying markers for Storage Manager stubs in SharePoint. Refer to the methods below to determine if the data is a stub or real content.

- If EBS Provider is used, navigate to the content database of the site collection in Microsoft SQL Server Management Studio. In the **AllDocs** table, if the value of the specified item in the **docflags** column is larger than **65533**, this item is stub. Otherwise, it is real content.
- If RBS Provider is used, navigate to the content database of the site collection in Microsoft SQL Server Management Studio. In the **AllDocStreams** table, if the value of the specified item in the **RbsId** column is not **Null**, the item is a stub; if the value of the specified item in the **Content** column is not null, the item is real content.

Customizing the Number of Storage Manager Job Threads

You can customize the number of threads to be used in the **Scheduled Storage Manager** jobs and **Convert Stubs to Content** jobs. Refer to the detailed steps below:

1. Access the installation path of DocAve Agent. The default path is ...\\AvePoint\\DocAve6\\Agent\\data\\SP2010\\Arch.
2. Locate **AgentCommonStorageEnv.cfg** and open it with Notepad.
3. Find the node named **MaxThreadCount** and change its value to a bigger number. The default value is **1**, which means **1** thread will be used in the **Scheduled Storage Manager** jobs and **Convert Stubs to Content** jobs. The maximum value of this attribute is **10**.
4. Save the modification and close the file.
5. Repeat the steps above on each DocAve Agent server.

After the modification is saved, all the newly started **Scheduled Storage Manager** jobs and **Convert Stubs to Content** jobs will be run using the specified number of thread. This will improve the efficiency of the corresponding job.

Checking the Job Status

In the **Realtime Storage Manager** tab or **Scheduled Storage Manager** tab, in the **Statistics** group, click **Job Monitor** to navigate to the job monitor of the corresponding module.

You can then view the status of the corresponding Storage Manager jobs.

This is the first step to take when monitoring jobs or troubleshooting for errors. For the detailed information of each option, refer to the [DocAve 6 Job Monitor Reference Guide](#).

Collecting Logs for Troubleshooting

If an error occurs within the job, refer to the steps below to collect the logs for troubleshooting.

1. Navigate to **DocAve > Control Panel > Common Application Settings > Log Manager**.
2. Click the **Log Manager** link to go to Log Manager.
3. Click **Log Settings** in the **Settings** group.
4. In the pop-up window, in the **Control Service** tab, set the log level to **Debug** in the **Log Level** column of each record.
5. Repeat step 4 in the **Media Service** tab, **Report Service** tab, and **SharePoint Agent** tab.
6. Click **OK** in the **Commit** group to save the settings.
7. Rerun the job to reproduce the issue.
8. Navigate back to Log Manager and select all of the items in each tab in the **Log Collection** field.
9. Click **Collect** in the **Actions** group to collect the logs.
10. Navigate to **DocAve > Job Monitor** to view the status of the collection job. Once the job is finished, select it and click **View Details** in the **Manage** group.
11. In the pop-up window, click **Download** in the **Download** group and save the downloaded logs to a local path.
12. To upload logs to our support team, [click here](#).

For further information on the Log Manager, refer to the [DocAve 6 Control Panel Reference Guide](#).

DocAve Storage Manager Use Cases

The following use cases apply the procedures described in this guide to several real-world situations. These use cases are intended to give the user an idea of when and why one would use common DocAve Storage Manager functionality. Note that these situations assume that the DocAve platform and applicable modules have been installed successfully, and that the prerequisite conditions are met (appropriate permissions for the Agent account, creation of logical devices, etc.).

DocAve Storage Optimization Suite in Conjunction with Tiered Storage

The DocAve Storage Optimization Suite enables SharePoint administrators to better manage data storage using hierarchical storage management (HSM). Moving data to an appropriate tier of storage lowers storage costs and improves SQL Server performance. With DocAve Storage Manager, the data can be extended to an appropriate tier, depending upon how frequently data is accessed or modified in SharePoint. An example of a tiered storage structure as it relates to extended SharePoint content is below:

- Content modified or accessed frequently is extended to Tier 1 storage (most expensive, exceptional system performance).
- Content that has not been modified for 6 months or more is extended to Tier 2 storage (less expensive than Tier 1, moderate system performance).
- Content that has not been modified for over 2 years is extended to Tier 3 storage (the cheapest of the three tiers).
- When there is no longer an immediate need for content to be accessible, use DocAve Archiver to move the data to NAS or Cloud storage.

To set up this tiered storage structure, Joe (the SharePoint administrator) must first determine the threshold size of data to leave in SQL. As [DocAve's whitepaper](#) explains, BLOBs greater than 512 KB or 1 MB should be externalized. Doing so improves performance for both reads and writes (given similar performance characteristics of the underlying storage platform), whereas access to a document smaller than 256 KB is faster with the BLOB stored in the SQL Server content database.

