



Usage Analysis Tools in SharePoint Products and Technologies

Date published: June 9, 2004

Summary:

Usage analysis allows you to track how websites on your server are being used. The Internet Information Services (IIS) log and usage analysis logs from Microsoft Windows SharePoint Services and Microsoft SharePoint Portal Server can provide a realistic picture of how a portal site is used in each company. The viewing and managing of these logging mechanisms will be discussed in this chapter.

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This white paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2007 Microsoft Corporation. All rights reserved.

Microsoft, FrontPage, SharePoint, SQL Server, and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Usage Analysis Tools in SharePoint Products and Technologies

This is a sample chapter from the Microsoft SharePoint Products and Technologies Resource Kit. You can obtain the complete resource kit (ISBN 0-7356-1881-X), which includes a companion CD-ROM, from [Microsoft Press](#).

Usage analysis allows you to track how websites on your server are being used. The Internet Information Services (IIS) log and usage analysis logs from Microsoft Windows SharePoint Services and Microsoft SharePoint Portal Server can provide a realistic picture of how a portal site is used in each company. The viewing and managing of these logging mechanisms will be discussed in this chapter.

Tying this information to records collected from performance measurements of the server running SQL Server, the index server, the search server, and the Web server provides a full picture of architecture behavior. Peak usage analysis in conjunction with records obtained from performance monitoring logs can be used for capacity planning. Based on this information, you can decide if you need an additional server in the front-end Web server farm. Diagnostic logging provides another related logging mechanism, but this is not discussed in this chapter.

Managing Usage Analysis

If you want to view usage data, you can use the site administration pages or Microsoft Office FrontPage 2003. Before you're able to perform this usage analysis, you need to configure usage analysis processing. You control settings for usage analysis processing from SharePoint Central Administration. You must be an administrator of the local server computer or a member of the SharePoint administrators group to configure usage analysis settings.

From SharePoint Central Administration, you can control the following:

Whether or not to log usage data.

- Usage analysis is not enabled by default. If you want to use the usage analysis features for your server, you must enable the usage analysis logging process. Windows SharePoint Services log files are created daily to track usage information. When the log files are processed, a flag is added to indicate that the log file has been processed. You can check whether the log file has been processed by opening the file using Notepad or any other text editor. If you see an ampersand (&) after the top-level site Uniform Resource Locator (URL), the log file has been processed. Log files are not automatically deleted. They are preserved in the following path: %WinDir%\system32\LogFiles\STS. This is the default path, but it can be customized. If you do not want to track usage analysis data and you want to conserve disk space, you can turn off data logging for usage analysis.


Usage analysis data is generated by Windows SharePoint Services and gathered into the Windows SharePoint Services log files, which are stored on the front-end Web servers. The name of the default log files folder is called STS. When usage analysis processing takes place, the Windows SharePoint Services logs are merged into the content databases on the computer running SQL Server or the Microsoft SQL Server Desktop Engine (Windows) (WMSDE). The log data is a summary record of transactions on your website. The IIS logs contain a complete overview of Web traffic, including calls to Web services, requests for images, and so on; the summary record contains only the information necessary to provide a usage overview.

Usage data is kept for a limited time in the content database for historical purposes. Daily information is stored for 31 days and monthly information for 31 months. All usage data is

stored as a binary image in the Webs table. Daily information is stored in the DailyUsageData column, and monthly information is stored in the MonthlyUsageData column.

Where the log files are stored and how many log files to create.


- By default, the log files are in %WinDir%\system32\LogFiles\STS. You can specify any other location you prefer for this root log folder.

 If you choose a different log file location, you must be sure to give the STS_WPG user group **Read**, **Write**, and **Update** permissions to the directory. Without these permissions, the usage log files cannot be created or updated by Windows SharePoint Services.

Inside this folder is a folder for every virtual server named using the Windows SharePoint Services virtual server globally unique identifier (GUID), and under those folders, a different folder for each day will be generated. The path of these folders cannot be modified.

