

Managing SharePoint Products and Technologies Performance at Microsoft

Note on IT

Published: November 2004

Microsoft deployed Windows® SharePoint® Services and Microsoft® Office SharePoint Portal Server 2003 to create an integrated enterprise-wide portal and team collaboration solution. This Note describes Microsoft IT's post-deployment experience managing the performance of this large SharePoint Products and Technologies environment, lessons learned, and their new best practices.

Document Definition

A Note on IT is a short, technically focused presentation of a specific topic related to Microsoft IT and is usually associated with an existing IT Showcase document. A Note might illustrate how Microsoft IT performs a specific operational task step by step or configures a hardware device or software application. It might also relate technical or operational details of a best practice or other key information regularly requested by customers.

Intended Audience

SharePoint Portal Server 2003 and Windows SharePoint Services administrators, backup operators, and operations support personnel.

Products & Technologies

- Microsoft Windows Server 2003 (including Windows SharePoint Services)
- SharePoint Portal Server 2003
- Microsoft SQL Server 2000
- Microsoft Operations Manager 2005

Introduction

In 2003, Microsoft IT deployed Windows SharePoint Services and Microsoft Office SharePoint Portal Server to create a global enterprise portal and team collaboration solution that supports more than 60,000 users, 250 group and division portals, 50,000 team sites, and manages more than 3 terabytes of information.

This Note presents lessons learned and the best practices Microsoft IT has implemented to address the performance monitoring and management needs of its enterprise portal and team collaboration solution. This Note will be helpful to Microsoft customers deploying SharePoint Products and Technologies in a similar environment. The specific topics covered in this Note include:

- Monitoring Growth and Performance
 - Microsoft Operations Manager 2005
 - SharePoint Report Tool (SPReport)
 - Windows Server™ 2003 Performance Console
- Optimizing Server Performance
 - Web Front-end Server Lessons Learned and Best Practices
 - Back-end Database Server Lessons Learned and Best Practices
 - Server Farm Network Optimization
- Appendix A - SharePoint Report Tool (SPReport)

The detailed design, implementation, and deployment of the centrally hosted collaboration platform is described in the Microsoft IT Showcase white paper *Deploying SharePoint Products and Technologies for Enterprise Collaboration*. Microsoft IT's deployment of SharePoint Shared Services is described in the white paper *Deploying SharePoint Portal Server 2003 Shared Services at Microsoft*; while the creation of Microsoft's internal enterprise portal is documented in *Microsoft Web Enterprise Portal*. All three white papers are available at <http://www.microsoft.com/technet/itshowcase>.

Note: For security reasons, the sample names of forests, domains, internal resources, organizations, and internally developed applications and files used in this document do not represent actual names used within Microsoft and are for illustration purposes only. In addition, the contents of this document describe how Microsoft IT runs its enterprise data center. The procedures and processes included in this document are not intended to be prescriptive guidance on how to run a generic data center and may not be supported by Microsoft Customer Support Services.

Monitoring Growth and Performance

Microsoft IT uses a combination of software tools to measure performance, monitor events, produce operational reports, and respond to issues as they arise. These are summarized in Table 1.

Table 1. SharePoint performance management tools used by Microsoft IT

Management Tasks	Management Tools
Event monitoring Alert notification Problem management	Microsoft Operations Manager 2005 Microsoft Operations Manager Management Packs <ul style="list-style-type: none"> • Microsoft Office SharePoint Portal Server 2003 • Windows SharePoint Services • SQL Server™ 2000 • Microsoft Windows Internet Information Server (IIS) 6.0 • Windows Network Load Balancing (NLB) • Microsoft Cluster Service • Windows Server 2003
Problem determination	SharePoint Portal Server Central Administration diagnostic logs
SharePoint operations reporting	SPReport reporting tool
Real-time performance monitoring Performance logging	Windows Server 2003 Performance Console Microsoft Management Console (MMC) Snap-ins <ul style="list-style-type: none"> • System Monitor ("perfmon") • Performance Logs and Alerts

Microsoft Operations Manager 2005

Microsoft IT deployed Microsoft Operations Manager (MOM) 2005 using the MOM management packs listed in Table 1. In addition to the SharePoint Portal Server 2003 and Windows SharePoint Services MOM management packs, the other key MOM packs used by Microsoft IT were those that helped to identify concerns and resolve issues with the Web front-end and database back-end servers: SQL Server 2000 and Internet Information Server MOM Packs. Microsoft IT used MOM 2005 to monitor performance monitor counter out-of-bounds conditions as well as selected events written to the Windows event log by a Windows SharePoint Services, SharePoint Portal Server 2003, Windows Server, or SQL Server 2000 running on the front-end or back-end servers.

The key performance counters that Microsoft IT uses to monitor the enterprise-wide SharePoint Products and Technologies infrastructure are listed in Table 2. Their interpretation is described in the section *Optimizing Server Performance*.

Table 2. Key performance counters used by Microsoft IT

Server Role	Counter	Category
All servers	Memory\Available Mbytes	Memory*
	Memory\% Committed Bytes in Use	Memory*
	Memory\Pages/second	Memory
	System\Processor Queue Length	Processor
	Processor\% Processor Time	Processor*
	Network\Network Queue Length	Network
Web front-end server	Process (w3wp)\% Processor Time	Processor*
	Process (w3wp)\Private Bytes	Processor*
	Web Service\Connection Attempts/second	Application load*
	ASP.NET\Worker Process Restarts	Application load
	Web Service\Current Connections	Application load*
	Web Service\Bytes Sent/Sec (all instances)	Network
	Web Service\Bytes Received/Sec (all instances)	Network
	.NET CLR Data\SQL Client: Pooled and NonPooled Connections (instance _global_)	Application load
	.NET CLR Data\SQL Client: Failed Pooled Connections (instance _global_)	Application load
Back-end database server	Process (sqlserver)\% Processor Time	Processor*
	Process (sqlserver)\Private Bytes	Memory*
	SQL Server\SQL Server: General Statistics: Current Connections	Database*
	SQL Server\SQL Server: General Statistics: Logins/Sec	Database*

***Note: Microsoft IT has used this counter as a "baseline performance indicator" for this category. These counter values are tracked over time to determine typical values and out-of-bounds limits.**

MOM 2005 was used to monitor, track, and respond to a selected set of critical events based on the above list of key performance counters as well as specific application and systems events. The list of critical events appears in Table 3.

