

***Microsoft Office
SharePoint Server (MOSS) 2007
Backup Strategies***



Table of Contents

Introduction	3
MOSS 2007 Recycle Bin	4
MOSS 2007 Native Backup Tools	7
Central Administration Backup	7
Command-line Backup	11
AvePoint Inc's Backup Solutions for MOSS 2007.....	13
Conclusion.....	22
Other AvePoint Solutions for MOSS 2007	23
Contact Information	23

Introduction

Microsoft SharePoint Products and Technologies have become widely popular and have helped everyone—from entire organizations to virtual teams to communities of users—collaborate and communicate more efficiently and effectively than ever. Because of this, Windows SharePoint Services and SharePoint Portal Server have been brought to the forefront in the list of mission critical tools.

Unfortunately, very few organizations have well documented SharePoint backup and recovery plans. Those that do have such plans generally focus on full disaster scenarios that involve the loss of an entire environment. These plans generally do not focus on item level backup and recovery. This is despite the fact that item level problems are significantly more commonplace and much more likely to occur than a full disaster. When item recovery does become necessary, it is usually accomplished by the re-uploading of content. However, such a re-upload causes a loss of content history, along with other meta data issues, and destroys many of the benefits of that formed the basis of putting the content within a SharePoint environment in the first place.

There has been much buzz about the enhanced features available in Microsoft Office SharePoint Server 2007, particularly in the areas of backup and recovery. While enhanced backup and recovery tools like a Recycle Bin and differential backups have been introduced, they do not completely satisfy the needs of SharePoint administrators. Administrators require tools that allow for fast restore capabilities at the portal, site, and item level. Because of these requirements, a third party backup and restore tool is still a necessary item in most MOSS 2007 environments. This paper examines the backup enhancements native to MOSS 2007, discusses shortfalls, explains why a third-party full fidelity item level backup and restore solution is required, and details how DocAve by AvePoint, Inc. addresses this need.

MOSS 2007 Recycle Bin

One of the main enhancements to Microsoft Office SharePoint Server 2007 is the addition of a Recycle Bin. For customers who use Windows SharePoint Services 2003, and who have not added a third-party recycle bin tool to their SharePoint 2003 environment, this is a very welcome addition. Without such a tool, a restore of an entire content database would be required to bring back a single deleted item, and as those who have tried such a restore know, this method creates a myriad of other issues.

The MOSS 2007 Recycle Bin is a two-stage recycle bin strategy designed to capture deletion events and allow for some retention of content before it is discarded. This Recycle Bin is part of the Web Application settings. These settings are accessed through the Central Administration Application Management, as shown in Figure 1.

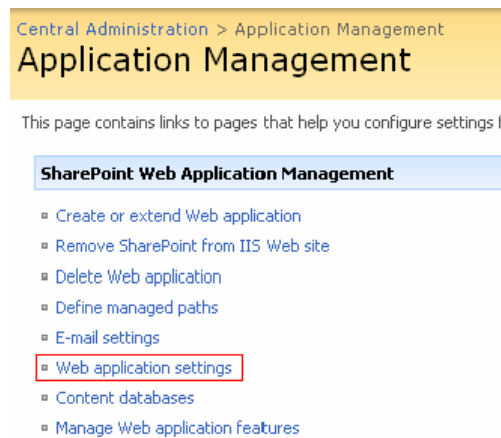


Figure 1: Accessing Web Application Settings

Clicking on the Web Application settings link, as shown in Figure 1 above, will open the Web Applications settings page. Recycle Bin settings, as shown in Figure 2, are on the bottom of this page.

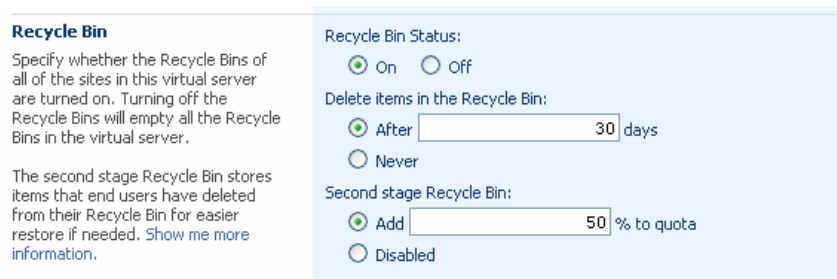


Figure 2: Recycle Bin Settings

Because the Recycle Bin is a Web Application setting, it can only be enabled or disabled for all of the site collections served by the Web Application; more granular, site level control is not possible. The first stage of the Recycle Bin is the user level or site level Recycle Bin, as shown in Figure 3. It is accessible by users and provides a site level view of deleted content—it will contain all items deleted from a particular site. The only adjustment on this bin, as shown in Figure 2, is retention time or the number of days to keep deleted content. Any changes to the retention time will apply to all site level Recycle Bins associated with that Web Application.




Figure 3: First Stage Recycle Bin

The second stage of the Recycle Bin is the administrator level or Site Collection Recycle Bin, as shown in Figure 4. It is only accessible by the administrator, and presents a site collection level view of deleted content. If there are multiple site collections supported by the Web Application and Recycle Bin is enabled in the Web Application settings, all of these site collections will have an independent Site Collection Recycle Bin. The only adjustment allowed on this recycle bin, as shown in Figure 2, is to disable or add to the quota size, and again changes will apply to all recycle bins served by that Web Application.



Figure 4: Site Collection Recycle Bin

When an item is deleted from a SharePoint site, it is moved to the Recycle Bin. It is held in the first stage Recycle Bin and visible to the end user for the retention time specified by the administrator. During this time it is also visible to the administrator in the Site Collection Recycle Bin. A user with permissions to that item can delete it from the first stage Recycle Bin. If this occurs, the item is still recoverable from the Site Collection Recycle Bin by the administrator. The site collection Recycle Bin has a view option that will only display items that have been deleted from end user Recycle Bins in that site collection.

One minor issue with this Recycle Bin solution is that it does not allow any granularity when it comes to restore options. If a user deletes an entire document library, but a single document needs to be restored, the entire document library must be restored.

