Compaq ActiveAnswers

Performance Brief

November 2001
Corporate Intranet Solution
Performance Brief
Prepared by:
Industry Standard Server Group (ISSG), Enterprise Solutions Engineering
Compaq Computer Corporation

Contents

Solution for Intranets from Compaq and Microsoft

Abstract: This Performance Brief is intended to provide information to Compaq field personnel, partners, and customers regarding the performance of the *Microsoft Solution for Intranets* running on Compaq *ProLiant* systems. It discusses the performance observed when running typical user workloads on appropriate Compaq *ProLiant* servers. The solution configurations described herein represent Corporate, Divisional and Team Site deployments, supporting a mix of representative portal activities and user tasks.

The solution software comprises Microsoft SharePoint Portal Server 2001, Microsoft SharePoint Team Services, Microsoft SQL Server 2000, and Microsoft Office XP components running on Microsoft Windows 2000. This document details the solution components, guidelines for installation on Compaq servers, workload descriptions and performance results, along with practical sizing and deployment examples.

This document is for people who will be, proposing solutions (professional services and sales), providing installation services or consulting, and who may be assisting the customer in deploying the Intranet Solution on Compaq systems. It focuses primarily on guidelines developed during testing in the Compaq Enterprise Solutions Engineering labs regarding deployment and performance on Compaq *ProLiant* servers. Compaq recommends the information provided herein be used in conjunction with the product solution information contained in additional white papers authored by Microsoft, noted in the Appendix.

Of these, the reader is especially encouraged to obtain Microsoft's White Paper *Microsoft Solution for Intranets – Prescriptive Architecture Guide,* which provides important product installation and tuning information for the solution configurations.

Notice

The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE, OR OTHER DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION, OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state or local requirements.

Compaq, Compaq Insight Manager, Deskpro, FASTART, NetFlex, NonStop, PaqFax, ProLiant, Prosignia, QuickFind, QVision, ROMPaq, SmartStart, and Systempro/LT are registered with the United States Patent and Trademark Office.

ActiveAnswers is a trademark and/or service mark of Compaq Information Technologies Group, L.P.

Microsoft, Windows and Windows NT, SharePoint, Team Services, and Portal Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Intel, Pentium and Pentium® III Xeon are trademarks and/or registered trademarks of Intel Corporation.

Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Solution for Intranets from Compaq and Microsoft
Performance Brief prepared by Industry Standard Server Group (ISSG), Enterprise Solutions Engineering

First Edition (November 2001)
Document Number Corporate Intranet Solution Performance Brief

Overview

This Performance Brief describes results from performance characterization testing being carried out at Compaq's Enterprise Solutions Engineering Performance Labs. It specifically discusses results obtained when running the Microsoft Solution for Intranets components on a range of Compaq *ProLiant* servers.

The following topics are discussed in this performance brief:

- Solution overview
- Performance results summary
- Mapping business needs to user task throughput
- Workload descriptions
- Characterization methodology
- Acceptable performance goals and system capacity limits
- Results tables and sizing examples
- Current product limitation guidelines
- Test system configuration details

Solution Components

Software

The key solution software components used and reported on in this document are:

- Microsoft SharePoint Portal Server 2001, V1.0 SP1 Beta 1 (BL4423)
- Microsoft SharePoint Team Services V1.0
- Microsoft SQL Server 2000
- Microsoft Office XP V1
- Microsoft Windows 2000 Advanced Server SP2
- Additional pre-requisite software and hot-fixes as noted in Microsoft's product installation guides and the solution Prescriptive Architecture Guide.

Note: Microsoft Windows Media Services is being promoted as a part of the Intranet Solution. The results and information reported in this document are focused on the products and services defined in the Corporate Intranet Solution Offering Prescriptive Architecture Guide (PAG). Future Compaq testing plans include the evaluation of Windows Media Services as it is defined in the Microsoft Solution for Intranets, Broadcast Component documents.

It should be noted that testing was performed on a Beta version of SPS SP1, which provided very stable and predictable performance. Evaluation of the released version will also be performed, when available, and any updated information will be published as a revision to this document.

Solution Architecture

The solution architecture comprises three main components that work together to provide a total corporate-wide solution. Large corporations typically include business divisions, which, in turn, include both large and small teams of people. These different organizational units share information in a variety of ways – within each team, between teams, across business divisions, and throughout the entire company.

- The business purpose, and thus the functional needs of each hierarchical group are different. Teams need to document, discuss, and share the research information. The Team Services component thus enables the hosting of portal sites that provide document repositories, discussion groups, and other capabilities that support team activities. Divisions typically unite teams within organizational or geographic boundaries. The SharePoint Portal Divisional site provides the following:
- A secure portal that facilitates division-wide news and announcements
- Indexing and cataloging of division team generated (and other internal) information
- Search and subscription capabilities
- Document management, workflow and approval features.

For organizations, the SharePoint Portal Corporate site provides information at the enterprise level and access to aggregated information from the divisional and team levels.

Each component therefore provides the required functionality to support each particular portal audience and facilitates construction of a corporate "federated model" such that users at all levels can obtain the appropriate information. The server architecture for this is illustrated below (see Figure 1).

