

VisualInfo for AS/400 Performance and Capacity Planning



Last Updated : August 10, 1999

Preface

This document is a consolidation of performance and capacity planning information related to ImagePlus VisualInfo for AS/400. Some of the information contained within this document was excerpted from the redbook Image and Workflow Library: Capacity Planning and Performance Tuning for VisualInfo and Digital Library Servers (SG24-4974-01). This redbook is available from <http://www.redbooks.ibm.com/>.

Preface	2
Chapter 1. VisuallInfo for AS/400	5
Introducing VisuallInfo for AS/400	5
VisuallInfo for AS/400 Components	5
Library Server	5
Object Server	7
Workflow Server	8
Client	9
Optical Systems	11
VisuallInfo for AS/400 Implementation	11
Chapter 2. VisuallInfo for AS/400 Performance Tuning	13
Communications	13
System Network Architecture (SNA)	13
Transmission Control Protocol / Internet Protocol (TCP/IP)	14
Client	15
Operating System	15
Workstation Configuration	15
Maximum Table of Contents Count	16
Workbasket Count	17
System-Assigned Workbasket	17
Server	17
AS/400 Performance Tuning	17
Logon	17
Search	18
Reorganizing Physical Files	19
File System	19
Security	19
Storage Management	20
Optical Subsystem	20
General Performance Considerations	20
Performance Considerations--Directly Attached Optical Libraries	21
Performance Considerations--LAN-Attached Optical Libraries	22
Capture Subsystem	23
Scanning	23
Compression	23
High-volume Image Capture	24
Chapter 3. VisuallInfo for AS/400 Capacity Sizing	24
Understanding Customer Requirements	25
Defining the Workload	26
Data Model	26
Input Processing	27
Item Processing	28

Storage Management 28

Modeling Workload 29

Additional Capacity Sizing Considerations 30

Estimating Storage Requirements 30

Disk Arm Requirements 31

AS/400e Servers 31

Appendix A. Creating a Model from Simulated Workloads .. 33

Simulating a Workload 33

Measuring the Workload 36

Creating a Model from Measured Data 37

Appendix B. Predefined Workloads 45

Creating a Model Using Predefined Workloads 45

Input Processing Workloads 52

Item Processing Workloads 54

Storage Management Workloads 55

Appendix C. Sizing Worksheet 56

Chapter 1. VisualInfo for AS/400

Introducing VisualInfo for AS/400

ImagePlus VisualInfo for AS/400 (VI/400) is a client/server software product designed to replace paper document processing with image processing for greater efficiency, reliability, and security. The VI/400 product provides capabilities to scan, import, store, secure, index, retrieve and route user-related information. VI/400 can manage document images and word-processing documents, as well as spreadsheets, voice and video.

VI/400 is a complete solution for managing multimedia objects. These objects are cataloged in a library server, providing a common repository of indexes for easy search and retrieval. The objects, which are stored in object servers, may be located centrally or decentralized and close to the users. In all cases, with appropriate user security, any object located anywhere on the network can be retrieved and presented to the user using the VI/400 client.

When you install VI/400, you are installing a combination library server, workflow server, and object server. VI/400 supports one library server and workflow server with one or more object servers. If you determine that you need the support of a second object server, you may install the VI/400 secondary processor on AS/400 system servers other than the library server.

VisualInfo for AS/400 Components

Library Server

The VI/400 library server is the single point of control for the system. It receives and processes requests from the VI/400 client, manages the storage and retrieval of VI/400 objects, and is responsible for controlling security and access.

The library server also contains the VI/400 database, or library. The VI/400 data model, implemented for the library, is reflected through the VI/400 client and the VI/400 application programming interfaces (APIs). The data model provided by VI/400 consists of the following concepts:

- Items
- Key fields
- Index classes

Items

An item is a document or a folder. Associated with each item are system- and user-defined attributes. System-defined attributes are stored once for each item. Examples of system attributes would be item creation date/time and item description. User-defined attribute values (key fields) are stored once for each instance of an item in an index class. End users create items.

Documents

Documents are items that can be stored, retrieved, and exchanged as a single unit among systems and users. A document is any object entered into VI/400 through scan or import.

Folders

Folders are items that provide a way to group related documents and other folders together. Grouping items in folders makes them easier to retrieve. A folder can contain zero to many items, but cannot contain itself as a subfolder. End users create folders. A folder can be placed in an index class and/or placed in another folder.

Item Relationships

Items can have the following relationships:

- When using the VI/400 client, each item is associated with one and only one index class.
- Each item can be located in zero to many folders. However, an item cannot be located in the same folder twice.

Key Fields

A key field is a label for standard information that you can use to index, identify, and retrieve items in index classes and folders. Users can assign values to key fields for each item in an index class.

Key Field Relationships

Each key field can be associated with zero to many index classes.

Index Classes

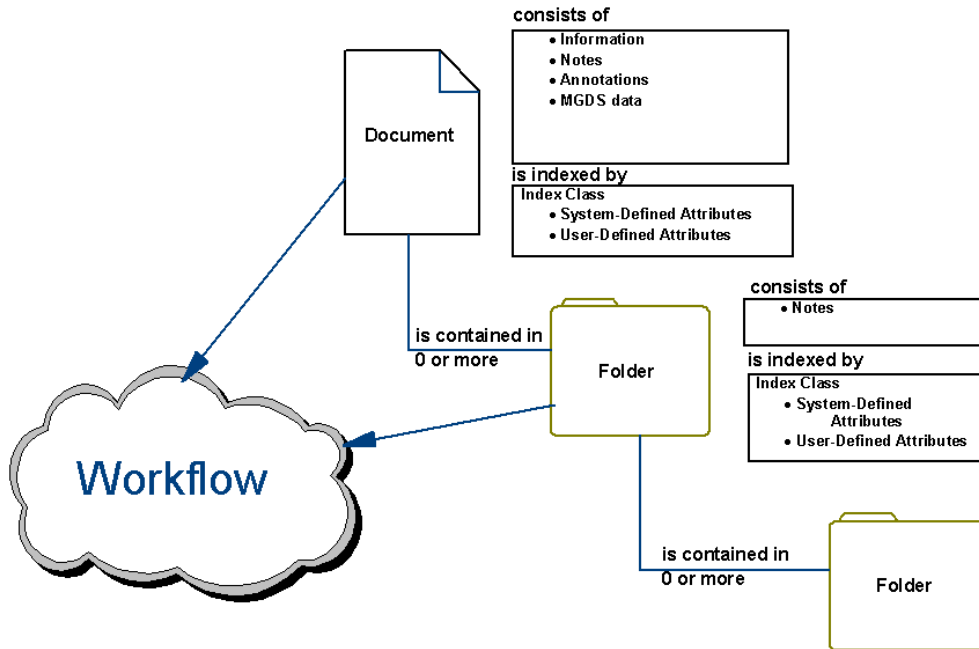
An index class is a way to group items, documents and folders with similar attributes for the purpose of storage and retrieval. An index class also gives you a way to group items for security reasons. Each item in an index class can have different key field values.

Index Class Relationships

Each index class can be a group of 1 to 8 key fields.

The following diagram shows the VI/400 data model.

VisualInfo for AS/400 Data Model



Object Server

A VI/400 object server is a repository for objects stored in the system. Each object server is associated with one library server. The VI/400 system can support multiple object servers. The architecture is entirely flexible with the best location for objects being determined based on system requirements and VI/400 system capabilities.

Object servers also provide system-managed storage (SMS) to automatically move objects from one storage medium to another, such as from DASD to optical. VI/400 manages the storage and migration of objects without user intervention. The objective is to assure a low cost of storage, yet maintain the best possible response time for users. This is usually accomplished by storing incoming documents on magnetic storage during their active life cycle and then automatically migrating them to optical storage when the retrieval frequency is much lower.

All objects stored in VI/400 are initially stored to AS/400 DASD. The objects are stored in directories as defined by the VI/400 system administrator. The system administrator has the option to specify which file system (i.e., "Root", QDLS, QOpenSys) the object directories are to be contained within. The VI/400 optical store process will also copy the object to optical. Given this, the object will exist on AS/400 DASD and optical for a specified period of time. After the

active life of an object on DASD has expired, it will automatically be removed from DASD. If subsequent to its removal from DASD a request is made to access an object, VI/400 will automatically retrieve the object from optical back to AS/400 DASD, if necessary. If the object exists on LAN-attach optical, it is possible to access the object directly from optical without staging it back to AS/400 DASD.

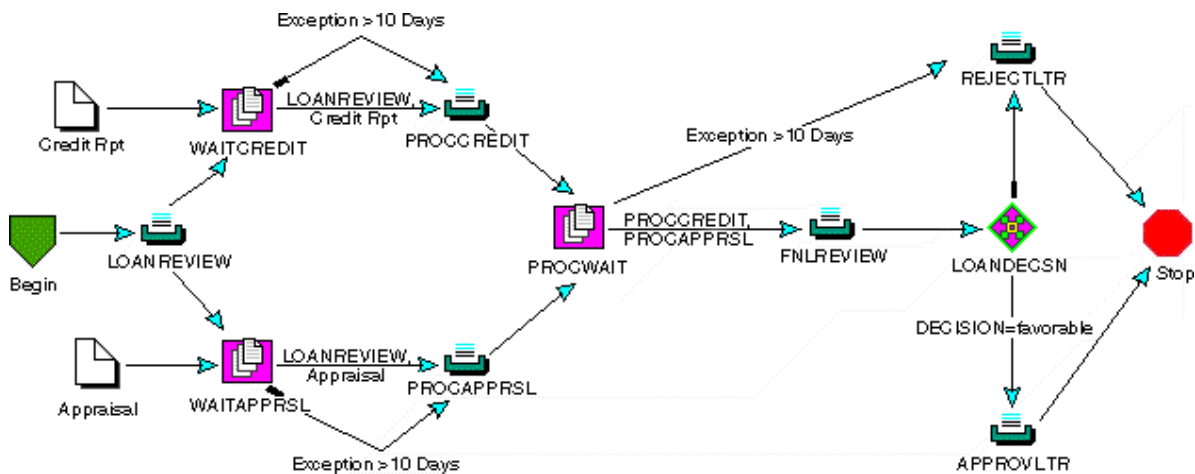
Optical store, optical retrieve, optical distribution, and object deletion are background processes which are started by the system administrator.

