

Administrator's Guide

 **NOVELL**®

For WINDOWS

USING ENVOY

TABLE OF CONTENTS

INDEX

SoftSolutions®

DOCUMENT MANAGEMENT SOFTWARE V 4.1

SoftSolutions is a registered trademark, and SpeedSearch and SuperSearch are trademarks of SoftSolutions Technology Corporation, a wholly-owned subsidiary of Novell, Inc.
All other brands and product names are trademarks or registered trademarks of their respective holders.

Version 4.1
© 1993-1995 Novell, Inc.
All Rights Reserved

167-000068-001—4/95

Novell, Inc.
1555 North Technology Way • Orem, Utah 84057-2399 U.S.A.
Telephone: (801)225-5000 • Fax: (801)229-1566

Quick Contents

Introduction	1
Setup Issues	5
Setup Procedures	37
Ongoing Procedures	181
Report Writer	245
SoftSolutions Utilities	275
Integrating Applications with SoftSolutions	397
Open Database Connectivity (ODBC)	407
SoftSolutions Error & Confirmation Messages	A-1
Integrations Error & Confirmation Messages	B-1
Integration Files	C-1
SoftSolutions for Windows Diagnostics Program (SSWDIAG)	D-1
Customer Support Services	E-1
Novell International Sales Centers	F-1
Index	

Table of Contents

Introduction	1
About This Manual	1
Documentation Conventions	4
Setup Issues	5
Operating System & Hardware Requirements	5
Operating System Interaction	5
Windows and Memory Requirements	6
Network Support Issues	6
SEMs & Multiple Protocol Support	6
File Services	7
*.MAP Files & OAMs	8
File Services & SEMs	9
Network Support Summary	9
Network Configuration Examples	9
NetWare Transaction Tracking System (TTS)	14
Indexer Workstations	14
Disk Space	14
Dedicated vs. Nondedicated Indexers	15
Using Multiple Dedicated Indexers	16
Optimizing a Dedicated Indexer's Performance	16
Configuring the File Server to Optimize an Indexer	17
Importance of a Balanced Configuration	18
Using the File Server as the Indexer	18
Additional Key Indexing Concepts	19
Organization of SoftSolutions Files	20
Archived Documents	20
Datasets	22
Documents	23
Drive Letters vs. "Logical" Paths	24
Mapping Utilities	26
Universal Filing System	31
Echoed Documents	32
Programs	33
Temporary Files	34
SoftSolutions Security System	35
Document Security	35
Screen Security	36
Screen & Document Security Interaction	36

Setup Procedures	37
Performing Operations in SoftSolutions Setup Windows	39
Applications	45
Control File	76
Custom 1-16	92
Custom Prompt Definitions	96
Dataset Groups	100
Datasets	103
Document Security Groups	107
Document Types	112
Enterprise Administrator (including Global Replicator)	118
EZLaunch Options	129
Initialize Portable Mode	131
Object Access Modules	135
Remote Document Server (RDS)	138
Screen Access Groups	149
Stop Words List	157
Strings Compiler Utility	159
Users	165
Virtual User (including Task Manager)	172
Workstation	177
Ongoing Procedures	181
Archival & Deletion of Documents	183
Archiving	184
Deleting	184
Archived Document Restoration	186
Deleting a Restoration Request	188
Restoring a Document	188
Searching for a Restoration Request	189
Backup of the SoftSolutions Database	191
Billing	193
Custom Fields Maintenance	194
Index Compaction	201
Index Verification & Reconstruction	204
Index Verification	204
Index Reconstruction	209
Types of Reconstruction	210
Determining Which Reconstruction To Use	212
Performing Reconstruction	214
Multiple Workstation Reconstruction	219

<i>Ongoing Procedures continued</i>	
Indexing & Indexing Errors	221
Indexer	223
Indexing Errors	226
Indexer Requests	227
Changing an Indexing Request	229
Deleting an Indexing Request	230
Searching for an Indexing Request	230
Mass Import of Documents Created Outside of SoftSolutions	232
Defining a Mass Import Configuration	233
Executing Mass Import	240
Temporary File Cleanup	243
Report Writer	245
Report Writer Window	245
Adding a Report Writer Definition	246
Changing a Report Writer Definition	252
Copying a Report Writer Definition	253
Deleting a Report Writer Definition	253
Executing a Report	254
Searching for a Report Writer Definition	255
Report Writer Examples	257
Archive Override Report By Project	258
Archive Report By Author	260
Billing Report	262
Document Listing By Author	264
Document Report By Project	266
Document Report By Type	268
System Usage Report By User	270
Report Writer Files	272
SoftSolutions Utilities	275
Terms Used for the Utilities	275
Utilities Shipped	276
Utilities by Function	277
Compatibility	278
Using the Utilities in a Windows Environment	278
Command Line Terms	280
Arguments	280
Document Specification Lists	281
Document Specification List Examples	283

<i>SoftSolutions Utilities (Command Line Terms) continued</i>	
Switches	283
Common Switches	285
Response Files	286
Utilities Listing	288
Utility Documentation Headings	289
Paths in Utility Examples	289
AddPro	291
ChgCre	297
ChgNum	299
ChgPro	301
Copy	306
CrePro	309
DelPro	316
FixCst	319
FixLF	323
FSize	327
GetPro	328
ImpDir	331
ImpPro	341
ImUser	353
IndTxt	356
LdUser	359
Move	362
NewVer	365
NoDoc	367
NoPro	369
OamCpy	371
PrnPro	375
ProExt	380
Prof	383
ResDoc	387
Wait	390
WinBil	394
Integrating Applications with SoftSolutions	397
SoftSolutions Integrations	397
Binary Compatibility of the SoftSolutions Datasets	399
Diagnosing Integration Problems	400
Customizing with the SoftSolutions Integration Manager	401

Open Database Connectivity (ODBC)	407
ODBC Overview	407
ODBC Terminology	408
Maintenance Options	409
Install a New SoftSolutions Data Source	409
Change a SoftSolutions Dataset's Information	410
Microsoft Query Example	411
SoftSolutions Database Tables for ODBC	415
Modifying the Tables	415
Date and Timestamp Literals	416
ACCESS	416
ACTIVITY	417
APPLICAT	417
BILLING	418
CONTROL	419
CUSTOM	423
DATASETS	423
DOCTYPES	424
DSGLOBAL	425
DSGROUPS	425
JOBS	425
PROFILES	426
REPORTS	429
RESTORE	431
SCREENS	431
SECURITY	431
SECUSER	432
USERS	432
VERSIONS	435
WORKSTATION	435
SoftSolutions Error & Confirmation Messages	A-1
Message Strings	A-2
Troubleshooting Disk I/O Errors	A-29
Disk I/O	A-29
Corruption	A-29
Integrations Error & Confirmation Messages	B-1