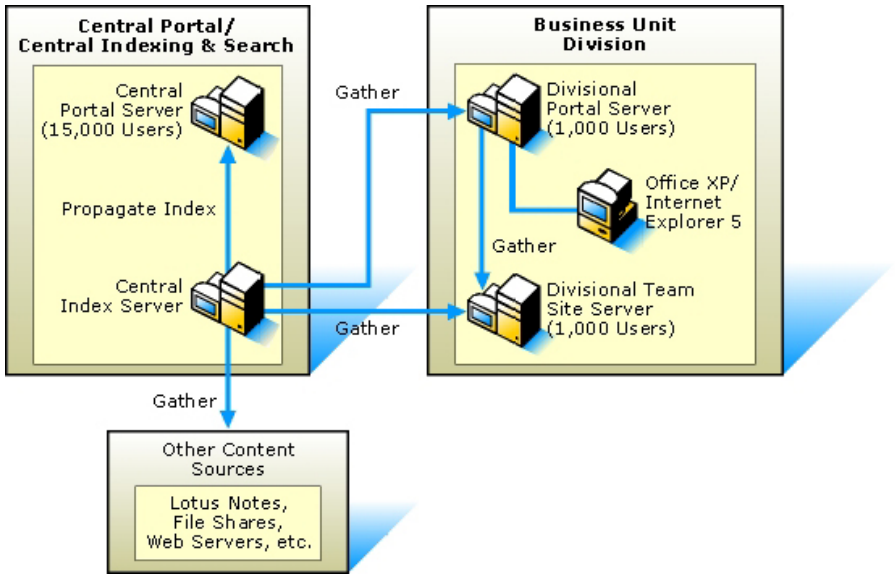

Figure 1: Server Architecture - Federal Model

The central portal server hosts the corporate portal. It is a single physical server that runs Microsoft SharePoint Portal Server. This server hosts the portal home page, all prescribed Web Parts, corporate-wide search, and subscriptions. This server does not crawl any content sources and does not provide any document management services. It uses the indexes propagated to it by the content index server to handle users' search requests. Additionally, some of the Web Parts in the corporate portal provide links to the division workspaces and team sites crawled by the content index server. This server is configured to support 15,000 users performing a variety of tasks.

The content index server crawls content and propagates the index to the corporate portal. It is one physical server that runs Microsoft SharePoint Portal Server. This server crawls the intranet and selected Internet sites, and then builds the necessary indexes. It can be configured to crawl file shares, Web sites (secure and anonymous), Microsoft Exchange public folders, SharePoint Team Services sites, and Lotus Notes databases. The indexes created by this server are propagated to the corporate portal server to facilitate search requests on the corporate portal. Although the server is kept busy, all of its activity occurs in the background relative to the activities of the other servers in this scenario. Users never contact directly with this server.

The divisional portal server hosts the divisional portal. It is one physical server that runs Microsoft SharePoint Portal Server. In addition to hosting the dashboard site for the portal and processing search requests for the divisional portal, this server provides document management services and search of division and team site content. This server is configured to support 1000 users performing a variety of tasks. One or more of the Web Parts in the divisional portal provide links to the team sites crawled by the content index server. Additional Web Parts and sub-dashboards can be customized after deployment to provide division-specific application functionality.

The divisional team site server hosts the team sites for the division. It is one physical server that runs Microsoft SharePoint Team Services and Microsoft SQL Server 2000. This server stores the content for all the team sites. It is configured to support as many as a total of 1000 users spread across as many as 100 team sites.

Determining Performance Requirements

Solution Platforms

Microsoft's White Paper *Microsoft Solution for Intranets – Prescriptive Architecture Guide* (hereinafter referred to as the *Microsoft PAG*) describes an example deployment for a typical corporate intranet solution, as illustrated in the previous section. It comprises a Corporate Portal, a Central Index Server, Divisional Portals, and Divisional Team sites. Microsoft provides suggestions in the PAG for approximate system configurations that could support each of the following:

- Corporate Portal (15,000 Users) – 4 x 700MHz CPUs, 2GB RAM
- Central Index Server – 2 x 700MHz CPUs, 1GB RAM
- Divisional Portal (1,000 users) – 2 x 700MHz CPUs, 1GB RAM
- Divisional Team Sites (1,000 total users) – 1 x 1.2GHz CPU, 1GB RAM

Compaq's recent testing in the ESE Labs has focused on the Corporate Portal, the Divisional Portal, and Divisional Team sites. Further testing is scheduled and will provide information on the Central Index Server and updated versions of the solution components, documented in an update to this guide.

Note: Compaq's memory and disk guidelines for deploying the portals, index and team sites on specific *ProLiant* servers differ in detail from those shown above. Memory utilization and tuning, discussion of disk logical volume design, and network utilization are discussed further in later sections, illustrating how these can affect the performance of both SPS and STS.

Mapping Requirements to Server Capacity

In order to size a solution for specific business needs, one must first understand the purpose of the solution, what specific tasks the supported users will be accomplishing, and their work patterns among other things. The intranet solution is presented to the user as a portal via a browser GUI, thus sizing a solution based simply on "number of users", is inappropriate. As a user's interaction with a browser-based application is essentially "stateless", system sizing is therefore most concerned with the following:

- Which transactions (or business functions) the users are employing, and
- Frequency or transaction arrival rate, more commonly referred to as *throughput rate*.

With a browser GUI, this also results in such metrics as HTTP Hit rate and, most importantly for sizing purposes, *HTTP Page rate*.