An additional issue, and the reason why the Recycle Bin is not considered a true backup solution, is that it only captures deletion events. Errors, data corruption, or disaster events are still possible. These problems can occur at the entire farm level or down to the individual item level, and they will not be covered by Recycle Bin. To address this, robust backup is required, and this backup needs to provide granular coverage down to the individual item level.

MOSS 2007 Native Backup Tools

Central Administration Backup

Microsoft Office SharePoint Server 2007 does offer some native tools for backup. The first is Backup and Restore through the Central Administration site, as shown in Figure 5.




Figure 5: Backup and Restore Tool in Central Administration

Using this tool, a backup can be performed at various levels, the highest level being the entire SharePoint farm and the lowest level being an entire content database. A screenshot of this is shown in Figure 6. It should be noted that this tool does not offer a great deal of granularity and does not provide coverage at the site, subsite, or item level.

Select	Component	Type	Description
<input checked="" type="checkbox"/>	Farm	Farm	Content and configuration data for the entire server farm.
	SharePoint_Config	Microsoft.SharePoint.Administration.SPConfigurationDatabase	Configuration data for the entire server farm.
<input checked="" type="checkbox"/>	Windows SharePoint Services Web Application	Windows SharePoint Services Web Application	Collection of Web Applications
<input checked="" type="checkbox"/>	SharePoint 8000	Web Application	Content and configuration data for this Web Application.
<input checked="" type="checkbox"/>	WSS8000_Content	Content Database	Content for the Web Application.
	WSS_Administration	Central Administration	Collection of Web Applications
	Web Application	Web Application	Content and configuration data for this Web Application.
<input checked="" type="checkbox"/>	SharePoint_AdminContent_5d417680-4e89-46cc-8847-964d23d7cc1a	Content Database	Content for the Web Application.
<input checked="" type="checkbox"/>	SPSearch	Windows SharePoint Services Help Search	Index user and assistance content and serve search queries

Figure 6: Central Administration Backup Process Builder

Besides allowing you to pick what to backup up, to the extent described above, this tool also lets you choose a full or differential backup. A differential backup only captures changes since the last full backup, and therefore is quicker to run. Best practices generally call for executing a full backup weekly and incremental backups daily, though this can be adjusted depending on the sensitivity of the data in the site collection or farm.

The tool also allows the administrator to specify the backup location and provides an estimate of the disk space required. These settings are shown in Figure 7.

Start Backup

Use this page to configure a backup job.

Backup Content
Specify content to include in the backup.

Backup the following component:
Farm

Type of Backup
Specify which type of backup to start:

- Full - backs up the selected content with all history.
- Differential - backs up all changes to the selected content since the last full backup.

Type of backup:
 Full
 Differential

Backup File Location
Specify where to store the backup files. Each backup job is stored in a separate subdirectory at the location you specify, and each object is backed up to an individual file in that subdirectory.

Backup location:

Example: \\backup\SharePoint
Disk Space Required: 61.65 MB.

OK Cancel

Figure 7: Central Administration Backup Settings

Besides only providing granularity to the content database level, one serious shortcoming with this tool is that it does not natively contain any scheduling options. Thus an administrator will need to manually kickoff the backup processes each time it needs to be run or write a customized program to perform backups. An example of a completed backup job is shown in Figure 8.

Backup	
Requested By	DEMONET04\administrator
Phase	Completed
Start Time	8/14/2006 3:00 PM
Finish Time	8/14/2006 3:00 PM
Top Component	Farm
Backup ID	1a895bd0-5321-4ef1-a876-bb2b580cb47a
Directory	c:\demo 2007 backup\spbr0000\
Backup Method	Full
Warnings	0
Errors	0

Figure 8: Completed Central Administration Backup

When it comes to restore, this tool walks you through a four step process. The first step is to specify the location of the backup files, as shown in Figure 9. The next step is to choose which backup file to perform the restore from. Once a file is chosen, the content of the file is displayed.

Figure 9: Specifying Location of Backup Files

In step 3, the content of the backup file to be restored can be chosen, as shown in Figure 10. However, the granularity in choosing what to restore is very limited, as it was when choosing what to include in the backup.

Select	Component	Type
<input checked="" type="checkbox"/>	[-] Farm	Farm
	SharePoint_Config	Microsoft.SharePoint.Administration.SPConfigurationDatabase
<input checked="" type="checkbox"/>	[-] Windows SharePoint Services Web Application	Windows SharePoint Services Web Application
<input checked="" type="checkbox"/>	[-] SharePoint 8000	Web Application
<input checked="" type="checkbox"/>	WSS8000_Content	Content Database
	[-] WSS_Administration	Central Administration
	[-] Web Application	Web Application
	SharePoint_AdminContent_5d417680-4e89-46cc-8847-964d23d7cc1a	Content Database
<input checked="" type="checkbox"/>	SPSearch	Windows SharePoint Services Help Search

Figure 10: Selecting Content to Restore

When performing the restore, the tool allows two options— Overwrite and Create New, as shown in Figure 11. The Overwrite option will overwrite all existing content with the content from the backup file, while the Create New option will create a separate duplicate copy without affecting existing content. When choosing Create New, name changes can also be applied. A restore job is kicked off right away when selected, as again there are no scheduling options.

Restore from Backup - Step 4: Select Restore Options

Use this page to specify options for restoring from a backup.

Warning: this page is not encrypted for secure communication. User names, passwords and other information will be sent in clear text.

Restore Component
Select the specific content to be restored. Click Change to select a different component.

Restore the following content:
Farm

Restore Options
When you restore a backup, you can either overwrite over the existing data or create a new copy of the data.

Type of restore:
 Create new
 Overwrite

Login Names and Passwords
If using SQL Authorization a login name and password are need to access this information. If no user name and password are provided then windows authorization is used.

Farm\Windows SharePoint Services Web Application\SharePoint 8000
Login name: demonet04\administrator

New Names
If you chose to create copies for the restored content server name, file location on the server, and the new databases.