Integration Files	C-1
SoftSolutions General Interface DLLs	C-1
Application Integration Files	C-3
Ami Pro for Windows 3.0	C-3
Lotus 1-2-3 for Windows 1.1, Release 4	C-4
Lotus Notes 3.0	C-4
Microsoft Excel for Windows 4.0, 5.0	C-5
Microsoft Word for Windows 2.0, 6.0	C-7
WordPerfect for Windows 5.1, 5.2, 6.0	C-11
SoftSolutions for Windows Diagnostics Program (SSWDIAG)	D-1
Running SSWDIAG.EXE	D-1
SoftSolutions Diagnostics Program Menu Options	D-2
Exiting the Diagnostics Program	D-2
DLLs	D-3
DLL Status	D-5
Viewing DLLs	D-6
Reporting on DLLs	D-7
Directory Searching Order	D-7
Viewing Searching Orders	D-8
Reporting Searching Orders	D-8
Reporting on a Combination of Diagnostics	D-9
Customer Support Services	E-1
Novell International Sales Centers	F-1
Index	

Introduction

Welcome to SoftSolutions® 4.1 for Windows, the advanced document management software for local and wide area networks (LANs and WANs) from Novell, Inc.

About This Manual

The SoftSolutions *Administrator's Guide* is divided into eight sections, six appendices and an index:

1. Introduction
 2. Setup Issues
 3. Setup Procedures
 4. Ongoing Procedures
 5. Report Writer
 6. SoftSolutions Utilities
 7. Integrating Applications with SoftSolutions
 8. ODBC Administrator
- A. SoftSolutions Error & Confirmation Messages
 - B. Integration Error & Confirmation Messages
 - C. Integration Files
 - D. SoftSolutions for Windows Diagnostics Program (SSWDIAG.EXE)
 - E. Customer Support Services
 - F. Novell International Sales Centers

The *Administrator's Guide* is written for the system administrator (the person in charge of setting up and maintaining SoftSolutions), and the SoftSolutions *User's Guide* is written for the everyday user (a person performing daily operations in SoftSolutions).

NOTE

If the software was purchased as a stand-alone product, the licensee should assume both the system administrator and user roles.

The *User's Guide* documents windows and procedures that apply to everyday users. Some of these windows can also be used during setup in customizing the appearance and contents of SoftSolutions windows (e.g., Screen Definitions, Document Desktop, Fonts, Colors, SpeedButtons, Workstation). Therefore, please use both manuals as needed in setting up SoftSolutions.

Although the *Administrator's Guide* includes brief information on using SoftSolutions in Windows (adding, deleting, changing, finding, etc.), the *User's Guide* also includes information on using SoftSolutions in Microsoft Windows and the fundamental operations in SoftSolutions.

Setup Issues

Setup Issues explains such topics as:

- Configuring the Indexer workstation
- Developing a logical directory structure
- Recognizing the types of files accessed by SoftSolutions
- Formulating the system security
- ▶ **These items should be addressed before performing the actual setup of SoftSolutions.**

Setup Procedures

Setup Procedures contains information on performing the actual setup, including the purpose for and use of each of the Setup windows.

Ongoing Procedures

Various files and indices must be maintained by the system administrator on an ongoing basis. *Ongoing Procedures* documents maintenance issues and procedures, including the utilities designed to assist in the maintenance of SoftSolutions.

Report Writer

The SoftSolutions Report Writer facilitates the creation of reports. Information is provided to aid in using the Report Writer to generate reports on data contained in SoftSolutions. Also discussed are archiving and deleting of “mature” documents using the Report Writer.

SoftSolutions Utilities

The utilities enhance operations with the SoftSolutions database.

Integrating Applications with SoftSolutions

Many applications can be integrated with SoftSolutions, providing document management of files created and edited by them. This section explains integrations with SoftSolutions, the binary compatibility of SoftSolutions datasets, diagnosis of integration problems, and customizing the integrations.

Open Database Connectivity (ODBC)

SoftSolutions' compliance with Microsoft's ODBC specification allows administrators to create user applications for querying the SoftSolutions database.

Appendix A

SoftSolutions error and confirmation messages are provided, with explanations.

Appendix B

Error and confirmation messages for integrations with SoftSolutions are listed here.

Appendix C

The various files and DLLs used in integrating applications with SoftSolutions are listed in this appendix.

Appendix D

This appendix documents the SoftSolutions 4.1 for Windows Diagnostics Program (SSWDIAG.EXE).

Appendix E

The appendix provides Novell GroupWare customer support information.

Appendix F

This appendix lists Novell's international Sales Centers, including telephone and fax numbers.

Documentation Conventions

A few documentation conventions are used to improve readability and enhance understanding.

bolded text	Field names (<i>e.g.</i> , Name field, Copy out document field, Document Security field, Document Number field).
ALL CAPS	Window and dialog box names (<i>e.g.</i> , DOCUMENT PROFILE window, CHECK OUT dialog box); filenames and keyboard entries (<i>e.g.</i> , WORD %P%F).
{Key Name}	Key references are enclosed within <i>Lenticular Brackets</i> , usually as they appear on an IBM® PC keyboard. For example, the “a” key is represented as {A}, the “escape” key as {Esc}, and the “plus” key as {+}.
bold and <u>underline</u>	SoftSolutions <i>options</i> selected from the menu bar in Windows. The bolded and underlined letter represents the mnemonic letter used to choose an option using the keyboard. When using mnemonic letters to perform a step, the step is always preceded by pressing the {Alt} key. (<i>e.g.</i> , {Alt}{X} chooses the Search Full- <u>Text</u> option).
“monospaced text”	SoftSolutions error and confirmation messages (usually within quotes), and example texts.
<i>[Italicized in Brackets]</i>	“Place holders,” which are words that identify information to be supplied by the user (<i>e.g.</i> , C:[<i>Filename</i>]).
{Items in Braces}	Options or switches that are optional for the command (<i>e.g.</i> , SSWAIT [<i>Minutes</i>]{ <i>Seconds</i> } {/STAT}, where the /STAT switch and <i>Minutes</i> are optional, but the <i>Seconds</i> must be given.)