Given this information, Joe [selects the target Web application](#) and sets up the following Realtime Storage Manager rule.

***Note:** Joe can also apply Storage Manager rules to content that was connected with [DocAve Connector](#).

Realtime Storage Manager

Storage Manager: Keep data connected to SharePoint

Rules and Settings
http://finweb.contoso.com/

<input checked="" type="checkbox"/>	Order	Rule Name	Criteria	Storage	Status
<input checked="" type="checkbox"/>	1	Active Content Tier 1 Storage	Size >= 512KB;	T1 LD	Enabled

Details:

Criteria: File size >= 512KB
 Storage: T1 LD
 Compression: None
 Encryption: None
 Database Selection: WSS_Content_6b8e9cd10bde40398f3b0f5f2bb6ef05
 Include New Content Databases: Yes

Figure 3: Extending active content larger than 512 KB.

Joe then uses Scheduled Storage Manager to configure rules for less-frequently accessed data:

Rules and Settings

Storage Manager: Keep data connected to SharePoint

Rules and Settings
http://finweb.contoso.com/

Rules

<input checked="" type="checkbox"/>	Order	Rule Name	Scope	Status
<input checked="" type="checkbox"/>	1	Active Content Tier 2 Storage	http://finweb.contoso.com/	Enabled

Details:

Criteria:
 1.Document, Size, >=, 512KB;
 2 Document, Modified Time, Older Than, 6Month(s)
 Criteria Filters: (1 And 2)
 Storage: T2 LD
 Compression: None
 Encryption: None

Figure 4: Moving content older than 6 months to tier 2 (T2) storage.

Lastly, Joe configures a rule for Scheduled Storage Manager to extend old content:

Rules and Settings

Storage Manager: Keep data connected to SharePoint

Rules and Settings
http://finweb.contoso.com/

Rules

<input type="checkbox"/>	Order	Rule Name	Scope	Status
<input type="checkbox"/>	1	Active Content Tier 2 Storage	http://finweb.contoso.com/	Enabled
<input checked="" type="checkbox"/>	2	Old Content Tier 3 Storage	http://finweb.contoso.com/	Enabled

Details:

Criteria:
 1.Document, Size, >=, 512KB;
 2 Document, Modified Time, Older Than, 2Year(s)
 Criteria Filters: (1 And 2)
 Storage: T3 LD
 Compression: None
 Encryption: None

Figure 5: Moving content older than 2 years to tier 2 (T3) storage.

Now that he’s created rules for active data, Joe wants to configure a rule that manages “old” data (data that has not been modified for 3 years or later). For this content, Joe configures an [Archiver rule](#) that moves the old data to NAS storage.

By using a combination of Storage Manager and Archiver (and even incorporating “connected” content into his storage management scope), Joe is able to increase SQL Server storage space and improve SQL Server performance. He uses Storage Manager to disseminate content to various tiers of storage, dependent upon how frequently the data is modified. Using information gleaned from AvePoint’s whitepaper on externalizing BLOB storage, Joe chooses to extend BLOBs that are 512 KB or greater. Finally, Joe configures Archiver to move the data to NAS (or Cloud) storage at the end of the content’s lifecycle, ensuring proper and timely archiving of documents within his scope.

DocAve Storage Optimization and Data Protection Suites in Compliance Environments

DocAve Storage Manager is not only used to better manage SharePoint data storage, it is also advantageous in a compliance environment. When using SharePoint in a compliance/regulated environment, it is essential that:

- All versions of active, editable content are kept in a compliance storage device where information, once written, cannot be modified (WORM storage or Dell DX Storage, for example).
- Active documents that should not be edited are preserved and immutable.
- Non-active documents are archived properly and remain preserved and immutable, but only for as long as necessary.

While the latter requirement can be accomplished using DocAve Archiver (to be discussed later in this section), the two former requirements can be accomplished using Scheduled Storage Manager. To that end, Scheduled Storage Manager can be configured to run routine searches on SharePoint content databases for data that matches specified criteria. When such data is found, Scheduled Storage Manager moves the real data to compliance storage while creating a data stub in SharePoint. So, in addition to increasing SQL storage space, Scheduled Storage Manager ensures that there will always be a record (or “breadcrumb”) of every document version.