Microsoft Internet Explorer
Warning: All selected items will be overwritten. Do you want them to be overwritten?
OK Cancel

Figure 11: Restore Options

It should also be noted that this backup tool is not very robust. Errors or corruption in the backup file will cause the restore process to fail. Thus, it is good practice to backup the backup files and to consider a more frequent backup schedule.

Command-line Backup

In addition to this tool, Microsoft Office SharePoint Server 2007 also contains the command-line backup tools smigrate and stsadm. These appear to be the same tools contained in SharePoint 2003, and thus have the same issues. Stsadm must be executed on the local server by a properly permissioned administrator. In function, it provides the same coverage as the Central Administration backup above. Smigrate does not require administrator permissions on the local server, but it is not a full-fidelity backup and will lose some customizations and security settings. Stsadm export and import are also available, but this is no more than a combination of the SharePoint 2003 stsadm and smigrate commands mentioned above. General usage information for stsadm export is shown in Figure 12. General usage information for stsadm import is shown in Figure 13.

```
C:\Program Files\Common Files\Microsoft Shared\web server extensions\12\BIN>stsadm -help export

stsadm.exe -o export
            -url <URL to be exported>
            -filename <export file name>
            [-overwrite]
            [-includeusersecurity]
            [-haltonwarning]
            [-haltonfatalerror]
            [-versions <1-4>
              1 - Last major version for files and list items <default>
              2 - The current version, either the last major or the last minor
              3 - Last major and last minor version for files and list items
              4 - All versions for files and list items]
            [-cabsize <integer from 1-1024 megabytes>]
            [-quiet]
```

Figure 12: Stsadm Export General Usage Info

```

C:\Program Files\Common Files\Microsoft Shared\web server extensions\12\BIN>stsadm -help import
Missing operation name or the operation name is invalid.

Usage:
    stsadm.exe -o <operation> [<parameters>]
    stsadm.exe -help [<operation>]

Examples:
stsadm.exe -o addpath -url http://server/sites -type wildcardinclusion
stsadm.exe -o adduser
    -url http://server/site
    -userlogin DOMAIN\name
    -useremail someone@example.com
    -role reader
    -username "Your Name"
    -siteadmin
stsadm.exe -o backup -url http://server/site -filename backup.dat -overwrite
stsadm.exe -o backup -directory c:\backup -backupmethod full
stsadm.exe -o createadminvs
    -admapcreateneu
    -admapidname StsAdminUs
    -admapidtype configurableid
    -admapidlogin DOMAIN\user
    -admapidpwd Password
stsadm.exe -o createsite -url http://server/site
    -ownerlogin DOMAIN\name
    -owneremail someone@example.com
stsadm.exe -o createweb -url http://server/site/web
stsadm.exe -o deletesite -url http://server/site
stsadm.exe -o deleteweb -url http://server/site/web
stsadm.exe -o enumsites -url http://server
stsadm.exe -o enumsubwebs -url http://server/site/web
stsadm.exe -o enumusers -url http://server/site/web
stsadm.exe -o extendvs -url http://server:80
    -ownerlogin DOMAIN\name
    -owneremail someone@example.com
stsadm.exe -o renameweb -url http://server/site/web1 -newname web2
stsadm.exe -o restore -url http://server/site -filename backup.dat
stsadm.exe -o restore -directory c:\backup -restoremethd overwrite
stsadm.exe -o setconfigdb -databaseserver server
stsadm.exe -o unextendvs -url http://server

For information about other operations and parameters,
use "stsadm.exe -help" or "stsadm.exe -help <operation>"

```

Figure 13: Stsadm Import General Usage Info

These tools are cumbersome manual command-line processes. While they are useful for performing various administrator tasks by a skilled administrator, they do not provide a true backup solution.

Thus the native backup tools in Microsoft Office SharePoint Server 2007 lack the needed automation, robustness, and granularity to provide full coverage of a SharePoint environment. With the above tools, the administrator has to manually control the backup processes by initiating various backups on a frequent basis or writing and managing a customized program to initiate these backups. When error or corruption events occur within the SharePoint environment at the item level, the administrator has to locate the proper backup and determine the best restore strategy. If the backup being used is from the Central Administration interface or from an stsadm backup, there is some limited granularity possible in the restore. This granularity helps, but there are still issues. If an overwrite is performed, any new content or changes since the last backup may be lost. If the Create New option is used, additional space will be required, content will need to be manually moved from the duplicated sites, and meta data and security issues can result. If the restore is being performed from an smigrate backup, customizations and security settings will be lost. Thus each available option has serious shortcomings. As with previous versions of SharePoint, the only complete backup solution available is to utilize a third-party tool.

AvePoint Inc's Backup Solutions for MOSS 2007

AvePoint's item level backup and restore product, DocAve 4.0, addresses all of the shortcomings inherent in the native backup and restore functionality in Microsoft Office SharePoint Server 2007 and further extends the administrators ability to properly manage their SharePoint environment. DocAve has been in existence since SharePoint 2001, and was the first item level backup and restore solution available for both SharePoint 2001 and 2003. The latest version, DocAve 4.0, expands DocAve from a backup and recovery tool to a full management platform. This is because DocAve 4.0 utilizes a scalable client-server deployment strategy that allows for deployment across multiple SharePoint instances, while enabling management and control from a single centralized browser-based interface. This architecture is depicted in Figure 14.