Table 3. Critical events monitored by Microsoft IT using MOM 2005

Event Source	Event
Windows SharePoint Services	Cannot connect to database
Windows SharePoint Services	Database capacity warning reached
Windows SharePoint Services	Database capacity reached
Web front-end servers	W3WP error

Web front-end servers	IIS start and stop times
Web front-end servers	Windows SharePoint Services stop time
Web front-end servers	Web Service\Connection Attempts/second > 500
Back-end database servers	Standard SQL Server 2000 events
All servers	Memory\Available Mbytes < 50 MB
All servers	Processor\% Processor Time: _Total(CPU Utilization) > 80%
All servers	Disk usage < 10%
All servers	Memory\Pages/second

Additional information on addressing the above events can be found in the section *Optimizing Server Performance*.

SharePoint Portal Server Central Administration Diagnostic Logs

The Configure Diagnostics Settings Web administration page in SharePoint Portal Server Central Administration Web console can be used to configure and view the application log files for the following SharePoint Portal Server 2003 and Window SharePoint Services processes:

- WWW Worker Process (W3WP)
- Notification Service
- Single Sign-on Service
- Administrative Service
- Search Service
- Search Filter Process
- Backup and Restore
- Audience Compilation
- Portal Creations
- Portal Deletions

SharePoint Report Tool (SPReport)

SPReport is a command-line tool developed by a project manager in the SharePoint product group to gather detailed statistics about the number of SharePoint sites, the number and types of lists created in each site, and the number and types of documents stored in a SharePoint server farm. SPReport was originally used by Microsoft IT and selected Microsoft customers to gather statistical data of interest to the SharePoint Product group. Microsoft has since made the SPReport source code available as a community-supported tool on the GotDotNet .NET Framework Community Web site at

<http://workspaces.gotdotnet.com/spreports>. The only technical support available for SPReport is the community-based support available on the GotDotNet Web site.

The SPReport data is can be saved to an Extensible Markup Language (XML) file and optionally formatted as more readable Hypertext Markup Language (HTML) report. More

information on obtaining and using SPReport can be found in *Appendix A - SharePoint Report Tool (SPReport)*.

Microsoft IT uses SPReport to identify "unused sites" - sites that have been created but not used. Technically, an unused site was defined as a SharePoint site if it contained only "unused lists" - lists that had not been accessed after they were created. Based on data gathered by SPReport, Microsoft IT estimates that 20 to 40 percent of the team sites in a server farm had never been used after being created.

To address its SharePoint site lifecycle management needs, Microsoft IT uses the *Configure Site Collection Use Confirmation and Auto-Deletion* feature in Windows SharePoint Services to automatically send notices to site collection owners requiring them to confirm that their sites collections are in use. This feature can be configured to delete unconfirmed site collections automatically.

Windows Server 2003 Performance Console

Performance Console is a Microsoft Management Console (MMC) application configured with two MMC snap-ins: System Monitor ActiveX control, and Performance Logs and Alerts. It can be found on the Administrative Tools program menu. System Monitor (previously known as Performance Monitor, or "perfmon") provides a real-time display of performance monitor data. *Counter Logs* provide capabilities for recording performance data based on a specific time interval. *Trace Logs* in the Performance Console enables recording of performance data based on specific events such as a performance counter value reaching, exceeding, or falling below a defined limit. Alerts can be used to send notifications to operations staff when an event occurs. Organizations that have deployed MOM 2005 would use the appropriate MOM management packs to capture and log this information.

Optimizing Server Performance

Globally, Microsoft IT deployed three regional server farms to support Microsoft employee collaboration needs. Each regional server farm is based on a similar design consisting of:

- Load-balanced Web front-end servers running Windows SharePoint Services
- Clustered back-end SQL Server 2000 database servers
- Shared storage area network (SAN) for database storage.

The Redmond, WA server farm, depicted in Figure 1, included additional servers to host Microsoft Web, the central enterprise portal used by Microsoft employees. Specific details on how Microsoft Web was designed and deployed can be found in the IT Showcase white papers listed in the *For More Information* section at the end of this Note.

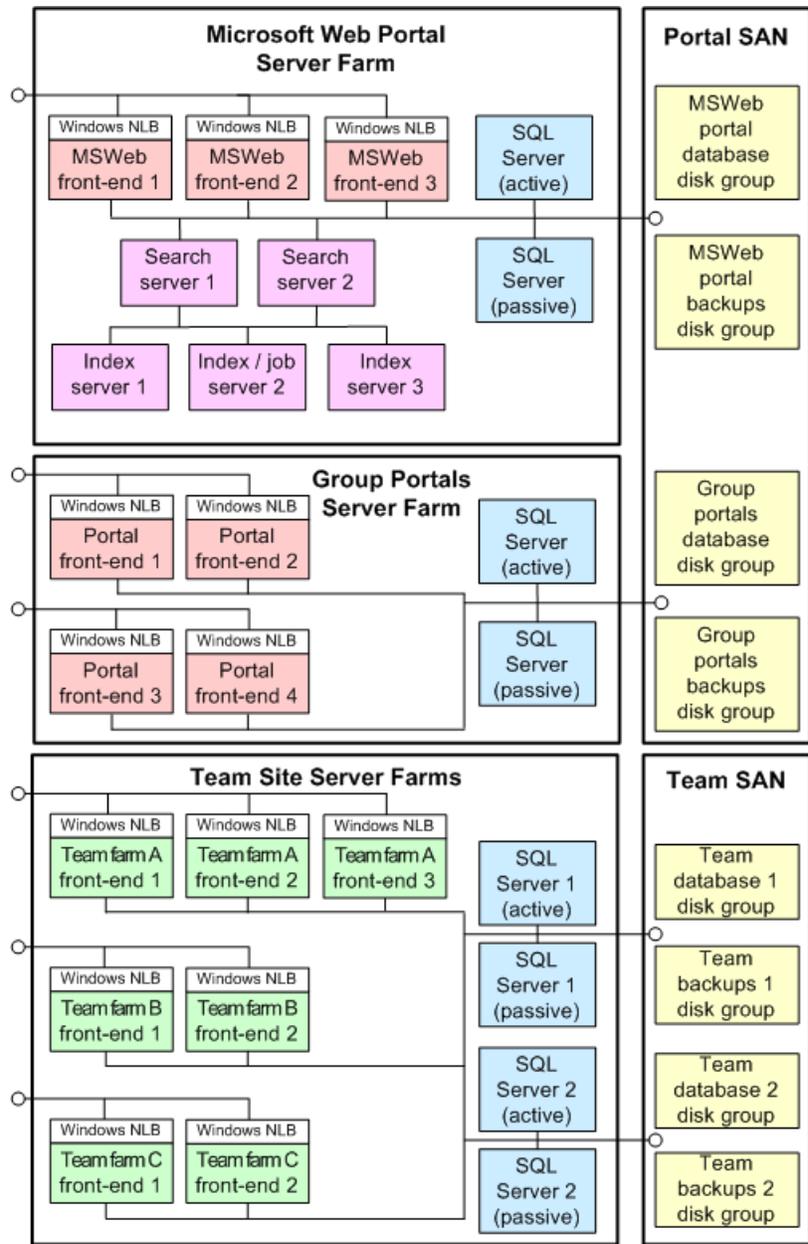


Figure 1. Redmond, WA portal and team site farm design

To increase the overall performance of the server farm, Microsoft IT monitors and manages each server based on its role. Overall performance of the SharePoint infrastructure at Microsoft is affected by three key system components:

- SharePoint Web front-end servers
- Clustered back-end SQL Server 2000 database servers
- Server farm network architecture.