When a user executes a portal function (for example - browse, search, or upload) or makes a selection from a list, they will almost always be presented with a new page (screen) of information. Some more complex functions, such as Document Management, are more resource-intensive, and can require multiple actions (or pages of information) to accomplish. However, in general one can equate a user function (or *operation*) to be equivalent to an HTTP page being returned.

We have therefore based our throughput results (and sizing methodology) on the premise of *HTTP Pages/second* as being the key throughput metric and being generally equivalent to the *operations/second* calculated as the required system throughput. Some portal functions are also known to be more complex (or resource-intensive) than others. This is accounted for by applying a *weighting factor* to these, as will be seen later in the example sizing scenario sections.

Workload and Methodology Development

Emulating the results of users interacting with this system requires that we define:

- A way to map the overall solution business needs to specific server *per-function throughput rate* and
- Prototypical workloads (specific transactions and mix ratios) that reflect the expected uses of the solution portals.

Further, we must decide on the performance goals of the systems (acceptable performance limits) so as to determine the systems *capacity limit* (measured in HTTP Pages/sec, as noted above).

Prior to beginning testing, Compaq and Microsoft performance engineers defined and agreed on a set of workload definitions and performance limits (for both SPS and STS) by which to apply representative workloads and judge the system performance and capacity. Microsoft also documented a set of business scenarios in the PAG that depict typical user populations and user activity for each scenario. Should the example values not exactly match a specific customer business need, any variable in the equation can be modified as needed and the resultant required *peak transaction rate* can be calculated.

The workload descriptions are shown below. These are followed by a summary of the key sizing metrics and the example scenarios relating to the Corporate, Divisional and Team sites, upon which Compaq's performance tests were based.

Workload Descriptions

The following are the various workload components utilized to characterize the three different portal types, running either Microsoft SharePoint Portal Server or Team Services as appropriate:

- Corporate Portal Server Workload
- Divisional Portal Server Workload
- SharePoint Team Services Workload

Corporate Portal Server Workload

- Home page related (75%) – Home page, News, Announcements, and so on.
- Search (5%) – Full text search to find information.

Home Page Related

User accesses the portal home page and emulates reading some of the content (for example, access to home page items, such as announcements or corporate news).

Simple Search

User enters a simple search request using the search scope capability. An example would be to search "Published Documents" (scope) for "portal" (text string). A list of related documents is returned. The user picks a random document near the top of the list (best results match percentages), opens the document and emulates reading the content for a short while. On completion, the user returns to the home page.

Divisional Portal Server Workload

- Frequently used (inherent as part of functions below, see details) – Home page, News, Announcements, and so on.
- Search (30%) – Full text search to find documents containing a specific text string.
- Browse Categories (65%) – Browse Taxonomy topics to locate documents containing relevant information.
- Document Management (5%) – Periodic check-in, check-out and publishing of user authored documents.

Frequently Used

User accesses the portal home page and emulates reading some of the content (For example, access to home page items, such as announcements or corporate news).

Search

User enters a simple search request using the search scope capability. An example would be to search "Published Documents" (scope) for "corporate portal" (text string). A list of related documents is returned. The user picks a random document near the top of the list (best match percentages), opens the document and emulates reading the content for a short while. On completion, the user returns to the home page.

Browse Categories

User picks a random category from the taxonomy. A list is returned containing documents and/or sub-categories. If documents are found, the user opens a random document and emulates reading it. If not, the user selects a random sub-category. If the returned list contains documents, the user selects and reads a document. This workload therefore emulates a user traversing a category tree. On completion, the user returns to the home page.

Document Management

User selects a random folder from a defined set of folders. User then selects a document and performs an appropriate Document Management (DM) operation, depending upon the current state of the document. For example, if document is checked out, check it in and optionally publish; else if it is published or checked in, then check it out). Additionally, on checkout there is a low percentage chance of also downloading the document to the client desktop. On check-in, there is also a low percentage chance of also publishing it. On completion, the user returns to the home page.

SharePoint Team Services Workload

- Contributing new material (25%) – Uploading documents or revisions to a document library.
- Reading material (30%) – Browsing Document Libraries and reading documents.
- Threaded discussions (45%) – Browsing, reading and contributing topics/replies to discussions.

Upload Documents

This workload uploads a random .doc, .xls, .ppt, .pdf, or .zip document from the client computer to the SharePoint Team Services server. The user enters into the homepage, selects "Document", selects a random Document Library from the Document Library list, and then uploads a document to the selected Document Library. The document sizes for upload ranges from about 50Kb to 10Mb and an average of 1Mb. This includes a random mix of file types (doc, ppt, xls, pdf, and zip).

Browse Documents

The user starts at the homepage, selects Documents, and selects a random Document library from the Document Library list. From the Document Library List, the user has a 25% probability of opening a randomly selected Document and reading the document. The user loops through the Document library twice before returning to the homepage.

Add Discussions

The user starts at the homepage, selects Discussions, and then selects a random discussion board. There is then a 50:50 chance of either composing a new topic or (if a thread exists) replying to an existing topic thread at a random point in the thread. If no topic exists in the discussion board, a new main topic is created. A total of 25 discussion boards are available. The topic or reply content averages about 5 lines in length.

Browse Discussions

The user starts at the homepage, selects Discussions, and then selects a random Discussion board from the 25 boards available. The user then selects a random discussion article to display. There is a 50% probability that the user will open a specific discussion topic/reply and read it. The user loops through the Discussion lists twice before returning to the homepage.