You can configure the limit of logs to be created on a daily basis, with a maximum number of 30 log files. If you set a maximum number of n log files, this applies to each virtual server. This means that the log folders of each of the virtual servers will contain at most n log files. In other words, if this number is set to 1 there will be 1 log file for each virtual server. Having many virtual servers and having many log files might reduce performance during logging. You should consider increasing the number of log files if a front-end Web server (standalone or member of a Web farm) has log files with more than a million entries.

The front-end Web server might not have enough memory to memory-map a really large log file, which leads to a situation in which the log file might not get processed. Each hit that a front-end Web server receives uses approximately 200 bytes (B) in a log file. As a result, approximately 200 megabytes (MB) of RAM are used to memory-map a log file that contains a million hits. Memory mapping occurs only for several minutes during usage processing. Because log files are processed serially, when you have several log files, a smaller memory footprint results when a log file is processed.

 Windows SharePoint Services logs HTTP 2.x information to the log files and does not log HTTP 3.x or HTTP 4.x information to the log file. When you use multiple log files, the log files are created at the same time, and all hits from one website are contained in the same log file. Usage data for a website is updated one time each day.

The number of requests that are sent to the back-end server during usage processing is proportional to the number of websites on the server. However, the memory footprint on the back-end server is not affected by the number of websites. The additional load that usage processing generates does not significantly affect the performance of the back-end servers on the server farm.

Whether or not to process the usage logs and when to do so.

- The tracking of usage data takes up a lot of disk space on the Web front-end server and the server where the content database is located. If you want to conserve disk space, you can stick to the default option and leave usage analysis processing disabled. In this scenario, you could use third-party tools to analyze the IIS logs. However, if you want to use the site administration pages to view usage data, you need to enable the processing of the usage logs.

By default, the log files are set to be processed every day at 1:00 A.M. During this process, the content of the Windows SharePoint Services logs is aggregated, summarized, and written to the content database. You can schedule the usage log to be processed at a more convenient downtime for your websites or for your SQL Backup schedule. You can specify an interval, by specifying the start time and end time of the interval, during which the usage processing job starts. The job is guaranteed to start at some time within the interval, but you cannot control when the job ends.

If your websites are primarily used by internal employees, for example, you might schedule the log to be processed at night, when demand on the sites is lower than during working hours.

Because usage analysis processing runs only once a day, when you enable usage analysis processing, you will not see any data until the next day. Log processing is done only for a single day's worth of data. If you turn off the log processing for a week but leave the data logging turned on, the next time you turn on processing, it will process only the previous day's log files. The log files for all the days before that will remain unprocessed.

When you configure usage analysis processing for a server, it takes effect for any existing virtual servers. If you later add a virtual server, you must configure usage analysis processing again to enable usage analysis for the new virtual server.

To configure usage analysis processing for a server

To configure usage analysis processing (shown in Figure 29-1), perform the following steps:

1. Click **Start**, point to **All Programs**, point to **Administrative Tools**, and then click **SharePoint Central Administration** or **Windows SharePoint Services Central Administration**. The configuration takes effect for any existing virtual server, so it does not matter which of these pages you use.
2. Under **Component Configuration**, click **Configure usage analysis processing**.
3. In the **Logging Settings** section, select the **Enable logging** check box.
4. In the **Log file location** box, type the location to store the log file.
5. In the **Number of log files to create** box, type a number between 1 and 30.
6. In the **Processing Settings** section, select the **Enable usage analysis processing** check box.
7. Under **Run processing** between these times daily, specify the range of times to start the usage analysis log processing. In the **Start** box, select the earliest time of day to begin running log processing. In the **End** box, select the latest time to end running log processing.
8. Click **OK**.