Web Front-end Server Lessons Learned and Best Practices

The following list of lessons learned and best practices are based on Microsoft IT's experience in managing several high-volume, load-balanced Web front-end server pools.

Causes of sustained high CPU utilization

When a Web front-end server has sustained CPU utilization greater than 80 percent, Microsoft IT has found the most common cause is a worker process that is hung or is waiting for a back-end server to respond; and is thus unable to cycle itself. The fix is to cycle the worker process.

A sustained non-zero value for *System\Processor Queue Length* counter also indicates a problem with a worker process. On a Web front-end server, a non-zero processor queue length counter value means that there are processes waiting to run. If processor CPU utilization also remains high, it may be a single worker process that needs to be cycled. To identify the problem process, use the Windows Performance Console to monitor the *Process (w3wp)\% Processor Time* counter.

Disk space usage

If disk space on a Web front-end server drops below 10 percent, Microsoft IT observed that the most common symptom was that Web Parts failed to render. The most common cause is the operating system volume is too small for a large, high-usage Windows Server 2003 server running ASP.NET and Windows SharePoint Services.

Microsoft IT typically configures the Web front-end servers with 4 gigabytes (GB) of memory requiring a page file of at least 6 GB. In addition, the ASP.NET Web server cache on a large, high-usage server can use several additional GB of disk space. Hence, the total disk space requirement for the operating system files on a Windows Server 2003 server running ASP.NET and Windows SharePoint Services is larger than the typical operating system partition size of 5 to 10 GB. To address this issue, Microsoft IT moved the page file to a second disk volume. In addition, the Web front-end servers were configured to use a larger operating system partition (18 GB of mirrored disk storage).

Increase the amount of physical memory and the available address space on Web front-end servers

For best performance, Microsoft IT configured its Web front-end servers with 4 GB of physical memory and used the /3GB switch to increase the address range of the server. More information can be found in the Microsoft Knowledge Base article *Available switch options for the Windows XP and the Windows Server 2003 Boot.ini files* available at <http://support.microsoft.com/default.aspx?scid=kb:en-us:833721>.

Internet Information Server (IIS) virtual server and application pool configuration

In an ideal situation, each virtual server is configured with its own application pool, but this is not always possible when large numbers of high-usage sites need to be configured on a Web front-end server.

To configure its Windows SharePoint Services Web front-end servers, Microsoft IT uses a *Web garden* design. A Web garden is an application pool that is configured with more than one worker process. (Web gardens are to be distinguished from Web farms, which use multiple servers for a Web site.) More information on Web gardens can be found in *Configuring Web Gardens* in the *IIS 6.0 Technical Reference* available at

http://www.microsoft.com/resources/documentation/IIS/6/all/techref/en-us/IISRG_WAS_31.msp.

As a general rule, for best performance Microsoft IT prefers to minimize the number of virtual servers configured on a Web front-end server. This reduces the size of the IIS metabase as well as the amount of server CPU and memory utilization. Based on its experience, Microsoft IT has established maximum limits for key configuration parameters on its high-usage Web front ends. These include:

- Number of worker processes per server
- Worker process virtual memory
- Number of application pools per server
- Maximum number of virtual servers per server
- Number of virtual servers per application pool.

The parameter settings are based on the following Microsoft IT server hardware configurations.

Table 4. Server Hardware Configurations

Server Role	Configuration
Web Front-end Server	Dual Intel Xeon 3.0 GHz processors 4 GB RAM 1 GB Network Interface Card (NIC) 100 MB NIC for Network Load Balancing communication
Index Server	Dual Intel Xeon 2.3 GHz processors 2 GB RAM 1 GB NIC
Search Server	Dual Intel Xeon 2.3 GHz processors 2 GB RAM 1 GB NIC
Back-end Database Server Node (medium server)	Quad Intel Xeon 2.3 GHz processors 4 GB RAM 1 GB NIC 100 MB NIC for cluster communication
Back-end Database Server Node (large server)	Quad Intel Xeon 3.0 GHz processors 8 GB RAM 1 GB NIC 100 MB NIC for cluster communication

Server memory configuration

The Web front-end servers were configured to use 4 GB of physical memory. One GB is assumed to be allocated to the operating system. This leaves up to 3 GB of memory available for use by ASP.NET worker processes (and other system services).

Worker processes

Using the Internet Information Server (IIS) Manager, Microsoft IT configures each worker process to use a maximum of 800 MB of physical as well as virtual memory. These new IIS 6.0 settings are set using the application pool's cycling properties in IIS Manager (illustrated in Figure 2). If a worker process exceeds these limits, the process is automatically recycled. The maximum limit of 800 MB is based on the desire to support a relatively large number of virtual servers with a single worker process.

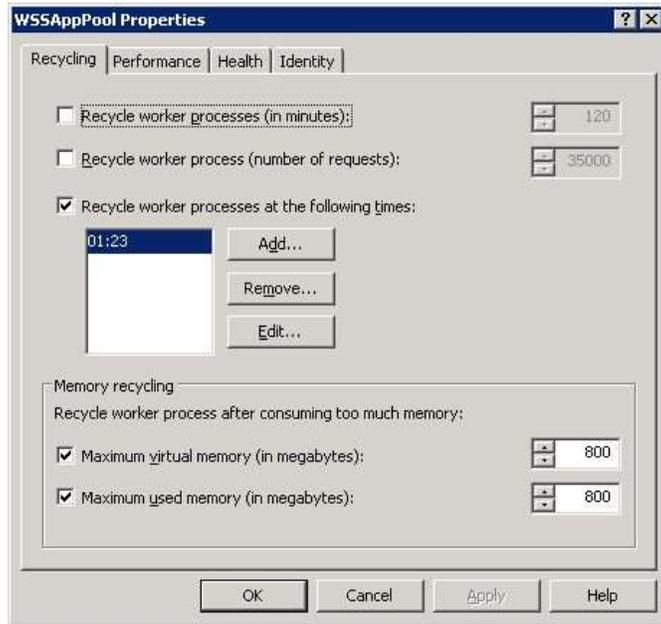


Figure 2. IIS Manager application pool cycling properties

With the default IIS application pool configuration, a worker process is shutdown if IIS cycles it five times in a row (*auto-rapid failure*). Microsoft IT has disabled auto-rapid failure to prevent employees from experiencing unexpected service outages. In this situation, MOM 2005 is used to notify Microsoft IT operations when cycling of work processes occurs repeatedly. Disabling auto-rapid failure coupled with automated monitoring is preferable to having a Web front-end being shut down in the middle of the night. Auto-rapid failure is configured using the application pool's health properties in IIS Manager (see Figure 3).