Workflow Server

The workflow server is not a separate installable component. When VI/400 is installed, the workflow server is installed with the library server. However, conceptually, the workflow server is a separate component which is integrated with the library server component. This integration allows for the routing of document and folder items.

The VI/400 workflow component provides for the movement of documents and folders from one workbasket or queue to another. Users can simply choose, on an ad hoc basis, where to send a document or folder. VI/400 also has a very powerful work management capability called Advanced Workflow.

Advanced Workflow provides a fast, efficient way to customize and automate business processes and it includes functions for building a work process and for automatically routing documents and folders through a business. The Work Management Builder of VI/400 provides a versatile, easy-to-use graphical tool which professionals can use to define and optimize the work process of a workgroup, department or enterprise. Using work step icons, a user can construct a comprehensive, easy-to-understand process diagram. Below is a diagram which was created with the Work Management Builder.



The process diagram defines:

- Where work is routed once it arrives
- Whether work must wait for certain events to occur or for certain documents or information to arrive
- When work is routed to several users or departments simultaneously

In addition, the business professional may use the Work Management Builder to define:

- What choices the user has at each step in the process
- When to exit to line-of-business applications
- What authority levels and work functions are available to individuals, workgroups, or entire departments
- How work should be prioritized, presented, and selected

All of these functions are defined graphically. No code needs to be written to create or modify the diagrams and there is no need to learn a programming language for general workflows.

The routing aspects of workflow can be handled in several ways, depending on business needs. Routing can be sequential, conditional, parallel or ad hoc:

- Sequential routing allows one person to work on a document or folder at a time.
- Conditional routing moves the document or folder to specified workbaskets when a predetermined condition is reached, such as a mortgage loan ceiling that requires audit review.
- Parallel routing allows different departments to work simultaneously; the work waits for all groups to finish before proceeding. Ad hoc routing allows users to manually route documents or folders to other users.

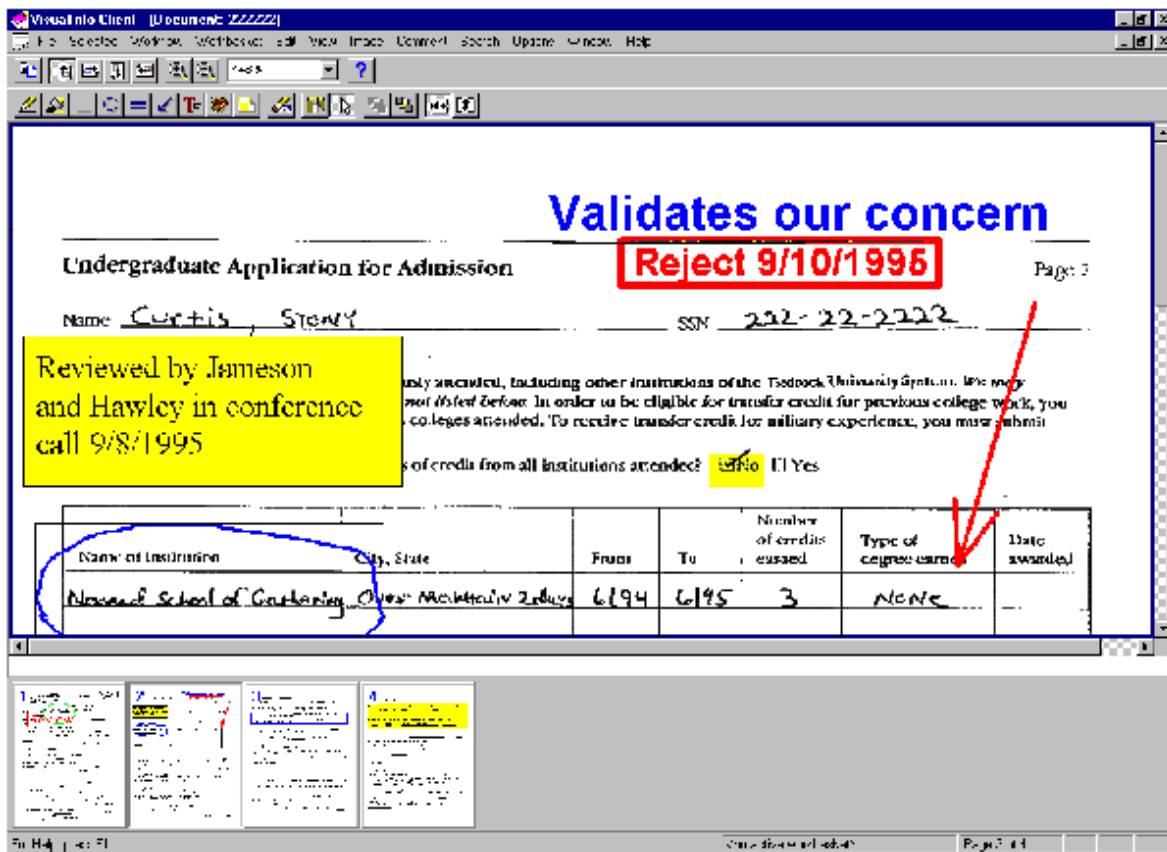
With advanced workflow, you can identify elements and decisions that are critical to a business, along with complex relationships and outcomes. As a result, users can implement policies and procedures consistently, resulting in increased productivity, enhanced quality, and improvement in the way an organization does business.

Client

The VisualInfo for AS/400 client provides a graphical user interface (GUI) that lets you bring documents into the system view so you can work with, store, and retrieve them. With the Windows NT and 95 client, the basic flow of documents and information is simple. To store a document in the system, you can import it from a directory or capture it by scanning its pages. You can then index the document with information that later helps you identify and retrieve it. At that point, you can electronically route it to other users, who can do the following:

- Organize documents into electronic folders. Documents can be placed into folders and folders can be placed into other folders. Documents and folders can reside in one or more folders.
- Retrieve groups of related documents and folders. The basic search capability allows users to query the system and retrieve documents and folders. An example would be entering an account number and retrieving the customer folder. An advanced search capability provides for the simple creation of complex queries which can then be saved for future use.
- Send a document or folder to other users for additional processing, using ad-hoc routing.
- Add a note to a document or folder.
- Annotate a document using a wide range of markups, such as highlighting, "type on" text, note, circle, square, line, arrow or stamp.
- Print or fax a document.
- Place a document or folder on a predefined workflow process.

Following is a screen shot of the VI/400 client image viewer.



The VI/400 client provides an extensive number of application programming interfaces and function calls that can be called from any programming tool supporting a C interface. Or, you can use the OLE automation interfaces. All of the API functions are available as methods to several different OLE objects. Microsoft Visual Basic, Powersoft PowerBuilder and other programming tools can be used to build custom applications. User exits are also available to enable integration. These capabilities provide tremendous flexibility when integrating with line-of-business applications.

Optical Systems

An optical storage device is an optional component in a VI/400 system. Optical libraries can be attached through a local area network (LAN-attach) or directly to an AS/400 server (direct-attach).

Lan-attach optical requires a separate optical controller which runs OS/2 Warp. Optical support and management functions are performed from the optical controller. Objects are stored and retrieved through a communications session between the VI/400 object server and the optical controller. VI/400 supports SNA communications between the AS/400 object server and the optical controller. Using LAN-attach optical, objects can be directly accessed from optical by a client device, without staging the object to AS/400 DASD.


Direct-attach optical may be attached to the primary library/object server or to secondary object servers. With direct-attach optical, the AS/400 is the optical controller and provides optical management support. Objects are stored and retrieved directly to and from optical using the Hierarchical File System APIs. A client request for an object which resides only on direct-attach optical will result in the retrieval of the object to AS/400 DASD for access by the client. The primary library/object server will communicate with secondary object servers using SNA communications.

VisualInfo for AS/400 Implementation

There are many ways to design and implement client/server applications. The VI/400 solution distributes logic and data to the clients. Using this approach offers the following advantages:

- Enhanced performance by having logic on the client and server.
- Enhanced performance by downloading, and caching in memory, relatively static data to the clients.
- Since the data downloaded to the client is read-only, VI/400 maintains a centralized database for the purpose of backup and recovery.
- Display processing is off-loaded from the server.

Distributed Logic and Data



Chapter 2. VisualInfo for AS/400 Performance Tuning

Communications

Whether you are using SNA or TCP/IP for communications, there are several values which can improve performance when properly set. These values should be configured on both the client workstations and the AS/400 server.

System Network Architecture (SNA)

Within the VI400 system, the SNA Advanced Program-to-Program Communications (APPC) protocol may optionally be used for communications between a VI/400 client and the VI/400 library server. The APPC protocol is always used for communications between:

- A VI/400 object server and a LAN-attach optical controller
- The primary VI/400 object server and a secondary object server

In APPC, a mode defines the characteristics of a session. The APPC compression and conversation buffering parameters within the mode description can be important in terms of performance.

Compression

If the mode is configured to allow for compression, performance can be improved. However, compression is very CPU-intensive. SNA supports two compression algorithms, Run Length Encoding (RLE) and Lempel-Ziv (LZ). In general, LZ compresses better than RLE, but costs more in terms of storage and processing.

