Realtime publishers

The Shortcut Guide™ To

Centralized SharePoint Administration

sponsored by

AvePoint®
Unleashing the Power of SharePoint™

Wendy Henry

Introduction to Realtime Publishers

by Don Jones, Series Editor

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We've made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book's production expenses for the benefit of our readers.

Although we've always offered our publications to you for free, don't think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you $40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the "realtime" aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We're an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I'm proud that we've produced so many quality books over the past years.

I want to extend an invitation to visit us at http://nexus.realtimepublishers.com, especially if you've received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you're sure to find something that's of interest to you—and it won't cost you a thing. We hope you'll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Copyright Statement

Chapter 1: Common Challenges in SharePoint Administration

Microsoft Office SharePoint Server 2007 and Windows SharePoint Services v3.0 (MOSS2007 and WSSv3) are among the leading network service applications available from Microsoft. More than just a communication platform, SharePoint has seized the market as a single management system for data storage, dissemination, manipulation, and business logic. As the product matures in established networks, a greater percentage of mission-critical business information is stored in SharePoint, requiring a higher degree of administration and maintenance. From documents to compiled files or custom list items to calendar appointments, today's SharePoint is used to organize, track, maintain, and retrieve the information imperative to corporate health.

Yet despite the capabilities and popularity of these products, the administrators of scaled MOSS2007/WSSv3 environments remain challenged by the infancy of the native administration utilities available out of the box. With such onus on one-stop-shopping for all mission-critical data from a single system, the degree to which that system must be monitored greatly increases. Even if business data is stored on disparate systems yet simply exposed on SharePoint sites, MOSS2007/WSSv3 is just as responsible for satisfying user needs as the alternative storage system itself. There is no way around it—SharePoint will become a top administrative concern in your network if it hasn't already.

SharePoint empowers business information workers by offering content contribution, collaboration, communication, and vast data retrieval opportunities. Therefore, like any organic system, SharePoint expands and contracts at the will of user behavior, which makes perfect planning of your deployment a challenge. Over time, scaling SharePoint to accommodate growth is inevitable, yet WSSv3 and MOSS2007 offer only a few native content-migration and feature-synchronization utilities. From installation through testing to production, deployment, and content control, there are many concerns that must be addressed to ensure SharePoint's ability to serve users. However, the logging and analysis features in WSSv3 and MOSS2007 are some of the more shallow capabilities in the entire product. From a disaster-recovery and fault-tolerance standpoint, SharePoint's native backup and replication offerings are cumbersome and prone to human error.

This guide will examine the important administration concerns in SharePoint and the tools available for centralized management. Specifically, the following topics will be discussed in detail:

- Administrative concerns in a scaled SharePoint environment and the limited native tools available for addressing them.
- A breakdown of SharePoint logical hierarchy and administration levels
- Scaling SharePoint across multiple locations while synchronizing content and security
- Best practices for archiving content and reporting on SharePoint

Deployment Considerations

Planning a deployment of SharePoint usually takes place prior to installation. But even if you are working in an existing SharePoint environment, your deployment design may need updating at any time. Deployment objectives include scaling the server farm, synchronizing WFE servers, migrating content, and connecting to distributed or legacy systems. Unfortunately, there are not very many native tools in WSSv3 or MOSS2007 that accommodate data replication and relocation.

According to the Microsoft Operations Framework (MOF), any smart SharePoint enterprise should consist of both a test and production environment. If nothing else, software updates should be thoroughly tested before being installed. Moreover, enhancements such as custom pages, custom Web parts, programmatic additions, and branding should be thoroughly developed and tested on a system that is isolated from the users to avoid unexpected interruptions to critical data access. Regrettably, if a simple mechanism is not in place to port the changes from the development system to the production environment, quickly and easily implementing a test environment may prove impractical.

Resource

MOF 4.0 is a collection of documentation designed to help IT professionals plan, deploy, manage, and support Microsoft networks. MOF can be downloaded from MS TechNet at http://technet.microsoft.com/en-us/library/cc506049.aspx.

Initial Installation of SharePoint Farm

According to Microsoft, there are four possible physical layouts of a SharePoint environment: single server, small farm, medium farm, and large farm. And although one would hope that most commercial SharePoint implementations are meticulously planned before installation, the reality is that inadequate planning or simple user behavior miscalculation can result in the need to grow a SharePoint environment from one layout to another. The SharePoint servers can play myriad roles in each layout (see Table 1.1). The content database(s) may need to be scaled to provide more growth opportunity. This could be accomplished by the DBA inside SQL Server by employing multiple secondary data files and filegroups to stretch the database across multiple hard disks within a single SQL instance.

Resource

For more information about SQL Server database optimization see http://www.microsoft.com/sqlserver.

Layout	SharePoint Server Roles	SharePoint Services
Single Server	All	All
Small Farm	Web Front End (WFE) Content Server	All None (SQL Server)
Medium to Large Farm (depending on combination of offloaded roles)	WFE—Central Administration Additional WFEs Application Server (Forms / Excel) Index Server Query Server Content Server	Central Administration WSS Web Application Forms or Excel Calculation WSS/Office-SP Server Search Query Web Service None (SQL Server)

Table 1.1: Examples of role distribution in SharePoint farm physical layouts.

SharePoint administrators can assist their SQL DBA counterparts by managing Web application distribution across multiple content databases in SQL Server. When a new site collection is created into a Web application, it is possible to instruct SharePoint to place the content into a separate content database if the site collection's anticipated content is expected to cause the existing content database to exceed a manageable size.

The method used to accomplish this depends on whether the separate content database already exists. If not, the STSADM.exe CLI utility's *createsiteinnewdb* operation will generate the new site collection and create the new content database on the fly. If the separate content database already exists, Central Administration can be used to render all other content databases offline except the desired separate content database prior to creating the new site collection (see Figure 1.1). Marking other content databases offline will interrupt hierarchy management in the site collections stored there. Luckily, creating a new site collection doesn't take long and you can return the other content databases to their Ready states soon.