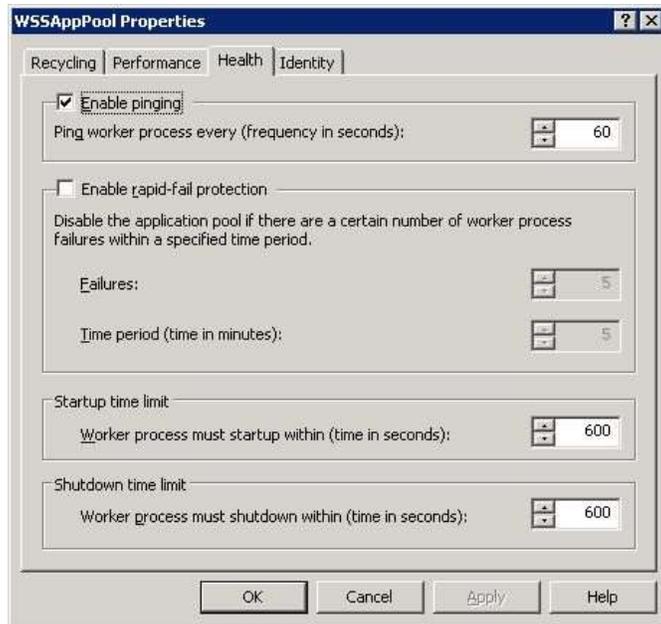


Figure 3. IIS Manager application pool health properties

Microsoft IT also found that worker processes on a heavily used Web front-end server would take more than 30 seconds to shutdown. To reduce the number of event log messages (and MOM 2005 notifications), Microsoft IT increased the worker process startup and shutdown time limits from 30 seconds to 5 minutes.

Number of application pools

With a maximum of 3 GB of virtual memory available to host worker processes and a maximum worker process size of 800 MB, Microsoft IT found that it can configure a Web front-end server with an average workload to support a maximum of three application pools with one worker process per application pool, or a single application pool with a maximum of three worker processes.

On a low-usage Web front-end server, Microsoft IT found the maximum number of application pools can be increased four, with one worker process each; or a single application pool with a maximum of four worker processes. The number of worker processes per application pool is configured using the application pool's performance settings in IIS Manager (illustrated in Figure 4).

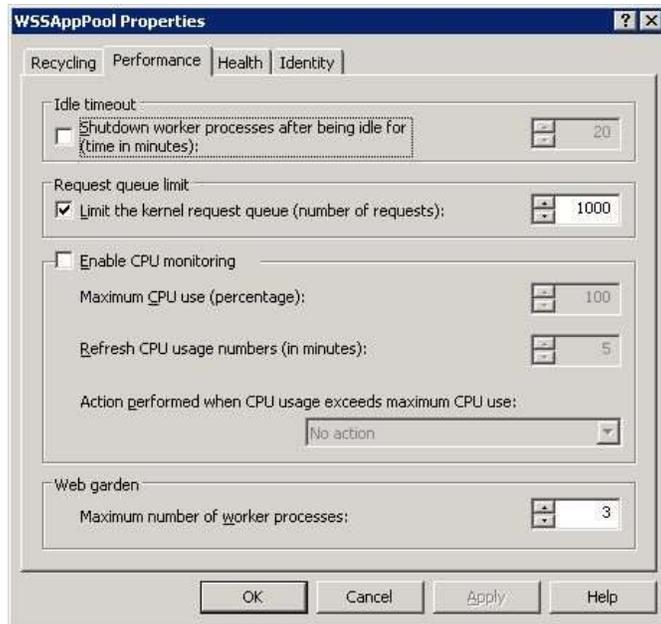


Figure 4. IIS Manager application pool performance properties

One virtual server or 50 virtual servers?

Microsoft manages several large, internal Windows SharePoint Services server farms hosting 10,000 sites or more. A majority of these deployments use a "Web garden" approach in their design: a single virtual server, single application pool, and multiple worker processes (the actual number is based on the amount of Web front-end server memory).

In contrast, Microsoft IT also manages a Windows SharePoint Services server farm of similar size that was created to consolidate 10,000 global sites into a single Windows SharePoint Services server farm. This farm is configured to use 50 virtual servers, four application pools and 100 content databases.

Based on this experience, Microsoft IT has found that a server-farm design based on Web gardens had:

- Better overall performance
- Lower total cost of ownership because of fewer databases to manage
- Lower support costs when it came to managing individual SharePoint sites

Number of virtual servers

SharePoint Portal Server 2003 supports up to 100 portal sites per server with each portal site requiring its own IIS virtual server. Microsoft IT also uses this value for the maximum number of virtual servers to be configured on its Windows SharePoint Services team farms (with each virtual server supporting a maximum of 50,000 site collections).

Theoretically, if a Web front-end server was configured to support 100 virtual servers with one application pool per virtual server, and at least one worker process per application pool, the high-usage server would require as much as 80 to 100 GB of memory. Given this, Microsoft IT worked to find a more optimal configuration that offered high-performance and scalability with more reasonable memory requirements.

Microsoft IT configured its Web front-end servers with 4 GB of physical memory. Given that a Web front-end server with 4 GB of memory can support up to four application pools (with one 1 GB worker process per application pool), Microsoft IT chose to configure a maximum of 25 virtual servers per application pool to meet the goal of supporting up to 100 virtual servers per Web front-end server (4 application pools per server with 25 virtual servers per application pool). Microsoft IT chose to deploy the following two Web front-end server configurations:

- In high-usage Windows SharePoint Services server farms, Microsoft IT found the optimal configuration for their Web front-end servers was one virtual server using one application pool and 3 work processes (see Figure 4).
- In the SharePoint Portal Server 2003 server farms, the Web front-end servers are configured with four application pools collectively supporting a maximum of 100 virtual servers (100 portal sites).