Setup Issues

Although SoftSolutions can be used immediately upon installation, additional setup may be required to take full advantage of SoftSolutions' features. The setup issues which should be understood and addressed before setting up SoftSolutions are documented in *Setup Issues*. These issues include:

- Operating system interaction *p. 5*
- Memory requirements *p. 6*
- Network support issues *p. 6*
- NetWare Transaction Tracking System (TTS) *p. 14*
- Configuring the Indexer workstation *p. 14*
- Recognizing the types of files accessed by SoftSolutions *p. 20*
- Developing a logical directory structure *p. 23*
- Understanding SoftSolutions security at the document and screen levels *p. 35*

Once the issues have been read, continue with the actual setup of SoftSolutions. This is documented in *Setup Procedures*.

Operating System & Hardware Requirements

Operating System Interaction

SoftSolutions has been designed to be network *independent*. All communications between SoftSolutions and the network file servers are handled through the networking features built into recent DOS operating systems, such as MS-DOS versions 3.1 and later.

SoftSolutions does *not* depend on any particular network protocol or topology. The only requirement is that the network must support DOS file sharing and locking. Most popular networks support this type of access. Your network vendor should be contacted if there is any doubt.

SoftSolutions can also operate on a “stand-alone” PC, with either the full licensed software loaded on the PC, or with a subset of the software (using SoftSolutions' Portable Mode) loaded from the network.

SoftSolutions supports Ethernet, Arcnet, StarLAN, Token Ring, etc., as long as the network operating system can operate on the topology. Obviously, the communications bandwidth for WANs influences the overall speed of SoftSolutions. Factors such as bridges, line transmission speed, routers, etc., can affect overall performance.

Windows and Memory Requirements ---

SoftSolutions for Windows is designed to run on any machine set up to run Microsoft Windows 3.1 or higher in standard or enhanced mode. SoftSolutions will run on a 386SX PC with 4.0MB of RAM. But, Novell recommends at minimum a 486 PC with 8MB of RAM. Additional RAM may be required, depending on the number of applications opened simultaneously in Windows.

Network Support Issues ---

SoftSolutions can be implemented under a true client/server architecture. The *server* side of SoftSolutions is available under NetWare, as well as SunOS, NCR SVR4, LAN WorkPlace, IBM AIX, SCO Unix, and Hewlett-Packard HP/UX. (*Contact Novell, Inc. for current availability.*)

SEMs & Multiple Protocol Support

NOTE

SoftSolutions for Windows will run in client/server mode on any TCP/IP product that supports the Window's Winsoc standard. This includes such products as the newest versions of BW/NFS and PC-NFS.

The *server* software of SoftSolutions is called a "Server Enhancement Module" (SEM). The SoftSolutions SEM runs as a NetWare 4.x NLM. SoftSolutions SEMs are network-specific. It is *not* possible to move a SoftSolutions SEM NLM and place it under an IBM Unix server. Each SEM has its own model number for ordering. SoftSolutions' SEMs are *not* interchangeable.

Being a client/server model, SoftSolutions also has a *client* side. The client software is the portion of SoftSolutions that runs on the workstations. Under DOS (including Windows), the SoftSolutions client software is fully interchangeable. This is true because support for NetWare and TCP/IP has been placed in the SoftSolutions DOS client. This feature is called "Multiple Protocol Support."

A strong benefit from Multiple Protocol Support is the capability of running SoftSolutions under a *heterogeneous network*. If a corporation has NetWare and Unix servers, for example, workstations running SoftSolutions can search both the NetWare dataset and the Unix dataset—with the same ease (from the user's point of view) as a *homogeneous network*. From the system manager's point of view, the heterogeneous network must be set up properly, which may require the workstation to have “stacked” protocols.

SEMs are network-specific because of the need to implement server-side software under different network operating systems. SoftSolutions DOS Client software is *not* network-specific because of the availability of the Multiple Protocol Support, which gives users increased searching flexibility on networks with dissimilar servers.

File Services

The SoftSolutions client-side software supports the protocol, *File Services*. This is intended to provide workstations with the ability to gain access to a dataset located on a server for which a SEM has not been specifically developed. With this feature, SoftSolutions client software can have access to a dataset on a server from either Banyan Vines, LAN Manager, LANtastic, NetWare Lite, OS/2 LAN Server, and most other implementations of PC networks.

Obviously, because a SEM with file services does not exist on the server side, the implementation is not as elegant as with NetWare or Unix. However, because of file services support, SoftSolutions is still very usable and powerful under almost all PC networks.

The previous paragraphs describe SoftSolutions running under a single dataset serviced by Banyan Vines, LAN Manager, LANtastic, NetWare Lite, OS/2 LAN Server, etc. What happens if a configuration has *multiple* datasets under these network operating systems? SoftSolutions without SEMs can, under certain circumstances, have access to multiple datasets located on different servers—*assuming that the workstations are attached to the desired multiple dataset servers on a permanent basis with uniquely-mapped drive letters*.

The networking issues to be addressed, and the complexity of the solutions, depend on the specific network and the local implementation methodology employed. Consider a simple scenario:

You have a network with a few workstations and two servers (two datasets) without SEMs. Each server has one single disk drive. Assume that the workstations have permanent drives mapped to both servers. Under this scenario, SoftSolutions *without* SEMs can search both servers, with the same look-and-feel as SoftSolutions *with* SEMs. The workstations will use file services to search each dataset.

***.MAP Files & OAMs**

Two SoftSolutions features exist for helping users under certain configurations:

- The *first* of these features is the *.MAP file, which allows the document path string on the DOCUMENT PROFILE window to be replaced with the correct drive letters. This assumes that the Profile's location points to physical servers and drives, or some other unique identifier.
- The *second* feature concerns the OAMs, which allow programmers to code customized algorithms for mapping drives dynamically, and then detaching after they are no longer needed. This feature is dependant on an individual network's capabilities.

In a complex configuration, it may not be possible to be permanently attached to all servers with unique drive letters mapped to all documents. The OAMs are the main tools for resolving this limitation. Even though SoftSolutions does not have OAMs for certain network implementations, SoftSolutions has made the architecture open enough to allow third party developers and end users to make available OAMs that compensate for the lack of availability of SEMs. As this availability grows, global searches for networks that are not supported via SEMs will be possible through file services.

You should use caution in assuming that larger Banyan Vines, LAN Manager, and OS/2 LAN Server networks can be implemented immediately in a global environment. Some work will need to be done, depending on the implementation adopted. Drive mappings, for instance, may need to be modified to more fully take advantage of this capability. SoftSolutions has many refined tools to aid in the management of the many types of networks that exist.