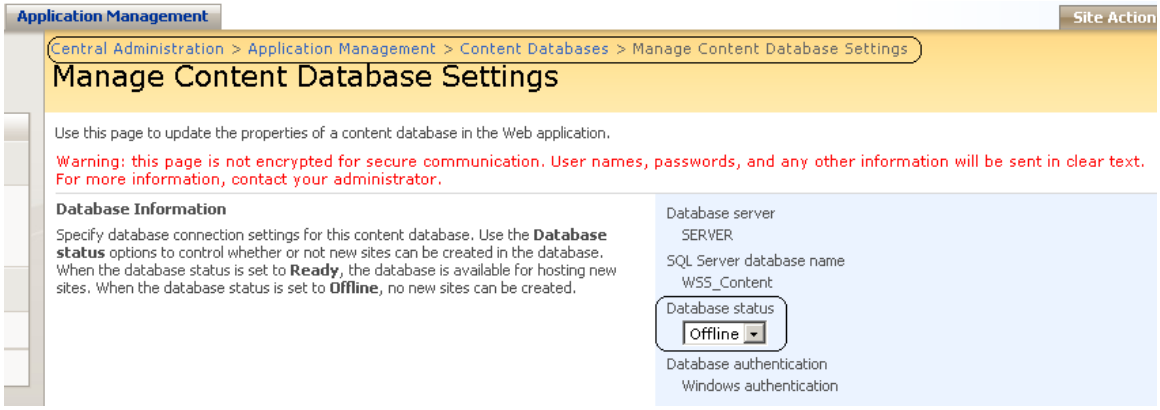

Figure 1.1 Selection of a specific content database's settings from Central Administration.

The disadvantages of using these methods to scale content databases are the time and expertise required. Either the SQL Server DBA or a highly authoritative SharePoint administrator with experience writing STSADM.exe syntax need to be involved and the SharePoint interfaces are not very intuitive. And what if these solutions only mask storage insufficiencies for awhile? Once you run out of room, moving SharePoint off of an inadequate SharePoint content server to another SQL Server is a whole different ball game.

Plan the SharePoint content server carefully because using the native Microsoft tools to move all SharePoint databases to an alternative server after SharePoint has been employed is not easy. The lengthy process Microsoft recommends results in:

- **Interruption**—Disconnecting all users and preventing further connection to SharePoint until the entire migration process is completed
- **Complexity**—There are distinct rules for migrating to a server of the same OS and SQL Server version (just a different host name) versus migrating to a server of newer OS and/or SQL Server version and/or 64-bit chipset (with the same host name)
- **Multiplicity**—OS steps require a member of the Windows OS Administrators group, SharePoint steps require a member of the SharePoint Farm Administrators group, and SQL Server steps require elevated SQL Server server-level role memberships; in a segregated IT department, this may require the involvement of multiple administrators

Resource

Moving SharePoint databases is outlined in the TechNet article available at http://technet.microsoft.com/en-us/library/cc512723.aspx.

The Microsoft procedures previously noted involve backing up SharePoint databases and restoring them to the new SQL Server content server, so keep in mind that the size of the SharePoint databases being moved will affect the amount of time it takes to create the .bak backup files, the duration of the .bak file copy/move operation, and the length of the restore process. What looks like a simple 8-step process on paper could actually take several hours or even longer if the necessary administrators have difficulty coordinating their schedules.

Another consideration for single-server or small farm is user throughput into the SharePoint enterprise. As the user base increases, say from a pilot group to all employees, asking one server to handle 100% of the user HTTP sessions becomes inefficient with increased wait times and retransmission requests (a performance benchmark on throughput per site collections is outlined in the TechNet article available at http://technet.microsoft.com/en-us/library/cc262787.aspx). The result is a degraded perception of SharePoint's performance. One solution is to add multiple WFE servers to the farm in an effort to spread out the HTTP traffic (see Figure 1.2). Yet scaling the farm with multiple WFE servers demands they be kept synchronized with the same configuration information, security, features, and files. Unfortunately, the native Central Administration and STSADM.exe utilities in SharePoint are laid out to configure each server independently, which means something as simple as copying a new feature solution or custom file to each WFE server demands repetitive administration.

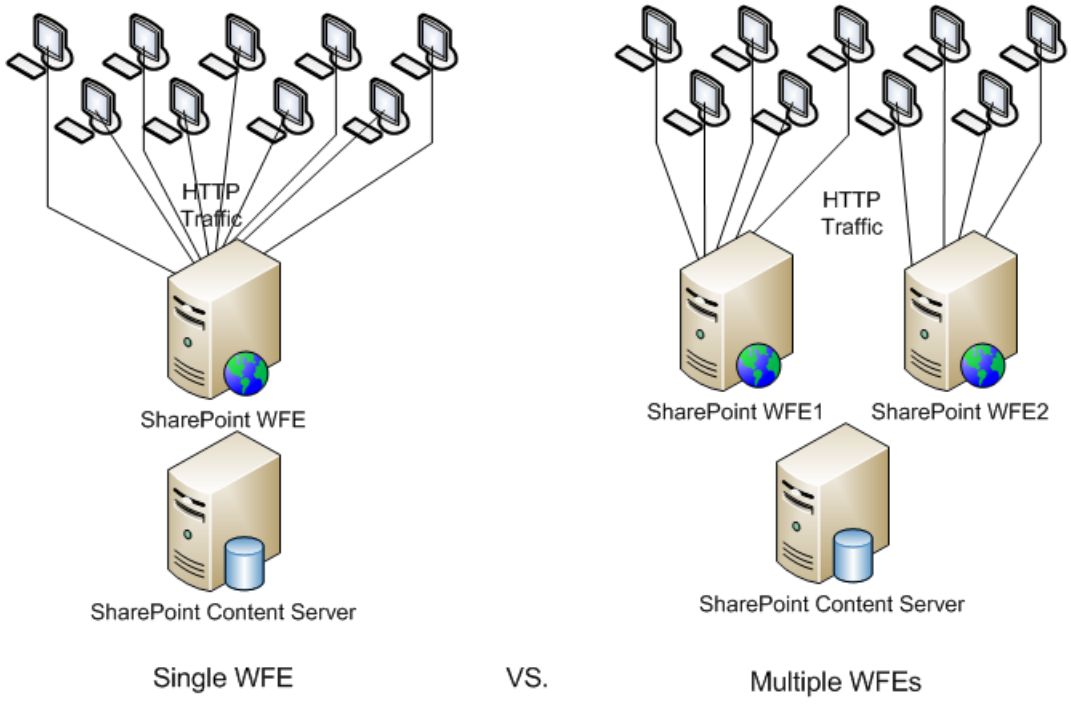

Figure 1.2: HTTP traffic scaled to multiple WFE servers.