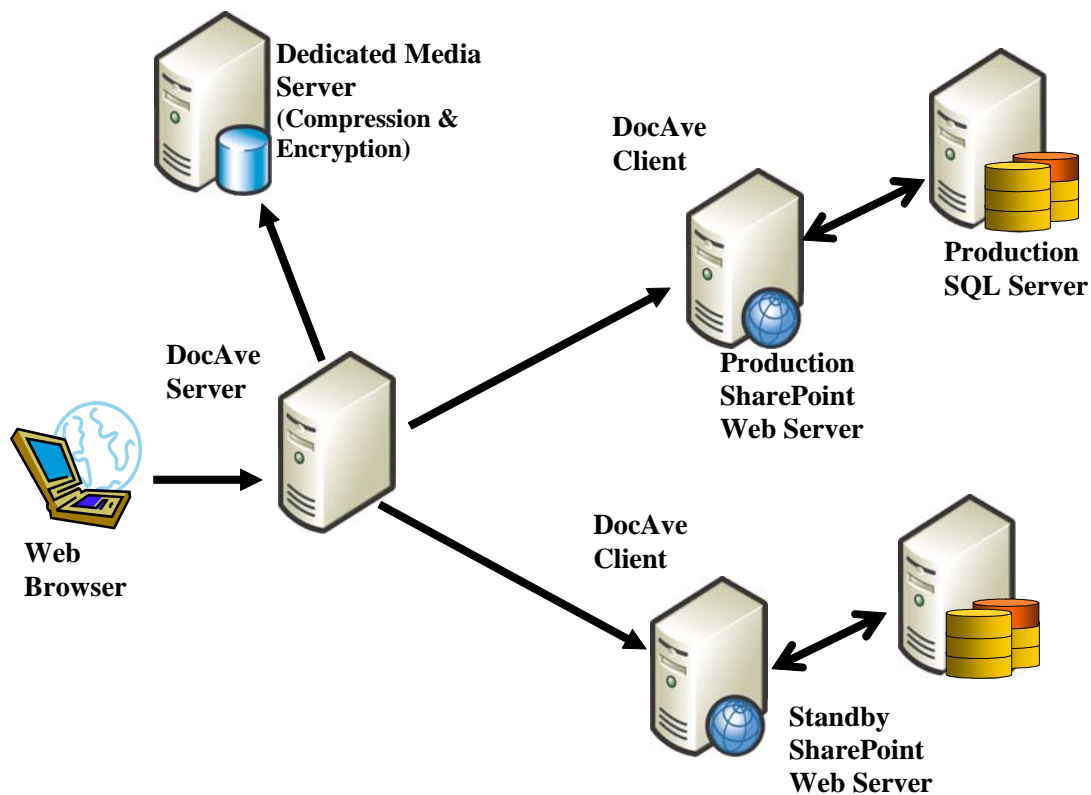


Figure 14: DocAve System Architecture

As shown in Figure 14, DocAve 4.0 Clients can be deployed on a front-end web server in each SharePoint environment. These clients are a light systems process, and do not use a significant amount of resources. The clients communicate with the DocAve 4.0 Server, which can be accessed via a web browser. While the figure depicts the DocAve 4.0 Client and Server installed on separate machines, they could be installed on the same machine.

In addition, a DocAve 4.0 Media Server can be designated to provide dedicated computing resources for DocAve 4.0 Server. Each DocAve 4.0 Server installation includes installation of DocAve Web Service, DocAve Network Service, DocAve Database Service, and DocAve Media Service. Each of these services can be run on separate dedicated machines, and thus there can be multiple dedicated Media Servers. The Media Servers performs the actual SharePoint data backup, thus effectively distributing the backup work load across multiple physical servers. Backup data on the Media Server supports full text indexing and search, and can also be encrypted and compressed.

DocAve 4.0 Server can also direct two DocAve 4.0 Clients to communicate with each other, allowing the transfer of SharePoint contents between SharePoint environments. Special care was taken in the DocAve 4.0 software architecture to ensure that this communication can survive very noisy or even intermittent data channels. Strenuous quality assurance tests on data packet fault tolerance included performing back ups and restores of SharePoint data across the Pacific Ocean for a period of three days, where Ethernet cables were intermittently unplugged to simulate disrupted data communications. Successful results from such data transfer tests showcased the strength of DocAve 4.0's data packet level fault tolerance features.

To address the shortcomings of the native backup and restore options in SharePoint, DocAve 4.0 provides granularity down to the item level in both backup and restore. The DocAve 4.0 interface is shown in Figure 15. It should be noted that DocAve 4.0 also includes modules covering disaster recovery, content management, archiving, auditing, and migration, though we will focus on the backup and restore functionalities. The manager window contains modules for performing either a site, subsite, or item level backup. Each of these modules will allow you to perform a full backup of the entire SharePoint environment, but the differentiation here is on the granularity of the restore. For example, if a full backup of a SharePoint farm were performed utilizing the item level module, the granularity on a restore from that backup would be down to a single individual item, such as a single version of a document, within that farm. If a full backup of a SharePoint farm were performed utilizing the site level module, the granularity on a restore from that backup would be available down to an individual site within that farm.