Workload Implementation

The Compaq developed workloads were implemented as separate "workload scripts" utilizing the Segue Silk Performer performance characterization toolset. Each workload component performed a mix of appropriate functions, operating on randomly picked search strings, categories or folders and documents. Realistic "think times" were also employed to emulate user activity such as perusing a list before selecting an item, entering form data or reading a page of a document. Each workload therefore seeks to emulate typical user behavior, rather than simple random activity.

Determining Capacity Limits

Each workload component (for example - Browse, Search, Document Management, Upload document, Threaded discussion, and so on.) was first employed in isolation so as to determine the peak server throughput possible for that specific activity. This peak was determined when the emulated user load applied caused either average CPU consumption to reach 85%, or key function response time average to exceed 3.0 seconds.

Experience has shown that utilizing these limits results in stable and predictable performance. In no tests did any other system resource (disk volume I/O bandwidth, memory utilization or network utilization) approach saturation. The actual utilization of the other resources observed during testing is provided in a later section.

Developing Sizing Models

Having determined the system capacity limits for each workload component, the expected results for various mixes of these were predicted via a sizing model. Each specific mixed workload was then run at an aggregate throughput rate expected to yield close to 85% average CPU consumption. In every case results confirmed the throughput prediction within 2%, thus validating the sizing model. This is well within acceptable experimental error, although it is possible that the combinations of functions may have caused small variations in actual system CPU consumption compared to the model (for example, increase due to context switching, or decrease due to cache efficiencies). The actual results versus the sizing model predictions are included in following sections.

Key Sizing Metrics

Specific business-related information, in the form of relevant sizing metrics, must be obtained from a customer in order to calculate peak throughput rates that the solution system(s) must support. Note that these are metrics that a customer should be able to estimate reasonably, without requiring an extensive or overly detailed needs analysis. Ranges and typical values are also shown as guidelines to sanity-check the validity of each metric value. These will differ depending upon the specific solution portal. For example, percent of active use of (or frequency of access to) the Corporate portal is likely much lower than that of a Team Site portal for a typical user. The following lists the key sizing metrics:

- **Number of users** – the number of potential users (or subscribers) who have access to the portal.
- **Percent of active users per day** – percentage of the above population who may access the portal during the day. Usually between 10% and 100%.
- **Operations per user per day** – accesses to the portal home page, searches, category browses, document retrievals, and so on. Depending upon the user's role (reader, author, reviewer, contributor, and so on.), this will usually vary between 1 and 10 per day.
- **Hours per day** – hours in the business day. Note that portal site access may cross time zones; therefore 10 hours per day may be a typical number.
- **Peak factor** – a ratio of how much the peak usage exceeds the average usage. Systems should always be sized for peak usage such that performance does not degrade unacceptably during busy periods. A ratio of 3:1 is typical.

These quantitative descriptions of each portal's usage characteristics can be used to estimate the required peak throughput. The following formula yields the peak throughput in *operations per second*:

$$\frac{\text{Number of users} \times \text{Percent of active users per day} \times \text{Number of operations per active user per day} \times \text{Peak factor}}{360{,}000 \times \text{Number of hours per day}}$$

Note: Peak factor is an approximate number that estimates the extent to which the peak dashboard site throughput exceeds the average throughput. This number typically ranges from 1 to 4. The number 360,000 is determined by:
100 (percent conversion) X 60 (minutes per hour) X 60 (seconds per minute).

The following sections provide some throughput requirement calculation examples for each portal (Corporate, Divisional and Team Site).

Example Portal Characteristics

The following are the different types of portal sites:

- Corporate Portal
- Divisional Portal
- Team Site Portal

The characteristics of each portal are described in the sections below.

Corporate Portal

The following table reflects characteristics suggested by Microsoft that could represent a corporate portal site intended to support approximately 15000 users. The workforce will typically make infrequent accesses to this site – these largely comprising browsing home page information (such as Corporate Announcements or News), in addition to searches to find information (located on the divisional and team servers and other information stores) that have been indexed to facilitate full-text searching. Note that the portal page activities (home page, announcements, news, and so on.) are considered less resource intense and are weighted accordingly compared to the Search function.

Table 1. Characteristics of Corporate Portal

Characteristic	Value
Number of users	15000
Percent of active users per day	50
Number of search operations per active user per day	2 (weight = 1.0)
Number of portal pages per active user per day	6 (weight = 0.5)
Total operations per active user per day	5
Number of hours per day	12
Peak factor	5

$$\frac{15000 \times 50 \times 5 \times 5}{360000 \times 12} = 4.34$$

As shown above, these characteristics yield a predicted peak throughput requirement of 4.34 operations per second.

For reference, testing has shown that a Compaq *ProLiant* DL580 (4 x 700MHz) running a mix of 75% Home page/news/etc. and 25% Search, achieved 5.43 operations/sec at approximately 85% average CPU consumption, with key function average response times of less than 3 seconds. Further details are provided in the Results section of this document.

Divisional Portal

The following table reflects characteristics suggested by Microsoft that could represent a divisional portal site intended to support approximately 1000 users. The workforce is reasonably diverse; utilizing a mix of common portal functions and a small percentage of uncommon functions, that is, document management activity. Most users are typically accessing the site to locate information (by browsing categories or searching for documents matching their needs), or to read divisional news or announcements on the home page. Note that the Document Management activities are considered more resource intensive than the other functions used and are weighted accordingly.