From Testing to Implementation

Scaling a SharePoint environment is not the only scenario where user content on the content server may need to be relocated to a different content server. Another occasion conducive to copying or relocating content would be from an intranet SharePoint enterprise to an extranet or internet-facing SharePoint environment. Additionally, updating production sites with newly tested and approved customizations requires content movement. With the advent of virtualization, establishing a development/test SharePoint enterprise has become affordable for most networks.

A primary concern regarding content replication or movement is maintaining original metadata. Creation timestamps, GUID, author identity, and object ownership can be easily overwritten if the content is simply redistributed from one SharePoint realm to another. MOSS2007 ships with options on the Operations page of Central Administration for deploying content (see Figure 1.3), but WSSv3 does not. Maintaining some metadata during a WSSv3 content deployment can be conditionally achieved using the STSADM.exe export and import operations but leaves little opportunity for transforming metadata or maintaining hidden metadata.

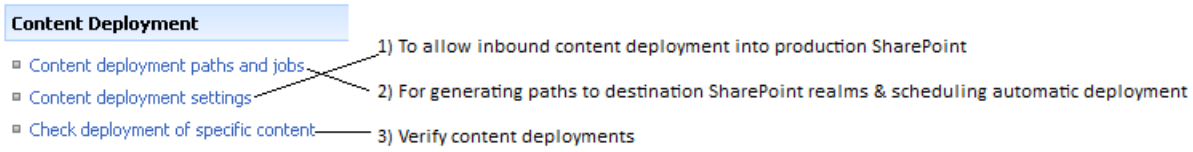

Figure 1.3: Central Administration / Operations / Content Deployment.

MOSS2007 farms do not accept inbound content delivered by the Content Deployment feature by default. Before you can engage in Content Deployment, you must configure the Content Deployment Settings. These settings can be accessed from the *Content deployment settings* link in the Content Deployment section of the Operations tab in Central Administration (see Figure 1.3 option 1). Once Content Deployment has been allowed and the settings configured, the two logical components of a deployment operation must be created: a path and a job.

A path defines origin and destination site collections per Web application for content deployment. Multiple paths can be created to identify multiple site collection relationships. Jobs can then be created for specific paths that designate which sites, lists/libraries, folders, or items in the path origin site collection are to be deployed and the frequency of automatic deployment operations. Since paths are defined at the site collection level, if you have several site collections to deploy from, you will need to create multiple deployment paths. Similarly, if you require granular deployments on differing schedules, you will need to create multiple deployment jobs. To create a deployment job, you will use the *Content deployment paths and jobs* link from the same section of Central Administration as the preceding course of action (see Figure 1.3 option 2).

Finally, you must create one or more jobs for each path that will specify the scope of content to be deployed (entire site collection, specific site, specific list/library, specific folder or item). These jobs can be scheduled to deploy automatically and monitoring of deployed jobs can be performed on the Check deployment of specific content link of the same section of Central Administration as the preceding course of action (see Figure 1.3 option 3).

Content Deployment in MOSS

Various blogs and magazine articles have been written on taming the MOSS2007 Content Deployment beast. Such detail falls beyond the scope of this guide, but you should be aware that setting up Content Deployment is granular and time consuming. Also, paths and jobs can become corrupted easily, so monitoring is a must! Essentially, content deployment must be enabled or not for the entire farm (default is not) then configured with myriad settings that include:

- Import & Export server designation in farm
- SSL, Temporary File location & Report recycling for farm
- Source Web Application & Site Collection
- Destination Central Administration server, Web Application & Site Collection
- Authentication method & credential into Destination
- Choice to deploy user information & security structure
- Schedules for automatic incremental deployments per job

A Quick Word About Quick Deploy

MOSS2007 Content Deployment offers a Quick Deploy feature via the Publishing Resources Feature that allows initiating deployment from the site itself rather than Central Administration for users. First, a Content Deployment job must be created and its Quick Deploy Settings enabled in Central Administration (see Figure 1.4).

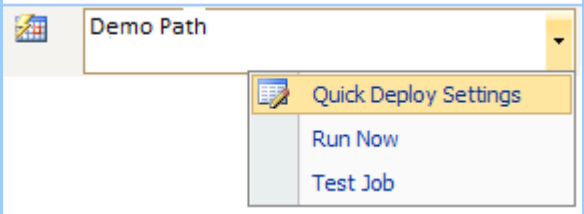

Figure 1.4 Quick Deploy Settings option for a path.

The default time interval SharePoint checks for Quick Deploy requests and distributes them is 10 minutes, but you can change this value in the Quick Deploy settings. Once Quick Deploy is configured, a page editor or library owner need only select Quick Deploy from the contextual drop-down menu of their object to submit a deploy request into the deployment queue, which will then be distributed at the next 10-minute interval per job settings.

MOSS2007 Content Deployment has several features, such as smart incremental deployments once the initial deployment has succeeded and automatic link webbing to deploy dependency content. But it can be cumbersome to configure (multiple paths by site collection, multiple jobs by granular content and different schedules) and it is not available in WSSv3. The STSADM.exe SharePoint Administration CLI utility can be used to export/import WSSv3 sites between enterprises and even retain security information in the process. And although STSADM.exe can be scripted and scheduled to be executed automatically via the Windows Server OS, it is still repetitive administration to configure multiple exports and imports.