Request Unit (RU)

An SNA RU is a defined block of data that contains application data and APPC related information. The AS/400 mode description should contain a value of 16384 or *CALC for the RU size. Specifying *CALC enables the AS/400 to calculate the RU size based on the communication line frame size. Assuming the frame size is set to 16393, set the Default RU size on the workstation to 16384.

SNA Pacing

The SNA pacing value specifies how many RUs can be sent or received before a pacing response is required. Set the inbound and outbound pacing values to 7.

In addition to the conversation-level parameters associated with the mode, there are several link-level parameters which can significantly effect performance.

Frame Size

The frame (path information unit (PIU)) size should be set to the maximum supported value. The maximum value will vary based on the use of token-ring versus Ethernet, LAN adapter type, line speed, and Ethernet standard. In a 16 Mbps token-ring environment the maximum value should be 16393. The maximum frame size for Ethernet will be between 1493 and 1500, based on the Ethernet standard.

LAN Acknowledgment Timer and LAN Acknowledgment Frequency

The LAN acknowledgment timer specifies the time interval to delay sending acknowledgments for received frames. The LAN acknowledgment frequency parameter specifies the maximum number of frames the system can receive before acknowledgment is sent. Personal Communications defaults the LAN acknowledgment timer to a value of 100 milliseconds. This value, in the AS/400 controller description, is defaulted to *CALC. Set this value to 30 milliseconds. In Personal Communications, the receive window count is used to determine the LAN acknowledgment frequency $((\text{receive window count}/4) + 1)$. A LAN acknowledgment frequency value of 3 should work best for VI/400.

LAN Response Timer and LAN Frame Retry Count

The LAN response timer and LAN frame retry parameters are used to determine how soon and how often a frame of data is transmitted again. The LAN response timer defines the time interval that the system waits before requesting a response for a frame that was sent. The LAN frame retry value specifies how many times the frame is transmitted again. Personal Communications defaults the LAN response timer to 3000 milliseconds. This value, in the AS/400 controller description, is defaulted to *CALC. If the AS/400 and workstation are on the same local area network, this value may be lowered to the minimum value of 500 milliseconds. The LAN frame retry value should be 10 or less.

LAN Maximum Outstanding Frames

Specifies the maximum number of frames that can be sent before waiting for an acknowledgment. The LAN maximum outstanding frames value on one station should never be less than the LAN acknowledgment frequency on the other station. This can significantly degrade performance. A value of 7 should work best for VI/400.

Transmission Control Protocol / Internet Protocol (TCP/IP)

For TCP/IP communications, the key performance parameters are the frame size, maximum transmission unit size, and the window size.

Frame Size

TCP/IP can use the line description frame size to determine the number of data characters to send in a single physical transmission. Use the largest frames size supported by the network.

Maximum Transmission Unit

The maximum transmission unit is the maximum size (in bytes) of IP datagrams that can be transmitted. A datagram is a basic unit of information passed over an internet network. The minimum size of any maximum transmission unit value is 576 bytes. The maximum MTU that can

be specified depends on the type of physical connection to the network. For Token Ring (16 Mbps), the maximum is 16388. For Ethernet, the maximum is 1500. Using the maximum values should improve performance.

Window Size

The CHGTCPA command may be used to specify a TCP/IP window size for sending and receiving. The minimum TCP buffer send and receive sizes should be 8192. The AS/400 performance improves with values up to 64,386.

Client

Operating System

To avoid potential memory resource problems, use Windows NT versus Windows 95, if possible.

Workstation Configuration

The VisualInfo Client is responsible for the presentation functions of the application, and therefore is an important component. Although there are minimum requirements for the client, informal testing and customer experience have shown that minimum configurations may not always provide good performance in all environments.

Even just a couple of years ago, hardware costs were enough to make us carefully consider a small processor or memory upgrade. However, as hardware costs continue to drop, and the cost of our people who use these systems continues to climb, the hardware cost differential between an adequate PC and a good PC just doesn't seem to be worth worrying about.

Processor Speed

The faster the processor in the client, the more quickly it can present the information it receives from the client. In general, the Pentium-class processors with speeds of 166MHz and above will generally provide better throughput. If the workstation will be performing only VI client functions, the speed of the processor can be slower than if the workstation is concurrently performing other application tasks.

Memory Requirements

The recommended memory configuration is 64MB. Memory in the client can make a significant impact on performance. While this may seem obvious, the following real-world example drives home this fact. At one customer site, it took 17 seconds to print a 4 page MO:DCA document on a Windows/95 client with 32MB memory.

When the customer upgraded the client to 64MB memory, the same document printed in 4 seconds.

This improvement is consistent with what is generally expected when you add memory to any computer. If the processor is memory constrained, nothing will run well. Memory is getting less expensive every day, so it may be a wise use of money to add some extra memory.

In the Windows/NT operating system, you can easily find out how much memory is being used. Open the Task Manager by clicking the right mouse button on the task bar. Select the Performance tab, and the charts show you CPU and memory usage. If the amount of memory used is more than the physical memory in the workstation, the processor will be paging into and out of memory, and that will slow performance.

If the workstation will be used for applications in addition to VI/400, it is important to check memory usage when other applications are running as well, to get an accurate picture of what the user will experience.

Windows/NT also provides a performance monitor as an administrative tool. To check that out, select Start, Programs, Administrative Tools (common), Performance Monitor.

Disk Space

The general rules for disk space on the client apply. There should be enough space available on the drive so that there is room for memory paging space, as well as work space for the client application to store temporary data. Of course, you must also include any space required by any other applications that run on that workstation.

Maximum Table of Contents Count

When you perform a search, open a folder, or open a workbasket, the resulting items are presented in a Table of Contents. The client application allows you to specify a limit for the number of items displayed in a Table of Contents. By default, Table of Contents windows display up to 100 items. If an index class or search result contains more than 100 items, only the first 100 are listed.

To control the number of items that display in a table of contents, do the following:

1. Select Preferences from the Options menu. The Operator Preferences window opens.
2. Click the Advanced tab.
3. In the Maximum TOC Count field, enter the desired number of items.
4. Click OK.

There is no upper limit to the number of items that display in a table of contents, but the higher you set your maximum number, the slower your system's performance can be when you retrieve a table of contents.

If you enter 0 as your maximum count, your tables of contents list all items.

Workbasket Count

Displaying the number of items in a workbasket slows processing speed, so if speed is an issue, you might want to disable this option. To disable the item count, do the following:

1. From the Options menu, select Preferences. This opens the Operator Preferences window.
2. Under Workbasket, unselect the check box for Display the number of items in each workbasket.

System-Assigned Workbasket

If a workbasket is defined as system-assigned, the next available item to be worked is automatically presented to the user. If a workbasket is not system-assigned, the user is presented a list of items in the workbasket to select from. There is a cost in terms of performance to display and maintain the workbasket TOC. If the workbasket contains a large number of items, consider making it a system-assigned workbasket.

Server

AS/400 Performance Tuning

Any effort to improve performance should begin with ensuring the AS/400 is properly tuned. Performance adjustments to the system may be made automatically or manually. In general, making performance adjustments automatically is the best approach. This may be accomplished by setting the QPFRADJ system value to indicate that the system should automatically adjust storage pools and activity levels. Refer to the document AS/400 Work Management (SC41-5306) for additional information regarding AS/400 performance tuning.

Logon

Logging on the VisualInfo for AS/400 is a heavy processing activity. When a user logs on to VisualInfo/400, the server sends a number of lists back to be cached in the client. At logon, a search (SQL query) returns index class information and attributes for all index classes, content classes, privilege strings for all tables, a list of all work baskets, a list of all workflow processes, and so forth.

This information is intended to save the client from having to search the server tables every time one of these items is needed. However, this can make the logon activity appear to be very slow to the user.

The only way to really "tune" this activity is to control the size of the tables to be returned to the client. Therefore, any unnecessary index classes, privilege strings, and processes that are not

required could be deleted from the system, and this would have a positive impact on the logon performance. Of course, this is all relative. Deleting one process from a set of 100 would not make a large difference to the logon, but deleting 20 of those 100 processes would.

The other basic way to improve this performance is to have the user stay logged on. As long as the cost of concurrent licenses is not an issue, this will provide the most consistent benefit.

Search

A key component of VI/400 is the search capability. Both Basic and Advanced Search can be used to select items from the VI/400 library server database, and there are a number of things that affect the performance of the searches. Every search that is requested runs a DB2/400 dynamic SQL query across the library server database. Therefore, whatever can be done to improve database query will improve the VI/400 searches. Note that a search against all index classes does not include FLRCLASS index class (this is used for legacy WAF and there are no entries in EKD0312 for this index class). There is also no search on DOCCLASS (all documents in EKD0310).

The VI/400 library server database supports up to 8 user-defined key fields for each index class. To support the search capability, there is a view across the physical data organized by each of the key fields on the index class. This view is an AS/400 logical file. However, these views are on the primary key only, so that a search using more than one key might still have to search a large number of records.

If there were an access path that contained each combination of key fields, the SQL query optimizer would be able to choose the access path that matched the key fields used in the search. However, this would require more than 40,000 views, and the system overhead to support these views would negate any query savings.

The best approach to improving the search performance is to build only those access paths that are most frequently used. If a large number of searches always use the same three key fields in the same order, having an access path with those key fields would improve those searches. This approach would be useful for any highly used combination of key fields searched. The problem is how to find out which ones are used. The solution, provided in all RISC-based AS/400 operating systems is a command called STRDBMON (Start Database Monitor). This command will log all file accesses for jobs that you request, and indicate whether an access path was used, and if not, make recommendations for access paths that would improve access.