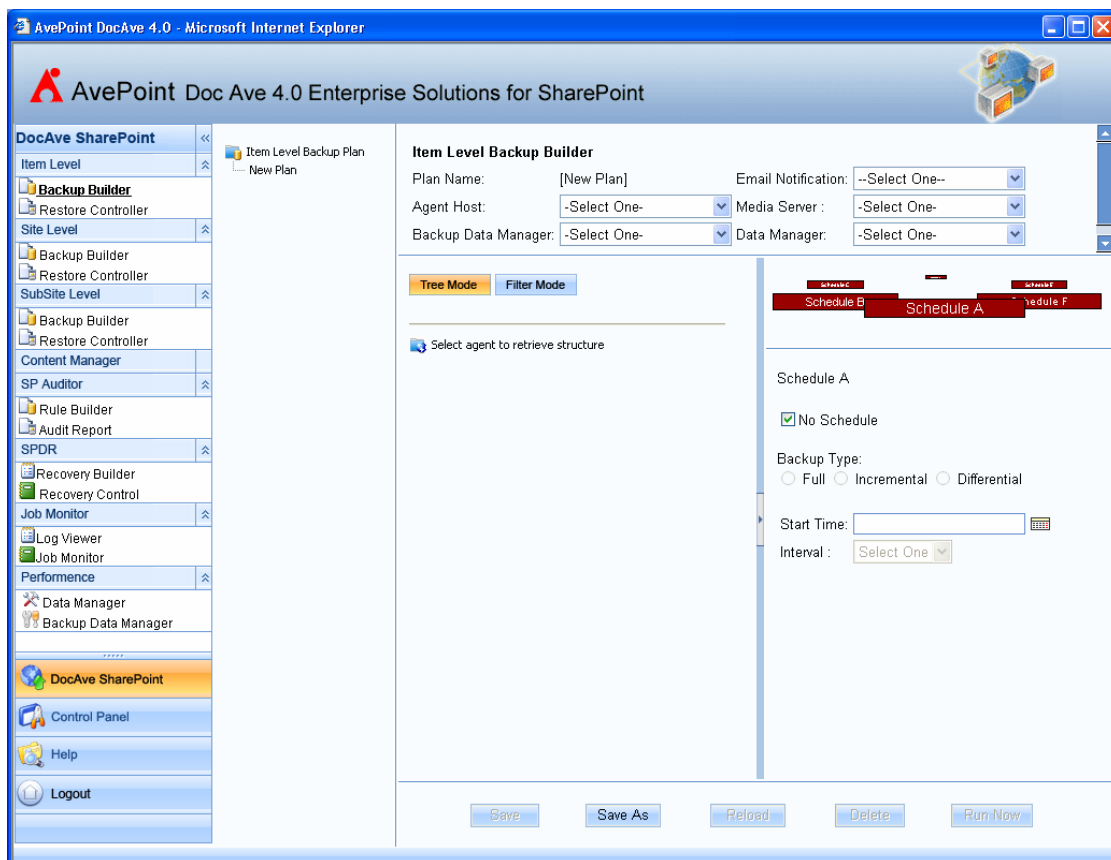


Figure 15: DocAve 4.0 Interface

A natural question that arises is why use the site level module, if the item level module provides more functionality. One reason is speed. While the item level module has been tested to perform at approximately 30 gigabytes per hour per plan, a site level backup is slightly faster at 35 gigabytes per hour per plan. Another common use of site level backup is the moving or promoting of a site. This is accomplished by performing an out of place restore. An administrator might also choose to only perform a site level backup on less mission critical sites or site collections. In addition, the site level module can be used to supplement the item level module. For example, site level backups on content could be performed every few hours, while item level backups might be performed once per day. The subsite level backup provides even more flexibility in developing backup strategies.

The backup and restore modules in DocAve consist of two units, the Backup Builder and the Restore Controller. Backup plans are created utilizing the Backup Builder. Two modes are available to select content to include in a backup plan, tree mode and filter mode. In the tree mode, shown in Figure 16, content is selected by drilling down through the tree structure of the SharePoint environment. The tree can be drilled down to individual folders. Content is selected by placing a check next to the tree item.

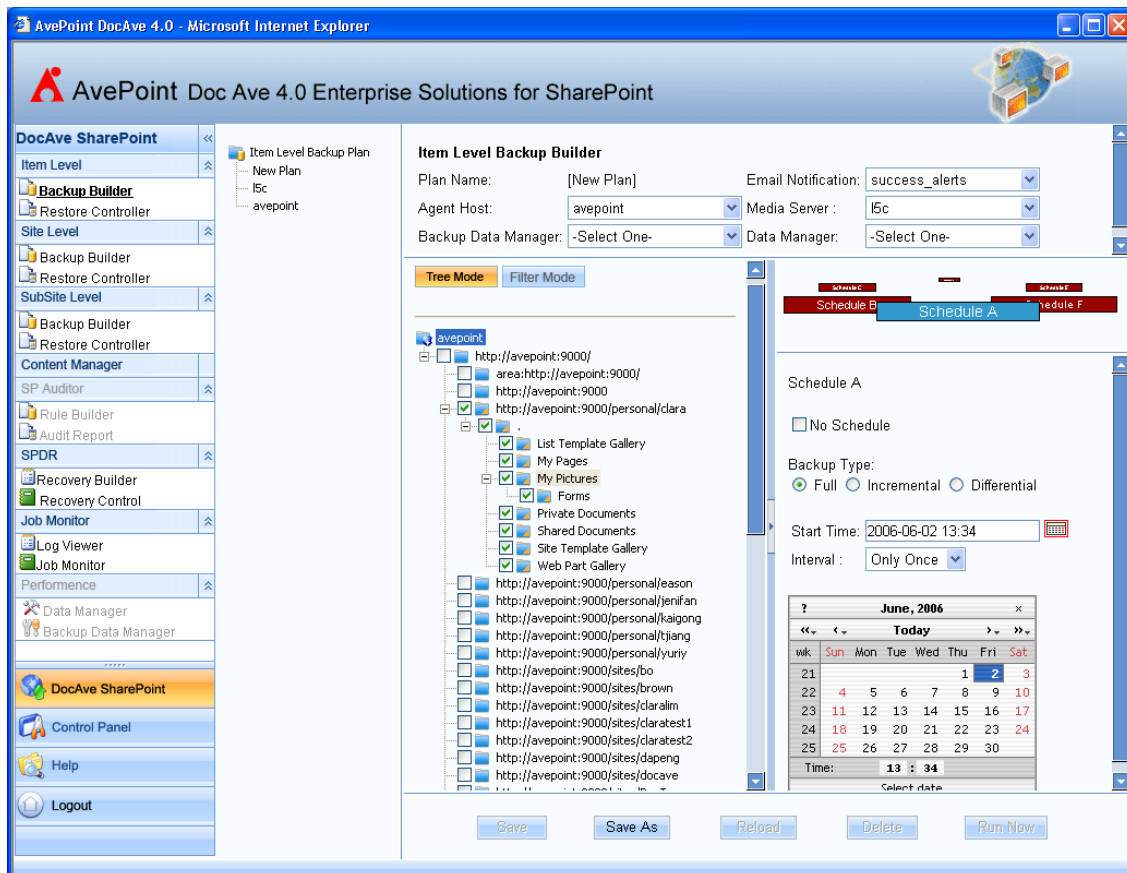


Figure 16: DocAve Backup Builder, Tree Mode

In filter mode, shown in Figure 17, content can be selected using simple and/or/negate logic. Available filters include sites, webs, lists, folders, files, and list items. Using this mode, administrators can quickly develop business rule aware backup strategies. For example, a plan can be created that performs a backup of the HR and Marketing SharePoint sites on a daily basis, but excludes any mp3 files on those sites. These two modes of designing backup plans offer significant advantages over the native backup options. Instead of being limited by the restraints inherent in the native backup tools, administrators can design plans to meet specific business needs.