When SEMs are available for any of these networks, the above items will not be a concern.

File Services & SEMs

File services can be mixed with SEMs. SoftSolutions can exist in a network with multiple servers, with some servers running SEMs, and with other servers running file services. A workstation connected to a mixed network can perform a global search and gain access to datasets via the SoftSolutions SEM NLM, SoftSolutions SEM for Unix, and file services—all because of the Multiple Protocol Support.

Network Support Summary

All of the above examples assume that issues are resolved relating to network connectivity, memory availability, Network Interface Card driver support, and the workstation's configuration (these are network integration issues which are *not* related to SoftSolutions).

Since the networking possibilities of SoftSolutions are so flexible, SoftSolutions invites SoftSolutions users and integrators to carefully experiment with different possibilities before committing to complex configurations. To discuss specific implementations, please call Novell GroupWare Customer Support Services (*see Appendix E*).

Network Configuration Examples ---

The following examples illustrate sample hardware configurations that can be set up to make use of certain SoftSolutions features: Multiple Protocol Support, Client/Server support using Server Enhancement Modules (SEMs), and non-Client/Server support (with global searching capabilities). (*Contact Novell, Inc. for current availability of SEMs.*)

These configuration illustrations consider the location of the SoftSolutions programs, data, and documents, as well as where the Indexer is placed.

Some of these illustrations depict mixed environments to highlight SoftSolutions' flexibility and to demonstrate what is important when configuring a system for SoftSolutions.

The following terms are used in the illustrations:

File Services

A direct communications link to information via a drive mapping. File services allows data to be accessed and transferred from a server to a workstation.

Client/Server

Two software components that permit splitting the execution of tasks between the server's CPU and the workstation's CPU. With SoftSolutions, this is accomplished using SEMs, such as NetWare NLMs or Unix daemon processes. Because the program tasks can be split up, the client/server approach generally yields better overall performance and provides some extra capabilities.

Server Enhancement Module (SEM)

A SEM is a separate software module for SoftSolutions that allows SoftSolutions to have access to data using the client/server technology. The SEMs are loaded on the server and are executed and run from the server with communication to and from workstations. Currently, a SEM is now available under NetWare 4.x. SEMs are available for SunOS, NCR SVR4, LAN WorkPlace, IBM AIX, SCO Unix, and Hewlett-Packard HP/UX. SEMs allow the network to split up the work between the Server and the Client, and therefore will generally give better performance and a more seamless method of searching across datasets.

Multiple Protocol Support

This is SoftSolutions' capability of accessing datasets via IPX/SPX, TCP/IP, or file services without the need for a different set of programs. As the location of each other dataset to be accessed is set up in SoftSolutions, the protocol used for accessing the data is specified. Theoretically, this means sequential searches can be made through several datasets that are accessed via different protocols, as long as the network can be properly configured to do so.

Data

A SoftSolutions dataset that consists of all Profile and setup records, including the indices of the databases and document full texts. When using a SEM for accessing the data, file services are not needed. But, if no SEM is used, file services are required for accessing the data. Datasets using a SEM can be mixed with datasets not using a SEM.

Documents

For the purpose of these configuration illustrations, *any files* that are managed by SoftSolutions are considered "documents." Documents can reside on the same file server as the data, on a separate server, or both. Documents managed by SoftSolutions require access via some method of file services.

Indexer

A dedicated workstation or server process that spends its time indexing documents that are managed by SoftSolutions. An indexer must have access to both data and documents in order to properly complete its function.

When preparing a SoftSolutions configuration, it is important to know the answers to the following questions:

1. On which servers will the users reside?
2. Where will the documents be located? Are file services available from each workstation to the location of the documents?
3. Where will the datasets be located? Will a SEM or file services be used for accessing the data? Which datasets will each user require access to for searching or for accessing the documents via retrieving, previewing, etc.?
4. Where will the Indexer reside? Will the Indexer be run from a workstation or as a process on the server?
5. Will terminals be used along with PCs in the configuration?
6. What communications software is currently available and what purchases may be necessary to allow file services where required?
7. If a workstation is to have access to multiple platforms, will the necessary software co-exist on a workstation (*e.g.*, BW/NFS & NetWare, NetWare & Banyan Vines, LAN WorkPlace, and LAN Manager)? Will there be enough memory to successfully run the planned configuration?

As is evident, many of the issues under consideration cannot be properly answered by SoftSolutions alone. Some issues may need to be referred to the network manufacturer or a systems integrator. The illustrations on the following pages should help to provide the questions that must be answered in order to configure a successful network for SoftSolutions usage.

