

Windows[®] IT Pro

Metalogix Content Migration and Upgrade for Microsoft SharePoint Server

By Michael Noel

Published: October 2008

metalogix

→ Contents

Why Choose the Content Migration Approach?	2
Out of Box Migration Approaches	2
In-Place Upgrade	2
Overview of the In-Place Upgrade Approach	2
Challenges with the In-Place Upgrade Approach	2
Metalogix comparison with the In-Place Upgrade Approach	2
Gradual Migration	3
Overview of the Gradual Migration Process	3
Challenges with Gradual Migration	3
Metalogix Comparison with Gradual Migration	4
Database Upgrade Approach	4
Overview of the Database Upgrade Approach	4
Challenges with the Database Upgrade Process	4
Metalogix comparison with the Database Upgrade Process	5
Benefits of Content Migration with Metalogix	5
Allow content owners to migrate (Migrate based on permissions)	5
Granular Migration	6
Content Filtering Synchronization/ Incremental Migration	6
Reorganize your sites during migration	6
Manage Site Collections and Content Databases	6
Server Consolidation	6
Server Virtualization	6
Ongoing Value of SharePoint Site Migration Manager	6
Site Structure Management	6
Incremental Updates	6
Conclusion	7

metalogix

Metalogix Content Migration and Upgrade for Microsoft SharePoint Server

Many organizations have deployed Enterprise Content Management and Collaboration solutions on Microsoft's SharePoint Products and Technologies in recent years. The latest SharePoint 2007 product line, Windows SharePoint Services (WSS) 3.0 and Microsoft Office SharePoint Server (MOSS) 2007, provide for out-of-the box collaboration, document management, and advanced search capabilities that have further accelerated the move to these technologies.

The popularity of Windows SharePoint Services 2.0 and SharePoint Server 2003 had led many organizations to deploy these versions to various business units and, in some cases, on a large scale across their organizations. While the capabilities of SharePoint 2003 were robust, with the release of SharePoint 2007, key functionality improvements in the software—such as vastly improved search, better document management and workflow, and restore capabilities with the Recycle Bin—made many of these organizations put plans into place to upgrade their environments to the new version.

Unfortunately, the migration path provided by Microsoft does not always address the needs of these organizations. The out-of-the-box process, outlined in this white paper, suffers from some fairly significant limitations and cannot be used in several common scenarios. Organizations looking to upgrade to SharePoint 2007 have subsequently been stuck in a situation where they can't make the move to the new technology without jeopardizing the existing knowledge built into the existing platform.

Fortunately for these organizations, Metalogix has developed the SharePoint Site Migration Manager (SSMM) tool to assist with the migration of SharePoint content from legacy SharePoint 2003 technologies to SharePoint 2007, allowing them to take advantage of the features within the new platform without unnecessarily risking their

data and intellectual property. SSMM has some rather significant advantages over existing migration tools and techniques, allowing for flexibility of migration between versions and farms, granular migration, reorganization of sites during the process, and many other enhanced capabilities. Unlike some content migration options, SSMM uses only fully supported features of SharePoint, so using the tool will not affect existing support agreements with Microsoft.

This paper discusses both the Microsoft migration approaches as well as migration using Metalogix's SSMM tool. Particular emphasis is placed on describing common migration scenarios and how those scenarios can be addressed using the different migration approaches. The paper itself can be used not only to better understand the capabilities of SSMM, it can also be used by SharePoint administrators specifically tasked with the role of migrating SharePoint content.

Why Choose the Content Migration Approach?

For organizations with a significant investment in an existing SharePoint 2003 environment, the answer is simple: documents and other vital intellectual property must be transferable to SharePoint 2007 if the project is to be successful. In many cases, significant customization of SharePoint 2003 sites has also been performed, tailoring the environment to the needs of the organization. These changes can make it difficult to migrate content successfully.

Out-of-the-Box Migration Approaches

Microsoft has created three unique out-of-the-box migration approaches for migration from SharePoint 2003 to SharePoint 2007. Each of these approaches is discussed in this paper, and comparisons to Metalogix SSMM functionality are noted for each section.