This command can be run for one or all jobs, and will report its findings. Based on the number of times it built an access path, you can determine what access paths would be most useful.

Reorganizing Physical Files

Another potential area for performance improvement is in file reorganization. As records are added and deleted from various tables in the database, the space occupied by the deleted records becomes unusable. Reorganizing the tables will reclaim that space and improve file access times. .

File System

Objects may be stored or accessed using a network drive mapped to the VI/400 object server. Alternatively, if there is not a network drive, VI/400 will store and retrieve objects through the command interface. A network drive is required for access to objects on a secondary object server.

Objects are stored in directories as defined by the VI/400 system administrator. The system administrator has the option to specify which file system (i.e., "Root", QDLS, QOpenSys) the object directories are to be contained within.

A series of test were conducted, varying the file system and the network drive option, to determine their effect on performance. These tests consisted of calling the VI/400 APIs SimLibCreateItem and SimLibCatalogObject repeatedly for a period of 1 hour to import a 70K object into VI/400. The results of the tests are reflected in the following table:

File System	Network Drive	Documents Created/HR
Root	No	870
Root	Yes (Client Access)	856
Root	Yes (NetServer)	858
QDLS	No	765
QDLS	Yes (Client Access)	766

The tests do indicate that there is some performance benefit associated with creating objects in the "Root" file system. The network drive option does not appear to have a significant effect on the rate of object creation. For all tests, there was not a significant variance in CPU or DASD arm utilization.

Security

Security configuration can also affect performance within VI/400. Since the server will only return items to which a user is authorized, there can be some performance benefit realized by limiting user access to items within VI/400. On the other hand, there is additional processing incurred to determine user authorization when securing items using groups and access lists.

Storage Management

Within VisualInfo for AS/400, there are several storage management jobs which operate in background. These jobs perform object deletion, interactive retrieve, optical distribution, optical retrieve, and optical store. If expired documents can be removed during off-shift, run the object deletion process then, instead of during prime shift activity.

Optical Subsystem

General Performance Considerations

Several factors can affect the optical performance of both LAN-attached and directly attached optical media libraries.

Volume Mounting and Dismounting

Volume mounting and dismounting is one of the most important factors affecting optical performance. It takes approximately 8 to 15 seconds to remove a volume, store it in a slot, retrieve a new volume, and mount it. If you can minimize the number of volume mounts and dismounts that your application requires, optical performance will improve.

Drive Contention

Performance can be severely affected by drive contention. The following conditions increase drive contention and should be avoided:

- Having only one drive available for use by applications libraries.
- Having many optical processes running that attempt to use different optical volumes at the same time.

Number of Directories and Files

Performance can be affected by having too few directories with too many files. Directories group related information to provide a means of quicker access. Typically, you get better performance from more directories with fewer files. Although there is no enforced limit on how many files there can be in a directory, you probably should not have more than 6000 for performance reasons.

Note: The number of files and directories does not affect the performance of the DUPOPT command.

File Size

The size of a file has a direct effect on the amount of time it takes to read, write, or copy the file. In general, the larger the file, the longer the operation can be expected to take.

Performance Considerations--Directly Attached Optical Libraries

Effects of File Attributes on Performance

When a file has extended attributes, they are stored separately from the data. When the data is written or copied, the attributes must also be written or copied. If file attributes are not required, attribute copying can be suppressed when copying between the QOPT and QDLS files systems by using the Change Optical (CHGOPTA) command. Setting the copy attributes (CPYATR) value on the CHGOPTA command to *NO suppresses the copying of attributes between the QOPT and QDLS file systems.

Volume Mounting and Dismounting Schedule Timers

This topic provides information on how the AS/400 system manages work requests to the directly attached IBM 3995 Compact Optical Library Dataserver. It is a high-level view and does not include all program logic.

You can change the queuing and scheduling logic used by the AS/400 system for directly attached optical media libraries by using the Change Device Description (Media Library) (CHGDEVMLB) command.

There are two timer values associated with optical media libraries that affect the scheduling of volume mounting and preemptive dismounting. Both can be changed using the CHGDEVMLB command. The first timer value (UNLOADWAIT) is the unload wait time. This value determines how long the system waits for a new request for a mounted volume before removing it. The second timer value (MAXDEVTIME) is the maximum device wait time. This value determines how long a volume with active requests can remain in a drive when there are requests waiting for other volumes that need to be mounted in a drive.

By using these two timer values, you can adjust the volume mount scheduling that is used by the optical media library to match your application's use of optical volumes.

You can change these timer values at any time; however, the new timer values will not become effective until the next time the device is varied on.

System job priority and limit timers are used to schedule volume mounting. The maximum number of volumes the 3995 can have mounted is equal to the number of drives in the optical media library. Keep the following points in mind as you schedule volume mounting:

A volume can remain mounted in an optical drive for the maximum device wait time if work requests with the same or lower job priority for a different volume have been received. An exception to this is when you are initializing a rewritable volume or using DUPOPT; the volume remains in the drive until the operation completes.

Work requests for mounted volumes are serviced before requests of the same or lower job priority for volumes not mounted.

For a multiple job environment, volumes are mounted based on the job priority for the work request. A work request from a job with a higher system job priority causes the required volume to be mounted to handle that request. The volume remains mounted for the maximum device wait time if work requests continue, dismounts after unload wait time inactivity, or is pre-empted by a work request from a job with higher system priority. If you are initializing a rewritable volume or using DUPOPT, the volume remains mounted until the operation completes.

If the work on a drive is interrupted because of a higher priority request, the maximum device wait time timer for the currently mounted volume is canceled. All future requests for that volume are queued for normal processing by priority.

If the volume needed for a work request is not mounted within 30 minutes, the job fails due to a time-out.

Performance Considerations--LAN-Attached Optical Libraries

This topic describes the ways you can control the performance of applications using an optical LAN server.

Controlling the Block Size for an Optical LAN Server

Optical LAN support is optimized to receive and send the largest possible block of data across the network. Currently, optical LAN support communicates with servers using data frames containing a maximum of 31 KB bytes of information. This does not directly affect the HFS user, but to minimize communications time, applications using large files should avoid read and write operations of less than 31 KB. Read or write operations larger than 31 KB are not a problem because they are automatically broken up into the optimum size.

Controlling Block Size for Optical LAN Server--Examples

Consider an example application that has 64 KB bytes of information to be written to an optical volume. The application could be designed to write blocks of 1 KB, 31 KB, or the entire file. If 1 KB blocks were written, the application would make 64 write requests. These 64 write requests would result in 64 communications requests.

If 31 KB blocks were written, the application would make 3 write requests (2(31 KB) + 1(2 KB)). These 3 write requests would result in 3 communications requests.

Finally, if the application requested that the entire file be written, there would be 1 write request. The file would automatically be broken up into 3 communications requests (2(31 KB) + 1(2 KB)).

All three approaches work, but the best performing one is the single write request of the entire file.

Capture Subsystem

The document capture subsystem refers to the components used to enter the documents into the imaging system. This is one of the most critical subsystems, since scanning documents accurately, reliably, and quickly is at the core of the ability to perform work from the digitized image.

Scanning

Scanners are used to capture and translate the paper document into digital data or image. This "cost of capture" includes the scanning of the source documents and indexing the images into the system. The index information includes at least one unique primary key and secondary keys that may be used to search and locate the object once it is stored into the system. Some solutions include auto-indexing the documents by using a bar code or by performing some type of Scanner performance is measured in pages per minute (ppm), and in impressions per minute (ipm) for duplex scanners, which is the number of pages per minute times two. There are desktop scanners (10-20 ppm), mid-range scanners (40-60 ppm), and high volume scanners (80-120 ppm) on the market today.

All scanners have rated speeds, the speed at which the scanner mechanism can capture the documents. However, the rated scanner speed provided by the manufacture can only be achieved under ideal conditions. Actual scanner throughput is considerably lower than the rated scanner speed. The only way to determine the actual scanner speed in your environment is to run a test using sample documents.

Some documents may need to be scanned in grayscale, which is slower than bitonal (black and white) scanning. Many documents can be scanned in at a resolution of 200 dpi but, for OCR or forms recognition, scan resolutions of 300 dpi or more are usually required. Depending on the scanner, scanning at 200 dpi can be twice as fast as scanning at 300 dpi. Selecting grayscale or bitonal scanning and the resolution (dpi) has a dramatic impact on the performance of the entire imaging system. For example, an 8.5x11 document captured at 200 dpi has a file size of 467 500 bytes, while the same document captured at 300 dpi is 1 051 875 bytes (uncompressed).

Compression

Another factor that affects the size of the generated file is the compression algorithm used to compress the bitmap file. Compression algorithms are classified as either lossy or lossless. An image viewed after lossless compression is identical to the way it was before being compressed. An image viewed after lossy compression is different from before it was compressed because some information has been lost. Commonly used compression formats include CCITT Group III,

Group IV, and JPEG. Compressed image size is typically one tenth of the original or less. The table below shows the file size resulting from scanning one page using TIFF6 format for different DPI and Grayscale values.

File Size after Compression (K Bytes):

DPI	100	200	400
Bitonal (B/W)	23.7	42.6	92.1
256-level Grayscale	104.2	292.5	796.8

High-volume Image Capture

The VisualInfo client includes the capability to scan documents with desktop scanners or to import files directly into VI/400. The client includes Pixel ISIS drivers, which support a variety of scanners. Once scanned, documents can be routed to an appropriate workbasket (queue) for indexing.