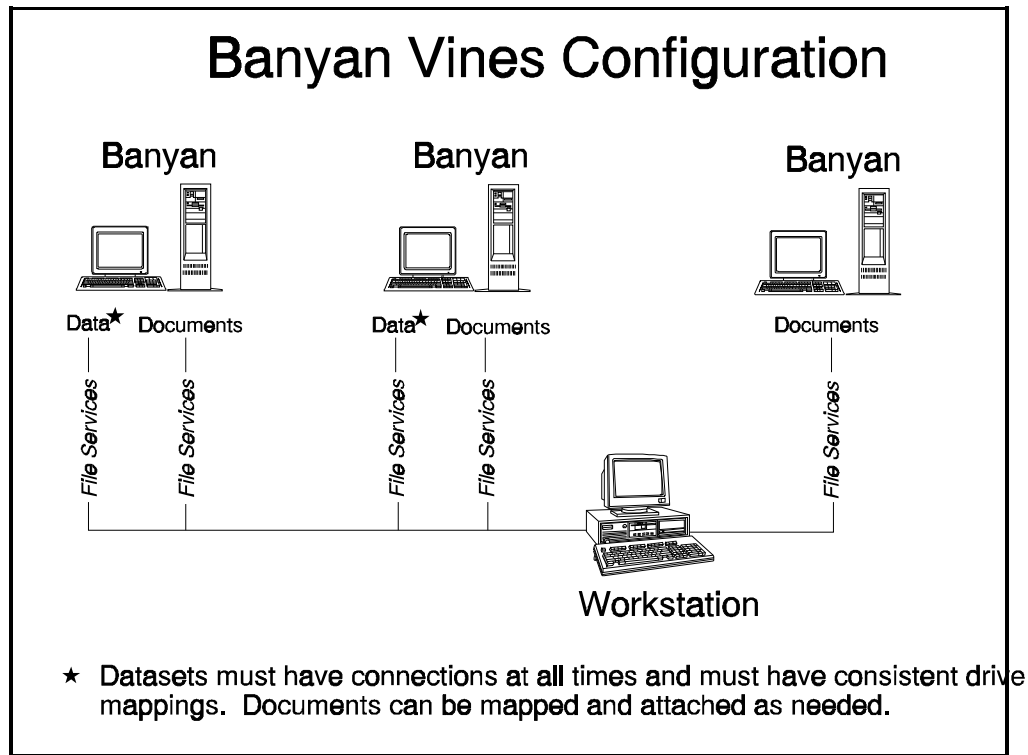


Figure 1

This illustrates a small Banyan Vines configuration that is set up for accessing all data files for searching without the use of a SEM. In order for the workstation to have access to all data files and documents, a file services connection must be made from the workstation to each server. In this case, a drive must be mapped to each dataset at all times because dynamic mapping cannot occur between datasets without the use of a SEM. The documents may either be mapped to different drive letters with a constant attachment or, a different set of Object Access Modules (OAMs) can be created to make the dynamic attachments for accessing documents. Using this configuration may not be as useful for large networks where SEMs are not available, because many drive letters will be used to make permanent attachments from each workstation to the servers where the data resides.

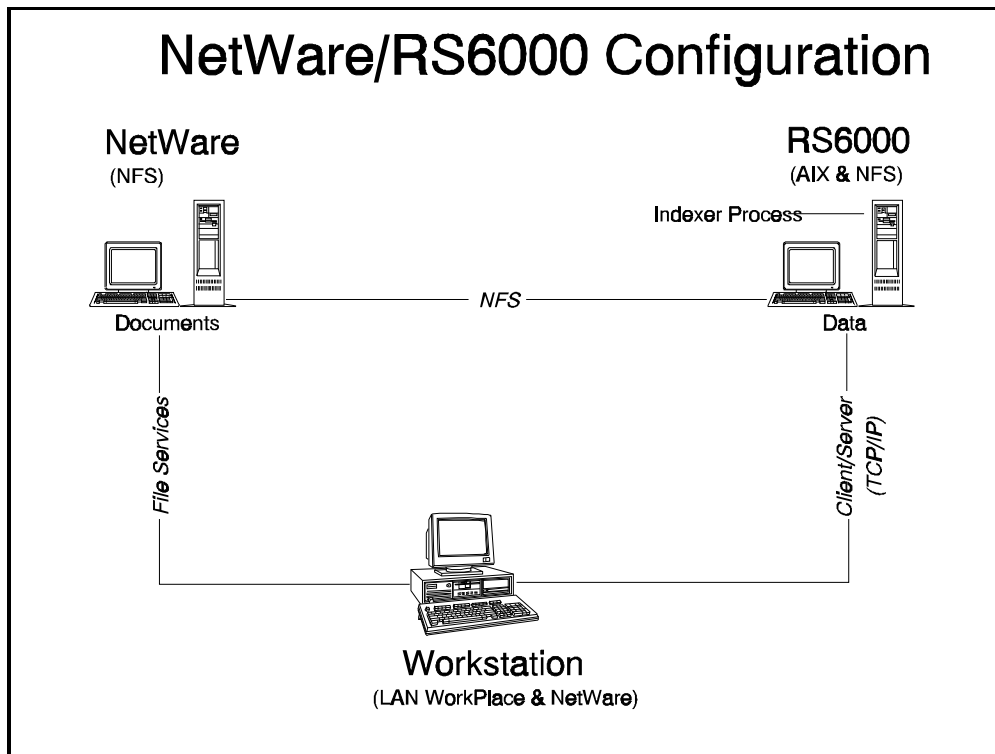


Figure 2

This configuration has the data placed on a RISC-based system running Unix. The data is to be accessed via a client/server connection using the SoftSolutions SEM for Unix. Since the RS6000 and NetWare systems both store documents, the workstation must have access to each system using file services. This is why LAN WorkPlace, and the NetWare shell are loaded on the workstation. The Indexer is running as a daemon process on the Unix machine to take advantage of the fast processing capabilities of the RISC system. To give the Indexer access to the documents on NetWare, an NFS (Network File Services) connection must be made between the RS6000 and the NetWare server.

NetWare Transaction Tracking System (TTS) _____

When using NetWare 3.11 (or later) and the SoftSolutions SEM, NetWare's Transaction Tracking System (TTS) can be used. TTS helps to preserve the integrity of the SoftSolutions data in the event of a file server crash. Novell, Inc. recommends that TTS be used. (*See the SoftSolutions Server Enhancement Module for NetWare documentation for additional information on TTS.*)

Indexer Workstations _____

The “Indexer” is a workstation connected to the system running SoftSolutions. The Indexer updates the SpeedSearch™ index files located in the SoftSolutions data directory with “words” from the documents' texts.

If SoftSolutions is fully implemented on a single-user PC (no network), the PC can also perform any required indexing.

Following is a discussion of disk space requirements for indexing and issues to be considered when determining your indexing configuration.

Disk Space

Temporary files are created when the Indexer is being run. Temporary files created by the Indexer can be as large as the largest document to be indexed. Therefore, the Indexer workstation must have at least as much hard disk (or RAM disk) space as the size of the largest document to be indexed. The location of these temporary files is defined in the WORKSTATION window for the Workstation ID used in starting SoftSolutions on the Indexer workstation.

Using the Indexer's fast hard disk or RAM disk as the temporary file location can increase the speed of the indexing process. The Indexer might fail when being run on a workstation with insufficient CPU RAM and/or hard disk space. For example, if a document contains too many unique words for the amount of memory available, an “Insufficient memory” error is reported.

Dedicated vs. Nondedicated Indexers

The Indexer can run continuously on one or more workstations dedicated only to indexing, *or* it can be run at night on one or more user's workstations. Optimum usage can be determined by indexing activity.

Dedicated Indexer

Technically, SoftSolutions does not require an additional workstation to be dedicated for indexing. But, a dedicated Indexer is the *most efficient* indexing configuration.

The advantages of the dedicated Indexer configuration are:

- 1) **Data Searching Availability:** Newly created or updated *documents* are indexed automatically and transparently, requiring no intervention by a user or system manager.
- 2) **Security:** The Indexer workstation and file server can be locked in a room where general users cannot gain access to them.

Nondedicated Indexer

A nondedicated Indexer configuration is useful for smaller networks. A particular user's workstation might be set up to index at night, during lunch, or during other periods of nonuse.