These configurations represent the trade-offs that can be made between the number of application pools and the number of virtual servers hosted by each application pool given

specific assumptions about the amount of server physical memory, the memory requirement per worker process, and the number of worker processes per application pool.

Configure worker processes to cycle once each day for problem sites

By default, IIS configures worker processes to cycle once each night (see Figure 2). For a virtual server with a problem site or a worker process that consistently uses large amounts of virtual memory (800 MB or more), Microsoft IT configures the worker processes in the application pool to cycle more frequently (for example, two or three times per day). Cycling more frequently is not recommended. The primary reason is that the ASP.NET cache is flushed each time a work process is cycled, which reduces performance.

Configure an extra application pool for isolating and troubleshooting virtual server issues

For troubleshooting purposes, a SharePoint site or a custom Web Part may create performance or reliability issues for a virtual server. One approach used by Microsoft IT to locate the problem is to host the suspect virtual server in its own application pool until the problem has been fixed.

Separate application pools can also be used for a portal site or team farm virtual server that have a significant amount of customization or that need to run under a separate service account.

Back-end Database Server Lessons Learned and Best Practices

In the Microsoft Web and team site server farms at Microsoft, the back-end database servers are deployed as two-node, active-passive SQL Server 2000 clusters connected to a shared storage area network (SAN) solution. This represented a trade-off in favor of simplified configuration and ongoing operation support versus using an active-active-passive multi-node cluster configuration that would have reduced server hardware costs.

The following lessons learned and best practices are based on Microsoft IT's direct experience managing server farms capable of supporting tens of thousands of sites consisting of several terabytes of physical storage.

Storage area network design

To lower server hardware and maintenance costs Microsoft IT configured the SAN environments to be used by four SQL Server clusters supporting five large portal and team site server farms (see Figure 1). Each SQL Server 2000 cluster was allocated 2 terabytes of database storage.

The three team farms which collectively host more than 20,000 SharePoint sites recently grew to exceed 2 terabytes in storage. Hence, Microsoft configured the three team farms to share a second SQL Server cluster and a second set of team SAN disk groups. The overall design is illustrated in Figure 1.

Use disk group pivoting

Disk group pivoting is a term that Microsoft IT uses to refer to the automated daily failover of a secondary SAN disk group from an active SQL Server node to the passive node so that the passive node can be used to backup database dumps on the disk group during non-business hours.

In the context of the Microsoft IT SharePoint infrastructure, disk-to-disk SQL Server database backups are written to a backup disk group by the active SQL Server node each evening.

The backup disk group is then *pivoted* (scheduled to fail over) to the passive SQL Server node to enable it to backup the disk backup files to streaming tape. The backup disk group is scheduled to pivot back to the active SQL Server node each morning. The typical configuration used at Microsoft is illustrated in Table 5.

Table 5. SQL Server disk group allocations for disk pivoting

SAN disk group	Usage	SQL server node this disk group is allocated to	
		Days/evenings	Nights
Data disk group	Database and log file disk storage	Active node A	Active node A
Backup disk group	Database dump and backup disk storage	Active node A	Passive node B

The data disk group was used for primary database storage, and during normal operations it was allocated to the active SQL Server node A. If node A was unavailable for a scheduled or unscheduled event, the disk group would automatically fail over to the passive SQL Server node B to maintain high availability of the SharePoint SQL Server databases.

The backup disk group was managed in a different manner. During the evening, the backup disk group was allocated to the active SQL Server node A, which used it to create online backups of the SharePoint databases found in data disk group. At night, however, the backup disk group was scheduled to automatically fail over to the passive SQL Server node B. During normal operation, this enabled the passive SQL Server node B to be used for disk-to-tape backup operations (rather than remaining idle as a standby SQL Server node). When the backup was complete, the passive SQL Server node B would then cause the backup disk group to fail back to the active SQL Server node A.

In the case where the active SQL Server node A became unavailable, the backup disk group in addition to the data disk group would fail over to the passive SQL Server node B.

Target moderate-sized content databases

Microsoft IT chose 25 GB as the target content size for content databases. This ideal database size was based on Microsoft IT's experience in managing the backup, restore, and recovery processes for large numbers of SharePoint sites in multiple server farms, as well as for the server farm design they deployed.

For the *Database capacity warning reached* and *Database capacity reached* limits for Windows SharePoint Services sites, Microsoft IT uses settings of 200 sites and 250 sites respectively. The design was based on a historical average at Microsoft of 60 MB of SQL Server storage per SharePoint site collection (projected to be 100 MB per site in three years) and the target maximum database size of 25 GB.

In addition, the maximum size of an individual SharePoint site can also be controlled. Microsoft IT set the storage limit value for individual sites to 500 MB. Microsoft IT recommends 2 GB as the maximum size of a SharePoint site.

Limit the maximum number of databases per SQL Server instance

For Windows SharePoint Services team sites, there must be at least one SQL Server database allocated per site collection. For SharePoint Portal Server 2003 portal sites, there are three databases per portal site and one portal site can be configured per virtual server.

Based on this, Microsoft IT's internal guideline is to limit the number databases per SQL Server instance to 300 (equal to three times the maximum number of virtual servers per Web farm).

Deploy additional content databases when the size of a database approaches the target content database size

At Microsoft, new SharePoint sites are being created at an average rate of 500 team sites per month. Given this growth rate as well as the additional storage requirements of its largest sites (that is, sites with more than 100 MB of content), each month Microsoft IT needs to allocate approximately 50 GB of new SQL Server database storage to support the average expected storage requirements for these sites. To accomplish this, in a typical month Microsoft will allocate two new 25 GB databases to the team site server farms.

When the size of a content database is approaching the target database size, Microsoft IT reduces the *Database capacity reached* limit on the existing database to a value less than the current number of sites in the database. This prevents new sites from being created in the previous database and forces them to be created in a database that is not at its *Database capacity reached* limit.

Ensure Available Memory is greater than 50 MB for back-end database and Web front-end servers

If the *Memory\Available MB* performance counter drops below 50 MB on a sustained basis on either the Web front-end or database servers, the most common scenario is that users will experience failures while uploading large files into a SharePoint document library. If the *Maximum Upload Size* is increased from the default value of 50 MB, the acceptable amount of available memory should be similarly increased.

For more information on large file support in Windows SharePoint services, see *Configuring Large File Support* in the Windows SharePoint Services Service Pack 1 version of the *Windows SharePoint Services Administrator's Guide* available at http://www.microsoft.com/resources/documentation/wss/2/all/adminguide/en-us/std05_SP01121517.msp.

Monitor disk space usage

If free disk space on a back-end database drops below 10 percent, file uploads are more likely to fail. In the case of a low SAN storage situation, the short-term solution is to extend the SQL Server database to use an additional disk volume. The longer term solution is to move the database to a new disk volume.