Limitations of Native SharePoint Content-Connection Tools

MOSS2007 Enterprise ships with only a few legacy and disparate system connectors (though more can be written as application definition files for the Business Data Catalog) for connecting to and migrating legacy or disparate system data. WSSv3 ships with only one connectivity tool: a Data View Web part configurable in SharePoint Designer 2007 that can connect to a limited number of data sources. Both the Business Data Catalog and SharePoint Designer 2007 Data View Web Part require a fair amount of connection string authoring experience and knowledge beyond most information workers' skills. Table 1.2 summarizes the limitations of native scaling, deployment, and connection tools in SharePoint.

Tool	Purpose	Limitations
Central Administration—Create/Extend Web Application	Scale a Web app across multiple content databases	Web app by Web app scope requires repetitive administration; provides only a temporary solution to increased storage needs
Central Administration—Services on Server	Configure SharePoint services per server in the farm	Server by server scope requires repetitive administration; advanced knowledge of service dependencies necessary when scaling a farm
STSADM.exe	Scaling Export/Import sites	Requires elevated OS privilege; granularity requires repetitive administration; automation depends on OS
Central Administration—Content Deployment *MOSS2007 Ent	Publishing content	Cumbersome setup, repetitive administration for deploying same set to multiple destinations; full deployments must be manually initiated; control of Quick Deploy
Data View Web Part	Connecting to external data sources	Limited providers: .NET, OLE DB, XML files, Web services, and SharePoint
Business Data Catalog (BDC) *MOSS2007Ent.	Connecting to external data sources	Requires extensive programming knowledge to construct necessary application definition files

Table 1.2: SharePoint limitations.

These limitations have opened the market for many successful third parties to provide robust centralized administration and data migration utilities. Features of such products include single-screen viewing of all servers, services, and logical structure of a SharePoint farm as well as centralized service and settings management that configure multiple servers in a single operation. Some applications offer granular metadata management both on the fly during data migration and in a centralized settings configuration for automated data deployment. And remember to choose a utility that offers external system connectors to the legacy systems you own. Doing so will make migrating from or interfacing into the external system with SharePoint much easier.

Controlling Contributed Content

The whole point of an organic collaboration platform like SharePoint is to empower information workers. After all, no one knows better how to structure and organize data than the folks who actually use it. And although some small SharePoint environments may cling to a centralized IT administration model by forcing users to submit structure and content changes to a select few individuals for processing, in medium-to-large enterprises, such a strategy would quickly prove overwhelming and unproductive. Decentralized administration via delegated authority is a better strategy for large dynamic collaboration platforms. But structural changes still need to be monitored to assure sufficient server resources are available, and content changes need to be monitored for approval, corporate/regulatory compliance, malice, and duplicity.

Monitoring Contributed Structure

SharePoint's logical structure of Web applications, site collections, and sites should be carefully planned and governed to ensure availability and content protection. Delegating structure maintenance to information workers increases the need for diligent governance. SharePoint includes its own security architecture that can be used to delegate logical structure authority. System-supplied SharePoint group and Permission Level objects are generated for each new top-level site in a new site collection based on the site template used (see Figure 1.5). It is then the default behavior of all new sites created beneath to inherit their parent site's security architecture. It is also the default behavior of each list/library item to inherit the security structure of its parent folder, the behavior of each folder to inherit the security structure of its parent list/library, and the behavior of each list/library to inherit the security structure of its parent site. Inheritance allows centralized management of authority delegation by restricting permission maintenance to the top of the site collection and letting the same permissions flow down the hierarchy to eventually protect each item. But inheritance can also be disabled at any logical level (site, workspace, page, list/library, folder, or list/library item) to configure a more granular permission structure. Inheritance and granular permission structures can be mixed throughout the site collection hierarchy.

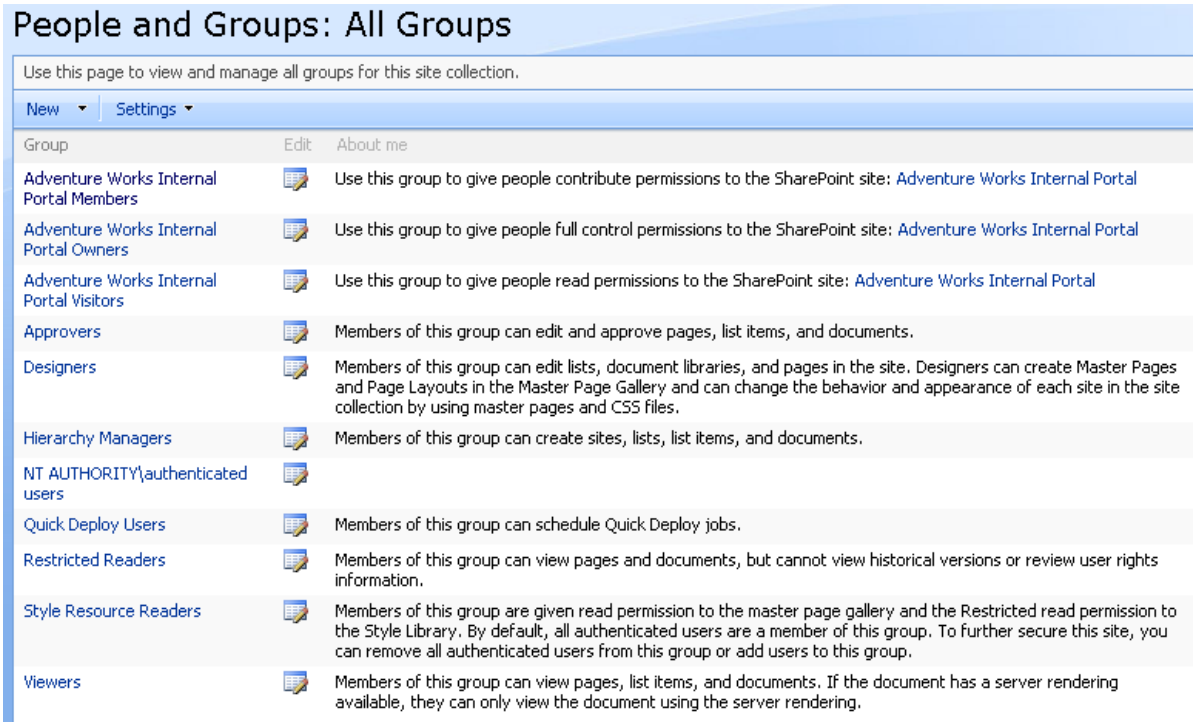

Figure 1.5: System-generated SharePoint groups of MOSS Collaboration Portal site template.