The advantages/disadvantages of the nondedicated Indexer configuration are:

- 1) **Advantage—Lower Equipment Costs:** Production workstations can do the indexing without the additional cost of a dedicated Indexer workstation.
- 2) **Disadvantage—Limited Searching Data:** The *disadvantage* is that until the Indexer is run on *new documents*, their full texts are not available to users for searching.

Very large network installations may require more than one dedicated Indexer workstation, or nondedicated Indexers can be used at night in addition to just one dedicated Indexer.

Using Multiple Dedicated Indexers

SoftSolutions performs DOS-level “record locking” to enforce database integrity. Under SpeedSearch, a single *unique* word is a *whole record*. When multiple Indexers are running, the system performs record locking to prevent different Indexers from manipulating the *same word* simultaneously. Indexing throughput can be increased by adding more dedicated Indexers, but this number should *not* exceed four or five. When the number of dedicated Indexers exceeds roughly four or five, the overhead for record locking increases significantly and performance is not increased. It is better to increase the *speed* of one, or a few, dedicated Indexers than to increase their number.

Optimizing a Dedicated Indexer's Performance

Each of the following helps to optimize a dedicated Indexer's performance:

- The faster the Indexer workstation, the better the performance.
- The Indexer should be at least as powerful (processor speed) as the file server.
- It is better to increase the speed of one Indexer CPU than to add a second or third Indexer.
- There should be one fast 486 or Pentium PC Indexer for every cluster of 100 to 125 workstations.
- Multiple dedicated Indexers are possible, but more than four or five are *not* recommended.
- The Indexer and the file server should have 16-bit or 32-bit network cards, if possible.
- The Indexer should *not* be a 286 or 386SX PC, unless the number of networked workstations are few.
- For large networks, the file server can possibly take advantage of two 16-bit or 32-bit network controller cards—one network card for all the regular workstations used in the network, and a separate network card for the dedicated Indexers. This avoids network collisions.
- Ensure that the file server has plenty of memory. If the Indexer's disk access to the file server is “cached” to RAM, performance for indexing increases significantly.
- The DOS-based Indexer is faster than the Windows-based Indexer. Therefore, SoftSolutions for Windows includes a DOS-based Indexer program. Both SoftSolutions 4.1 for DOS and SoftSolutions 4.1 for Windows use the same database, so the DOS-based Indexer should be used.

For dedicated Indexer workstations, a monochrome monitor is sufficient and cost-effective.

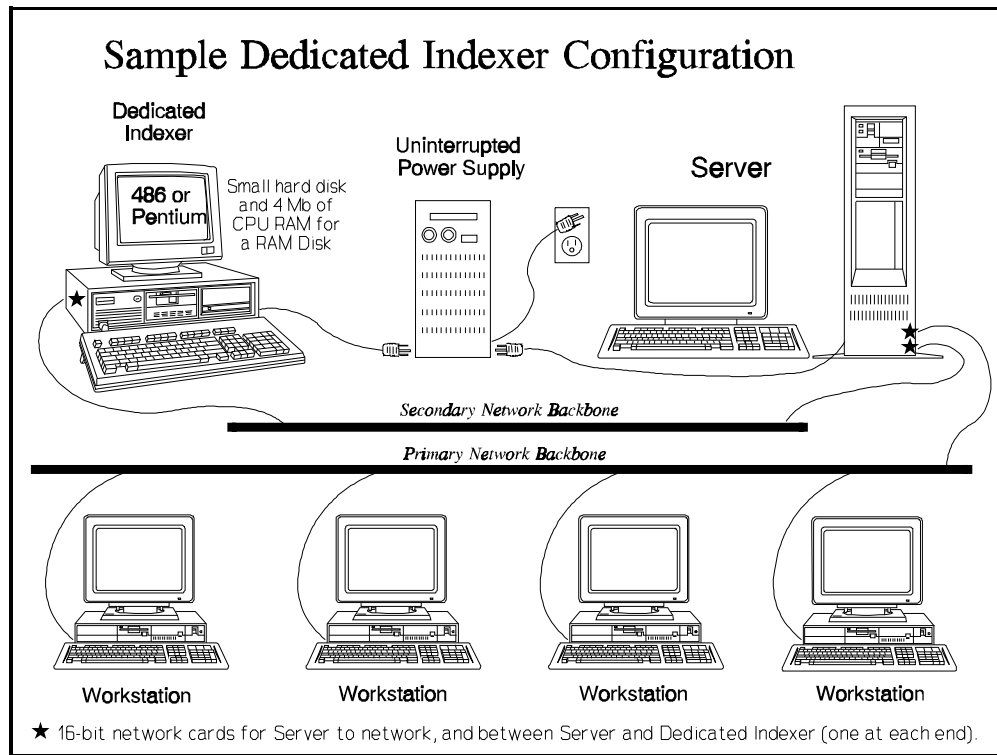


Figure 3

Configuring the File Server to Optimize an Indexer

When the Indexer is running, it accesses the dataset stored in the file server as quickly as the hardware can deliver it. But, the Indexer's performance degrades if the file server slows down indexing accesses because the file server's CPU is constantly accessing physical disk blocks. Following are two ways to optimize the Indexer:

- 1) **Add memory.** The greater the file server's memory, the more the network operating system allocates RAM for disk caching. Disk caching on the *file server* is imperative. In systems such as NetWare, the operating system allocates all available memory (not being used by regular programs) for caching. Although network performance monitoring utilities may indicate that sufficient memory is available for the regular programs, additional RAM might improve the Indexer. Memory is an inexpensive way to significantly boost the performance.

- 2) **Add a second network controller card.** If the file server has two network cards, one can be assigned to the normal network traffic for disk and print services, and the second can be dedicated to the Indexer traffic. This strategy decreases the potential of network collisions for normal users. A second network controller card is even more important when multiple dedicated Indexers are used. To improve the bandwidth between the Indexer and file server, configure a 16-bit or a 32-bit controller card (8-bit network cards are adequate for users' workstations).

Importance of a Balanced Configuration

In any aspect of systems optimization, the key is a “balanced” design. The system always performs at the level of its *weakest* component. If the Indexer is too slow, replacing it with a Pentium machine does *not* improve performance if there is also a *memory* bottleneck at the file server. (In fact, the bottleneck is worse!) Increasing the performance of a single critical network resource raises the overall throughput to its next resource bottleneck.

