

February 2011

The Essential Guide to Migrating SharePoint Content

By Colin Spence



SPECIAL ADVERTISING SUPPLEMENT TO SHAREPOINTPRO CONNECTIONS

SharePoint migrations provide many challenges, but if done properly they can yield significant benefits to an organization. This Essential Guide provides an overview of the options you should consider when migrating to SharePoint 2010, including assessing readiness to migrate, an overview of the standard Microsoft options, and the advantages of using AvePoint's DocAve SharePoint Migrator tool. The guide also covers DocAve compatibility with non-SharePoint data repositories.

The term migration summons forth images of birds or other animals moving in an orderly fashion across great distances, and is quite apt when talking about the data that needs to be moved. You can argue that an In-Place Upgrade isn't technically a migration, but I've found that these are few and far between and generally don't meet the needs of most organizations (I discuss this below). While many different types of SharePoint migrations are possible, due to the large number of SharePoint-branded products that have been released by Microsoft over the past decade, this guide focuses on the most typical migrations occurring today, which tend to be from SharePoint 2007 to SharePoint 2010.

Although SharePoint 2003 (2.0) and even SharePoint Portal Server 2001 (1.0) environments are still in production, they are less common and pose additional challenges due to the age of the software. For example, no direct migration path from the 1.0 or 2.0 versions of SharePoint technologies, including SharePoint Portal Server 2001, SharePoint Portal Server 2003, SharePoint Team Services, and Windows SharePoint Services 2.0 to SharePoint 2010 is supported using Microsoft out-of-the-box tools. Using Microsoft techniques, the only way to migrate these older SharePoint environments to 2010 is to upgrade the servers and sites to SharePoint 2007 first, and then follow one of the valid migration paths. AvePoint offers tools to migrate SharePoint 2003 content to SharePoint 2010, but these tools aren't reviewed in detail in this guide. However, the interfaces and capabilities of these tools are similar to the SharePoint 2007 to 2010 migration tools discussed.

Assessing Reasons to Migrate

Without lingering unduly on the topic, an organization should have clear reasons to move from one SharePoint platform to another. Just because a new version of SharePoint (or any other software) is on the market doesn't mean that it makes solid business sense for the organization to migrate. However, once the decision has been made and before too many commitments are made in terms of the timeline, you should review the existing SharePoint environment (assuming 2007 for this guide) to predict the challenges that might lie in the way. By getting a better understanding of the current state of the SharePoint 2007 environment, IT can better decide on the upgrade path and process.

You should review the overall health of the SharePoint 2007 environment. A quick checklist follows:

- Check the server logs on the SharePoint front-end server(s) and the SQL back-end server(s). Experienced server administrators will quickly be able to differentiate between the nuisance warnings and alerts and the errors that could indicate configuration or performance issues.
- Review the sizes of the SQL content databases to have a sense for how much data there currently is, and how much storage space will be needed in the new environment.

- Review the SQL maintenance plans to ensure that log files are being shrunk, and that the database consistency checker is running.
- Test a recent backup of the SharePoint 2007 environment. Just because the backup looks like it has been completing doesn't mean it can actually be restored, and many (most?) organizations don't test restores often enough.
- Compare the production farm to a new, uncustomized farm with a tool such as WinDiff and look for differences.
- Review the web.config files for the production IIS web sites (typically located in the C:\inetpub\wwwroot\wss\VirtualDirectories folder), look for safe controls that aren't Microsoft, look for application keys (for example the text string "PublicKeyToken").
- Look in the GAC (typically the C:\Windows\Assembly folder); third-party DLLs are a sure sign that there might be something custom in the environment.
- Look in the solution store (Central Administration > Operations tab > Global Configuration > Solution Management), because many customizations are deployed through solutions.
- Look in the SharePoint hive, especially in the layouts, images, and themes folders (generally in C:\Program Files\Common Files\Microsoft Shared\Web Server Extensions\12).
- Review the programs from the Control Panel to look for third-party applications.