Imagine you have made several users members of the Hierarchy Managers SharePoint group (which is assigned the default Manage Hierarchy Permission Level), giving them the ability to create new sites, lists/libraries, and items. If these users have been creating structure rampantly and without a change control doctrine in place, you've lost control of what logical structures exist and where they are. Unfortunately, SharePoint offers little in the way of structure discovery to keep tabs on what is out there. MOSS2007 ships with a Site Map (Figure 1.6) in the Site Directory site template that is dynamically populated with links to all sites in the site collection. And although this may prove useful for discovering actual sites, it does not list pages, lists/libraries, or items. Worse, in a multiple site collection environment, you would have to visit each site collection's site directory separately! SharePoint Designer 2007 offers simple reports that can enumerate site hierarchies including the pages and lists/libraries (but not items) in WSSv3 or MOSS. Actually, this functionality is programmatically equivalent to STSADM.exe operations such as *enumsites* that would produce the same information if executed from the CLI by a site collection administrator. But again, these SPD2007 reports and STSADM.exe operations are site collection specific.

Figure 1.6: Typical site map.

Monitoring Contributed Content

The authority to create logical structure will likely be limited to a few talented information workers and therefore is probably one of the lesser focuses of SharePoint governance. The crux of monitoring SharePoint will almost certainly center on information worker content contributions because more users will be granted the ability to edit pages and add list/library items than those who can add sites. Diligently monitoring content goes beyond setting permissions to allow only certain users to post information. The content posted by authorized users must adhere to corporate and regulatory standards, and may require approval by editors and/or managers. Furthermore, according to 2009 statistics gathered by Ontrack data recovery, up to 32% of data corruption is caused by well-intentioned authorized users!

Content approval has long been a priority in multi-author environments. The process of content approval in SharePoint can be as simple as a list/library setting that allows one of any members in the Approvers SharePoint group to pass a newly submitted or changed item. More sophisticated serialized approval workflows are supported in SharePoint by the Windows Workflow Foundation of the .NET Framework prerequisite software:

- **Workflow Templates—**System-supplied templates configurable in the browser GUI provided per active Features
- **SharePoint Designer 2007**—Workflow Wizard GUI offers stepped workflows of finite conditions and resultant actions with optional If/Else conditional logic and parameterization but that are linked to only a single list
- **Visual Studio—**Programmatic workflows based on the SharePoint object model

Unfortunately, SharePoint Designer 2007 and Visual Studio require extensive training beyond the scope of most information workers and the only Approval workflow template installed by default has limited configuration. Furthermore, workflows require an initiating user action to kick off automatically or manual invocation, making them reactionary by nature.

But controlling content is about more than reactively approving or rejecting already-submitted items. To ensure data availability, malicious content must be thwarted *before* being saved into SharePoint and proactive protection of critical data from deletion is also essential. Unfortunately, SharePoint does not include a free native antiviral solution, though a separate antivirus product such as Microsoft Forefront Security for SharePoint can be additionally installed and offers only simplistic file name extension blocking as a means of preventing the upload of undesirable content. However, blocking all files of a specific file name extension without considering other mitigating circumstances may be too broad of a limitation.

Also, SharePoint supplies only a user-centric two-stage Recycle Bin for the salvaging of erroneously-deleted content. By default, deleted items are purged from the user's Recycle Bin after 30 days and from the Site Collection Administrator Recycle Bin in 30 days. Meanwhile, users are unable to access the deleted critical data. Wouldn't it be better to prevent the deletion of the necessary critical information in the first place?

Limitations of Native SharePoint Discovery Tools and Workflows

Monitoring information worker-generated structure and content can prove challenging when using only native SharePoint tools. In addition to the lack of structure discovery utilities and proactive content control tools, site usage information is difficult to produce. Usage Analysis in WSSv3 and MOSS2007 can be employed from Central Administration to provide user action details about a site but requires exhaustive logging that is resource intensive. Usage Analysis only records user activity during the time range configured for it, and most networks can only suffer the high resource utilization during non peak hours (when you need to monitor information workers the most). The XML log files generated are poorly formatted and contain unresolved GUID numbers rather than friendly object names, making them difficult to understand.

Third-party SharePoint add-on software has opened up the possibility of monitoring SharePoint not only from a centralized IT perspective but also in a decentralized model. Most applications offer secure, understandable GUIs that can be customized by information workers to serve them best. The SharePoint object model is hidden by most utilities under a GUI form-based hood that allows the information worker to prevent content containing keywords from being uploaded or created into SharePoint as well as requiring Web parts be digitally signed before being added to a page. Content validation, not approval, should be the main objective.

Maintaining Data Availability

Despite the best of monitoring and control mechanisms, eventual failure or exhaustion of consumable resources is an imminent probability. Entire volumes have been written about maintaining SharePoint; an exhaustive list of all possible calamities and proposed solutions here is beyond the scope of this guide. But I would be remiss in not addressing the topic, especially because anything from power outages to hardware malfunction to software corruption can render a SharePoint enterprise useless. Disaster Recovery plans should include not only fault tolerant redundancy solutions but also backup strategies and location failover policies. But what about when the SharePoint enterprise is not experiencing any failures yet the users are reporting degraded performance?