To establish this system using the DocAve Storage Optimization Suite, Joe (the SharePoint administrator) first [selects the Web application](#) to which he wants to create and apply a rule. Joe must decide what criteria he will use to identify the data that should be “extended” from SharePoint to external BLOB storage. Knowing that his company classifies sensitive legal contract documents using a **Contract** Content Type in SharePoint, Joe creates the following criteria rule in the [Scheduled Storage Manager rule creator](#):

	Level	Rule	Condition	Value	
1	Document	Content Type	Matches	Contract	

With this criteria rule, Scheduled Storage Manager will externalize all SharePoint content that is classified with the **Contracts** Content Type. In the [Rules and Settings](#) pane, Joe configures Scheduled Storage Manager to run on a schedule of every evening. This configuration ensures that:

- All **Contract** documents (including newly-uploaded ones) are extended to compliance storage.
- The Scheduled Storage Manage plan will run during business off-hours, so the plan will not affect SharePoint performance during work hours.

In addition, Joe configures Data Protection Suite module [Granular Backup and Restore rules](#) to back up the Web application content on a nightly basis. By configuring rules for both the Storage Manager and

Granular Backup and Restore modules, Joe is now guaranteed to have a record in compliance storage of all active content versions.

Lastly, to further manage compliance data in storage, Joe wants to remove old **Contract** Content Type documents from SharePoint. Joe configures an [Archiver rule](#) to archive these documents if their modification date is older than 6 months. Joe chooses to run this Archiver job on a weekly basis.

By leveraging the DocAve Storage Optimization Suite and Data Protection Suite functionality, Joe is able to manage compliance data storage appropriately. Scheduled Storage Manager extends **Contract** Content Type documents to compliance storage. Regular Granular Backup and Restore backups ensure a record of all content versions, even if the content is deleted from SharePoint. Finally, DocAve Archiver removes the data stub from SharePoint, removes the real data from compliance storage, and archives the content in compliance storage.

Appendix A: Stub Database Inheritance

In some cases (records management, for example, where there will be millions of files), it may be necessary to configure stub databases down to the Web application or content database level. Should this be the case, refer to the information below.

Note the following stub database inheritance rules for newly-added Web applications and content databases:

- By default, the stub database of the parent node will be used by the lower level.
- If there is no stub database configured for the parent node, the stub database of the grandparent node will be used, and so on.

If the stub database is not configured for a specified SharePoint object when running a Storage Manager job, the object’s stub database will also be automatically configured using the rules above. The stub database configuration will then be saved and displayed in DocAve. The corresponding stub database will be used in all of the later Storage Manager jobs performed on the specified SharePoint object.

For example, assume that you only configured a stub database for a Web application and set a Storage Manager rule on it. You did not configure the stub database for any of the content databases under the Web application. When the Storage Manager rule is triggered on a specified content database in the Web application, the content database inherits the stub database of the Web application. The stub database configuration is saved and is used in all of the later Storage Manager jobs performed on this content database.

The x in the table below indicates that the corresponding component remains not configured. If you do not wish for a lower level to inherit the stub database of a higher level, expand the tree to the specified level and configure a stub database for the lower level separately.

Selected Nodes	Configuration Status of the Stub Database			
	Existing Web Application	Existing Content Database	Newly-Added Web Application	Newly-Added Content Database
Only the Farm Node	x	x	Inherit the farm’s stub database.	Inherit the farm’s stub database.
Only a Web Application Node	Only the stub database of the selected Web application node will be configured.	x	x	If the newly added content database is in the selected Web application node, it will inherit the Web application’s stub database. Otherwise, it will remain not configured.

Only a Content Database Node	×	Only the stub database of the selected content database node will be configured.	×	×
Farm Node and a Web Application Node	Only the stub databases of the farm node and the selected Web application node will be configured.	×	Inherit the farm's stub database.	If the newly added content database is in the selected Web application node, it will inherit the Web application's stub database. Otherwise, it will inherit the farm's stub database.
Farm Node and a Content Database Node	×	Only the stub databases of the farm node and the selected content database node will be configured.	Inherit the farm's stub database.	Inherit the farm's stub database.
Web Application Node and a Content Database Node	Only the stub databases of the selected Web application node and content database node will be configured.	Only the stub databases of the selected Web application node and content database node will be configured.	×	If the newly added content database is in the selected Web application node, it will inherit the Web application's stub database. Otherwise, it will remain not configured.
Farm Node, a Web Application Node and a Content Database Node	Only the stub databases of the farm node, the selected Web application node and the content database node will be configured.	Only the stub databases of the farm node, the selected Web application node and the content database node will be configured.	Inherit the farm's stub database.	If the newly added content database is in the selected Web application node, it will inherit the Web application's stub database. Otherwise, it will inherit the farm's stub database.

Appendix B: Coexistence of EBS and RBS

When both the RBS provider and the EBS provider are enabled, RBS will be used first. Refer to the following table for more information.