Service Pack 2 for SharePoint 2007 includes a pre-upgrade check utility that lets administrators check the readiness of their environment for upgrade to SharePoint 2010. This pre-upgrade check runs as an extension to the STSADM command-line tool. Assuming the SharePoint 2007 environment is completely upgraded to SP2, you can run this tool with minimal risk because it's read-only and makes no modifications to any of the files on the server. Run the pre-upgrade check by typing

```
stsadm -o preupgradecheck
```

The pre-upgrade check tool runs through a number of tests and checks the environment for compliance with SharePoint 2010 requirements. It produces a detailed report that outlines which areas of the existing farm are ready for upgrade, and which ones are in need of remediation before they can be upgraded.

Another best practice is to create an Excel grid of the different site collections and sites and then record information of interest that will be useful during the migration. For example, a key piece of information is whether the site or site collection will be migrated. Similar to the process of moving to a new house, a migration is a good time to hire a dumpster and get rid of "stuff" that is no longer of value. Many old sites or even whole site collections haven't been modified in months or years, and IT may determine that this content does not need to be migrated. Instead, IT might choose to use SharePoint backup tools to back up the whole site collection to drive storage or tape, or use the stsadm -o export command to export individual sites.

As mentioned in the checklist above, testing a restore of a SharePoint 2007 site collection or even the whole farm is always a good idea. True, it will take a number of hours to configure one or more new virtual or physical servers, and then install SharePoint and restore the site

collections or databases, but then IT will know it works and hopefully overcome any issues encountered along the way. A test environment is often recommended when organizations are migrating so the production environment isn't impacted.

In-Place Upgrades

It's a fairly rare situation where an In-Place Upgrade is the best approach. That being said, it might be a logical step to take for certain SharePoint 2007 implementations, where the server hardware meets the rigid requirements of SharePoint 2010 and is new enough to be judged as offering the reliability and performance that will meet the collaboration and document management needs of the organization for the next three to five years.

To consider an In-Place Upgrade, the basic requirements for SharePoint 2010 technologies need to be met by the existing hardware. These are

- Windows Server 2008 x64 or Windows Server 2008 R2 x64
- SQL Server 2005 x64, SQL Server 2008 x64, or SQL Server 2008 R2 x64 for the database
- Service Pack 2 for WSS 3.0 and MOSS 2007

In many cases, the existing SharePoint 2007 hardware doesn't meet these requirements, so the option for the In-Place Upgrade is off the table from the beginning. Note that SharePoint 2010 has higher recommendations for RAM and processor capabilities that can also provide challenges, as Figure 1 shows. In many cases where the SharePoint 2007 farm was configured based on recommended Microsoft specifications several years ago, these recommendations might not be met. The products supported by Microsoft for In-Place Upgrades include the previous versions only, as Figure 2 shows.

Finally, the risks involved in performing an In-Place Upgrade generally outweigh the benefits. The basic risk is a failure during the upgrade process, which can happen even in a relatively simple or small SharePoint 2007 environment. Even if the In-Place Upgrade goes perfectly, legacy code will still be on the server, which could cause future problems. For these and other reasons the In-Place Upgrade is generally not "the best" upgrade option for most organizations.

Database Attach Method

For many reasons, the Database Attach process is better suited for most organizations running SharePoint 2007 that want to keep all the data in the new environment and don't want to change the taxonomy significantly. Due to the relative ease of the process, this method is also well suited to organizations that have time constraints and that need to perform the upgrade over short periods of time, such as overnight or over a week-end. When the migration takes place it's done one content database at a time and it's performed by attaching

the content databases from the SharePoint 2007 farm to the new farm and then performing the upgrade.

IT has the opportunity to install newer, typically more powerful hardware to run SharePoint 2010 and thereby gain performance benefits. In addition, you can configure the SharePoint 2010 farm from scratch to meet best practices and leverage many of the new features in SharePoint 2010, such as managed accounts.

Also, the new SharePoint 2010 servers contain no legacy code, which many purists prefer. And knowing that the server has been freshly built and configured for SharePoint 2010 use helps many SharePoint farm administrators sleep better at night.

Note that in the Microsoft "official" steps (<http://technet.microsoft.com/en-us/library/cc262483.aspx>), the Database Attach Method involves configuring the new SharePoint 2010 farm, then detaching the content dbs in SQL Server Management Studio, and then taking the farm offline. Then the content dbs are attached in SQL Server Management Studio for the new SharePoint 2010 farm to upgrade the content. As the next section points out, one the Microsoft "hybrid" upgrade methods can reduce the down-time when compared to the pure Database Attach Method.