The most significant planning issue to balance the network under SoftSolutions is the Indexer configuration for larger networks. Everything else runs very efficiently with little network change. If the guidelines and recommendations outlined herein are followed, the network will be optimized. These suggestions apply equally for the initial network design as for improving an existing configuration.

You must evaluate each component in light of the overall configuration to obtain balance—namely: the file server (CPU, RAM for caching, number and speed of disks, controller cards); workstation nodes; dedicated Indexers; network topology; rate of document production; document mix; network operating system; wide area bridging; systems management; expertise level of the systems integrator and in-house support; training and expectation level of users; etc. Even two seemingly unassociated elements, such as network security setup and search performance, are interrelated.

Using the File Server as the Indexer

Implementing a dedicated Indexer *does* render significant advantages at a relatively low cost. The file server already has the significant task of managing the network. The burden of additional work can degrade its response time. Additionally, in order for the file server to perform the indexing task, a different program for each network implementation is required.

Today, the price of computers is very reasonable. Future technology will bring the prices even lower and provide a better and faster computer for the same cost.

Additional Key Indexing Concepts

Dynamic Documents

SoftSolutions is unique in its position as a full-text Indexer because its algorithms must manage “dynamic” documents. Most other full-text systems manage primarily “static” documents—documents which are no longer being changed, such as depositions, manuals, signed contracts, and litigation material.

The challenges of managing dynamic documents are much greater. SoftSolutions does not know in advance the number of documents to be indexed and searched, the relative sizes of documents, or how often changes will be made. The cumulative dynamic changes being made by all workstations to all files in a network can be overwhelming.

Transparent Indexing

All indexing activities must occur “transparently”—meaning that a user need not formally request that a document should “now” be indexed. This should happen automatically as the document is completed. Full-text indexing activities should not prevent a user's workstation from being used for its next activity. Indexing should occur in the “background,” leaving a user's workstation (and the file server) undisturbed.

The chosen network architecture to implement the management of “dynamic” documents in full text with the simultaneous requirement of “transparency” and “background processing” is through dedicated or background Indexers.

Indexer Security

Because the dedicated Indexer is such a part of the overall network, Novell, Inc. recommends that it be placed in a secure location where unauthorized personnel cannot have access to it. In fact, it should be treated similar to the file server in terms of its importance to the network.

The dedicated Indexer can be the hardest working workstation in the network. Most network managers take great care in planning the file server configuration—similar care should be taken for the dedicated Indexers. The dedicated Indexers and the file server should be connected to an uninterruptible power supply (UPS).

Organization of SoftSolutions Files

In addition to the SoftSolutions program files, SoftSolutions accesses many different types of files, including data files, documents (word processing documents, spreadsheets, graphs), echoed documents, archived documents, and temporary files.

To give maximum flexibility and control over the network, the system administrator must determine *where* each of these types of files should be stored. A description of each of these types of files is provided to aid in determining the best location for the various files. Information on the SoftSolutions program files and the directories in which these files are placed when initially installing SoftSolutions is found in the appendix of the *SoftSolutions Installation Guide*.

Archived Documents

Disk Space

The archive location may contain one or several archive “sets” or subdirectories. In the CONTROL FILE window, the system administrator may stipulate the amount of space to be set aside for each “set.” The maximum size for a set should equate closely to the capacity of the archival medium. For example, if 1.2MB floppies are being used, an archive set should be limited to about 1MB in size.

The size of an archive set should be such that *at least two* sets will fit in the archive location defined in the APPLICATIONS window. This is to ensure that when one set becomes full, there is enough space remaining so that another set can be created automatically and used.

For large-capacity archival mediums, ensure that the size of *one* archive set *divides evenly* into the size of the archival *medium*. This prevents wasted space. For example, if the archival medium's capacity is 100MB, then using 35MB for the size of an archive set allows only two archive sets to fit on the medium, wasting 30MB.

Because some types of archival media use some of their available space to record file storage data, make the *cumulative* archive set size 10% smaller than the space available on the archival medium. For example, a 100MB archival medium size can accommodate three 30MB archive sets with 10MB (the 10%) left for file storage data.

Considerations

Documents, according to their document types, can be archived if they are not accessed within a specified number of days. The archive location is specified in the archive Object Access Module in the APPLICATIONS window. The default is \ARCHIVE. The archive location may also be to an optical disk or a tape drive.

An “Archive Set” is a subdirectory with a name that follows a pattern established in the CONTROL FILE window (such as ARCHIVE\ARC000001). SoftSolutions uses this pattern to automatically create subdirectories under the archival directory for each subsequent archive set.

An archiving report is used to move documents into the current archive set until the accumulated size of the documents approaches the set's maximum size as defined in the CONTROL FILE. Then it creates another set and begins archiving documents to that set. For example, when ARCHIVE\ARC000001 becomes full, SoftSolutions automatically creates “ARC000002” and begins archiving to this subdirectory. Thus, the main archival directory must be located on a disk with at least as much free space as can be used by *two* archive sets.

Datasets

A SoftSolutions “Dataset” is a collection of SoftSolutions Profiles, Setup records, and SpeedSearch indices. This information is contained in multiple data files (*.DAT, *.X01, *.X02, *.X03, and *.KDT) and is directly maintained by SoftSolutions.

Files

ACCESS.DAT	DATASETS.DAT	PROFILES.DAT	SECUSER.DAT
ACCESS.X01	DATASETS.X01	PROFILES.X01	SECUSER.X01
ACTIVITY.DAT	DATASETS.X02	PROFILES.X02	SECUSER.X02
ACTIVITY.X01	DOCTYPES.DAT	PS.KDT	SORT.LST
APPLICAT.DAT	DOCTYPES.X01	PWL.KDT	STOP.LST
APPLICAT.X01	DOCTYPES.X02	REPORTS.DAT	STORE.DAT
APPLICAT.X02	DSGLOBAL.DAT	REPORTS.X01	STORE.X01
BILLING.DAT	DSGLOBAL.X01	REQUESTS.DAT	USERS.DAT
BILLING.X01	DSGLOBAL.X02	REQUESTS.X01	USERS.X01
COMMENT.DAT	DSGROUPS.DAT	RESTORE.DAT	USERS.X02
CONTROL.DAT	DSGROUPS.X01	RESTORE.X01	VERSIONS.DAT
CUSTOM.DAT	ICHAR.LST	SCREEN.DAT	VERSIONS.X01
CUSTOM.X01	JOBS.DAT	SCREEN.X01	WORDS.DAT
CUSTOM.X02	JOBS.X01	SECURITY.DAT	WORDS.X01
CUSTOM.X03		SECURITY.X01	WORKSTAT.DAT
		SECURITY.X02	WORKSTAT.X01

Disk Space

The total size needed for the data directory and its one subdirectory (STATUS) depends on such factors as:

- 1) the volume of documents entered into the system;
- 2) the percentage of these documents that are full-text indexed; and,
- 3) whether billing is used.