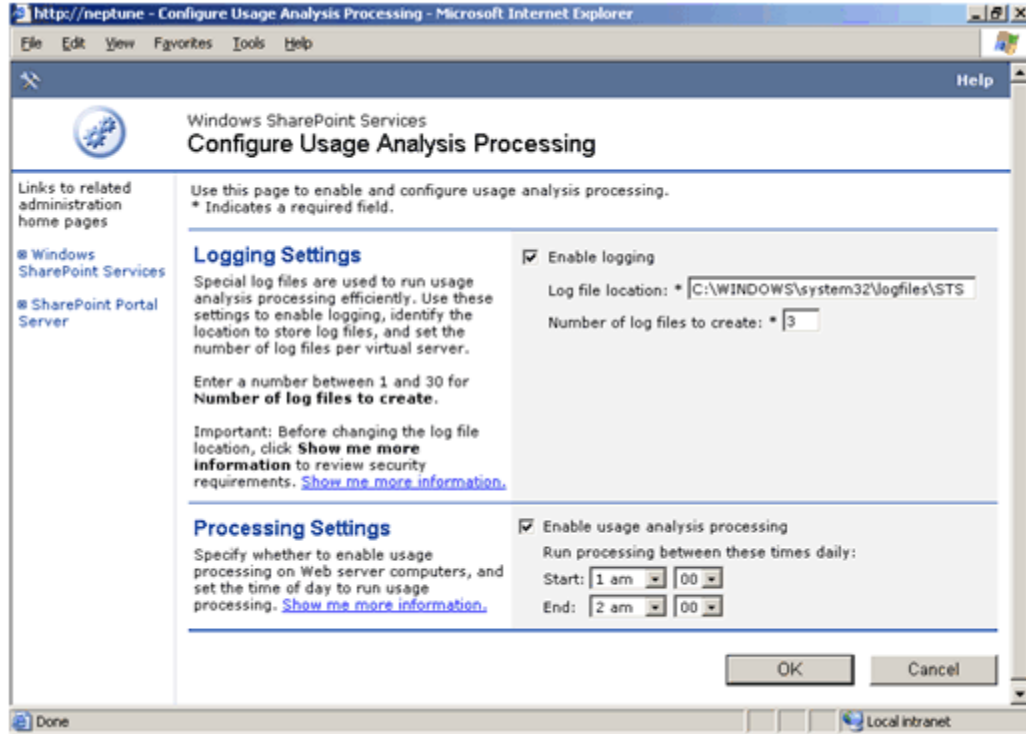


Figure 29-1. Configuring usage analysis using SharePoint Central Administration

You can also configure usage analysis processing at the command prompt by using the Stsadm.exe tool. This tool can be used to schedule the timed job responsible for starting the processing of log files.

You use the **getproperty** operation with Stsadm.exe to retrieve information about timed jobs. The **getproperty** operation takes the propertyname parameter as well as the -url parameter.

If you want to retrieve the usage analysis job settings, type the following command at the command prompt:


```
stsadm -o getproperty -pn job-usage-analysis -url http://<virtual server URL>
```

You use the **setproperty** operation with Stsadm.exe to schedule timed jobs. The **setproperty** operation takes the propertyname and propertyvalue parameters as well as the -url parameter. For a timed job, the propertyvalue parameter is the frequency and time when you want the job to be performed. If you want to schedule the usage analysis job, you can type the following command at the command prompt:

```
stsadm -o setproperty -pn job-usage-analysis -pv "daily between hh:mm:ss and hh:mm:ss" -url http://<servername>
```

For example, to schedule a job to be processed daily between midnight and 1:00 A.M., you would specify the propertyvalue parameter as

```
stsadm -o setproperty -pn job-usage-analysis -pv "daily between 0:00:00 and 01:00:00" -url http://woodgrove
```

 You can set the start and end time of the interval to be the same in order to process the log files at a precise time. Remember to set this time at least a few minutes in the future to allow the timer job to schedule the process.

⚠ When configuring usage analysis processing via SharePoint Central Administration, any changes in the processing schedule can be retrieved via the getproperty operation of stsadm.exe. Because of a bug, changes in the scheduling of the usage analysis job via the setproperty operation of stsadm.exe will not be visible in the SharePoint Central Administration pages. The changes are valid nonetheless.