Contention refers to competitive requests for the same resource causing delays. In SharePoint, contention occurs when multiple user client machines attempt to connect to the same WFE SharePoint server for HTTP services or when multiple WFE servers need sessions with a single index or content server. Scaling strategies such as employing multiple WFE servers to handle the client requests or clustering the content server for workload balancing (and fault tolerance) can alleviate such contention to a certain degree. But there is also hardware contention within each SharePoint server and network latencies that must be monitored and addressed to keep the enterprise functioning within acceptable Service Level Agreement (SLA) requirements.

Disaster Recovery / Fault Tolerance

From a fault-tolerance standpoint, a SharePoint single-server physical layout is a recipe for single point of failure. Scaling out the farm to a small farm by separating SQL Server onto a separate content server at least follows the Microsoft best practice for the database engine product that states SQL Server should be implemented on a dedicated server. But the small farm still leaves SharePoint services vulnerable to a single host disaster. The more you scale the physical layout, the more redundancy you can introduce for each SharePoint service. Unfortunately, the licensing model of MOSS2007 can make large farm scaling cost prohibitive for most small to medium businesses.

> **Resource**
>
> Microsoft TechNet provides an article regarding planning and monitoring SQL Server specifically for SharePoint support. It can be downloaded from http://technet.microsoft.com/en-us/library/cc287996.aspx.

From a backup and restore point of view, a scaled farm requires multiple backup procedures to protect all the servers. The content server must keep diligent backups of all SharePoint databases as well as the SQL Server system databases via SQL Server administration to recreate the entire SQL Server instance in the event of loss. And even though the remaining WFE servers can take up the slack in the event one WFE fails in a multi-WFE layout, the same workload across fewer servers will cause a performance degradation due to increased utilization of fewer resources unless a 1+n fault tolerance model was implemented with a hot spare server. Administrators must independently back up each WFE server's IIS instance in a manner that mitigates recovery time during degraded performance. In segregated administration, the backups in SQL Server and IIS are commonly the privilege of non-SharePoint administrators. Lastly, SharePoint administrators must perform backups of the SharePoint enterprise itself and often divide the logical structure to be backed up at different stages based on data volatility and size. Granular backup choices are limited to sites, no smaller.

Backups in SharePoint

There are three (3) possible interfaces for conducting backups in SharePoint: Central Administration, STSADM.exe from CLI, and SharePoint Designer 2007. Depending on which portion of the logical design is to be backed up, the appropriate tool should be used:

- Central Administration—Backup entire farm or select Web applications or chosen configuration databases
- STSADM.exe—Backup a specific site collection or entire farm
- SharePoint Designer 2007—Backup a specific site

There are complete books written regarding disaster recovery plans for SharePoint. Backup procedures not included are direct database backups in SQL Server and IIS configuration backups via IIS.

Contention

Due to SharePoint's architecture, contention vulnerability is lurking at every connection. Depending on the physical layout, contention for hardware resources may cause interruptions, such as in a single-server implementation. Many other factors must be considered in addition to server hardware. Network congestion resulting from bandwidth contention, refused client connections can be the result of inadequate TCP/IP resources on the SharePoint WFE servers running IIS, and delayed Search results are the product of competition for the SharePoint index servers or failed index builds when the index server is unable to contact a busy content server.

To alleviate individual server hardware contention, Microsoft best practices strongly encourage the use of multiple hard disks, multiple NICs (perhaps streamed together to appear as one host on the TCP/IP network), and multiple CPUs. Workload balancing WFEs and clustering the SQL Server content server are also strongly encouraged. Optimizing IIS and SQL Server products individually fall outside the scope of this guide, but Microsoft and others have published many technical books on just those subjects.

Limitations of Native Tools

SharePoint Central Administration is the native tool for configuring and starting the various SharePoint services. In a farm, only one of the WFE servers will have the necessary IIS structure to serve the Central Administration Web site. Already we have a single point of failure concern for accessing Central Administration not to mention a contention dilemma. And the pages for managing services prove inconvenient because they filter based on server (not service). Thus, each server in the farm must be configured separately, which results in time-consuming repetitive administration when a similar change needs to be made to all machines (see Figure 1.7).

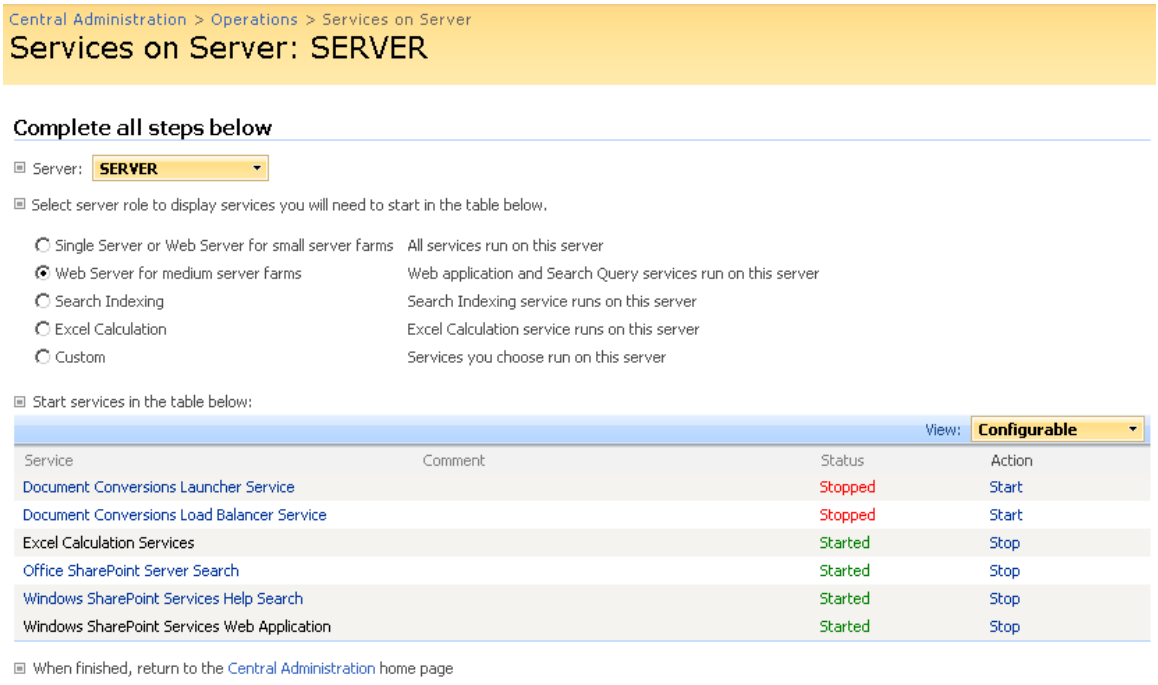

Figure 1.7: The Central Administration Operations Services on Server page.