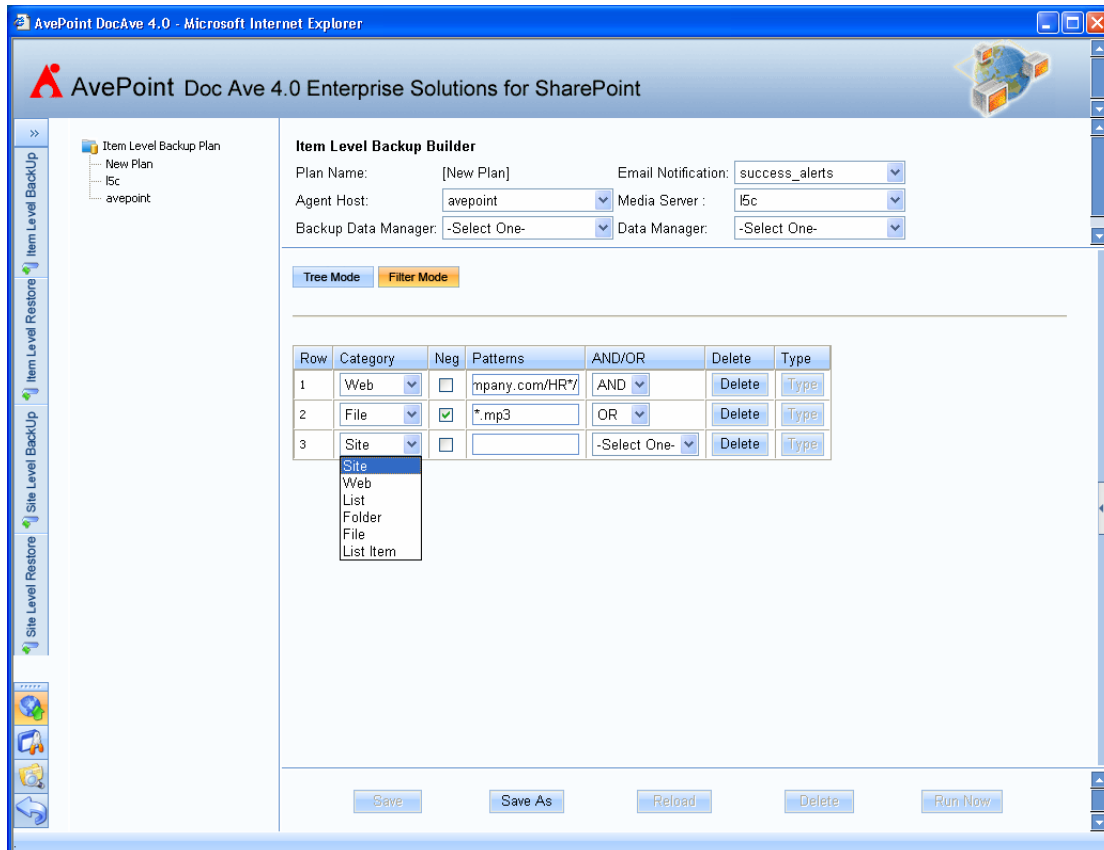


Figure 17: DocAve Backup Builder, Filter Mode

While the native backup tools require administrators to write custom programs to schedule backups, the DocAve Backup Builder contains a simple scheduling agent. This is shown in Figure 16. Each backup plan can contain up to 6 independent schedules within the plan, and numerous backup plans can be constructed. This allows for near limitless scheduling options.

In addition, while the native backup tool offers only full and incremental options, DocAve also offers a differential option. An incremental backup performs a backup of any changes since the last completed backup. If a full backup is done weekly followed by daily incremental backups, the full backup plus all the incremental backups are required. If an incremental backup in the chain is lost, then later incremental backups will lose their connectivity. With a differential backup, any changes since the last full backup are captured. Thus all that is required is the full backup and the last differential backup.

When performing a restore, the Restore Controller, shown in Figure 18, allows the administrator to search the backup files. Content to be restored can be selected either through a tree mode view of the backup contents or by using search mode. In tree mode, granularity is available down to the item level, for example a specific individual version of a single document can be selected to be restored. The restore can also be started immediately or scheduled to occur at a later time. The restore can also be done in place, or back to the location it was originally when the backup was performed, or out of place, which can be anywhere within any SharePoint instance.