For higher volume scanning, a batch capture subsystem is usually implemented. Kofax Ascent Capture and InputAccel from Input Software are production-level document capture applications that integrate high-performance batch scanning, image processing, Optical Character Recognition (OCR), Intelligent Character Recognition (ICR) and document indexing into a single package that significantly lowers the cost of capturing large quantities of documents into VisualInfo. These capture subsystems have "release" processes which integrate with VI/400. They provide for scanning in a disconnected mode so that you can determine the best time to upload batches of scanned documents to your VI/400 system. The ability to move the input processing workload off-hours provides a significant performance benefit. In addition, since indexing is performed at the workstation using the front end capture application, the indexing workload is eliminated from VI/400.

Chapter 3. VisualInfo for AS/400 Capacity Sizing

Capacity sizing is the process of estimating system resources required for a specified application workload, such as VisualInfo for AS/400. The capacity sizing process encompasses the following steps:

- Understanding customer requirements
- Defining the workload
- Modeling the workload

Understanding Customer Requirements

The initial step in a capacity sizing exercise is to understand the existing customer environment and identify ImagePlus requirements. Generally, requirements are gathered from customer site studies and walkthroughs. As a part of this effort, the existing system of paper documents and processing is analyzed to determine:

Workflow characteristics

- Where work or documents originate
 - ▲ External mail
 - ▲ Internal mail
 - ▲ Computer output
 - ▲ Fax
- How documents are routed
- How documents are processed
- Whether document processing follow regular, repeatable procedure
 - ▲ Automated workflow with predefined process
 - ▲ Ad hoc workflow
- Whether additional information is required before working a document
- The information that effects the employee's decisions when working a document

Filing characteristics

- Where are documents stored
- How are documents retrieved
- How are documents and folders indexed
- Retrieval rates
- Access controls

Document characteristics

- What are the different type of documents (e.g., invoices, claims)
- Document size
- Pages per document

- Life cycle
- How many documents are received daily

Defining the Workload

Once the existing customer environment and requirements are understood, it is possible to make fundamental system design decisions related to VI/400. Performing the design will allow for specific customer workloads to be determined, and eventually modeled.

Data Model

Based on the requirements analysis, you should have an understanding of the paper documents currently processed. You need to relate this information to the VI/400 data model concepts of documents, folders, index classes, and key fields.

Documents

A VI/400 document item will usually correspond to an existing paper document. Associated with the document, there may be notes or annotations. The base document, notes, and annotations are each separate objects (files). Collectively, they make up a VI/400 document. For each type of document, the following questions should be answered:

1. How is the document indexed?

Determine the index class and associated attributes that will be used to represent the document.

2. What is the document size?

Obtain real documents from the customer and scan to determine the document size. A typical single sided 8 ½ x 11 inch sheet of paper will result in a 50 KB image when scanned at 200 DPI.

3. How many days will the objects reside on DASD?

A document should reside on DASD for the number of days that is actively worked and frequently accessed. A common value for the life of a document on DASD is 15 days.

4. How many documents are received daily?

Determine how many documents of this type are created each day.

Folders

A folder is a container for document items and other folders. You may think of a folder as the logical equivalent of a folder in a file cabinet containing documents that have some sort of relation to each other. The following questions regarding folders should be answered:

1. How is the folder indexed?

Determine the index class and associated attributes that will be used to represent the folder.

2. What items will the folder contain?

Determine what documents and folders the folder will contain.

Input Processing

Input processing refers to the method in which documents enter the system and are indexed for subsequent retrieval. Documents can be scanned, imported, or faxed into VI/400. This may be accomplished using the VI/400 client or some alternative scanning subsystem. From a workload perspective, the method of input is not really significant. That is, internally the processing is the same, and therefore, the workload is the same regardless of how documents enter the system. More significant, in terms of workload, is the indexing process.

Using the VI/400 client, documents are scanned or imported in one step, while indexing is a separate process. Typically, scanned or imported documents are assigned to the NOINDEX index class and added to a workbasket where they are subsequently indexed. The same is true for documents arriving via fax. So, using the VI/400 client for input processing there will be two workloads, one for input (scan, import, fax) and another for indexing.

Since most document capture applications provide high-performance batch scanning and document indexing in a single package, when these capture subsystems “release” documents to VI/400, the documents are input and indexed in one step. This eliminates the separate indexing workload. Additionally, with the capture subsystems, you determine the best time to upload batches of scanned documents to your VI/400 system allowing you the ability to entirely move the input processing workload off-hours.

In addition to determining how documents are input and indexing, to characterize the input processing workload, the following additional questions need to be answered:

1. What is the arrival rate?

This is best expressed in terms of documents per hour. In many cases, documents will not be input evenly throughout the day. For example, scanning may occur for only a number of hours

during the day. For sizing purposes, it is important that the workload reflect the actual input rate per hour and not a rate for the entire period of operations.

2. Is “Autofoldering” performed?

During indexing, it is possible to automatically add a document to a folder based on a key field value. If a folder does not exist, it will be created and the document added.

3. Is the item started on a workflow?

During indexing, it is also possible to automatically start an item on a workflow.

Item Processing

Within VI/400, items may be accessed by performing searches or by opening a workbasket. Once an item is accessed, there are various functions that may be performed on the item. For example, adding a note, annotating, printing, faxing, or continuing along a workflow. Some of these actions will effect the AS/400 workload, others will not. How an item is accessed, the associated processing, and rate of item processing effects the AS/400 workload.

The decision to use the advanced workflow feature of VI/400 will also effect the AS/400 workload.

The workloads related to processing items in VI/400 can be expressed, in several ways. For example, the following parameters could be used to describe a workload:

- 100 searches per day
- 200 documents displayed
- 10 hours of operation per day

Alternatively, the same workload could be described with the following parameters:

- 5 users
- 2 searches per hour per user
- 2 documents displayed as a result of a search

Since Best/1 will be used for modeling the workloads, you want to eventually determine the hourly workload for each different type of user.

Storage Management

Within VisualInfo for AS/400, there are several storage management jobs which operate in background. These jobs perform object deletion, interactive retrieve, optical distribution, optical retrieve, and optical store. These jobs will create an additional load on the server. The following parameters are used to characterize this workload:

- Number of optical stores
- Number of optical retrieves
- Number of objects deleted
- Number of store processors
- Number of retrieve processors
- Number of delete processors

Modeling Workload

Once the VI/400 workload profiles are understood, it is possible to estimate system resources required to support these workloads. This stage of the capacity sizing process should be performed using BEST/1. BEST/1 is a capacity planning tool for the AS/400 that is a part of the AS/400 Performance Tools licensed program (5769-PT1). For a complete understanding of the BEST/1 tool it is recommended that you obtain the manual BEST/1 Capacity Planning Tool, SC41-3341.

To create a BEST/1 model from the workloads that have been defined involves the following steps:

1. Simulating the defined workloads
2. Collecting performance data
3. Creating a model from performance data

This process of modeling the workloads you defined will provide the most accurate sizing estimates. However, the above process is obviously a significant effort. Given this, there is an alternative method for estimating system resources that provides a lesser degree of accuracy, but is easier to perform.

This alternative involves using predefined workloads that represent typical VI/400 activities. Any method of sizing requires benchmarks and performance measurements based on a defined workload. In this case, these benchmarks and measurements of certain VI/400 workloads have been done for you. BEST/1 is still used to create the model, however, the predefined workloads may be used. Using predefined workloads, instead of the actual workloads you defined, will have an effect on the accuracy of the capacity sizing. The accuracy will be dependent on how closely the predefined workloads reflect the actual workloads. Also, it is very important to emphasize that the addition of the individual workload utilizations is not linear. That is, if the two workloads were run on the same system concurrently, it is unlikely that the total CPU utilization would be the sum of the utilizations obtained by running the workloads independently. This is a result the increased level of contention on the system. However, at low levels of utilization on a system,

reasonably accurate answers may be obtained by adding utilization figures. Additional information regarding the use of predefined workloads may be found in Appendix B.

If you do not have access to BEST/1 for the purpose of modeling the workload, you may provide the information detailed in Appendix C to the Image Competency Center to have a sizing performed.

Additional Capacity Sizing Considerations

Estimating Storage Requirements

Follow these steps to estimate DASD storage requirements for document objects:

1. Calculate average document size = average size of a page * average number of pages
2. Calculate daily storage requirement = number of documents scanned and imported daily * average document size (1).
3. Calculate DASD required for documents = daily storage requirement(2) * average number of days document is stored on DASD.

Note: If the number of days on DASD is different for different types of documents, you should perform this calculation separately for each document that has a different number of days on DASD, to be as accurate in your estimate as possible.

Follow these steps to estimate DASD storage requirements for note and annotation objects:

1. Determine average size of note and annotation part
2. Calculate daily storage requirement = number of new parts (note, annotation) * average part size (1).
3. Calculate DASD required for additional parts = daily storage requirement(2) * average number of days object is retained on DASD.

To approximate the storage requirements for the VI/400 index files, complete the following computation:

DASD required for index files = 1200 Bytes * Number of document items
+ 618 Bytes * Number of folders
+ 683 Bytes * Number of item parts (note, annotation)
+ 66 Bytes * Number of items in folders

Disk Arm Requirements

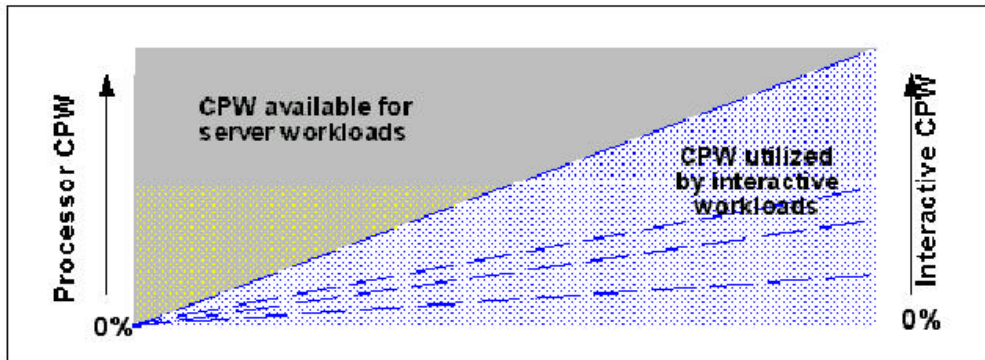
With the availability of higher capacity disk devices, fewer drives are required to satisfy storage requirements. However, since the VI/400 application is I/O intensive, configuring too few disk arms can create a bottleneck and constrain CPU performance. Therefore, it is important to not just consider storage requirements when determining the number and type of disk devices.