Table 2. Divisional Portal - Example Characteristics

Characteristic	Value
Number of users	1000
Percent of active users per day	80
Number of common operations per active user per day	15 (weight = 1.0)
Number of uncommon operations per active user per day	1 (weight = 3.0)
Number of operations per active user per day	18
Number of hours per day	12
Peak factor	4

$$\frac{1000 \times 80 \times 18 \times 4}{360000 \times 12} = 1.33$$

As shown above, these characteristics yield a predicted peak throughput requirement of 1.33 operations per second.

For reference, a Compaq *ProLiant* DL580 (2 x 700MHz) running a mix of 65% Browse, 30% Search and 5% Document Management activities, achieved 4.10 operations/sec at approximately 85% average CPU consumption, with key function average response times of less than 3 seconds. Further details are provided in the Results section of this document.

An alternate server choice would be a Compaq *ProLiant* DL380 (2 x 1.2GHz), which would provide slightly better performance, but cannot be expanded with additional CPUs.

Team Site Portal

The following table reflects characteristics suggested by Microsoft that could represent a Team Services deployment intended to support multiple team sites totaling approximately 1000 users. The workforce would likely be a number of project teams, each ranging from 10-20 users in size. Most users are typically accessing these sites to collaborate on team projects by uploading project-related documents, browsing document libraries, reading material or contributing to threaded discussions. Note that as team sites are targeted at project-focused groups, the percent active and operations per day are typically higher than for the other portals. Weighting has been applied to the document upload save and replace functions, as the resource demands are higher than the other STS functions.

Table 3 - Team Site Portal - Example Characteristics

Characteristic	Value
Number of users	1000
Percent of active users per day	80
Number of uploads and document saves per user per day	12 (weight = 2.0)
Number of browses and discussion activities per user per day	18 (weight = 1.0)
Number of operations per active user per day	42
Number of hours per day	12
Peak factor	3

$$\frac{1000 \quad X \quad 80 \quad X \quad 42 \quad X \quad 3}{360000 \; X \; 12} = 2.33$$

As shown above, these characteristics yield a predicted peak throughput requirement of 2.33 operations per second.

For reference, a Compaq *ProLiant* DL380 (1 x 1.2GHz) running a mix of 25% Uploading, 30% Browsing and 45% Discussions, achieved 50.00 operations/sec at approximately 90% average CPU consumption, with key function average response being sub-second. Further details are provided in the Results section of this document.

An alternate server solution would be a *ProLiant* DL360. The choice between these would likely be predicated on the number of internal and/or external disks required, optional Compaq *Smart Array* I/O controllers and other hardware factors relevant to the specific solution system needs.

Performance Results

The following section details actual results obtained in the ESE Labs, when applying the workloads to specific Compaq *ProLiant* servers suitable for each specific portal deployment.

The results show the maximum throughput achievable for each relevant portal function tested. Other metrics, such as CPU%, memory usage, and network bandwidth are also reported.

Additionally, several "mixed-function" workloads were run to both validate the sizing model and to provide actual results to compare to the predicted throughput needs. In every case, the results obtained exceeded the estimated throughput requirements for each portal server (as shown in the previous typical sizing calculations). Examples of these results are included for each portal.

Significant data collection and analysis was performed for each test. This included monitoring metrics such as CPU, memory, disk I/O, network traffic rates, function throughput and response times, HTTP hits and pages per second.

Corporate Portal Server

Compaq *ProLiant* DL580

4 x 700MHz CPUs, 2MB L2 cache

2GB main memory

These tests yielded a peak throughput rate of 5.43 operations per second, exceeding Microsoft's estimated throughput requirements example of 4.34 operations per second for a corporate portal deployment on a single 4-CPU server. User perceived response times (that is, HTTP Page times) for Searching were less than 3.0 seconds on average. Home page response times were typically sub-seconds. This is largely due to the new caching schemes employed in the SP1 release of SharePoint Portal Server. These caching parameters were set up as recommended in the Microsoft PAG. The Search DB cache sizes were also setup per the PAG to optimize search performance. The following table summarizes key performance metrics.

Table 4. Key Performance Metrics

	Avg. CPU%	Memory (MB)	Network Kb/sec	HTTP hits/sec	HTTP pages/sec
Search, Home page, etc.	84	1547	1,169	84.22	5.43

Notes:

1. Avg. CPU% - is the average portal server CPU consumption measured during the test.
2. Memory (MB) - is the peak memory utilization observed, in MB. Note that the Store cache max was set to a high value to maximize store access performance. The Search DB cache was also set to an increased value (per PAG) to improve search performance.
3. Network Kb/sec - is the aggregate bi-directional average network traffic between the emulated clients and the portal server, in Kb/sec. This mainly comprises the HTTP request/response packets.

4. HTTP hits/sec - represents the average number of HTTP packets required to generate the number of HTTP pages/sec. Each page may require more or less "hits" depending upon the page complexity and number of items to display.
5. HTTP pages/sec - is generally equivalent to user operations/sec – the key throughput metric. Each user action almost always results in a new page being displayed as a result of selecting an action or document.
6. The desired workload comprised 75% Home page related and 25% Search. As the Search workload itself included periodic home page access and was calibrated to deliver these overall percentages, we can therefore simply use the maximum measured overall workload throughput (HTTP Pages/sec, as reported above) to determine corporate portal performance.