NOTE: SSMM does not directly support migration from SharePoint 2001 versions, such as SharePoint Team Services and SharePoint Server 2001. The content within any existing SharePoint 2001 Sites must first be migrated to SharePoint 2003 using Microsoft's SPOUT/SPIN tool, downloadable from <http://www.microsoft.com/downloads/details.aspx?familyid=FEEA83B4-7F46-4040-BCE2-5341E5C1B107&displaylang=en>, before it can be migrated with the SSMM tool.

In-Place Upgrade

Overview of the In-Place Upgrade Approach

The In-Place Upgrade approach--provided by Microsoft--

is a method by which an individual server is upgraded in place with all content to SharePoint 2007. The process upgrades all site content on the server at the same time.

The following high-level steps are included in an in-place upgrade:

1. Run the Pre-Upgrade Scan Tool to uncover potential issues with Sites and Workspaces.
2. Fix any issues uncovered by the Pre-Scan Tool before migrating Sites. This can include removing unsupported web parts, re-ghosting sites, and other fixes identified in the PreUpgrade Report.
3. Install .NET Framework 3.0 on the single Front-end server.
4. Install the SharePoint 2007 binaries on the server, selecting In-Place Upgrade. Wait for all Site Collections to be upgraded.

Challenges with the In-Place Upgrade Approach

There are several significant limitations to the In-Place Upgrade approach. First and foremost, it is the riskiest migration strategy because there is no fallback strategy in the event of problems. Several other key challenges exist with this strategy:

- The In-Place Upgrade Process can be used only on a single-server farm.
- It can only be used to migrate from WSS 2.0 to WSS 3.0 or from SPS 2003 to MOSS 2007. No cross-platform migrations are allowed.
- The environment is completely down during the process.
- Unghosted sites (sites that have been modified with tools like FrontPage 2003) may not upgrade and could potentially fail. Other site customizations, which include things as simple as Discussions List and deprecated features such as bucketwebs, will not migrate.
- Running out of disk space during the in-place upgrade process can torpedo the entire migration and leave it in a half-migrated state. This is a fairly common scenario as the process requires more than double the disk space that is currently used.
- The process does not allow for changing or modifying existing database versions.

Metalogix comparison with the In-Place Upgrade Approach

SSMM offers some rather significant advantages over the In-Place Upgrade approach. For example, the ability to migrate site content in waves, rather than all at once, is vital. Additionally, SSMM does not suffer from the same limitations, hardware requirements, and database considerations as the In-Place Upgrade does. Indeed, flexibility is the key feature trait within SSMM, as SharePoint

Admins are not handcuffed by restrictions when migrating content.

SSMM migrates content using the native SharePoint API, FrontPage RPC, and SharePoint web services, so the source content is not modified in any way and Microsoft support agreements are not violated. This approach allows for a great degree of flexibility in migration approach, dramatically reducing the risk inherent in an In-Place Upgrade. Key differentiators between the In-Place Upgrade approach and migration using SSMM include the following:

- SSMM uses a familiar, easy-to-use tree-based interface for migration of content. The In-Place upgrade process requires the use of command line utilities, XML files, and other complexities.
- SSMM removes most of the risk associated with performing a migration because it is not data destructive. The In-Place upgrade process requires direct changes to be made to production data.
- SSMM allows for migration to new hardware, software, and different farm characteristics. The In-Place process must be run on the existing SharePoint 2003 server.
- SSMM can migrate the vast majority of sites that fail the In-Place Upgrade Scan Tool, while preserving security, metadata, document versions, and web part content. The In-Place Upgrade process is notorious for failing even when there have been small customizations made to a site.

Gradual Migration

Overview of the Gradual Migration Process

For the vast majority of customers, the in-place upgrade process is too risky and limiting. Indeed, Microsoft often recommends against using the in-place process on all but the smallest servers. For these organizations, Microsoft built a second, more thorough, migration approach that allows content to be migrated gradually, rather than all at once. While an improvement over the in-place option, there are still several serious limitations with this process that must be examined before implementing it as a migration strategy.

The gradual Migration process consists of the following high-level steps:

1. Run the Pre-Upgrade Scan Tool to uncover potential issues with Sites and Workspaces.
2. Fix any issues uncovered by the Pre-Scan Tool before migrating Sites. This can include removing unsupported web parts, re-ghosting sites, and other fixes identified in the PreUpgrade Report.
3. Install the .NET Framework 3.0 on any SharePoint Front-ends.