Increase the amount of physical memory and the available address space on back-end database servers

Microsoft IT usually configures its back-end database servers with 8 GB of physical memory and uses the /3GB switch in the Boot.ini file to increase the address range of the server. In addition, the /PAE switch in the Boot.ini and the AWE enable option in SQL Server were used to allow SQL Server 2000 to utilize more than 4 GB of memory. Without the /PAE switch SQL Server can only utilize up to 3 GB of memory. More information can be found in the Microsoft

Knowledge Base article *How to configure memory for more than 2 GB in SQL Server* available at <http://support.microsoft.com/default.aspx?scid=kb;en-us;274750#7>.

In situations when an existing pair of back-end database servers is configured with 4 GB of physical memory per server and the server memory cannot be easily upgraded, Microsoft IT often chooses to deploy a second SQL Server instance (active/passive SQL Server cluster).

Increase the amount of virtual memory available to back-end database servers when hosting a large number of databases

In addition, SQL Server 2000 was configured to use the -g512 startup parameter. The -g startup parameter leaves additional Memory-To-Leave (MTL), unreserved virtual memory available for database allocations. More information can be found in the Microsoft Knowledge Base article *PRB: There May Not Be Enough Virtual Memory with Large Number of Databases* available at <http://support.microsoft.com/default.aspx?scid=kb;en-us;316749>.

Use file servers, Windows Distributed File System (DFS) and storage area networks (SANs) for long-term archiving

Microsoft IT prefers to use file servers and DFS for archiving large numbers of documents that require online, non-collaborative access. For large document collections, SANs are used for physical storage. File servers are the preferred way for storing:

- Microsoft IT-managed software product distribution files
- Scripts and executable files
- Personal database files (such as Microsoft Access database (.mdb) files large Outlook personal store (.pst) files)
- Large audio and video media files
- Other files greater than 100 MB in size (for example, disk imaging files).

Blocked file type list changes

Microsoft IT made minor modifications to the list of file types that are blocked from being uploaded to a SharePoint document library:

- MP3 audio (.mp3) and Outlook personal store (.pst) were added to the blocked file type list because these files can be larger and there are very few business reasons for needing to upload MP3 files to a SharePoint site at Microsoft.
- Compiled online help (.chm), desktop shortcut (.lnk) and Internet shortcut (.url) were removed from the blocked file type list.

Additional Considerations

Occasionally, there were additional performance trade-offs that needed to be considered. One decision to be made was whether to leverage the "blocked file types" feature in Windows SharePoint Services versus enabling anti-virus scanning on the SQL Server database servers. The blocked file types feature can be used to prevent executable files such as .exe and script files from being uploaded and stored in a SharePoint document library. This helps reduce the risk of viruses being spread through the use of team sites, especially when used in conjunction with an effective desktop virus scanning strategy. Because anti-virus software increases overall server utilization, Microsoft IT chose to exclude the following database paths from being scanned by the anti-virus software (in addition to the defaults):

- MS Search full text indexes
- SQL Server 2000 databases
- SQL Server 2000 logs

Server Farm Network Optimization

When the network adapter utilization or network queue length counters are above normal, it is an indication that the server farm network design and/or network usage patterns need to be optimized.

Microsoft IT found it was very important to be aware of all the different types of network traffic on the server farm local area network. This includes front-end HTTP requests and responses

for SharePoint site, document library and list content, back-end database queries, database backup and restore operations (both disk-to-disk and disk-to-tape operations), user profile synchronization, and content indexing and search related traffic.

Optimizing network capacity and network design

In the earlier stages of the SharePoint Product and Technologies deployment, Microsoft IT observed that end users were occasionally experiencing slow Web server response times. The initial solution was to add a second 100 MB network interface card (NIC) to each Web front-end server. The first NIC handled incoming HTTP requests from users, while the second NIC was used by the Web front-end servers to connect the back-end database servers.

However, as additional Web front-end servers were added to a server farm to support additional user connections, performance as perceived by users did not improve. The root cause of the problem was the combination of the increased user HTTP requests and back-end database access traffic coupled with the fact that the Microsoft IT disk-to-tape backup service was also using the same network segments during peak-use hours. Microsoft ultimately upgraded the bandwidth of the server farm network to 1 GB. Network utilization dropped to approximately 10 percent - a decrease of 60 to 80 percent - and the simpler network topology was easier to manage. An alternative approach is to deploy a virtual local area network (VLAN) to interconnect a second NIC on the Web front-end servers with the NIC on the back-end database servers. More information on this configuration can be found in the white paper *Windows SharePoint Services Network and Load Balancing Design* available at <http://www.microsoft.com/downloads/details.aspx?FamilyID=cb0d258f-2815-4af6-bb6d-d6ac20b34173&DisplayLang=en>.

.NET Common Language Runtime (CLR) SQL client database connections

On the Web front-end servers, *.NET CLR Data\SQL Client: Pooled and NonPooled Connections* and *.NET CLR Data\SQL Client: Failed Pooled Connections* are good counters for measuring the relative performance of the connections between the Web front-end servers and the back-end SQL server database servers. If the failed pooled connections counter is high relative to the number of open Web front-end connections or overall server baseline performance, Microsoft IT has found that the server farm network design and available network bandwidth were often the causes. This problem was solved by optimizing the network capacity or network design using the approaches describing previously.

Search indexing service

If unexplained increases in network bandwidth utilization during specific periods of the day occur, Microsoft IT will use the Web front-end server IIS logs and Windows Network Monitor to determine if the source of the increased network traffic is from a SharePoint index server.

Microsoft IT uses several approaches to optimize peak and overall network bandwidth used for content indexing:

- Reduce the number of concurrent indexer threads that can be active. This is configurable on a site-by-site basis.
- Avoid indexing the same content more than once. Consider centralizing the deployment of index servers for content of enterprise-wide interest. (SharePoint Portal Server 2003 only)

- Use scheduled incremental crawls to reduce the indexing times. (SharePoint Portal Server 2003 only)
- Use a small number of content indexes. Use a separate (smaller) content index for managed content because smaller indexes produce higher relevancy search results. Use larger content indexes for unmanaged content (for example, for indexing file servers and Internet sites). (SharePoint Portal Server 2003 only)

Summary

While each resource in a SharePoint server farm has its own set of performance counters and events to monitor and manage, Microsoft IT, through its experiences operating a large integrated enterprise-wide portal and team collaboration solution, has found that there are four primary areas of focus for performance management:

- Physical and virtual memory configuration of the Web front-end and back-end database servers
- Configuration of the Web front-end server application pools, number of worker processes, and virtual servers
- Number and size of the back-end server databases
- Bandwidth and configuration of the server farm network connections.

