

# **SIZING AND CAPACITY PLANNING**

---

ENSURING YOUR SHAREPOINT ENVIRONMENT SUITS YOUR NEEDS

---



## NOTE:

This document is based upon and targeted at a SharePoint 2003 environment. Now that SharePoint 2007 has been released to manufacturing, Anexinet engineers are working on similar recommendations for a SharePoint 2007 environment. If you are anticipating a rollout of a SharePoint 2007 environment, please consider this document only a guideline and contact us for further information.

## WELCOME

This article represents the second installment in a series aimed at exploring Microsoft's SharePoint Products and Technologies (a term used to encompass both Windows SharePoint Services and SharePoint Portal Server) in real world scenarios.

SharePoint is one of the fastest growing products in Microsoft's history – and one that only gets more important as time goes on. In 2005, Microsoft CEO Steve Ballmer stated:

---

*SharePoint is the number one product at this point in the history of Microsoft. We expect \$400 million revenue from that product line - faster than any other product in Microsoft's history.<sup>1</sup>*

---

In addition, at the 2005 Professional Developer's Conference, Steven Sinofsky (Executive Vice President for the Office Division) stated:

---

*[Windows SharePoint Services] is at the core of every Microsoft collaboration and information-sharing web initiative*

---

With this kind of attention, a lot of people are asking what all of the buzz is about. To help meet some of this need and to help organizations understand the benefits of SharePoint, in this and subsequent articles, we'll explore many different facets of the products, including:

- Benefits and Drawbacks
- Features and Capabilities
- Getting Started
- Preparing for the Future
- Customization

➤ Integration with Other Systems

If you have specific questions or would like to discuss how SharePoint can help your organization be more productive, save time and money and manage all of your organization's content, please contact Paul Knight at [pknight@anexinet.com](mailto:pknight@anexinet.com) or David Mann at [dmann@anexinet.com](mailto:dmann@anexinet.com).

## INTRODUCTION

SharePoint was initially introduced as a product by Microsoft in 2001 as **SharePoint Team Services** and **SharePoint Portal Server**. Since that time, the products have morphed once (the current 2003 release) and is on the verge of another significant upgrade (as part of Office 2007). Anexinet's engineers have been working with the products since the early beta days (known then as Tahoe) and with even earlier semi-related incarnations – namely 1997's *Team Manager*. Most recently, we have been involved with the Office 2007 release of SharePoint (and the entire Office System) since the pre-Alpha phase as part of Microsoft's invitation-only Technology Adopter Program (TAP). This inside access to Microsoft's developers and Program Managers allows us to know the products earlier and better than anyone else in our market. It allows us to write, speak, train and publish books<sup>ii</sup> on the new releases while others are still installing the products for the first time.

This article is the second in our *SharePoint in the Real World* series. In this article, we'll examine capacity planning and server-sizing recommendations for SharePoint 2003.

## SERVER ARCHITECTURE

Server architecture in the SharePoint world is a multi-faceted construct. There are numerous ways to design and build a SharePoint infrastructure – your specific business requirements will drive the details of your final architecture. At a high level, the following architectures are available:

1. Single Server – all SharePoint and database components are installed on a single physical machine. Can make use of either SQL Server or WMSDE (the desktop version of Microsoft's database engine)
2. Small-Farm – All SharePoint components are run on a single server, database components are run on a separate dedicated SQL Server machine
3. Medium Farm – SharePoint components are split off onto dedicated servers, depending on needs; Database components are run on a single or clustered servers, depending on needs
4. Large Farm – Database components are running on multiple, clustered servers. SharePoint components (search, index, web front-end, site collection, etc) are all split off onto individual servers – clustered or load-balanced as need be.

Each architecture supports a different number of users and levels of processing power. The following should be considered *general guidelines* for average-use user counts. Information for these recommendations was taken from the SharePoint Products and Technologies Resource Kit by Bill English, et al.

Configuration	Maximum User Counts	Notes
<b>Single Server</b>	10,000	Not recommended – see below.
<b>Small Farm</b>	15,000	Recommended configuration for most installations, even for a user base less than 10,000 users.
<b>Medium Farm</b>	50,000	
<b>Large Farm</b>	∞	Network saturation occurs before servers are overloaded

Naturally, this table presents very approximate guidelines. Even 5000 power users could overwhelm a small farm environment if there are significant other demands (indexing, document load, searching, etc) placed on the server concurrently.

While it is supported by SharePoint, the Single Server environment is not one typically recommended by Anexinet. While this architecture is appealing to smaller companies and an initial foray into collaborative functionality, it is nonetheless not a recommended alternative for the following reasons:

1. Difficulty in upgrading to another architecture should it become necessary for business, performance or other technical reasons – such as a geographically distributed user base
2. Lack of Fault Tolerance – there is no clustering option for the single server approach
3. Lack of scalability – there is no load balancing option for the single server approach

All of these problems are resolved with a Small Farm configuration. Technically, the Small Farm environment can be set up on a single server initially (with the understanding of the lack of Fault Tolerance and Load Balancing) and then expanded out to multiple servers as needs increase. The main difference in this case is that the ability to upgrade to other architectures is not limited when starting from the Small Farm. This is often a great way to get started with SharePoint.