4. Install SharePoint 2007 onto existing SharePoint 2003 Front-ends, choosing the upgrade using the Gradual Migration Process.
5. Create a new farm and a new web application. The process works by maintaining two parallel Web Applications and farms so that client requests can be directed to the appropriate side, depending on if that particular site collection has been migrated.
6. After installing on all web front-ends, the Site collections can be migrated one at a time, starting with the root site collection. For the root site collection, IISReset must be run after the migration.
7. After all site collections have been migrated, the process is completed by clicking the Complete Upgrade option, which collapses the old SharePoint 2003 web application and infrastructure.

Challenges with Gradual Migration

While vastly superior to the in-place upgrade process, there are still limitations with this approach:

- After the migration is complete, the existing server performance is often slower than it was in SharePoint 2003. This is due to imperfections in the way that SharePoint 2007 is installed on top of SharePoint 2003. For this reason, Microsoft recommends exporting site collections into a new, freshly created farm after the migration is complete.
- Granularity of migration is limited to individual site collections. In other words, you can split the migration, but only by site collection. Because many SharePoint 2003 environments consist of a single site collection, this limits those environments to effectively migrating everything at the same time.
- You cannot migrate between the 32-bit version of SharePoint 2003 and the 64-bit version of SharePoint 2007. The 64-bit version of SharePoint 2007 is highly recommended for new farms, because new versions of SharePoint will be 64-bit only. This presents somewhat of a Catch 22 for organizations considering the Gradual Migration process.
- You cannot change database versions or swap out farm roles during the upgrade process; you can do this only after the upgrade is complete.
- Any existing backwards-compatible SharePoint 2003 Web Storage System Document Libraries (created in SharePoint 2001 and migrated to SharePoint 2003) must have their contents dragged onto a file server and then copied into new SharePoint 2007 document libraries before the migration process can be completed. All metadata is lost during this operation.
- The upgrade process often leaves site collections with broken navigation elements, strange formatting, and malfunctioning web parts. The process itself, while

designed to be flexible, cannot take into account all of the factors and variations in SharePoint 2003 sites, so migrated sites can often have a strange look and feel after the process is complete. This is particularly the case for 'Unghosted' sites.

Metalogix comparison with Gradual Migration

Once again, the key feature difference between SSMM and other options (such as the Gradual migration approach) have to do with the way data is migrated. Both Microsoft approaches use a process in which the data is directly modified on the original server. SSMM, however, assumes that a new farm, on new equipment, is setup and configured for the migrated content, and that a non-invasive process of reading the content and copying it to the new farm is accomplished. At the same time, once the content is migrated, links within the new server will be updated.

While the gradual approach provides for less risk than an in-place upgrade approach, it still suffers from some serious limitations, and often will not work for complex content or large farms. In addition to benefits already mentioned, the following are some of the advantages SSMM has over the Gradual Migration option:

- SSMM can migrate content in a more granular approach, rather than requiring an entire site collection to be migrated at one time, as the Gradual approach requires.
- Content owners can use SSMM to migrate sites as long as they have rights to the content itself. The Gradual migration process requires full farm admin rights to migrate a site collection.
- SSMM allows for migration to new 64-bit servers, new service packs, new hardware, new OSes like Windows 2008, new database levels—all without touching the existing SharePoint 2003 farm. The Gradual migration approach requires the farm to retain 32-bit status, database levels, and server OS versions.
- SSMM gives Admins the ability to reorganize their content, site collection structure, and content database layout, all of which must remain the same with the Gradual migration process.
- The migration of sites using SSMM can be performed from a remote machine or the administrators desktop, while the Gradual migration process must be run directly on the server.
- SharePoint content can be re-templated while it is migrated when using SSMM, something that is not possible with the Gradual migration process.
- SSMM does not require any type of pre-scan tool to be run, such as what is used with the Gradual migration.

Database Upgrade Approach

Overview of the Database Upgrade Approach

A third, often overlooked approach to migrating SharePoint content is to use the Upgrade Database approach. Using this approach, the SharePoint databases from the source SharePoint 2003 farm are backed up and restored onto new SharePoint 2007 farm servers, where they are upgraded in place by SharePoint 2007.