AS/400e Servers

AS/400e servers 720, 730, and 740 can be configured to meet a wide range of performance requirements. Whether the system is running mostly back-office applications, newer e-business applications, or a mixture of both, the performance can be customized on the new 7XX servers to match business needs. Each model includes a base processor and interactive performance feature. Optional processor and interactive performance features can be selected to balance the server's performance for a given workload. Increasing performance on installed servers can be done simply and with little disruption. Unlike their predecessors, the AS/400e servers 720, 730, and 740 now offer two CPW ratings to allow performance to be customized for a given workload:

- Processor CPW represents maximum relative performance for a processor feature running commercial processing workloads. This value can be used to compare relative performance between various processor features. Processor CPW was known as Client/Server CPW in prior releases.
- Interactive CPW represents the relative performance available to perform interactive workloads. The amount of Interactive CPW consumed reduces the available Processor CPW by the same amount. Interactive CPW for a specific processor is changed through the installation of optional Interactive Features in the Models 720, 730, or 740.

The interaction of the Processor CPW and Interactive CPW is illustrated by the following Chart.



This chart shows the CPW available for server workloads as the interactive workload increases. On the left side of the chart is a scale from 0-100% representing the amount of Processor CPW available for server workloads (non 5250-type workload). On the right-hand side of the chart is a similar scale reflecting the amount of Interactive CPW being utilized by Interactive workloads (5250 based). Each of the dotted lines dissecting the rectangle represents various levels of Interactive CPW that can be purchased when ordering an Interactive Feature for 7XX servers. (Interactive features are not available on the 170). At any point in time, the amount of CPW used to perform interactive workloads reduces the CPW available for server workloads by an equal and proportionate amount.

For example:

If 70% of the AS/400 is being used for Interactive workloads (that is, the Interactive Feature offers a CPW level equal to the Processor CPW), then 30% processing power is left on the machine to perform server workloads. No tuning or management is required. If on the other hand, none of the AS/400 is being used for Interactive workloads, then all of the Processor CPW is available to perform server workloads. No tuning or management is required.

For best performance, all critical system resources should be kept in balance by proper configuration. Most of the workload associated with VI/400 is server related. System administration is the only interactive processing within VI/400. Given this, the Processor CPW should be maximized.

Appendix A. Creating a Model from Simulated Workloads

As previously mentioned, the most accurate method of capacity sizing involves simulating the defined workloads, collecting performance data, and creating a model from the performance data.

Simulating a Workload

Workloads may be simulated using the VI/400 API test driver, VI400TST. The VI/400 API test driver provides the capability to interactively call the VI/400 APIs using a menu-like interface. As an alternative to interactively selecting APIs and specifying the appropriate parameter values, the test driver can also be driven using a response file. The VI/400 API test driver and response files provides an easy mechanism for simulating workloads. The following is a sample response file:

```
*****
* Create an item in the NOINDEX index class
*****
301          SimLibCreateItem
0           Expected return code
1           Document
6           Index class
3           Number of attributes
40          Source
CAPTURE
41          User
PERFUSR
42          Timestamp
1999-01-25-15.11.34.237000

*****
* Import the "scanned" document
*****
107         SimLibCatalogObject
0           Expected return code
null
C:\VIPERF\Perf70k.mda
1           Part number
1           Content class (MODCA)
0           Close
128        Affiliated type
+DFT
```

To create a response file which reflects the various workloads, perform the following steps:

1. Turn VI/400 logging on.
2. Use the VI/400 client application to run the function associated with the workload. For example, scan a document, index a document, or process an item.
3. From the VI400.LOG, identify the APIs and associated parameters used to perform the function.
4. Create a VI/400 test driver response file with the appropriate APIs and parameters as determined from the VI/400 log.

Input can be read from a file by using either VI400TST < user.input.file or option 1000, which will cause reading to switch to the file name entered next. When reading from a response file, prompts will not be displayed, and input will be echoed to the screen. Also, each function requested must be followed by the expected return code. After the function has been called, the expected return code is compared to the actual, and a message displayed if they do not match. It is also possible to nest response files.

Function number 1002 will cause processing to pause and you must press ENTER to continue. This allows checking results as the response file is run. If a response file contains several pauses, they can be skipped by using 1003. Entering 1003 again will toggle back to pause mode.

Lines beginning with "*" will simply be echoed. Numeric responses can be followed by comments. Alphanumeric responses must not have trailing blanks or comments.

While the VI/400 API test driver delivered with the product is acceptable for simulating a workload, there is a modified version of this program which contains some additional capabilities to simplify this task. The modified VI400TST, available from <ftp://ftp.software.ibm.com/as400/products/visualinfo/performance/>, contains the following enhancements:

- A delay function for simulating key/think time
- An iterate function which allows for the repeated execution of a response file
- A set item identifier function which can be used to set the current item identifier to a specific TOC entry
- The ability to read attribute values and search expressions from a data file
- The ability to open multiple attribute sets (SimLibOpenItemAttr).

The following response file demonstrates some of these additional enhancements.

```

*****
* Perform a basic search
* The search expression is set to null, so read the
* search expression from a data file
*****
801          SimLibSearch
0           Expected return code
1024        Both
3           SIM_SEARCH_USER_ATTR
101         ICDOC1
101         Return limit
null
9999

*****
* Change attribute
* Obtain the attribute value from a data file
*****
203          SimLibWriteAttr
0           Expected return code
101         Attribute ID
null

*****
* Set current (default) item ID to first TOC entry
*****
1009         Set current item ID
1

*****
* Delay for 5 seconds- Key/Think time
*****
1006         Delay
5000        5 seconds

*****
* Iterate through the response file 100 times
*****
1007         Iterate
100          Iteration count

```

The function number to specify a data file is 1008. Typically, this will be entered interactively before issuing the function to use a response file. The following data file contains an attribute value and a search expression. The data must not have trailing blanks or comments.

```
PERFORMANCE101
A101 == "PERFORMANCE101"
```

Measuring the Workload

Prior to measuring a workload, you will need to identify a capacity planning unit of measure. The workload will be measured with regard to this capacity planning unit of measure. From a VI/400 perspective, this will typically be a measurement in terms of documents scanned/hour, documents indexed/hour, or an items processed/hour.

After determining how the workload will be measured, the next step is to simulate the workload and collect AS/400 performance data. Use the VI400TST test driver to run repeatedly the response file which reflects a specific workload. While the workload is being simulated, collect AS/400 performance data. Use the OS/400 STRPFRMON command for performance data collection to start the performance monitor.

```

                                Start Performance Monitor (STRPFRMON)

Type choices, press Enter.

Member . . . . . *GEN      Name, *GEN
Library . . . . . > QPFRDATA  Name
Text 'description' . . . . . Performance Data for VI/400
-----
Time interval (in minutes) . . . 05      5, 10, 15, 20, 25, 30, 35...
Stops data collection . . . . . *ELAPSED *ELAPSED, *TIME, *NOMAX
Days from current day . . . . . 0      0-9
Hour . . . . . 1      0-999
Minutes . . . . . 0      0-99
Data type . . . . . *ALL      *ALL, *SYS
Select jobs . . . . . *ALL      *ALL, *ACTIVE
Trace type . . . . . *NONE     *NONE, *ALL
Dump the trace . . . . . *YES     *YES, *NO
Job trace interval . . . . . .5      .5 - 9.9 seconds
Job types . . . . . *DFT      *NONE, *DFT, *ASJ, *BCH...
                                + for more values
                                -----
                                More...

F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
F13=How to use this display   F24=More keys
```

Enter the STRPFRMON command parameters and collect performance data for a duration of one hour. The duration over which performance data is collected can be specified at the time

STRPFRMON is run. Alternatively, the ENDPFRMON command can be used to end the performance monitor and stop data collection.

After running the workload for one hour, end the VI/400 API test driver and determine the number of capacity planning units of measure. For example, if an input processing workload was run, identify how many documents were scanned or imported. The workload measurement can be obtained in a number of ways. If the iterate function is being used to repeatedly run a response file, the displayed iteration count may be helpful in terms of measuring the workload. It may also be possible to obtain the workload measurement by determining how many items are contained within a workbasket or index class. If the workload is simulating scanning and adding an item to an indexing workbasket, the number of items scanned can be determined by the number of items in the indexing workbasket.

Creating a Model from Measured Data

BEST/1 provides the facility to use measured AS/400 performance data in building a model for capacity planning.

BEST/1 for the AS/400

Use the STRBEST command to start BEST/1.

```
                                BEST/1 for the AS/400

Select one of the following:

    1. Work with BEST/1 models

    5. Create BEST/1 model from performance data

   10. Work with results

   50. About BEST/1

   60. More BEST/1 options

Selection or command
====> 1

-----
F3=Exit  F4=Prompt  F9=Retrieve  F12=Cancel
```

On the BEST/1 for AS/400 menu, select option 1, Work with BEST/1 models.