There are many tools available in Microsoft's arsenal of utilities for monitoring resource contention; they are platform specific. The OS can monitor TCP/IP and software performance counters (see Figure 1.8). SQL Server Management Studio will reveal locking contention for SharePoint databases. IIS can report on site statistics, and SharePoint diagnostic logging can shed light on application warnings or critical errors concerning availability. But viewing all these freely-included tools in the respective products is disjointed and doesn't produce a cohesive cause-and-effect picture of poor SharePoint performance.

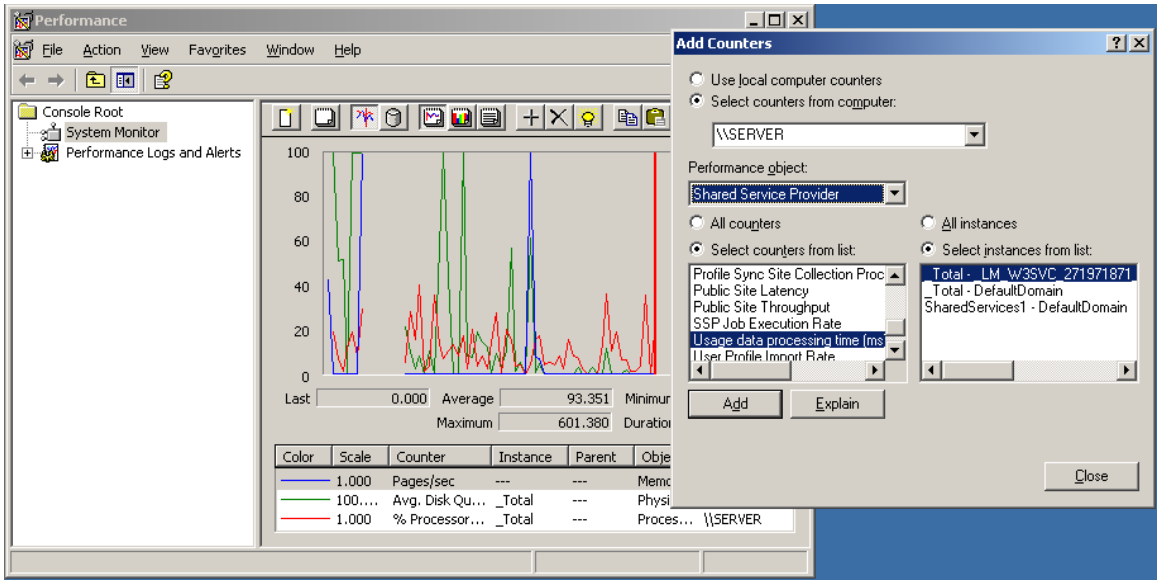

Figure 1.8: Windows Server OS Performance Console—MOSS SSP Performance Object counters.

Several third-party SharePoint administrative applications specialize in centralized monitoring tools that reduce the time and effort required to manage a SharePoint enterprise. By combining reports on the network, hardware, OS, IIS, SQL Server, and SharePoint into a single interface, these products add value to SharePoint through more efficient administration and quicker response time to potential issues. Furthermore, most of these centralized applications allow SharePoint service configuration changes to be applied to multiple servers in the farm simultaneously instead of the decentralized "per server" or "per Web application" filters available in Central Administration.

Corporate and Regulatory Compliance

The granular security architecture, varied Web service delivery, and highly customizable format of SharePoint make it a viable candidate for storing mission-critical and sometimes private data. It's reasonable to assume someone, at some point, may want to audit the data for compliance to government or corporate regulations. Ah, but as easy as it is for users to find and retrieve their contributed content, reporting on user activity and security changes prove a bit more challenging to the SharePoint administrator.

Microsoft touts SQL Server Reporting Services as a potential report tool for SharePoint. By gaining connection into the SharePoint Web services from SSRS, reports can be programmatically built around SharePoint statistics by using the SharePoint programming model. But for those who lack the time, money, or expertise to produce such a solution, there are a few limited native reporting tools in SharePoint or a third-party alternative application.

Auditing 101

Do not over-audit. Words to live by. When choosing what exactly to report on from your SharePoint enterprise, choose only that information being sought by an auditor or needed for compliance documentation. Unnecessarily auditing systems is resource intensive and counterproductive. Is the objective to simply be made aware of content additions? Workflows that create log files can help. Or must you keep track of SharePoint errors and warnings and the probable causes that preceded them? Diagnostic logging in Central Administration will happily generate many huge log files for you. Is visit tracking important? IIS W3C logs or Central Administration Usage Analysis can provide the details. Ah, but Usage Analysis also includes metadata regarding SharePoint sessions, which may be too much information if your goal is to avoid over-auditing.

SharePoint's Auditing Model

SharePoint can be broken down into four essential auditing branches: content metadata, logical structure, service architecture, and usage. Additionally, some SLA or contract agreements may require historical examples of informational, warning, or critical error messages that the SharePoint services have reported to the OS. Regardless of which or all of these categories your auditing needs stem from, SharePoint's native auditing tools are distributed, passive, and dependent. For example, each of the four branches create and store their log files separately resulting in a distributed set of data necessary to produce a single fully-arcing audit. Reporting on content metadata changes via reactionary workflows is passive. And usage analysis depends on special log files that are configured with a storage path and maximum file number per Web application (distributed), which makes usage analysis processing dependent on the file system and its security configuration (see Figure 1.9).