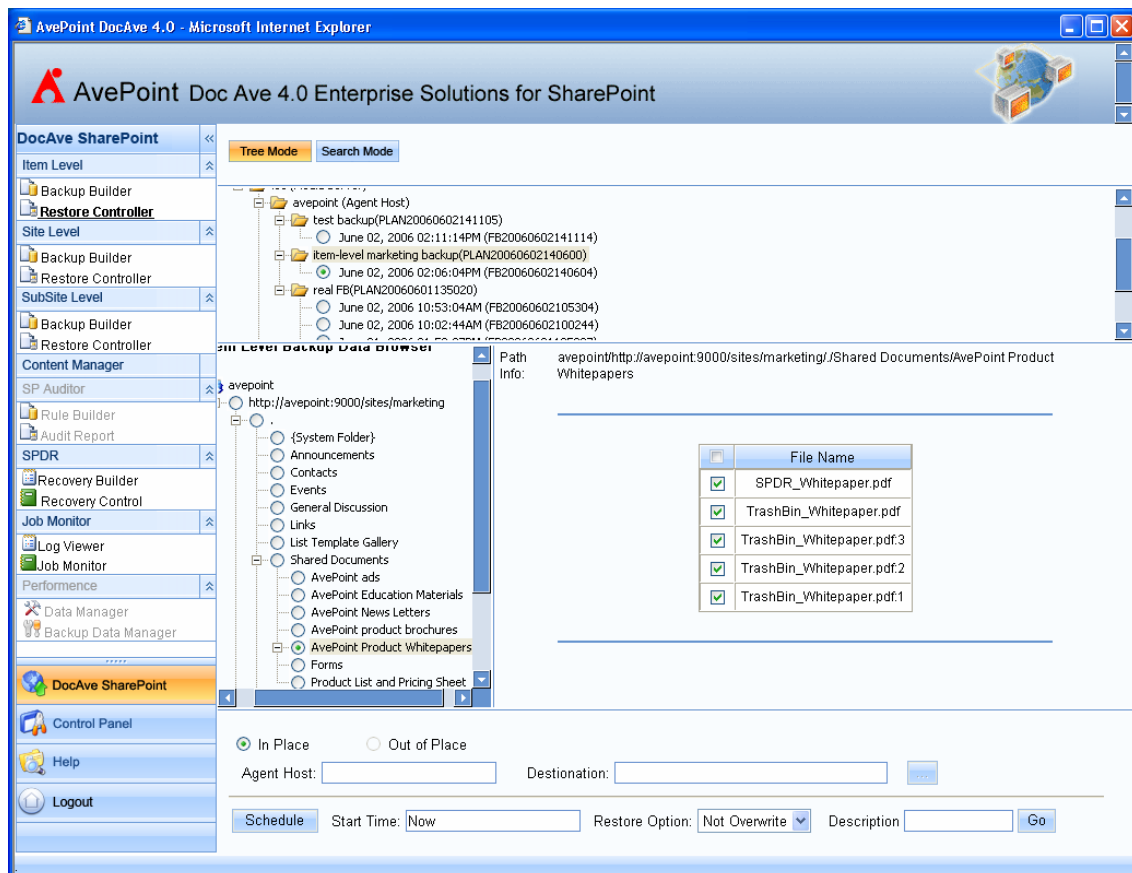


Figure 18: DocAve Restore Controller

Two further enhanced features that DocAve 4.0 offers over the native backup tools are notification and logging. The notification settings, shown in Figure 19, can be set up to notify administrators when jobs are completed. Either a summary or a detailed report can be chosen to be sent. Multiple recipients can be entered for each report.

The screenshot shows the 'AvePoint DocAve 4.0 Enterprise Solutions for SharePoint' web interface. The left sidebar contains a 'Control Panel' menu with options like Data Management, Control Services, Reporting, Email Notification, System, and License Management. The main content area is titled 'New' and contains the following form fields:

- Name: New
- Application Type: -Select One-
- Outgoing Mail Server (SMTP):
- Port: 25
- Sender (e.g. kirbyg@avepoint.com):
- User Name on SMTP:
- Password on SMTP:
- Confirm Password:
- Summary Report Recipients: For example: John@msn.com
- Summary Report Notification Level: -Select One-
- Detailed Report Recipients: For example: John@msn.com
- Detailed Report Notification Level: -Select One-
- Message Format: -Select One-

At the bottom of the form are buttons for 'Save As', 'Save', 'Clear', 'Delete', and 'Test Account Settings'.

Figure 19: DocAve Notification Settings

The DocAve Log Manager, shown in Figure 20, allows for logging of both DocAve itself and the backup and restore processes. Various logging options are available. The exact logs and their level can be selected through the Log Level Settings. The amount of log data to retain can be controlled via Log Data Pruning.