Visual Upgrade Option

Visual upgrade is an option in SharePoint Products Configuration Wizard for both the in-place or Database Attach upgrade, where the administrator can choose to keep the existing look and feel of SharePoint 2007 sites or upgrade to the SharePoint 2010 look and feel. In some cases, it makes sense to not immediately upgrade to the new SharePoint 2010 look and feel because end users might not have been trained on the new features and interface, or because IT wants to upgrade select sites to ensure that nothing breaks in the process. Note that My Sites immediately assume the new SharePoint 2010 look and feel, so the visual upgrade options don't apply to My Sites.

Hybrid Upgrade Methods

Microsoft also offers the option of "hybrid" upgrades, labeling them as the Read-Only Database approach and the Detach Databases approach (see the Technet URL referenced in the previous section). As this section discusses, the Read-only Database approach is really just a slight variation on the "pure" Database Attach method.

In the Read-only Database approach, the new SharePoint 2010 farm is configured, then the SharePoint 2007 environment is set to read-only, and the content databases are backed up from the old farm and restored to the new farm. The databases are then attached in SQL Server Management Studio and the upgrade process runs, upgrading the content dbs. In general, this is a preferred method for performing a SharePoint 2007 to 2010 upgrade. The risks are minimized because the

Deployment Size	Processor Minimum requirement
Small Deployments	64-bit, four cores
Medium Deployments	64-bit, eight cores
Combined size of content databases	RAM recommended for SQL Server
Minimum for small production deployments	8 GB
Minimum for medium production deployments	16 GB
Recommendation for up to 2 TB	32 GB
Recommendation for the range of 2 TB to a maximum of 5 TB	64 GB

Figure 1: SharePoint 2010 Processor Requirements and SQL Server RAM Recommendations

Product	SharePoint Foundation 2010	SharePoint Server 2010 Standard Edition	SharePoint Server 2010 Enterprise Edition
Windows SharePoint Services 3.0 with SP2	X		
Microsoft Office SharePoint Server 2007 Standard Edition		X	
Microsoft Office SharePoint Server 2007 Enterprise Edition			X
SharePoint Server 2010 Trial		X	X

Figure 2: Products Supported by Microsoft for In-Place Upgrades

SharePoint 2007 farm is left operational and is simply in read-only mode.

The second hybrid approach is the Detach Databases method, where the SharePoint 2007 farm is taken offline, the content databases are detached, and an in-place upgrade is performed on the SharePoint 2007 farm, which will upgrade the other databases but not the content databases that have been disconnected. The administrator then reattaches the content databases, upgrading them in the process. This essentially breaks the process into two steps, which will alert IT to any problems with the upgrade of the non-content databases prior to attempting to upgrade the content dbs.

This second hybrid approach is generally seen as overly complex without immediate apparent benefits, because if the upgrade fails after the content databases have been disconnected, the farm will be down until it can be fixed or rebuilt.

DocAve SharePoint Migrator

Many organizations take time to investigate third-party migration tools. There could be a number of reasons for this research; for example, some organizations test one or more migration methods discussed above, only to find that the upgrades or migrations don't complete successfully and then decide to investigate third-party options. Other organizations might have needs not met by the out-of-the-box options, and they look for these features in third-party tools. The following sections cover a number of capabilities provided by the DocAve SharePoint Migrator tools that provide value to organizations with more complex migration needs.

Support for Multiple Data Sources

A common requirement in many companies is the need to migrate content from sources not supported by the out-of-the-box migration tools in SharePoint Server 2010. As an example, some organizations have multiple versions of SharePoint implemented, and they want to consolidate and migrate to a single, high-performance SharePoint 2010 farm. Such a migration is complex without the aid of third-party tools. Take for example an organization with a WSS v2 farm (Farm1) and a SharePoint Server 2007 (MOSS 2007) farm (Farm2) that wants to consolidate content to a SharePoint Server 2010 farm (Farm3). An out-of-the-box upgrade would require first upgrading the Farm1 WSS v2 to WSS v3, then upgrading WSS v3 to MOSS 2007, and then upgrading that to SharePoint Server 2010. And that's just Farm1. Farm2 would need to be upgraded to SharePoint Server 2010, and *then* the content needs to be migrated from Farm1 and Farm2 to Farm3, which requires additional steps. Throw into the mix a need to migrate content from Exchange Public Folders, one or more file shares, and content in other third-party document management or enterprise content management (ECM) tools such as Documentum or LiveLink, and the process becomes even more complex.