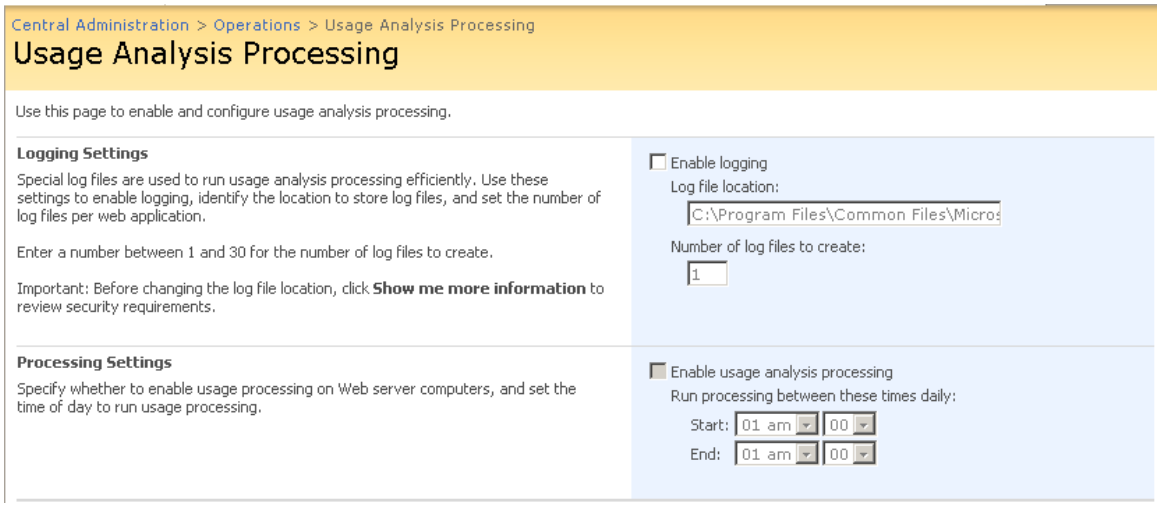

Figure 1.9: MOSS2007 Usage Analysis depends on WSSv3 usage logging.

Limitations of Native SharePoint Auditing Tools

Site administrators in both WSSv3 and MOSS2007 environments have access to usage information about their sites. In WSSv3, usage log files can be enabled via the STSADM.exe CLI utility. In MOSS2007, both WSSv3 usage log files and optional usage analysis can be enabled in Central Administration. Site administrators can then access the results of usage analysis via the Site Administration category Site usage reports link in their browser GUI. Similarly, Site Collection administrators can view the usage reports for an entire site collection via the Site Collection Administration category Site Collection usage reports link in their browser GUI. Because SharePoint provides the usage links per each site's Site Settings page, an administrator requiring usage information on more than one site in a Site Collection but not all sites therein would need to visit multiple pages and collate the reports.

Similarly, MOSS2007 offers Site Collection Audit reports per Site Collection that can be configured to automatically generate reports concerning anything from content access to security or audit configuration changes. The reports can then be downloaded per site collection as .XML file types and opened with any XML editor application (see Figure 1.10). The links to configure the Audit settings and download the Audit logs can both be found in the Site Collection category of Site Settings. But again, the reports must be accessed using repetitive steps to access the Site Settings GUI of each site collection. Also, the reports are secured to be accessible only by site collection administrators!

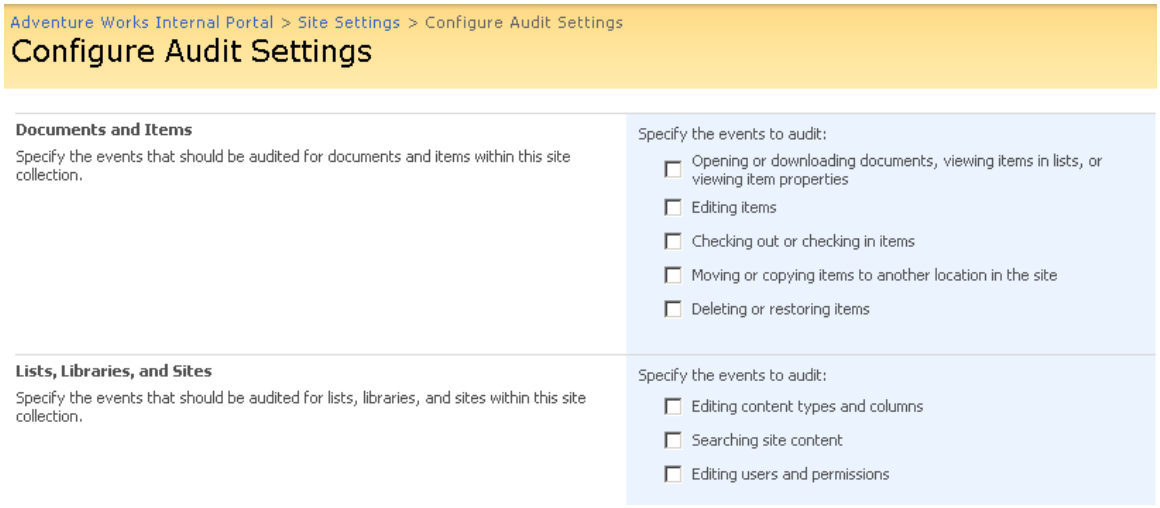

Figure 1.10: MOSS2007 audit settings for a site collection.

Summary

This chapter introduced the guide and outlined many administration challenges for Microsoft SharePoint administrators. It examined deployment concerns including physical layout and migrating content from development to production. It outlined the realistic need to control user-defined structure and uploaded content beyond simple approval methods. Microsoft best practice guidelines were mentioned for suggested data availability plans, and regulatory compliance by the SharePoint servers was discussed.

In the remaining chapters, this guide will step through the SharePoint Administration Hierarchy and explore managing multiple locations. Future chapters will also discuss archiving and reporting best practice solutions. Additionally, the concept of realistic auditing introduced in this chapter will be further examined. So stick around!