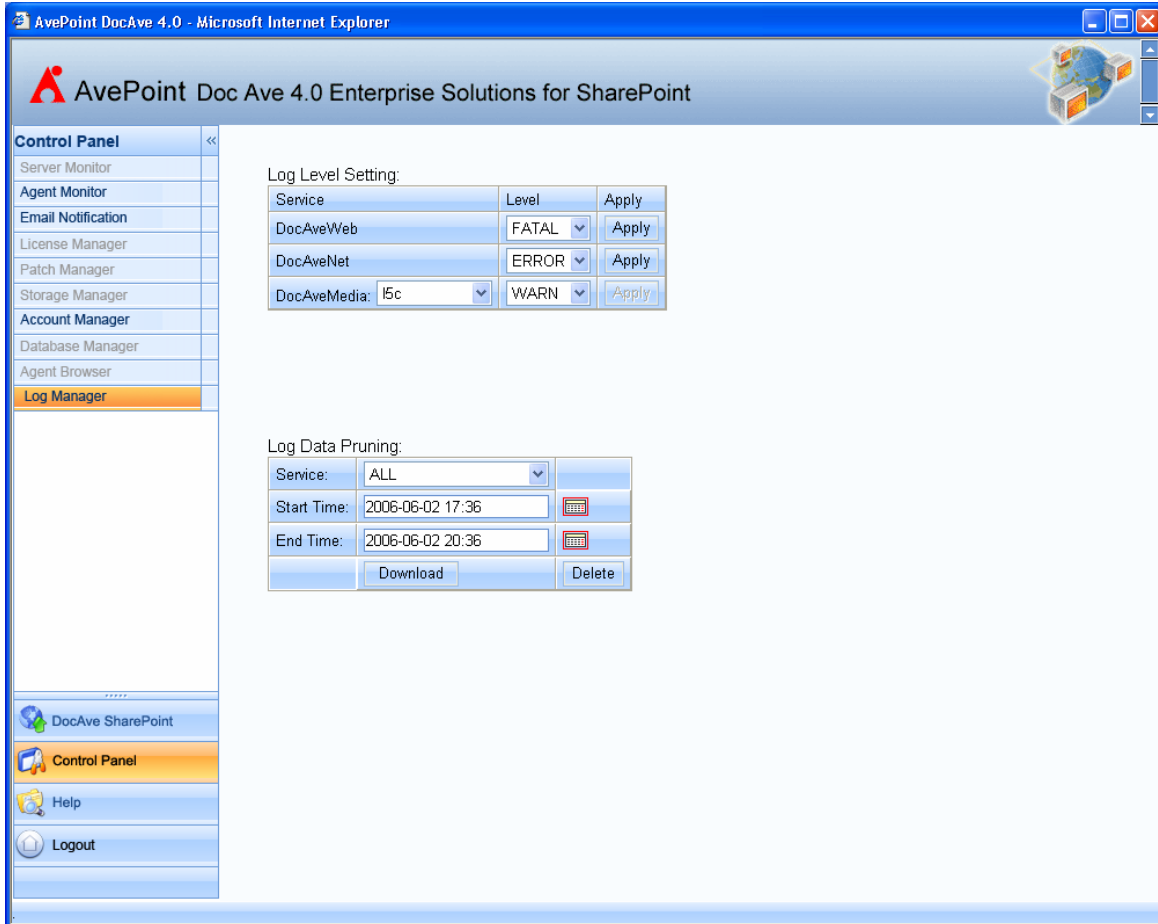


Figure 20: DocAve Log Manager

One further significant enhancement over the native backup tools is backup data retention management. With the default SharePoint backup tools, the administrator needs to manage the backup files manually, deleting files as they become out of date and obsolete or writing custom programs to do so. DocAve 4.0 offers the ability to manage backup data retention policies. By adjusting the Data Retention Settings, shown in Figure 21, administrators can specify exactly how much backup data to keep. Cleaning of old backup data can be triggered by either the start or completion of a specific type of backup job. Data Coalescing is also available through the Data Retention settings. Data Coalescing merges or collapses both incremental or differential backup data onto the previous full backup. This saves disk space needed to store backup data and can allow administrators to effectively stretch out backup cycles without the added risk of losing incremental data.

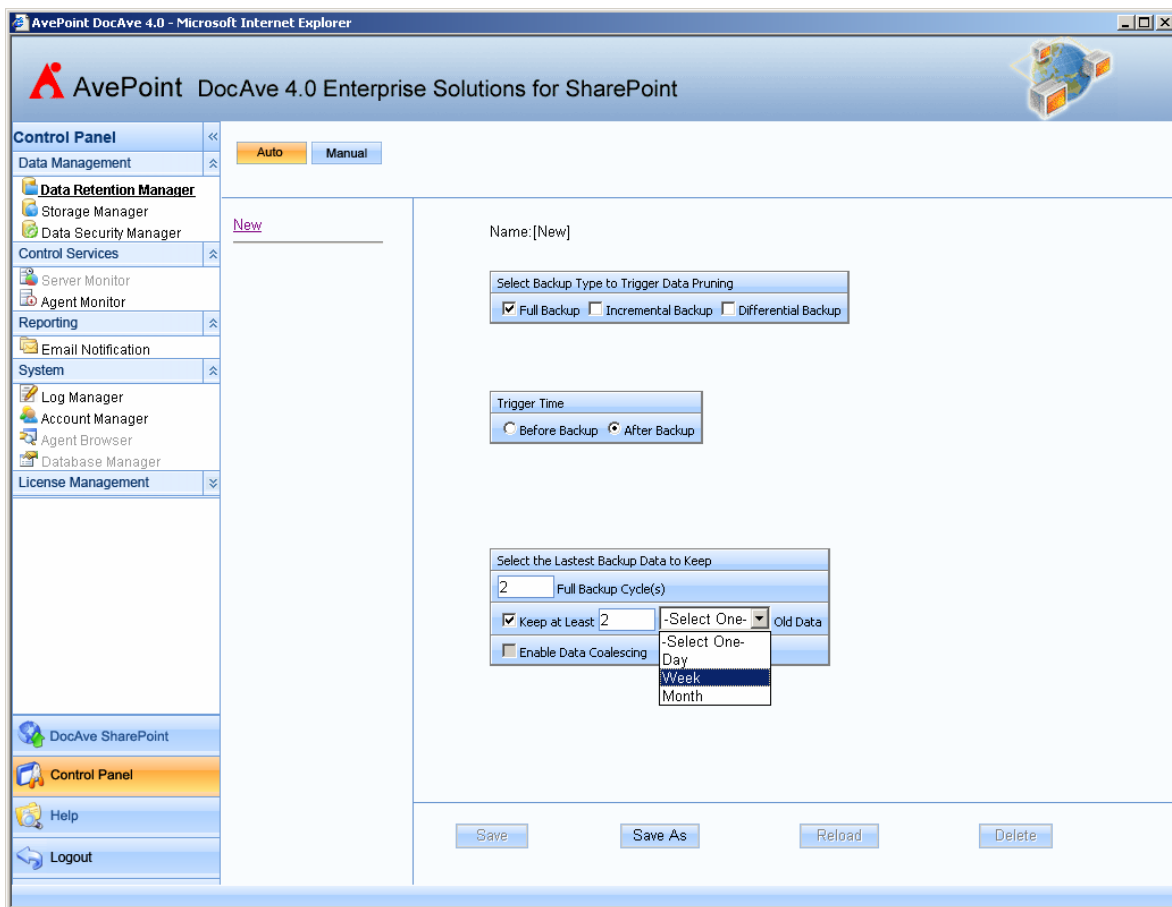


Figure 21: DocAve Data Retention Settings

Conclusion

In conclusion, while Microsoft Office SharePoint Server 2007 does offer some additional features and enhancements over Windows SharePoint Services 2003 when it comes to backup and recovery, in the end these improvements do not change the fact that a third-party tool is required to provide full robust automated backups at all the necessary levels. DocAve offers numerous enhancements over the native backup and restore tools and addresses all of the shortcomings inherent in these native tools.

Other AvePoint Solutions for MOSS 2007

Besides item-level backup and recovery for MOSS 2007, AvePoint offers other value-added solutions to address the SharePoint administrators' need for enterprise-strength SharePoint item-level archiving, auditing, and disaster recovery. Please visit us on the web or contact an AvePoint representative near you to find out more information.

- DocAve 4.0 Archiver
- DocAve 4.0 Auditor
- DocAve 4.0 Content Manager
- DocAve 4.0 Migrator
- DocAve 4.0 SPDR (SharePoint Disaster Recovery)
- DocAve 4.0 SharePoint TrashBin

A 30-day free trial of all AvePoint solutions is available for download via the web at <http://www.avepoint.com/download/download.htm>

Contact Information

Headquarters

AvePoint, Inc.
370 Campus Drive, Suite 104
Somerset, NJ 08873

Tel: (800) 661-6588

Fax: (732) 271-8655

Email: Sales@AvePoint.com

For other AvePoint locations, please visit us on the web at <http://www.AvePoint.com>