Work with BEST/1 Models

```
Work with BEST/1 Models

Library . . . . . VI400DTA   Name

Type options, press Enter.
  1=Create   3=Copy   4=Delete   5=Work with   6=Print   7=Rename

Opt  Model      Text                               Date      Time
  1  VI400MODEL

(No models in library)

Command
====> _____
F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F12=Cancel
F15=Sort by model   F16=Sort by text   F19=Sort by date and time

Bottom
```

Select option 1 to create a model and enter a name for the model to be built.

Create BEST/1 Model

```
Create BEST/1 Model

Select one of the following:

  1. Create from performance data
  2. Create from predefined and user-defined workloads

Selection
  1

F3=Exit   F12=Cancel
```

Select option 1, Create from performance data.

Create BEST/1 Model from Performance Data

```

Create BEST/1 Model from Performance Data

Model . . . . . : VI400MODEL

Type choices, press Enter. Use *SLTHOUR to select an hour-long time period or
use *SLTITV to select select first and last interval of a one to two hour
time period. The time period selected should be representative of your peak
processing activity.

Text . . . . . VI/400 Model

Performance data:
Member . . . . . Q991370837 Name, F4 for list
Library . . . . . QPFRDATA Name

Start time . . . . . *SLTITV Time, *FIRST, *SLTHOUR, *SLTITV
Start date . . . . . *FIRST Date, *FIRST

Stop time . . . . . *LAST Time, *LAST
Stop date . . . . . *LAST Date, *LAST

F3=Exit F4=Prompt F12=Cancel

```

Enter the member name and library containing the AS/400 performance data. Select an hour long time period. If you specify *SLTITV for the Start time, the following screen is displayed which allows you to select the start time and end time of the performance data to be included in the model.

Select Time Interval

```

Select Time Interval

Library . . . . . : QPFRDATA Performance member . . : Q991370837

Type option, press Enter. Select first and last interval.
1=Select

Opt  Date      Time          ---Transaction---  --CPU Util--  I/Os per Sec
      Date      Time          Count   Rsp Time  Total  Inter  Sync  Async
  1  05/17/99  08:43:48      0        .0       17     0     31    15
  -  05/17/99  08:48:49      0        .0       18     0     31    15
  -  05/17/99  08:53:50      0        .0       18     0     30    15
  -  05/17/99  08:58:50      0        .0       17     0     30    14
  -  05/17/99  09:03:51      0        .0       17     0     31    15
  -  05/17/99  09:08:51      0        .0       17     0     32    15
  -  05/17/99  09:13:52      0        .0       17     0     30    14
  -  05/17/99  09:18:52      0        .0       17     0     31    15
  -  05/17/99  09:23:53      0        .0       17     0     31    15
  -  05/17/99  09:28:53      0        .0       17     0     31    15
  -  05/17/99  09:33:53      0        .0       17     0     30    14
                                          More...

F3=Exit  F12=Cancel  F15=Sort by interval  F16=Sort by count
F17=Sort by rsp time  F18=Sort by total CPU util  F19=Sort by total I/Os

```

Classify Jobs

```

                                Classify Jobs

Select one of the following:

    1. Use default job classification
    2. Classify jobs into workloads
    3. Use existing job classifications

Selection
   2

F3=Exit  F12=Cancel

```

Select option 2, Classify jobs into workloads.

Specify Job Classification Category

```

                                Specify Job Classification Category

Type choice, press Enter.

Category . . . . . 1

1=User ID
2=Job type
3=Job name
4=Account code
5=Job number
6=Subsystem
7=Pool
8=Control unit
9=Comm line
10=Functional area

F3=Exit  F12=Cancel

```

This display allows you to select the job category you want to use in assigning the AS/400 server jobs to workloads. Options 1 and 3 work well for VI/400.

Assign the AS/400 server jobs to workloads. If you selected to categorize jobs by user ID, the user ID associated with the TCP/IP listener job can be used to assign the jobs to the appropriate workload. If you selected to categorize jobs by job name, you would identify the TCP/IP listener job in the list of jobs and assign it to the workload.

If you want to run and model several workloads at the same time, use multiple TCP/IP prestart jobs with different port numbers and user IDs. This will enable you to properly assign VI/400 jobs to the appropriate workload.

Specify Paging Behaviors

```
Specify Paging Behaviors

Type choices, press Enter.

Workload      Paging Behavior
              (F4 for list)
QDEFAULT      *GENERIC
VI400WORK     *GENERIC

F3=Exit  F4=Prompt  F12=Cancel

Bottom
```

Accept the default value of *GENERIC for all workloads.

Define Non-Interactive Transactions

The following display allows you to specify the number of capacity planning units of measure for each server workload. Change the value of the Type field to *NONE and enter the number of capacity planning units of measure in the Total Transactions field. The value entered is per hour. For example, if a scan workload was simulated and 300 documents were input in one hour, the Total Transactions field would contain a value of 300, representing the number of capacity planning units of measure.

Define Non-Interactive Transactions

Job classification category : User ID

Type choices, press Enter.

Workload	---Activity Counted as Transaction---	Quantity	Total Transactions when Type = *NONE
QDEFAULT	*LGLIO	100.0	0
VI400WORK	*NONE	100.0	300

Bottom

Type: *LGLIO, *CMNIO, *CPUSEC, *PRINT, *NONE

F3=Exit F12=Cancel

Save Job Classification Member

Save Job Classification Member

Change values if desired, press Enter.

Member	_____	Name
Library	VIDATA	Name
Text	_____	
Replace	N	Y=Yes, N=No

This display allows you to save the job classifications as a member in a library.

Confirm Creation of BEST/1 Model

Confirm Creation of BEST/1 Model

Type choices, press Enter.

Model	<u>VI400MODEL</u>	Name
Library	<u>VIDATA</u>	Name
Text	<u>VI/400 Model</u>	_____
Replace	<u>N</u>	Y=Yes, N=No
Job name	<u>CRTBESTMDL</u>	Name, *JOB
Job description	<u>QPFRJOB</u>	Name, *NONE, *USRPRF
Library	<u>QPFR</u>	Name, *LIBL, *CURLIB

F12=Cancel

Verify the fields and press Enter to confirm creation of the BEST/1 model.

Appendix B. Predefined Workloads

Predefined workloads can be used as an alternative to creating a model from measured data. The predefined workloads represent typical input processing, item processing, and storage management workloads in VI/400. Using predefined workloads, instead of the actual workloads you defined, will have an effect on the accuracy of the capacity sizing. The accuracy will be dependent on how closely the predefined workloads reflect the actual workloads.

The saved workloads are available from the following site:
<ftp://ftp.software.ibm.com/as400/products/visualinfo/performance/>

Creating a Model Using Predefined Workloads

BEST/1 for the AS/400

Use the STRBEST command to start BEST/1.

```

                                BEST/1 for the AS/400

Select one of the following:

    1. Work with BEST/1 models
    5. Create BEST/1 model from performance data
   10. Work with results
   50. About BEST/1
   60. More BEST/1 options

Selection or command
===> 1
-----
F3=Exit  F4=Prompt  F9=Retrieve  F12=Cancel
```

On the BEST/1 for AS/400 menu, select option 1, Work with BEST/1 models.

Work with BEST/1 Models

```
Work with BEST/1 Models

Library . . . . . VI400DATA   Name
Type options, press Enter.
  1=Create   3=Copy   4=Delete   5=Work with   6=Print   7=Rename

Opt Model      Text                Date      Time
  1  VI400MODEL

(No models in library)

Command
====> _____
F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F12=Cancel
F15=Sort by model   F16=Sort by text   F19=Sort by date and time

Bottom
```

Select option 1 to create a model and enter a name for the model to be built.

Create BEST/1 Model

```
Create BEST/1 Model

Select one of the following:

  1. Create from performance data
  2. Create from predefined and user-defined workloads

Selection
  1

F3=Exit   F12=Cancel
```

Select option 2, Create from predefined and user-defined workloads.

```

                                Create BEST/1 Model

Model/Text:  VIMODEL

Select one of the following:

    1. Work with workloads
    2. Specify objectives and active jobs

    5. Analyze current model
    6. Analyze current model and give recommendations
    7. Specify workload growth and analyze model

    10. Configuration menu
    11. Work with results

    50. About BEST/1

                                                                    Bottom

Selection or command
====> _____
F3=Exit   F4=Prompt   F9=Retrieve   F12=Cancel   F15=Save current model

```

Select option 1, Work with workloads.

Work with Workloads

```

                                Work with Workloads

Model/Text:  VIMODEL

Type options, press Enter.
  1=Create   2=Change   3=Copy   4=Delete   5=Display   6=Print   7=Rename
  8=Save workload to workload member   9=Edit transactions

Opt      Workload      Text
-----
(No workloads defined)

                                                                    Bottom

F3=Exit   F6=Add saved workload   F9=Add predefined workload   F12=Cancel
F13=Combine workloads

```

Select F6, Add saved workload.

Add Saved Workload

```

                                Add Saved Workload

Library . . . . . VIDATA      Name

Type option, press Enter.
  1=Select

Opt  Member      Text                                     Date      Time
  1  INPUT70WB   Input 70K document to WB                          05/24/99  15:27:09
  _  SEARCHDOC   Search/process document                          05/24/99  14:48:21
  _  SEARCHFLR   Search/process folder                            05/24/99  14:47:56
  _  INPUT70AF   Input 70K document; Auto-folder                  05/24/99  14:47:25
  _  INDX70AFAW  Index 70K doc; Autofolder; Autoworkflow          05/24/99  14:46:14
  _  INDEXT70AF  Index 70K document; Auto-folder                  05/24/99  14:44:34
  _  INDEXT70    Index 70K document                               05/24/99  14:35:06
  _  INPUT70AW   Input/index 70 document; auto-workflow           05/24/99  14:34:20
  _  INPUT70     Input 70K document to index class                 05/24/99  14:32:05
  _  INPUT200WB  Input 200K document to WB                        05/24/99  14:30:41

                                                Bottom
F3=Exit   F5=Refresh   F12=Cancel  F15=Sort by member  F16=Sort by text
F19=Sort by date and time
```

Use option 1 to select a saved workload. The Work with Workloads screen will be displayed again. To add additional workloads use F6=Add saved workload.