Divisional Portal Server

Compaq *ProLiant* DL580

2 x 700MHz CPUs, 2MB L2 cache

1 GB main memory

Compaq *ProLiant* DL380-G2

2 x 1.2GHz CPUs, 512Kb L2 cache

1GB main memory

A mixed-function workload was used to represent overall user activity. The workload comprised 65% category browsing and document reading, 30% searching and 5% document management. Using the appropriate percentages of maximum achievable function throughput resulted in a predicted mixed workload aggregate throughput of:

(0.65 x 4.72) + (0.30 x 2.90) + (0.05 x 2.37) = 4.06 operations/second

The mixed workload tests on a 2-CPU DL580 yielded a peak throughput rate of 4.15 operations per second, compared to the predicted 4.06 operations per second, thus reasonably validating the sizing model. The throughput achieved also exceeded Microsoft's estimated throughput requirements example of 1.33 operations per second for a Divisional portal deployment on a single 2-CPU server. User perceived response times (that is, HTTP Page times) were less than 3.0 seconds on average except for the Home page, which was typically sub-second. This is largely due to the new caching schemes employed in the SP1 release of SharePoint Portal Server, which were set up as recommended in the Microsoft PAG. The following table summarizes key performance metrics.

Table 5. Divisional Server Results

Test Performed	Avg. CPU%	Memory (MB)	Network Kb/sec	HTTP hits/sec	HTTP pages/sec
Browse	90	1535	897	66.90	4.72
Search	80	1544	547	42.95	2.90
Document Management	89	1562	113	10.29	2.37
Mixed Load	89	1554	658	57.91	4.15

Notes

1. Avg. CPU% - is the average portal server CPU consumption measured during the test.
2. Memory (MB) - is the peak memory utilization observed, in MB. Note that the Store cache max was set to a high value to maximize store access performance. The Search DB cache was also set to an increased value (per PAG) to improve search performance.
3. Network Kb/sec - is the aggregate bi-directional average network traffic between the emulated clients and the portal server, in Kb/sec. This includes both normal HTTP request/response packets, plus document downloads (open document to read) and also document upload/download resulting from document management check-in, check-out activity.
4. HTTP hits/sec - represents the average number of HTTP packets required to generate the number of HTTP pages/sec. Each page may require more or less "hits" depending upon the page complexity and number of items to display.
5. HTTP pages/sec - is generally equivalent to user operations/sec – the key throughput metric. Each user action almost always results in a new page being displayed as a result of selecting an action or document.
6. The test divisional portal deployment in the ESE Performance Lab included a custom Web Part on the Home Page that displayed a list of the last 10 published documents. It was felt that including a custom web part on this home page was more realistic than a "vanilla" installation.

Team Site Server

Compaq *ProLiant* DL380-G2

1 x 1.2GHz CPU, 512Kb L2 cache

1GB main memory

Compaq *ProLiant* DL360

1 x 1.2GHz CPU, 512Kb L2 cache

1GB main memory

A mixed-function workload was used to represent overall user activity. It comprised 25% document upload, 30% document browsing and reading, 25% browsing discussions, and 20% contributing to discussions. Using the appropriate percentages of maximum achievable function throughput, resulted in a predicted mixed workload aggregate throughput of:

(0.25 x 18) + (0.30 x 56) + (0.25 x 72) + (0.20 x 57) = 50.7 operations/second

The mixed workload tests on a 1-CPU DL380 yielded a peak throughput rate of 50.0 operations per second, compared to the predicted 50.7 operations per second, thus validating the sizing model. The throughput achieved also vastly exceeded Microsoft's estimated throughput requirements example of 2.33 ops/sec for a Team Sites portal deployment on a single 1-CPU server. User perceived response times (that is, HTTP Page times) were typically sub-seconds. The following table summarizes key performance metrics:

Table 6. Team Sites Server Results

Test Performed	Avg. CPU%	Memory (MB)	Network Kb/sec	HTTP hits/sec	HTTP pages/sec
Add Discussions	83	411	2,332	612	57
Add Documents	74	411	3,405	242	18
Browse Discussions	87	449	2,325	861	72
Browse Documents	83	470	6,325	769	56
Mixed Workload	93	412	4,056	617	50

Notes:

1. Avg. CPU% - is the measured average CPU consumption on the STS Server.
2. Memory (MB) - is the peak memory utilization on the STS Server.
3. Network Kb/sec - is the measured average 2-way network traffic in Kb/sec. Note that approximately 10,000Kb/sec is the absolute limit for a 100Mbit network segment. The browse document workload consumes more than 6,000Kb/sec of network bandwidth, even though a document is opened (and thus downloaded to the client) only 25% of the time when a user browses a Document Library.
4. HTTP hits/sec - represents the average number of HTTP packets required to generate the number of HTTP pages/sec. When compared to pages/sec, it shows the typical average number of hits to create the page content. Many of the hits (For example, those relating to static graphics) will be satisfied via a client cache hit.
5. HTTP pages/sec - is generally equivalent to user operations/sec – the key throughput metric. Each user action almost always results in a new page being displayed as a result of selecting an action or document.