## HARDWARE SPECIFICATIONS

The figures presented in the table above are based upon each server in the configuration consisting of the following hardware:

- Processor: (2) 2 GHz Pentium 4
- Memory: 2 GB

Obviously, more powerful hardware is better, but does not translate directly into increased user support. For example, doubling the processing power and/or memory will improve performance but will not necessarily double the number of supportable users.

## STORAGE NEEDS

All SharePoint content is stored inside SQL Server. Therefore, the majority of the disk-based storage will need to be dedicated to the SQL Server machine(s). In order to plan for current and future needs, there are a number of things to take into account:

- \* Current document storage
- \* Document versioning plans/requirements
- \* Future document growth

Depending on the server architecture, there are other storage considerations to be taken into account as well, as shown in the following table:

Element	Storage Needs	Notes
<b>Document Storage</b>	Two times the total file size of the documents to be stored – <b>including versions</b>	Versioning is typically overlooked when planning this need. For example, if you plan on storing 1000 documents with an average size of 500 KB, and plan on keeping 10 versions on average, you will need 10 GB of capacity:  1000 documents * .5MB * 10 versions * 2  This storage is required in the database
<b>List Storage</b>	---	Hard to estimate this as it depends on how heavily your organization uses Lists in SharePoint. As a very rough guideline, assume 1 MB per user.  This storage is required in the database
<b>Indexing</b>	50% of total indexed content: documents (excluding versions), List storage and any other	This storage is required on each Index server

	content (websites, file shares, Exchange public folders, etc) that will be indexed.	
<b>Search Catalogs</b>	50% of Indexing storage requirements	This storage is required on each Search server

## CAPACITY PLANNING

The final element of SharePoint that plays into designing a server architecture is the actual capacity of SharePoint to store and manage documents, List content and sites within the application itself. The following table, taken from the Windows SharePoint Services Administrator's Guide (available online at: <http://www.microsoft.com/resources/documentation/wss/2/all/adminguide/en-us/stsb07.msp?mfr=true>), lists the guidelines for maintaining acceptable performance within any SharePoint installation – regardless of the server architecture.

Object	Scope	Guideline for optimum performance	Comment
<b>Site collections</b>	Database	50,000	Total throughput degrades as the number of site collections increases.
<b>Web sites</b>	Web site	2,000	The interface for enumerating subsites of a given Web site does not perform well much beyond 2,000 subsites.
<b>Web sites</b>	Site collection	250,000	You can create a very large total number of Web sites by nesting the subsites. For example, 100 sites each with 1000 subsites is 100,000 Web sites.
<b>Documents</b>	Folder	2,000	The interfaces for enumerating documents in a folder do not perform well beyond a thousand entries.
<b>Documents</b>	Library	2 million	You can create very large document libraries by nesting folders.
<b>Security principals</b>	Web site	2,000	The size of the access control list is limited to a few thousand security principals, in other words users and groups in the Web site.

<b>Users</b>	Web site	2 million	You can add millions of people to your Web site by using Microsoft Windows security groups to manage security instead of using individual users.
<b>Items</b>	List	2,000	The interface for enumerating list items does not perform well beyond a few thousand items.
<b>Web Parts</b>	Page	100	Pages with more than 100 Web Parts are slow to render.
<b>Web Part personalization</b>	Page	10,000	Pages with more than a few thousand user personalizations are slow to render.
<b>Lists</b>	Web site	2,000	The interface for enumerating lists and libraries in a Web site does not perform well beyond a few thousand entries.
<b>Document size</b>	File	50 MB	The file save performance degrades as the file size grows. The default maximum is 50 MB. This maximum is enforced by the system, but you can change it to any value up to 2 GB (2047 MB) <b>if</b> you have applied Windows SharePoint Services Service Pack 1.

Note that these are *soft* limits – not hard coded into SharePoint. If you get much beyond these limits, however, performance begins to degrade rapidly.

## CONCLUSION

There is as much science as there is art to properly sizing a SharePoint server architecture. A large part of the equation, too, is planning for future needs. A properly planned and implemented SharePoint environment can grow to nearly limitless capacity as your organization’s needs change with time. SharePoint itself is a highly scalable, cost-effective application for delivering a strong document management and collaboration solution to nearly any size organization.

If you have questions or would like to discuss this or any other SharePoint topics in more detail, please feel free to contact Paul Knight at [pknight@anexinet.com](mailto:pknight@anexinet.com) or David Mann at [dmann@anexinet.com](mailto:dmann@anexinet.com).

Look for the next article in the *SharePoint in the Real World* series coming soon from Anexinet’s Portals and Collaboration Team.

<sup>i</sup> Reuters (<http://www.reuters.com/newsArticle.jhtml?type=domesticNews&storyID=6816430>)

<sup>ii</sup> [http://www.amazon.com/gp/product/1590597001/sr=1-1/qid=1156775098/ref=sr\\_1\\_1/002-2861238-5995211?ie=UTF8&s=books](http://www.amazon.com/gp/product/1590597001/sr=1-1/qid=1156775098/ref=sr_1_1/002-2861238-5995211?ie=UTF8&s=books)