For More Information

Microsoft IT Showcase White Papers

- Deploying SharePoint Products and Technologies for Enterprise Collaboration, <http://www.microsoft.com/technet/itshowcase>
- Microsoft Web Enterprise Portal, <http://www.microsoft.com/technet/itshowcase>
- Deploying SharePoint Portal Server 2003 Shared Services at Microsoft, <http://www.microsoft.com/technet/itshowcase>
- Monitoring Enterprise Servers at Microsoft, <http://www.microsoft.com/technet/itshowcase>

Microsoft SharePoint Products and Technology

- Microsoft SharePoint Products and Technology Product Web site, <http://www.microsoft.com/sharepoint/>
- TechNet: Microsoft SharePoint Products and Technologies, <http://www.microsoft.com/technet/prodtechnol/sppt/default.aspx>
- Monitoring Events for Windows SharePoint Services, Administrator's Guide for Windows SharePoint Services, <http://www.microsoft.com/resources/documentation/wss/2/all/adminguide/en-us/stsg04.aspx>
- Monitoring Performance, Administration and Scripting Tools Online Help for Windows Server 2003

Microsoft Operations Manager

- Microsoft Operations Manager Product Web site, <http://www.microsoft.com/mom/default.aspx>
- Microsoft Operations Manager Management Packs, <http://www.microsoft.com/management/mma/catalog.aspx>

- Microsoft Office SharePoint Portal Server 2003 Management Pack Guide, <http://www.microsoft.com/downloads/details.aspx?FamilyID=aebc81a9-53de-425c-a22d-d634a99ea5f4&DisplayLang=en>
- Microsoft Windows SharePoint Services Management Pack Guide, <http://www.microsoft.com/downloads/details.aspx?FamilyID=e0f42659-e9db-4dff-ae6a-59f6c4eb04e1&DisplayLang=en>

Microsoft Product Support

- How to configure memory for more than 2 GB in SQL Server, <http://support.microsoft.com/default.aspx?scid=kb;en-us;274750>
- Available switch options for the Windows XP and the Windows Server 2003 Boot.ini files, <http://support.microsoft.com/default.aspx?scid=kb;en-us;833721>

Microsoft Sales Information Center

For more information about Microsoft products or services, call the Microsoft Sales Information Center at (800) 426-9400. In Canada, call the Microsoft Canada information Centre at (800) 563-9048. Outside the 50 United States and Canada, please contact your local Microsoft subsidiary. To access information through the World Wide Web, go to:

<http://www.microsoft.com>

<http://www.microsoft.com/itshowcase>

<http://www.microsoft.com/technet/itshowcase>

For any questions, comments, or suggestions on this document, or to obtain additional information about Microsoft IT Showcase, please send e-mail to:

showcase@microsoft.com

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. Microsoft grants you the right to reproduce this White Paper, in whole or in part, specifically and solely for the purpose of personal education.

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.

Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, email address, logo, person, place or event is intended or should be inferred.

© 2004 Microsoft Corporation. All rights reserved.

This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY. Microsoft, Windows, and Windows Server 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.

Appendix A - SharePoint Report Tool (SPReport)

SPReport is a command-line tool developed by a project manager in the SharePoint product group to gather detailed statistics about the number of SharePoint sites, the number and types of lists created in each site, and the number and types of documents stored in a SharePoint server farm. SPReport was originally used by Microsoft IT and selected Microsoft customers to gather statistical data of interest to the SharePoint Product group. Microsoft has since made the SPReport source code available as a community-supported tool on the GotDotNet .NET Framework Community Web site at <http://workspaces.gotdotnet.com/spreports>.

The report data is can be saved to an Extensible Markup Language (XML) file and optionally formatted as more readable Hypertext Markup Language (HTML) report.

Background

SPReport is a command-line tool developed by the SharePoint product group to help Microsoft IT gather detailed statistics about the number of SharePoint sites created in a server farm, the number and types of lists created in each site, and the number and types of documents stored in each SharePoint document library.

Technical Support

Technical support for SPReport is the community-based support available on GotDotNet Web site. No formal support for SPReport is available from Microsoft Customer Support Services.

Using SPReport

SPReport is a command-line utility that gathers metrics about a site and its subsites, the document library and other lists in each site, and the documents in each document library.

SPReport uses the SharePoint object model to retrieve information about the current configuration and hence, the tool needs to be run locally from a SharePoint Web front-end server. The current version of SPReport cannot be run remotely.

Table 6 describes how to use SPReport and the command line options it supports.

Table 6. SPReport command line options

```

SPReport -g | -v <virtualserver> | -s <site> [-f <filename>] [-e]

Operations and Parameters:

-g Report for all sites on this farm.

-v Report for all sites on this virtual server.

-s Report for this site and all sites underneath it.

-f Output to this file name prefix - default is sprOut.

-e Exclude subsites from reporting.

-p Privacy mode - do not record URLs in output.

-t Test mode - include unused lists and sites in data reporting.

-n Number of threads to use - default is 5

-x Uses an existing xml and xsl file to create a new htm.

-m Merge results from several reports into one.

```

The following table is a list of SPReport command line examples.

Table 7. SPReport command line examples

Command line	Description
<code>SPReport -v http://myvs</code>	Gather metrics and report on all of the sites on the virtual sever myvs
<code>SPReport -s http://mysrv/site</code>	Gather metrics and report on this site and all subsites beneath it
<code>SPReport -g -f fileName</code>	Gather metrics and report on all sites on the local server
<code>SPReport -x -f fileName</code>	Generate a new HTML report using the existing SPReport XML data and XSL transform files
<code>SPReport -m sprOut1.xml;sprOut2.xml;sp rOut3.xml -f sprOutMerge</code>	Merge several SPReport XML data files together to create a single file called sprOutMerge.xml

SPReport Sample Statistics

The following tables are samples of the statistics gathered by SPReport. SPReport summarizes the SharePoint statistics at following levels of detail: virtual server, SharePoint site type, SharePoint list type, and document type.