Product Limit Guidelines

Several Microsoft white papers (such as *Capacity Planning for SharePoint Portal Server 2001* and *Microsoft Solution for Intranets PAG*) describe a number of recommended product limits. These involve three main areas:

- Maximum throughput
- Server Resource consumption
- Portal Corpus limits

The first two of these are also discussed within this Performance Brief with regards to the measured maximum throughput achievable, within recommended resource consumption limits. In order to achieve the performance levels reported, it is also important not to exceed the current product design limits with regard to the corpus of information being stored and other key variables. The reader is strongly encouraged to read the Microsoft's White Papers, noted in the Appendix, for detailed recommendations.

Test System Configurations

The following describes the test system configurations. Note that the quantity and size of disks shown was for characterization purposes in a lab environment. The number and size of these will vary for each customer deployment depending upon a number of factors including, but not limited to:

- Portal purpose (corporate, divisional, team site)
- Quantity and size of planned corpus content (on divisional site)
- Total of off-node crawled information and propagated indexes
- Quantity, size and retention period of information for team sites

Corporate and Divisional Portal Servers

- Compaq *ProLiant* DL580 server, 2 or 4 x 700MHz CPUs (100MHz bus, 2MB caches)
- 1GB or 2GB main memory
- Compaq *Smart Array* 5302 disk I/O controller
- 10/100Mbit network interface card (at 100Mbit)
- 4 internal disk drives (18.2GB each)
- External Ultra SCSI-3 Compaq *StorageWorks* drive bay containing 14 x 18.2GB disks

Notes:

1. The corporate site server was configured with 4 x 700MHz CPUs and 1GB memory. The divisional site server was configured with 2 x 700MHz CPUs and 2GB memory.
2. While configured with 2GB main memory, actual memory usage on the divisional server was observed to be about 1.5GB. Memory usage will vary depending upon system and application tuning. The PAG notes that the SharePoint Store maximum cache size and Search DB cache limits can be increased to increase cache capacity, and thus improve performance. Lab tests confirmed this; in fact, in some tests the Store process size was observed to approach 1GB. Compaq therefore recommends configuring 2GB of memory to take advantage of the increased cache sizes possible, and resulting increases in server performance.
3. Both the internal and external disk drive bays were connected to a Compaq *Smart Array* 5302 intelligent controller, rather than the default on-board SCSI controller. This provides both much higher potential performance (I/O bandwidth) and also facilitates easy setup of RAID Arrays via Compaq's disk array management utility.
4. The internal and external disks were formed into logical volumes (RAID arrays) such that the operating system and the different SharePoint data structures could be installed onto different logical volumes to suit their individual I/O needs. This is not only essential in a lab testing environment (so that detailed I/O analysis can be performed), but is also strongly recommended in real-life deployments in order to optimize I/O performance. The following discussion provides more information.

Employing Multiple Logical Volumes with SPS

Testing has shown that creating multiple logical volumes, each employing an *appropriate* RAID level, can

- Provide optimal performance and
- Fault tolerance for each type of data structure to be hosted on the volumes.

RAID levels provide different benefits, which should be matched to the needs of the data to be stored on a particular logical volume.

RAID-0 (simple striping) allows the creation of a single logical disk volume spanning multiple physical disks. This provides two key benefits –

- The ability to create large volumes, and
- The ability to spread the volume I/O operations across multiple spindles (high throughput).

No fault tolerance is provided, however.

RAID-1 (mirroring) allows the creation of a logical volume where the content is "mirrored" on two physical drives.

- This provides a measure of fault tolerance as data is not lost should one of the drives fail.
- Secondly, it has been shown that I/O performance is better (faster access) when the volume data is mostly read. Conversely, if the data is mostly written, performance can degrade slightly as writes have to occur to both physical disks.

RAID0+1 is a mirrored stripe set (thus combining the benefits of both these RAID levels) requiring a minimum of 4 physical disks (a 2-disk stripe mirrored onto a second 2-disk stripe). This should be used to increase the fault tolerance of a 2-disk stripe. If a larger stripe set is required (3 or more disks), then RAID-5 is likely a better choice.

RAID-5 (data guarding) provides a higher level of fault tolerance by writing parity data across all array members. Volume I/O is also aggregated across the disks. While a single large RAID-5 logical volume may be the easiest to set up, and may be adequate for most I/O performance needs, there is anecdotal evidence provided by the Compaq *StorageWorks* group that, RAID-5 volumes of less than 5 disks can perform better. ADG (Advanced Data Guarding also known as RAID-6) was used for the Team Services configuration, and is discussed in the section detailing that server's configuration.

SharePoint Portal Server requires and creates several different types of data storage (for example - database, logs, temp areas, indexes, and executables), for which different RAID level volumes should be formed. Compaq recommends that creating multiple logical volumes to hold each (as shown in the example below) will provide the optimum performance. If spare disks are also initially left in the *StorageWorks* cabinet (marked unused), then adding disks to increase specific volume storage at a later date will be much easier. The following table illustrates examples of using RAID-0 (simple stripe), RAID-1 (mirroring) and RAID-5 (data guarding) where appropriate.

Table 7. An Example of SPS Server Logical Volume Layouts

VOLUME CONTENT	DRIVE	DISKS	RAID LEVEL
Operating system	C	2 x 18.2GB	1
Swap file, kits, etc.	D	Partition on above volume	1
	E	CD drive	N.A.
SharePoint installation	F	2 x 18.2GB	1
Search temp, Search logs, DB logs	G	2 x 18.2GB	0
SharePoint Database	H	5 x 18.2GB	5
SharePoint Indexes	I	5 x 18.2GB	5

Note: Unless third party software products are used for backup/restore, space should also be allocated for backups. The size required for these will usually be the combined "in-use" sizes for the Database and Indexes. In the above example, it might be possible to use Drive "I", as the Index size is typically about one third that of the database size.