The DocAve Migrator product family supports the following data sources:

- SharePoint 2001
- Windows SharePoint Services (WSS) v2
- Microsoft SharePoint Portal Server (SPS) 2003
- Windows SharePoint Services (WSS) v3
- MOSS 2007
- Exchange Public Folders
- File Systems and Networked File Shares
- Documentum eRoom v6.0 and above
- EMC Documentum v6.5 and above
- Lotus Notes v.6.5 and above
- Open Text Livelink 9.5 and above
- Open Text Vignette v7.x and above
- Oracle/Stellant v7.x and above

DocAve products support the following migration targets:

- WSS v3, MOSS 2007
- WSS v2, SPS 2003 (for Exchange Public Folders and SharePoint 2001)
- Microsoft Exchange 2007/2010 (only for the Lotus Notes to Exchange Migrator tool)
- SharePoint Foundation 2010, SharePoint Server 2010

Note that the appropriate combination of DocAve clients must be purchased and installed and you need to budget for those costs. The DocAve administrative console must be installed, and the agents must be installed on one source farm server and one destination farm server. To offset this cost, typically the costs of outside consulting or increasing staffing levels during the migration project are reduced by use of DocAve migration tools. The different DocAve content migration modules are available for evaluation purposes to test with limited amounts of data (generally 1 GB), which is helpful, because "try before you buy" is a key step in product evaluation.

DocAve SharePoint Migrator Walk Through

This section provides an overview of the process you need to go through once the DocAve tools are installed on the source and destination servers. The basic options that are accessible to the administrator are SharePoint 2003 to 2007, SharePoint 2003 to 2010, and SharePoint 2007 to 2010 modules. For this guide, I'm concentrating on SharePoint 2007 to 2010 migrations. I'll review the tools needed to build a migration plan, as well as some of the more powerful features available from the Mapping and Filter tools.

Much of the "magic" happens in the Plan Builder (Figure 3). As the name suggests, the Plan Builder lets you configure the details of the migration. You can schedule the migration to run later, or immediately. DocAve allows for the definition of Filters and Mapping Settings separately; these definitions are then referenced from the Plan Builder. This approach lets

you configure these standards separately and reuse them in the Plan Builder, rather than having to define all these variables repetitively. The following section walks through the primary capabilities of these tools.

Figure 3 shows the Migration Settings tab in the Plan Builder. The interface is intuitive, with the source content defined in the left-hand pane, and the destination defined in the right-hand pane. You simply choose which site collections, sites, lists, or even items that are to be migrated. Then you select a destination for the items. Figure 3 shows an example where three sites from the source SharePoint 2007 farm are selected and will be migrated beneath a subsite (subsite1) in the destination SharePoint 2010 farm. Note that you can migrate the source sites to specific content databases, which you can also set from this interface.

You can then schedule when the job will run; for example, you can schedule the job to run late at night, or you can click the Run Now button to run the job immediately. End users won't be impacted because SharePoint object model is used for all activities, so downtime is typically not needed. Note also that you can schedule jobs to run at predefined intervals (such as weekly) on a Full or Incremental basis. Many organizations find value in being able to migrate a portion of content, test the results, and then migrate any changes after through testing; for example, a week later. The Incremental option allows for this level of migration.

Links are provided to the Filter and Mapping tools, as well as shown under the Migration Settings tab in Figure 3. Mapping setup configuration provides options in the following areas:

- Common Settings
 - Site Collection / Site
 - List
 - Permission
 - Alerts
 - Characters

- Permission Configuration
- List Level Configuration
- Site Level Configuration

While I don't have room to review all these options, some key options in the Mapping Setup let you make the following decisions or modifications:

- Preserve the look and feel of existing SharePoint sites, and allow end users to update their sites' user experience. This lets you set the Visual Upgrade settings from the DocAve interface.
- Restore all permissions, only site permissions, only list permissions, only folder permissions, only item permissions, or do not restore any permissions. These different options let you determine which combination of permissions are restored. In some cases, the use of unique permissions on items, folders, or lists can cause administrative headaches, and IT might want to simplify overall governance of SharePoint by no longer allowing or enforcing item- or folder-level permissions.
- Migrate alert of the list, library, folder, item, and document levels.
- Set maximum length of the SharePoint URL, folder name, or file name. Out of control URLs can cause problems and IT might use the migration process to truncate these.
- Map to a new destination domain name. Reorganizations or acquisitions might require a change in domain name, which is easily mapped during a DocAve migration.
- Map source users to new destination names. User names might also change when migrated to a new domain, which is also easily configured.
- Change source columns to a new column type, such as a Managed Metadata column on a list level or on a site level. A key list that is used by many users on a daily basis can be tuned or modified to comply with

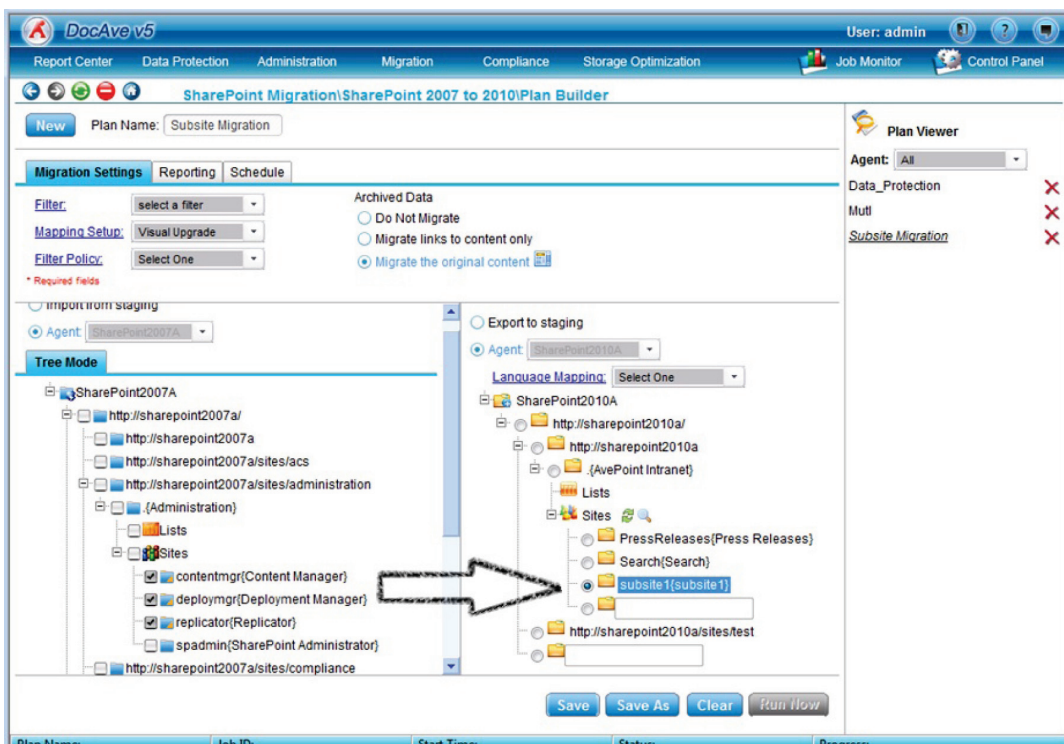


Figure 3: DocAve SharePoint Migrator Plan Builder Screen

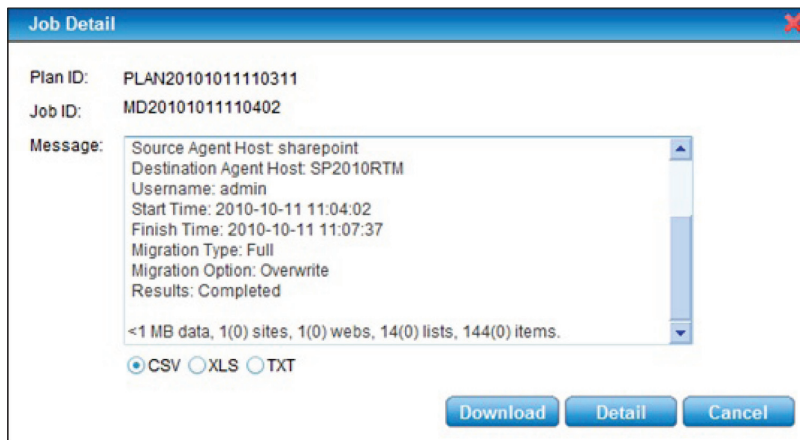


Figure 4: Job Detail Viewed from the Job Monitor

best practices (such as replacing a Single Line of Text column with a Managed Metadata column).