The Database Upgrade approach is similar to the In-Place upgrade method, in that all data is upgraded at the same time. What this approach adds is the ability to move the existing content to a new farm built on new servers, which can have the added advantage of allowing for new OS versions, SQL versions, and 64-bit architecture changes. The following high-level steps are involved with the Database Upgrade approach.

1. Build a new SharePoint 2007 farm for the databases to be restored to. Do not create content within the farm.
2. Run the Pre-Scan Upgrade tool on the SharePoint 2003 farm. Review the results and adjust as necessary.
3. Set existing SharePoint 2003 databases to Read-only status.
4. Backup the existing SharePoint 2003 databases.
5. Restore the SharePoint 2003 databases to the new SharePoint 2007 database role server in the farm.
6. Restore the content database that includes the root site collection in a farm to the new server farm either using the STSADM command line tool or the GUI. The database upgrade process will begin.
7. Restore the additional content databases to the new farm. Check the upgrade status by reviewing the upgrade log.
8. Review the migrated sites and fix errors, broken hard-coded URLs, and legacy document libraries.

Challenges with the Database Upgrade Process

The Database Upgrade migration approach has a big advantage over both previously mentioned approaches in that it allows for new server hardware and software to be utilized for the new SharePoint 2007 farm. In addition, since the source databases and farm are not directly touched, it reduces some of the risk associated with an upgrade, because they can be brought back online and fallback to the SharePoint 2003 farm is possible. That said, however, this approach suffers from many of the same limitations as the In-Place upgrade process, because the mechanisms it uses to upgrade are very similar. In addition, there are some added disadvantages and challenges with this approach:

- The amount of drive space required to perform this type of migration can be three to four times the amount that is currently allocated to the SharePoint 2003 database.

In addition, a large amount of log files are generated during the upgrade, which can have the effect of causing the logs or database drive volumes to run out of space during the upgrade, causing it to fail.

- The Pre-Scan Upgrade tool must still be run with this type of migration, or else the site collections will not upgrade properly.
- The entire SharePoint 2003 environment must be in read-only state during the migration process, and it is recommended that it be shut down.
- The new content databases will not be fresh, newly created SQL databases, but will inherit the structure, compatibility level, and other characteristics of the original SharePoint 2003 databases.
- Site collections and logical structure of SharePoint cannot be changed during this migration.
- All of the limitations with the previous approaches in regards to difficulties migrating customized sites still apply to this scenario.

Metalogix comparison with the Database Upgrade Process

While the Database Upgrade process allows for a new farm to be created, legacy database structure and site structure are still maintained. SSMM, on the other hand, copies content, permissions, versions, metadata, and views from a source farm to a unique and freshly created target farm. In addition, several other key advantages of SSMM exist over the Database Upgrade process:

- No unnecessary drive space is consumed on the target farm during the migration process.
- New SQL Databases can be created with compatibility levels and other settings that match the new version of SQL Server used on the target farm. For example, the new target farm could utilize SQL Server 2008 x64 Enterprise Edition for the Database role, and all new databases would be created with 2008 settings, rather than having migrated databases lifted from legacy SQL 2000/2005 servers.
- SSMM's granular migration approach gives the best of both worlds...the ability to migrate slowly over time, like with the Gradual migration approach, while

also giving the ability to build on new hardware, such as with the Database Upgrade Process.

- No pre-scanning of content is required, and site content owners can control the migration process, something which is definitely not possible with the Database Upgrade Process.
- Failure of the migration of a content database is difficult to troubleshoot, and even if 99 out of 100 content databases migrated, the final failed database could require fallback to the old server. SSMM does not have this limitation.