Table 8. Summary of all Site Types sample statistics

SharePoint Reports Summary

<p>Host Name: PSN-W23-012C, teamsites.contoso.ca Report Date/Time: 10/8/2004 2:52:50 PM Version: 7_13_2004 Total Site Collections Scanned: 5</p>
<p>Summary of all Site Types Total Number of In-Use Sites: 8 Total Number of Unused Sites: 15 Percent Large (250MB+): 0% Percent Medium (5-250MB): 0% Percent Small (5MB-): 100%</p>
<p>Percent Frequently Visited (100+ visits per day): 0% Percent Rarely Visited (< 1 visit per day): 50% Percent Increasing in Popularity: 100% Percent Decreasing in Popularity: 0% Average Number of Visits per Site (in the past Month): 34 Percent of Visits made to the Home Page: 32%</p>
<p>Percent with a Customized (Unghosted) Home Page: 0% Percent that are Top-Level Sites: 50%</p>
<p>Total Number of In-Use Lists: 12 Percent of In-Use Lists that are Stale: 17% Total Number of Unused Lists: 41 Total Size: 1MB Average Size of a Site: 0MB Average Number of In-Use Lists per Site: 2 Average Number of Unused Lists per Site: 5 Average Number of Items per Site: 15 Average Number of Alerts per Site: 0</p>
<p>Average Number of Users per Site: 2 Average Number of Domain Groups per Site: 0 Average Number of Cross-Site Groups per Site Collection: 0 Average Number of Readers (User) per Site: 0 Average Number of Readers (Domain Group) per Site: 0 Average Number of Contributors (User) per Site: 1 Average Number of Contributors (Domain Group) per Site: 0 Average Number of Web Designers (User) per Site: 0 Average Number of Web Designers (Domain Group) per Site: 0 Average Number of Administrators (User) per Site: 1 Average Number of Administrators (Domain Group) per Site: 0</p>
<p>Portal Personal Site Quick Links Average Number of Quick Links per Site: 1 Percent of Quick Links that are Public: 0% Most Visited (92 visits)</p>

Table 9. Site Type sample statistics

Site Type: Team Site Percent of all Site Types: 38% Percent of all Site Visits: 78% Total Number of In-Use Sites: 3 Total Number of Unused Sites: 2 Percent Large (250MB+): 0% Percent Medium (5-250MB): 0% Percent Small (5MB-): 100%
Percent Frequently Visited (100+ visits per day): 0% Percent Rarely Visited (< 1 visit per day): 0% Percent Increasing in Popularity: 100% Percent Decreasing in Popularity: 0% Average Number of Visits per Site (in the past Month): 69 Percent of Visits made to the Home Page: 30%
Percent with a Customized (Unghosted) Home Page: 0% Percent that are Top-Level Sites: 100%
Total Number of In-Use Lists: 8 Percent of In-Use Lists that are Stale: 25% Total Number of Unused Lists: 22 Total Size: 1MB Average Size of a Site: 0MB Average Number of In-Use Lists per Site: 3 Average Number of Unused Lists per Site: 7 Average Number of Items per Site: 20 Average Number of Alerts per Site: 0
Average Number of Users per Site: 2 Average Number of Domain Groups per Site: 0 Average Number of Cross-Site Groups per Site Collection: 0 Average Number of Readers (User) per Site: 0 Average Number of Readers (Domain Group) per Site: 0 Average Number of Contributors (User) per Site: 1 Average Number of Contributors (Domain Group) per Site: 0 Average Number of Web Designers (User) per Site: 0 Average Number of Web Designers (Domain Group) per Site: 0 Average Number of Administrators (User) per Site: 1 Average Number of Administrators (Domain Group) per Site: 0 Most Visited (92 visits)

Table 10. Summary of all List Types sample statistics

Total Number In-Use: 12 Percent Stale: 17%

Total Number Unused: 41
 Total Number Large (500+ items): 0
 Total Number Medium (25-500 items): 2
 Total Number Small (25- items): 10
 Average Number of Items: 10
 Average Number of Custom Views: 1
 Average Number of Custom Properties: 26
 Average Number of Custom Required Properties: 7
 Maximum Number of Custom Properties: 33
 Total Custom Properties (Single-Line Text): 77
 Total Custom Properties (Multi-Line Text): -2
 Total Custom Properties (URL): -2
 Total Custom Properties (Choice): 0
 Total Custom Properties (Multiselect Choice): 4
 Total Custom Properties (Ratings/Matrix): 0
 Total Custom Properties (Boolean): 0
 Total Custom Properties (Number): 0
 Total Custom Properties (Currency): 0
 Total Custom Properties (Date/Time): 1
 Total Custom Properties (Calculated): 0
 Total Custom Properties (Lookup): 0
 Total Custom Properties (User): 0
 Total Customized (Unghosted) View Pages: 0
 Total Customized (Unghosted) Form Pages: 0
 Percent with Custom Views: 42%
 Percent with Custom Properties: 25%
 Percent with Moderation: 0%
 Average number of Visits per List: 13
 Percent of Visits made to View Pages: 47%
 Percent of Visits made to Form Pages: 15%
 Percent of Visits made to Documents in the List: 37%
 Most Visited (31 visits)

Table 11. Single List Type sample statistics

List Type: DocumentLibrary
 Percent of all Lists: 15%
 Percent of all In-Use Lists: 42%
 Percent of all List Visits: 46%
 Total Number In-Use: 5
 Percent Stale: 0%
 Total Number Unused: 3
 Total Number Large (500+ items): 0
 Total Number Medium (25-500 items): 0
 Total Number Small (25- items): 5

Average Number of Items: 12
 Average Number of Custom Views: 1
 Average Number of Custom Properties: 12
 Average Number of Custom Required Properties: 12
 Maximum Number of Custom Properties: 12
 Total Custom Properties (Single-Line Text): 11
 Total Custom Properties (Multi-Line Text): 0
 Total Custom Properties (URL): 0
 Total Custom Properties (Choice): 0
 Total Custom Properties (Multiselect Choice): 0
 Total Custom Properties (Ratings/Matrix): 0
 Total Custom Properties (Boolean): 0
 Total Custom Properties (Number): 0
 Total Custom Properties (Currency): 0
 Total Custom Properties (Date/Time): 1
 Total Custom Properties (Calculated): 0
 Total Custom Properties (Lookup): 0
 Total Custom Properties (User): 0
 Total Customized (Unghosted) View Pages: 0
 Total Customized (Unghosted) Form Pages: 0
 Percent with Custom Views: 60%
 Percent with Custom Properties: 20%
 Percent with Moderation: 0%
 Average number of Visits per List: 14
 Percent of Visits made to View Pages: 22%
 Percent of Visits made to Form Pages: 0%
 Percent of Visits made to Documents in the List: 78%
 Percent with Versioning: 0%
 Average Size: 195KB
 Percent of size taken up with Old Versions: 0%
 Most Visited (26 visits)

Table 12. Document type sample statistics

Document Type: .doc
 Total Documents: 6
 Average number in Version History: 0
 Average Size: 12KB
 Average number of Visits per Document: 1
 Percent of all Documents: 10%
 Percent of all Document Visits: 7%
 Most Visited (1 visits)