Modify Model

A BEST/1 workload is made up of one or more functions. Since the saved VI/400 workloads were created using performance data, the workloads contain one function per hour. This function represents all of the work done by this workload.

Each function is made up of one or more transactions. The transactions per hour represents the workload throughput. Depending on the workload, the transactions per hour will be, the number of documents input per hour, the number of documents indexed per hour, the number of items processed per hour, the number of objects stored to optical per hour, or the number of objects deleted per hour.

You can make adjustments to the model to more accurately reflect the environment that is being sized. These adjustments may be made to the a workload, function, or transaction.

Suppose an environment being sized contains three workloads: scanning, indexing and folder processing. Assume scanning and indexing will be accomplished using the VI/400 client application. Given this example, the following display reflects the saved workloads which are required for the model.


```

Work with Workloads

Model/Text:  VIMODEL      VI/400 Model

Type options, press Enter.
  1=Create   2=Change   3=Copy   4=Delete   5=Display   6=Print   7=Rename
  8=Save workload to workload member   9=Edit transactions

Opt      Workload      Text
-        INDEX70AF     Index 70K document; Auto-folder
-        INPUT70WB     Input 70K document and add to WB
9      SEARCHFLR     Search/process folder

Bottom
F3=Exit   F6=Add saved workload   F9=Add predefined workload   F12=Cancel
F13=Combine workloads

```

Selecting option 9, Edit transactions, for a workload displays the following screen.

```

Edit Transactions

Workload . . . . :  SEARCHFLR   Search/process folder
Workload type . . :  *NORMAL           CPU architecture :  *RISC

Type changes, press Enter.

Function      Trans  Pool      Transactions      CPU Time      Permanent
              Type   ID   Priority  per Function      (Secs)        Writes
SEARCHFLR    2    2    20      22.00          73.909      99.1
SEARCHFLR    2    1    0       1.00           3.784       57.8

Bottom

Trans Type: 1=Interactive, 2=Non-interactive

F3=Exit   F11=Alternative view   F12=Cancel   F14=Read I/Os   F15=Write I/Os
F16=Paging values   F17=Paging behaviors   F18=Chars transferred

```

The Transactions per Function field on the Edit Transactions screens represents the workload throughput. In this example of the SEARCHFLR workload, this is twenty two folders processed per hour. This value may be modified to more accurately reflect the environment being sized.

If option 2, Change, is selected from the Work with Workloads display, the following screen is displayed.

```

Change Workload

Workload . . . . . : SEARCHFLR
CPU architecture . . . . : *RISC

Type changes, press Enter.
Workload text . . . . . Search/process folder
Workload type . . . . . *NORMAL      F4 for list
Usage mode . . . . . 4              1=Casual, 2=Interrupted, 3=Steady,
                                      4=N/A

Function Text          Functions  Avg K/T  -----Tns per Function-----
per User              (secs)   Inter    Non-inter
Function of SEARCHFLR 1.00      N/A     .00      23.00

Bottom
F3=Exit  F4=Prompt  F6=Work with functions  F9=Specify chars to comm lines
F10=Specify I/Os to ASPs  F11=Show all functions  F12=Cancel

```

Since the saved VI/400 workloads were created using performance data, the workloads contain one function per hour. This function represents all of the work done by this workload. The Functions per User value may be modified to more accurately reflect the environment being sized.

Specify System Configuration

Prior to analyzing a model, BEST/1 requires a valid AS/400 system configuration. Since the model was created from predefined workloads, a hardware configuration is not automatically selected. However, BEST/1 supplies a set of standard AS/400 hardware configurations that may be used. If you attempt to analyze a model that has no hardware configuration defined, the Select Configuration Member display is presented allowing you to select one of the supplied configurations.

Analyze Model

After making adjustments to the model and specifying the system configuration, you may use BEST/1 to analyze the model. From the Create BEST/1 Model or Work with BEST/1 Model screens, select option 5 to analyze the current model.

Create BEST/1 Model

Model/Text: VIMODEL

Select one of the following:

1. Work with workloads
2. Specify objectives and active jobs

5. Analyze current model
6. Analyze current model and give recommendations
7. Specify workload growth and analyze model

10. Configuration menu
11. Work with results

50. About BEST/1

Bottom

Selection or command

====> 5

F3=Exit F4=Prompt F9=Retrieve F12=Cancel F15=Save current model

F20=More BEST/1 options

Using BEST/1, you can make simulations of changes to the hardware configuration and additional modifications to the model itself.

Input Processing Workloads

Member	Description	Capacity Planning Unit of Measure
INPUT70WB	Scan or import a 70K document and add it to an indexing workbasket. This workload reflects scanning or importing from the VI/400 client.	Documents input/hour
INPUT200WB	Scan or import a 200K document and add it to an indexing workbasket. This workload reflects scanning or importing from the VI/400 client.	Documents input/hour
INPUT70	Import and index a 70K document. The document is not added to an indexing workbasket. The implication is that the appropriate index information is available when the document item is created. This workload could be used to represent the “release” processing associated with a capture subsystem.	Documents imported and indexed/hour
INPUT70AF	Import and index a 70K document. Based on index values, the document will be added to a folder. If the folder does not exist, it is created and the document is then added. Each folder will contain three documents, so every third document input will cause a folder to be created. This workload could be used to represent the “release” processing associated with a capture subsystem.	Documents imported and indexed/hour
INPUT70AW	Import and index a 70K document. The document is indexed into the appropriate index class. The document item is started on a predefined process. This workload could be used to represent the “release” processing associated with a capture subsystem.	Documents imported and indexed/hour
INDEX70	This workload simulates indexing a document from a system-assigned workbasket using the VI/400 client application.	Documents indexed/hour

Member	Description	Capacity Planning Unit of Measure
INDEX70AF	This workload simulates indexing and auto-folding from a system-assigned workbasket using the VI/400 client application. Based on index values, the document will be added to a folder. If the folder does not exist, it is created and the document is then added. Each folder will contain three documents, so every third document indexed will cause a folder to be created.	Documents indexed/hour
INDX70AFW	This workload simulates indexing, auto-folding and auto-workflow from a system-assigned workbasket using the VI/400 client application. Based on index values, the document will be added to a folder. If the folder does not exist, it is created and the document is then added. The new folder is started on a predefined process. Each folder will contain three documents, so every third document indexed will cause a folder to be created and started on a process.	Documents indexed/hour

Item Processing Workloads

Member	Description	Capacity Planning Unit of Measure
SEARCHFLR	This workload simulates processing a folder. A basic search is performed for a specific folder. The search results contain a single folder which satisfies the search criteria. The folder is opened and a note is added to the folder. The folder contains three documents which are displayed.	Folders processed/hour
SEARCHDOC	This workload simulates processing a document. A basic search is performed to retrieve a document item. The search results contain ten documents which satisfy the search criteria. A single document is selected and opened. A note is added to the document.	Documents processed/hour
ADHOC	This workload simulates ad hoc workflow. A workbasket is opened and a folder work package is selected and opened. The folder contains three documents which are displayed. A note is added to the folder. Processing is completed by routing the item to another workbasket.	Folders processed/hour
ADVWF	This workload simulates advanced workflow. A workbasket is opened and a folder work package is selected and opened. The folder contains three documents which are displayed. A note is added to the folder. Processing is completed by continuing the item on the process.	Folders processed/hour

Storage Management Workloads

Member	Description	Capacity Planning Unit of Measure
OPTSTORE	Store document objects residing in the “Root” file system, to LAN-attach optical.	Documents stored/hour
OBJDELETE	Delete expired documents from DASD. The documents reside in the “Root” file system.	Documents deleted/hour

Appendix C. Sizing Worksheet

If you are a business partner or IBM employee that does not have access to BEST/1 for the purpose of modeling a workload, you may provide the information below to the Image Competency Center to perform a sizing. After completing the worksheets, call the ICC Hotline at determine where to send the information.

		General Information Provide general information regarding system use.
1		Input processing hours per day Enter the number of hours a day objects will be loaded into the system through scanning, importing, or inbound facsimile processing.
2		Indexing hours per day Enter the number of hours during the day the customer will be indexing documents in the indexing workbasket.
3		Item processing hours per day Enter the number of hours during the day the system will be available for document and folder searching, opening, routing and other processing.
4		Searches per day Enter the number of daily ad hoc searches for document or folder items.
5		Workbasket items processed per day Enter the number of items processed daily from workbaskets.
6		Advanced workflow Is advanced workflow used?
7		Number of users Enter the number of concurrent users.
8		Business days per year Enter the number of days per year the system will be creating and processing documents.

The following information should be provided for each item or group of items with different characteristics.

		Item Information Provide details for each document and folder type in the system.
1		Item type Specify whether the item is a Document or Folder .
2		Items per day Enter the number of items of this type that will be created each day.
3		Size For document items, enter the size of the document in K bytes.
4		Days on DASD For document items, enter the number of days the item parts will remain on DASD.
5		To be indexed Is the document going to be placed in an indexing workbasket? With a scanning subsystem the document enters the system with indexing information and does not need to be processed from an indexing workbasket.
6		Number of folders added to Enter the number of folders the item will be indexed into.
7		Autofoldering Is the document going to be automatically added to a folder?
8		Workflow Is the item automatically started on a process?
9		Accesses per day Enter the number of times items of this type that are accessed daily?
10		How accessed How are items of this type accessed? 1- Search 2- Workbasket 3- Folder TOC