Benefits of Content Migration with Metalogix

Allow content owners to migrate (Migrate based on permissions)

A key feature with SSMM, illustrated in figure 1, is the ability to delegate the migration process to content owners, giving them the flexibility to choose when content gets migrated. Because SSMM uses the standard SharePoint API to move content from one location to another, the user of the tool simply needs to have rights to the content to be able to move it. For many organiza-

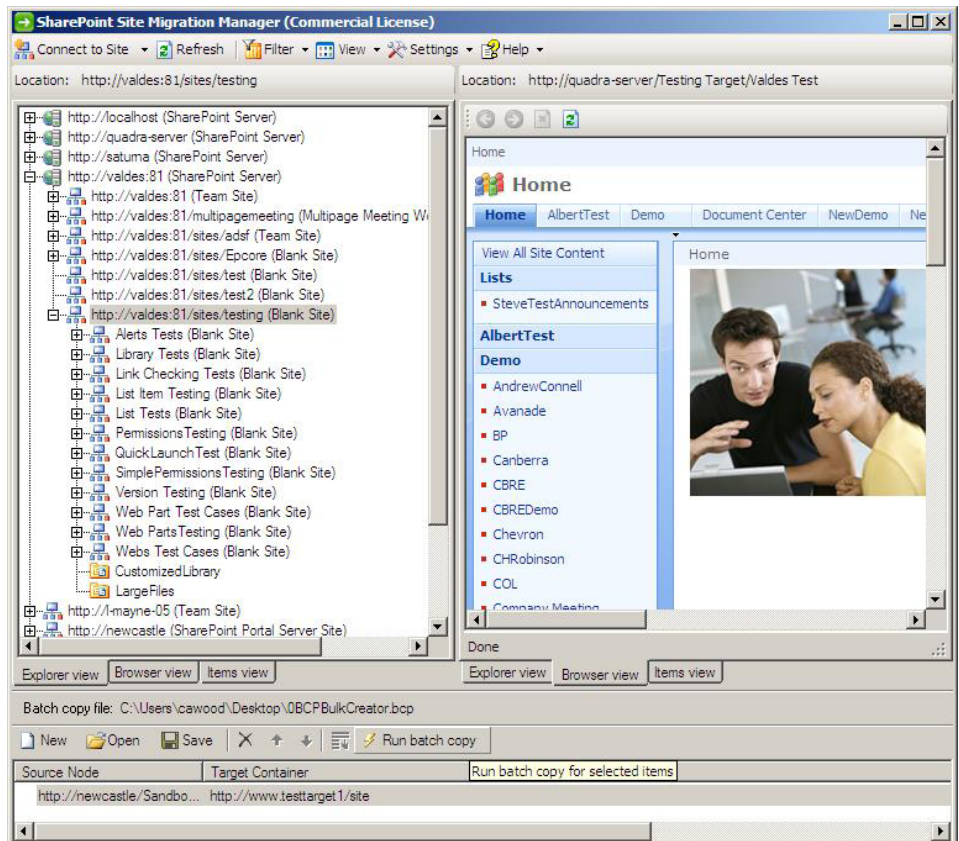


Figure 1

tions, the content owners may have unique schedules or may want to be in control of the process. This process provides for that type of flexibility, and lets them fit their own content migration into their schedules.

Granular Migration

One very useful feature in SSMM is its ability to perform granular migrations of content. Content owners can decide which portions of a site collection get migrated first, and then migrate the content in waves, rather than en masse.

Content Filtering Synchronization/Incremental Migration
The user interface of the tool allows for content owners and admins to define what type of content is migrated, and whether to filter out specific types of content or sites. This allows site admins to restrict certain types of files, list items, or other content to be abandoned in the old site, or migrated across. In some cases, it may make more sense for old, unused site content to be retired with the old platform. SSMM gives you that type of flexibility to decide what gets migrated and what stays behind.

Reorganize your sites during migration

Site structure can be redefined while migrating. Sites and Site collections can be split into multiple targets, and administrators and content owners can define what content is migrated. This allows for visual and navigational changes to be made to the content, giving a site architect the ability to reorganize and restructure the content.

Manage Site Collections and Content Databases

Many SharePoint 2003 Admins built their farms with a small number of content databases (in some cases only one), which limits the scalability, performance, and recoverability of the farm. Metalogix's SSMM allows for that content to be restructured and reorganized into a best practice approach of multiple site collections split among multiple content databases, without the need to change any of the logical structure of the existing sites.

Server Consolidation

Since SSMM can pull and push data from multiple sources and targets, it allows organizations to collapse multiple SharePoint 2003 farms into a single SharePoint 2007 farm. Many organizations are surprised to find out how many instances of SharePoint exist in their organization across multiple departments. In many cases, individual team sites were created on desktops and rogue servers—and content was decentralized across the organization. This content can be brought back under centralized control using SSMM, and multiple SharePoint sites can be consolidated onto a single managed farm.