- Replace a source template ID with a new template ID on a list or site level. This could be used to change sites created from a Team Site template to a customized template to add functionality to corresponding sites.

Filter options include:

- Time Range settings based on created time or modified time
- Version filters to set the number of versions captured, and whether major only or major and minor versions are captured.

Although this partial list of Mapping options and Filter options might seem overwhelming for organizations wanting to use the migration process as a means of cleaning up the data that ends up in the SharePoint 2010 lists and libraries, they can be very valuable. By using the filters, IT could migrate content that has been changed in the last two years to a new production site, migrate older content to an archive site, and then—to further reduce total data migrated—migrate only the last major version of documents.

As you'll also see in Figure 3, the Plan Builder provides several options for managing archived data. If the DocAve Storage Optimization tools were used to archive data, you could migrate this content without needing to first restore it to the SQL Server content databases. This is also true if you archived content by using the free DocAve Extender. The DocAve Extender offloads SharePoint BLOBs to file-based storage and utilizes Microsoft EBS and RBS interfaces. Once in SharePoint Server 2010, you can then migrate from EBS BLOB API technology to the newer RBS technology on SharePoint Server 2010 for content externalization.

Reporting on the results of the plan are very flexible, and the logs of the migrations are quite detailed. This allows IT to be alerted when migrations finish, and find complete details of the migration from the DocAve interface. Figure 4 shows the basic job detail available from the Job Monitor after a job has run. Note that the basics of the job are summarized, such as start and finish time, as well as type of migration and results in terms of site, web, lists, and items migrated. Clicking on the Detail button lets you peruse the results item by item, which can be very valuable when you're migrating mission-critical data.

Note that in these discussions of the DocAve Migrator tools, no mention was made of In-Place Upgrades, Database Attach Method, or Hybrid upgrades and migrations. DocAve

needs a source (which can be a variety of products, as discussed above) and a SharePoint 2010 farm (or other target, as discussed above) and then its engine takes care of migrating the content to its new home without forcing additional work on the part of the IT team responsible for the migration. The new source does need to be configured and ready to receive content, but the DocAve migration tools make this very easy and intuitive. Should the organization choose to cull or trim older data, or reduce the number of versions of documents migrated, or change templates used, or levels of permissions migrated, these options are all possible through the DocAve tools.

Going Forward

A basic set of preparation tasks to perform prior to an upgrade or migration to SharePoint 2010 are provided, along with the hardware and operating system requirements for SharePoint 2010 that IT should be familiar with before embarking on a SharePoint upgrade or migration. By performing these tasks and being familiar with the requirements for SharePoint 2010 IT can start to get a sense of the organization's best migration and upgrade path.

I've provided a solid grounding in the out-of-the-box upgrade and migration options. I've also discussed the "official" Microsoft methods and reviewed the DocAve SharePoint Migrator toolset with an eye toward the tools that you would have at your fingertips.

For organizations with smaller, simpler SharePoint 2007 implementations, the hybrid Read-Only Database approach is typically the recommended method from my experience because it's the least impactful on the source data. Should this process fail, or if an organization realizes that the more sophisticated toolset provided by AvePoint in the DocAve SharePoint Migrator product would ease the migration process, I recommend that the product be evaluated and tested.

Colin Spence, an MCP and a MCTS in SharePoint and a Partner at Convergent Computing, performs in the roles of Senior Architect, Practice Manager, and Technical Writer for the organization. He focuses on the design, implementation, and support of Microsoft-based technology solutions, with a current focus on Microsoft SharePoint technologies. He has been implementing SharePoint-based solutions since 2003 and has over 20 years of experience providing IT-related services to a wide range of organizations. He has worked with AvePoint products since 2007. Colin has authored several best-selling books on SharePoint products, including SharePoint 2010, contributes to numerous publications and speaks regularly on SharePoint technologies.