Team Site Server

- Compaq *ProLiant* DL380 server, 1 x 1.2GHz CPU (100MHz bus, 512Kb cache)
- 1GB main memory
- Onboard Compaq *Smart Array* 5I controller
- Compaq *Smart Array* 5302 Controller
- 10/100Mbit network interface card (at 100Mbit)
- 2 internal disk drives (18.2GB each)
- External Ultra SCSI-3 Compaq *StorageWorks* drive bay containing 14 x 18.2GB disks

Notes:

1. The Compaq *ProLiant* DL380 can be configured with 2 CPUs (1.2GHZ), however, a single CPU system was configured for testing so as to evaluate the Divisional Team Site Server configuration defined in the PAG.
2. Although configured with 1GB main memory, actual memory usage on the Divisional Team Site Server was observed to be typically less than 512MB. However, installing 1GB main memory is recommended in the Microsoft PAG, and is also endorsed by Compaq.
3. The internal disk drive bays were connected to the on-board 5300 5I SCSI controller. The external disk drive bays were connected to a Compaq *Smart Array* 5302 intelligent controller. This provides higher potential performance (I/O bandwidth) and also facilitates easy setup of RAID arrays via Compaq's disk array management utility.
4. In this case the internal disks comprised a RAID 1 mirror set and the external disks were formed into a single ADG logical volume. Testing showed that disk and controller I/O rates were well within the performance limits of this configuration.

Advanced Data Guarding (ADG)

The Team Services virtual directories and Microsoft SQL Server 2000 databases were stored on an external 14-member RAID ADG array. RAID ADG (also known as RAID-6) is similar to RAID-5 in providing fault tolerance by writing parity data across all members of the array, but allows for increased fault tolerance over RAID-5 by writing 2 sets of parity data across the array. Logical volumes can survive the failure of 2 members in an array configured using RAID ADG. There is, however, some degradation in write performance when using RAID ADG due to the overhead associated with writing 2 sets of parity data.

Table 8. An Example of SPS Server Logical Volume Layouts

VOLUME CONTENT	DRIVE	DISKS	RAID LEVEL
Operating system Swap file, kits, etc. Team Services installation SQL Server 2000 installation	C	2 x 18.2GB	1
CDROM	D	CD drive	N.A.
Team Services Virtual Directories SQL Server 2000 databases	E	14 x 18.2 GB	ADG

Note: If high levels of document uploading are anticipated, it may be beneficial to split the Virtual Directories and Microsoft SQL Server 2000 database onto separate logical volumes, as both these structures can experience heavy write activity during document uploads.

Conclusions

The following summarizes our key findings:

- The current range of Compaq *ProLiant* servers can provide appropriate configurations to support all of the Intranet Solution components. The servers best suited are the *ProLiant* DL360, *ProLiant* DL380, and *ProLiant* DL580.
- In each test case the Compaq *ProLiant* server employed exceeded performance requirements, by delivering significantly more throughput (operations/second) than the solution's typical workload was estimated to require. This implies that reserve capacity is available.
- The sizing models developed for each component provided throughput *predictions* that were consistently within 2% of the *actual measured throughput,* when running the prototypical mixed workloads. Validating the algorithms and values to this degree of confidence will thus allow us to complete the development of an Intranet Solution Planning & Sizing tool, planned to be available before end CY2001.
- The consistent and predictable results obtained are a testament to the Microsoft products' design and implementation. Despite the fact that STS V1.0 and a Beta version of SPS SP1 were utilized in the testing, product performance and stability was excellent.
- The performance improvements Microsoft has incorporated into SharePoint Portal Server SP1, particularly the new caching schemes, yield significant performance improvements in both maximum throughput achievable and in dramatic reductions in many function page response times. Note that the instructions in the PAG should be followed carefully in order to utilize these new caching schemes to their best.
- SharePoint team Services has proven to be both easy to set up and deploy, but also to deliver exceedingly efficient performance. Even when deployed on one of the *slowest* currently available Compaq *ProLiant* server, performance greatly exceeded that required to meet the prototypical Team Site workload needs.

Future Work

The following is a list of Solution-related "future" work currently planned for execution in the Compaq ESE Labs:

- Complete and field-test the Intranet Solution Sizing and Planning Tool.
- Evaluate updated versions of SPS SP1 and STS when available from Microsoft.
- Update this Performance Brief as new results are obtained.
- Characterize the Intranet Solution Central Index Server component and develop server recommendations. This will include crawling/indexing STS sites, Web sites, Exchange folders, etc.
- Characterize the Intranet Solution Windows Media Server Broadcast Edition component.

The latest information can always be found at:

- Compaq-SharePoint related web pages at http://compaq.com/sharepoint
- Compaq *ActiveAnswers* site http://activeanswers.compaq.com

Appendix

Related Documents

Microsoft Solution for Intranets Prescriptive Architecture Guide.

Capacity Planning for Microsoft SharePoint Portal Server 2001 White Paper.

Available from: http://www.microsoft.com/SharePoint/

Departmental Deployment of SharePoint Team Services.

Large-Scale Deployment of SharePoint Team Services.

Available from: http://www.microsoft.com/frontpage/sharepoint/default.htm