EBS Provider Status	RBS Provider Status	Case	Result
Enabled	Enabled on content database B, B is one of the content databases of Web application A and site collection C is in content database B	<ol style="list-style-type: none"> 1. Created a new rule on Web application A. 2. Break the rule inheritance on site collection C and create a new rule. 	<p>The new rule created on Web application A will be an RBS rule. The new rule created on site collection C will be an EBS rule.</p> <p>When a file is uploaded to C, the RBS rule configured for A will be checked first. If the file triggered the RBS rule, it will be extended.</p> <p>If the file does not trigger the RBS rule, the EBS rule configured for C will be checked. If the file triggered the EBS rule, it will be extended. If not, it will be uploaded to content database B in SQL Server.</p>
		Create a new rule on site collection C directly.	<p>The new rule created on site collection C will be an EBS rule.</p> <p>RBS is enabled for the content database where C is, but no RBS rule is configured, when a file is uploaded to C, the EBS rule will be checked.</p> <p>If the file triggered the EBS rule, it will be extended. If not, it will be uploaded to content database B in SQL Server.</p>

	Enabled on content database B, B is one of the content databases of Web application A and site collection C is not in content database B	<ol style="list-style-type: none"> 1. Created a new rule on Web application A. 2. Break the rule inheritance on site collection C and create a new rule. 	<p>The new rule created on Web application A will be an RBS rule. The new rule created on site collection C will be an EBS rule.</p> <p>Since only EBS is enabled for the content database where C is, when a file is uploaded to C, the EBS rule will be checked. If the file triggered the EBS rule, it will be extended. If not, it will be uploaded to the corresponding content database in SQL Server.</p>
		Create a new rule on site collection C directly.	<p>The new rule created on site collection C will be an EBS rule.</p> <p>Since only EBS is enabled for the content database where C is, when a file is uploaded to C, the EBS rule will be checked. If the file triggered the EBS rule, it will be extended. If not, it will be uploaded to the corresponding content database in SQL Server.</p>
	Enabled on Web application A and all the content databases of A	<ol style="list-style-type: none"> 1. Created a new rule on Web application A. 2. Break the rule inheritance on site collection C and create a new rule. 	<p>The new rule created on Web application A will be an RBS rule. The new rule created on site collection C will be an EBS rule.</p> <p>When a file is uploaded to C, the RBS rule configured for A will be checked first. If the file triggered the RBS rule, it will be extended.</p> <p>If the file does not trigger the RBS rule, the EBS rule configured for C will be</p>

			checked. If the file triggered the EBS rule, it will be extended. If not, it will be uploaded to the corresponding content database in SQL Server.
		Create a new rule on site collection C directly.	<p>The new rule created on site collection C will be an EBS rule.</p> <p>RBS is enabled for the content database where C is, but no RBS rule is configured, when a file is uploaded to C, the EBS rule will be checked.</p> <p>If the file triggered the EBS rule, it will be extended. If not, it will be uploaded to the corresponding content database in SQL Server.</p>

Index

- Agent Group**, 18
- AgentToolSP2010StorageEBS.exe, 15
- BLOB Provider**, 12
- Configure Stub Database**, 13
- Convert Stub to Content**, 30
- Copy from an existing rule**, 21
- Create a new rule**, 21
- Criteria**, 21
- Data Compression**, 20
- Database Selection**, 20
- Disable**, 26
- Edit**, 25
- Enable BLOB Provider**, 14
- Enable EBS for Farm**, 14
- Enable RBS for Farm**, 14
- Enabling EBS Using the Tool, 15
- Enabling RBS Using the Tool, 16
- Encryption**, 20
- General Settings, 17
- Include New Content Databases**, 14
- Inherit**, 24
- Install BLOB Provider Binaries**, 12
- Number of Threads**, 18
- Orphan BLOB Retention, 17
- Overview**, 15
- Processing Pool, 18
- Realtime Storage Manager, 7
- Remove**, 26
- Reuse an existing rule**, 21
- Rule Based**, 28
- Rule Management, 19
- Scheduled Storage Manager, 7
- Scope**, 9
- Scope Based**, 29
- Scope panel**, 9
- Stop Inheriting**, 24
- Storage Optimization**, 7
- View Rule**, 28

Notices and Copyright Information

Notice

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

Copyright

Copyright © 2012 AvePoint, Inc. All rights reserved. All materials contained in this publication are protected by United States copyright law and no part of this publication may be reproduced, modified, displayed, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written consent of AvePoint, 3 Second Street, Jersey City, NJ 07311, USA or, in the case of materials in this publication owned by third parties, without such third party's consent.

Trademarks

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

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

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

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

Changes

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

AvePoint, Inc.
3 Second Street
Jersey City, NJ 07311
USA