In addition, because a new SharePoint 2007 farm is created for the migration, new hardware, software versions, 64-bit architecture, and newer database servers can be provisioned, all with best practice architecture and without the need to design around previous SharePoint version design flaws.

Server Virtualization

Server Virtualization is a hot topic in today's business data centers. Many organizations are opting to virtualize the majority of their server infrastructure, including SharePoint servers. The fact that SSMM is target agnostic allows for full flexibility to build the new SharePoint 2007 farm on virtual machines if required. This applies to virtualization software from any one of the common virtualization platforms in use today.

Ongoing Value of SharePoint Site Migration Manager

Metalogix SSMM is not simply a migration tool. The software has a number of features that can be used to manage existing SharePoint lists and sites. Among these are several key features that make SSMM highly useful after a migration, or for organizations that simply need to manage their existing SharePoint 2007 investment.

Site Structure Management

Having access to the SSMM client allows SharePoint content owners to easily reorganize the SharePoint server if organizational mistakes are made. Instead of having to simply live with inefficient structural errors, they can instead be quickly rectified. The following are a list of several of the more popular management features in SSMM:

List splitting

Microsoft recommends that SharePoint lists contain less than 2,000 items. If you find that your lists are growing too large (e.g., more than 2,000 documents in a document library), or you simply want to organize your list items more effectively, you can use SSMM to split lists across multiple SharePoint sites, site collections, or farms. To split a list, you can use SSMM to create new lists and then copy the items from the original list to many locations.

List merging

If two lists have the same name, merging the lists is as simple as copying one list into another and choosing the Append items to list option. If the lists have different names, you can easily combine them using the SSMM item copy view.

Rename lists during copy

With this feature, you can move to better naming conventions by simply renaming the lists during the copy operation.

Copy a list schema

SSMM provides the ability to duplicate the columns of an existing list. This facility makes it easy to share column sets without copying data or using list templates.

Reorganize Folder content

Using the copy folder option, you can restructure a document library to use folders or collapse folders into a single list.

Split sites across multiple Content Databases

You can refactor your content to manage distribution to larger (or smaller) sets of content databases. If your structure has gotten out of hand, this will allow you to get back to a performance-oriented architecture.

Merge sites into a single site collection

Using SSMM, you can simply copy sites from one site collection to another. This process is as simple as copying and pasting.

Incremental Updates

It may be that a company would like to continue to roll out sites, lists, or other sections of their SharePoint server over time. Having SSMM allows these organizations to perform incremental copies and manage these rollouts. It is also possible to filter items by date, so only changes made in a certain period of time could be copied for roll-out or backup purposes.

Conclusion

Microsoft provides for three unique migration options for moving SharePoint content from SharePoint 2003 to SharePoint 2007, each of which has its various advantages and disadvantages as outlined in this document. Metalogix SharePoint Site Migration Manager is a tool that offers a fully supported and capable alternative to these approaches.

SSMM provides value to SharePoint administrators by giving them more flexibility in migration approaches, with features such as the ability for content owners to control the migration and the capacity to define at a granular level when specific content is migrated. In addition, SSMM also provides for ongoing value as a management tool for SharePoint sites, allowing for structural changes to be easily made to SharePoint sites and lists. These features have positioned SSMM as a valuable tool not just for SharePoint migrations but also for the ongoing site maintenance often required by content owners and SharePoint Site Admins.

Michael Noel is an MVP for SharePoint Server and an MCSE+I. He has been involved in the computer industry for nearly two decades, and has significant real-world experience helping organizations realize business value from Information Technology. Michael has authored several major best-selling industry books that have been translated into seven languages with a total worldwide circulation of over 150,000 copies. Significant titles include *SharePoint 2007 Unleashed*, *Exchange Server 2007 Unleashed*, *Windows Server 2008 Unleashed*, *ISA Server 2006 Unleashed*, *SharePoint 2003 Unleashed*, and many more. Currently a partner at Convergent Computing in the San Francisco Bay Area, Michael's writings and worldwide public speaking experience leverage his real-world expertise designing, deploying, and administering IT infrastructure for his clients.