Analyzing Website Usage

Usage analysis processing provides usage reports on Windows SharePoint Services sites. Site usage reports are useful for identifying which content on your Windows SharePoint Services sites is being heavily used or is used very little. This helps you to understand which sites are candidates for archiving and which sites should be kept online.

When you view a usage report in Windows SharePoint Services, the data is arranged into a tabular report. You must be a member of the administrator role (or have the **View Usage Data** right) for a site to view the site usage statistics. Summary and detailed usage reports supply information such as the following:

- Number of page hits for each individual page
- Number of unique users
- Browser and operating system information
- Referring domains and URLs

To view single site usage data

To view a site usage report (as shown in 29-2), perform the following steps:

1. Navigate to the Windows SharePoint Services site, and then click **Site Settings**.
2. Click **Go to Site Administration**.
3. Under **Management and Statistics**, click **View site usage data**. Data can be displayed by monthly summary or daily summary.

Use this page to view a detailed usage report for this Web site. The report does not include data for sites under this Web site. To see detailed data for these sites, see their corresponding usage reports. For usage information on all sites in this site collection see the Web site collection usage summary.

Select Report:	OS	Month Summary	Go!	
OS	Page	Recent Month	Most Recent Day	Most Recent Day Visits
Windows Server 2003	OS	12/12/2003		1
	Browser			
	Referrer URL			

Figure 29-2. Viewing single site usage data in the site administration pages

Usage data is processed for an entire site collection on one server at a time. The usage data is collected and stored per site, which is why it can be viewed only per site, not per site collection. Even though the data is logged and stored for an entire site collection, when you view the data in the site administration pages, you can see only the data for a particular website or subsite, not for the entire site collection. You can use the Site Collection Usage Summary page to see the total number of hits for a site collection; for detailed information, you must use the Site Usage Report page for the individual site or subsite. The site collection usage reports also help you track how much storage space your sites are using. This information is gathered as part of the quota tracking for Windows SharePoint Services sites.

To view a site collection usage summary

This report provides a summary of the entire site collection usage. You can view the usage data and storage information for an entire site collection at the same time. (See Figure 29-3.)

1. Navigate to the top-level website, and select **Site Settings**.
2. Select **Go to Site Administration**.
3. Under **Site Collection Administration**, select **View Site Collection Usage Summary**.



Figure 29-3. Viewing the site collection usage summary in the site administration pages

Internet Information Services (IIS) Logs

IIS logs include detailed information—such as who has visited sites and what was viewed—in terms of total visits, average visits, page views, and trends over time. The SharePoint Products and Technologies log files contain a subset of the data available in the IIS logs. This subset contains the information that is most relevant to the SharePoint Products and Technologies

administrator. The IIS logs, however, contain valuable information as well. Careful analysis of the IIS logging data helps you to discover how much traffic is going to portal sites, how much is going to Windows SharePoint Services sites, and how much is going to search operations. The IIS logs also contain hits to pages that are excluded in the Windows SharePoint Services logs, specifically pages in the `_layouts` folders. These folders contain Windows SharePoint Services application pages, which are not important to users trying to track the use of the content in their site.

Configuring the Log Files

IIS logging is enabled by default for each virtual server. The recommended log file format is W3C, which is the default. This enables you to specify which fields are included in the log file. By limiting logging to the W3C fields that are most important to your customer, you can limit the log file size and simplify the analysis. Depending on the amount of traffic to your sites, the size of your log files can begin to consume valuable memory resources and CPU cycles. You need to balance the gathering of detailed data against the need to limit IIS log files to a manageable size and number.

Setting Up Logging

Set the logs to be created on a daily basis, which creates a new log file each day for each virtual server. Ensure that the local Administrators group and the IIS_WPG group have the appropriate permissions to access to the log files directory. Within the Logfiles directory, IIS creates a separate directory for each IIS virtual server log with naming based on the virtual server instance ID.

To view the directory and log file for each virtual server through IIS Manager

1. Right-click a virtual server, click **Properties**, and then select the **Properties** check box next to **Active log format**.
2. The **Log File** name is displayed at the bottom of the screen.

Reading the Log File Data

The IIS logs are ASCII files that can be read using a text editor, but third-party utilities or FrontPage are typically used to analyze IIS and generate meaningful, formatted reports and graphical representations of usage data. How the IIS logging information is presented and accessed depends on the tool used to present the IIS log data. IIS log files can also be logged in a database that complies with Open Database Connectivity (ODBC), such as a Microsoft SQL Server database, and SQL Server Query Analyzer can be used to generate reports.

Viewing Website Reports Using FrontPage

In a previous section, you've seen how to view usage data using the site administration pages. Various website reports, including a report about usage data, is available using the Reports view from FrontPage 2003. The Reports view allows you to display a variety of reports that provide detailed information on the files, shared content, problems, and workflow status of your website, in addition to statistics on site usage. All reports except the usage reports are always available in FrontPage and will not be discussed in this chapter. The usage reports are available only if SharePoint Team Services, FrontPage Server Extensions, or Windows SharePoint Services is installed on a server. If FrontPage is used to view usage reports about Windows SharePoint Services sites, FrontPage retrieves the usage data from the content database.

To display a particular report, activate Reports View by clicking the Web Site tab at the top of the working area of the FrontPage window and then clicking the Reports View button at the bottom of the tab. Then select the particular report you want to see from the drop-down menu at the left end

of the Reports toolbar. (The menu button is labeled with the name of the currently displayed report, such as Site Summary.)

Alternatively, in any FrontPage view you can choose the particular report you want to see from the View, Reports submenu. (See Figure 29-4.) This will activate Reports View and show the selected report.

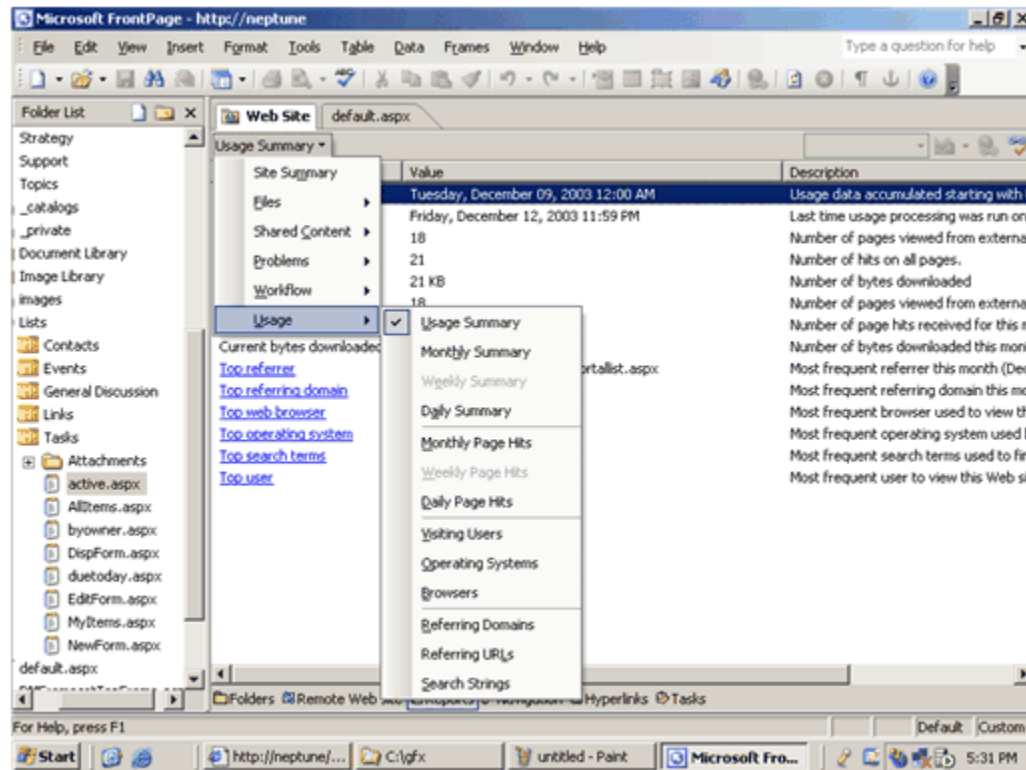



Figure 29-4. Viewing website reports in FrontPage 2003

The usage reports displays activity statistics collected by the Web server. The following table discusses all the reports that are part of the Usage view.

Table 29-1. Usage Reports

Report	Description
Usage Summary	Displays overall statistics for your site collected since inception. The inception date is one of the reported statistics.
Monthly and Daily Summaries	Displays total visits, total page hits, total hits of all kinds, and percentage of hits.
Monthly and Daily Page Hits	Displays, by period, the number of times that Web visitors requested each page in your site.
Visiting Users	Displays the identities of site visitors to your site. However, unless you require Web visitors to identify themselves by user name and password, this report will be blank. In intranet or extranet

Report	Description
	environments, this report provides valuable information.
Operating Systems	Reports how many visits came from computers running Windows 95, Windows 98, Windows NT, Windows 2000, Windows 2003, Macintosh, various forms of UNIX, and so forth.
Browsers	Reports how many visits came from various browsers, such as Microsoft Internet Explorer 5.0 and later, and various versions of Netscape.
Referring Domains	Reports the names of all websites—anywhere—that contain hyperlinks that Web visitors followed to your site.
Referring URLs	Reports the locations of all Web pages that contain hyperlinks that Web visitors followed to your site.
Search Strings	Reports a history of keywords that Web visitors entered on Search forms on one of the large search engines such as Yahoo! or AltaVista. This tells you how people are finding your site and what topics interest them. In intranet or extranet environments, this report will probably remain empty.


 Because weekly usage data is not stored in the content database, the Weekly Summaries and Weekly Page Hits reports are not available when FrontPage reports are used to analyze usage data on a server where Windows SharePoint Services is installed.

When you generate reports in FrontPage, you can save or copy the data to HTML or into other Microsoft applications, such as Microsoft Office Excel. This is useful for archiving report data or for sharing it with others. You can save report data from the following types of reports: Files, Shared Content, Problems, Workflow, and Usage. You do this by clicking **Save As** on the **File** menu. You cannot save the Site Summary and Usage Summary reports. You can copy website report data to another Office program, such as Excel or a Web page in your website, by right-clicking anywhere in the aforementioned reports and clicking **Copy Report**.

Troubleshooting

If you see a “no data” message when you try to view the Site Usage Report page, you should first check whether logging and usage processing is turned on. This is turned off by default. You should also check whether the site is new. If the site is new (created today) or has not been used before today, no data will appear until the usage log processing has been done (usually within 24 hours). After that, make sure there has been activity on your site within the last 31 days. Daily usage data is kept only for the past 31 days, so if there has not been activity you will see the “no data” message.

If you find you’re having trouble processing usage analysis data, you should check whether the SharePoint Timer Service is enabled or needs to be restarted. If you need to enable or restart the SharePoint Timer service on a Web server computer, you can do so by using the **Services** control panel.

 To enable or restart the SharePoint Timer service, you need to be a member of the local computer’s Administrators group or you need to have been granted permissions to manage services on the local computer.

Summary

This chapter shows how to manage usage analysis. The chapter also discusses how to analyze website usage using the site administration pages and website reports in FrontPage. In addition, it shows how to configure IIS log files to paint a complete picture of portal site usage. It also assists you in troubleshooting usage analysis and scheduling the timed jobs responsible for processing the Windows SharePoint Services log files.