An estimate is approximately 1-4MB for each 1,000 documents stored in SoftSolutions (10-40KB average document size), assuming billing and full-text indexing are used.

Considerations

The SoftSolutions dataset (several specialized data files) must be stored under a directory accessible to all SoftSolutions users. This directory is known as the “data” directory, which can have any directory name and reside in any directory location (*e.g.*, F:\SOFTSOL\DATA). SoftSolutions automatically creates one subdirectory (STATUS) under the data directory during installation.

Because the files stored in the data directory are accessed frequently by SoftSolutions, the overall performance of SoftSolutions is strongly affected by the performance of the file server holding the data directory. If the network has more than one file server, consider putting the data directory on the fastest server.

NOTE

All users require the operating system's "read" and "write" rights to the data directory. Users also need "read" and "write" rights to the STATUS directory for creating an audit trail or error logs during mass importing, and for running the SoftSolutions Indexer. Each user also requires the operating system's "sharable" rights to the data directory. If the SoftSolutions SEM (Server Enhancement Module) is in use, *all* rights to the data directory can be removed for users.

Documents

Documents tracked by SoftSolutions can be spread across many file servers. They can even be stored individually on workstations' hard disks (*not* recommended).

Two things need to be considered when deciding *where* to store documents:

- 1) Documents that are stored on a workstation's local hard disk cannot be accessed by other workstations. Also, they cannot be indexed, archived, or deleted unless the Indexer or archival/deletion report is run from the workstation where the documents are stored.
- 2) All users requiring access to a document must gain access by using the same *full* name. For example:
 - a) A particular user accesses file server #1 by the drive letter designation F:\.
 - b) The user creates a document on file server #1 and defines the document's location and name as "F:\WP\BHC\MEMO\[MYFILE]."
 - c) If "[MYFILE]" is to be retrieved by other users, they must also have access to file server #1 and use the drive letter designation F:\.

A file's full name includes the file server designation or drive letter, the directory names (path), and the file's name.

Novell, Inc. recommends that SoftSolutions be used to automatically determine the storage location of documents created through SoftSolutions based on data stored in the SoftSolutions Profiles. Such data includes information from the author, creation user, document type, document security group, custom fields, and application fields, or any combination.

The directory structure may be customized to the specific requirements of the organization, or of each user, department, or practice group in the organization. The directory structure is defined in the CONTROL FILE window. *(See the definition for the “Loc” field in the CONTROL FILE on page 83 for information and examples on creating a directory structure.)*

Drive Letters vs. “Logical” Paths

Novell, Inc. recommends using logical paths versus drive letters for defining document locations. For example, using FS1/VO1:\JAC\CORR instead of F:\JAC\CORR. SoftSolutions' mapping utilities can then be used to convert this path to a format supported by the application used for accessing the document. By doing so, the path is more meaningful to the user, and mapping for all workstations need not be the same in order for all workstations to have access to the necessary documents. *(See the definition for the “Loc” field in the CONTROL FILE on page 83 for information and examples on creating a directory structure.)*

As an example using NetWare with a word processing application, the following illustrates two workstations (WS1 and WS2) in a network with three file servers (FS1, FS2 and FS3):

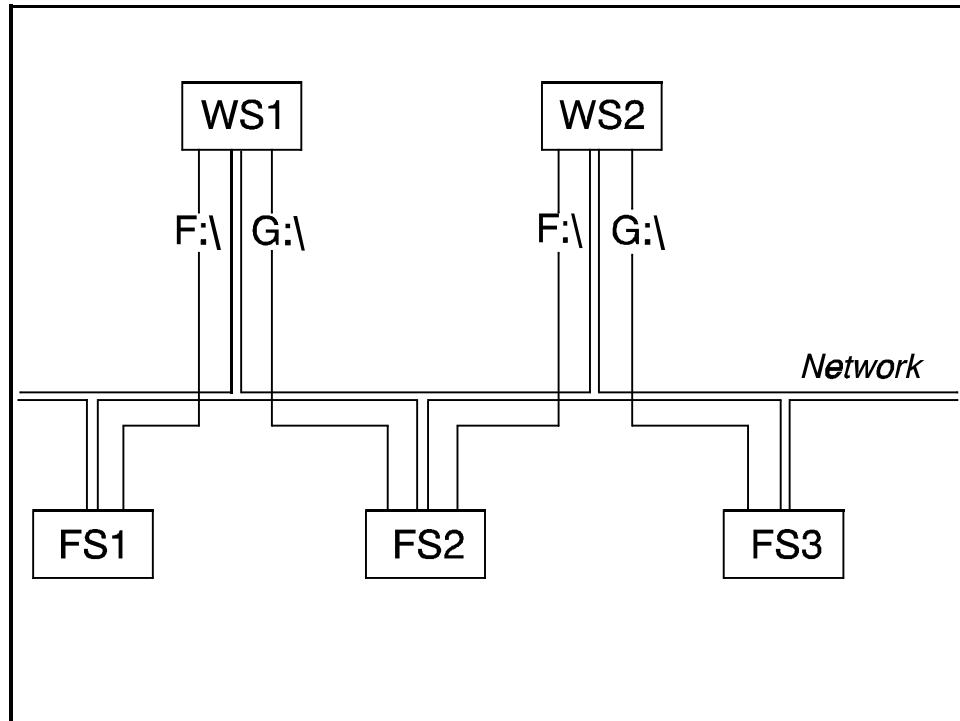


Figure 4

- The drive designations for WS1 may be mapped as F:\ for FS1 and G:\ for FS2; and for WS2 may be mapped as F:\ for FS2 and G:\ for FS3. (DOS requires these designations to be mapped individually per workstation.)

*Notice that each workstation has drive letter designations of F:\ and G:\. Also note that WS1 has **not** been mapped to FS3, and WS2 has **not** been mapped to FS1. Workstations only need to be mapped to file servers per individual user requirements.*

Using Only Drive Letters

- A user on WS1 creates a document on file server FS2 and defines its location as G:\WP\CLF\MEMO.
- A user on WS2 tries to retrieve the MEMO from FS2 and fails because WS2'S G:\ mapping points to the FS3 file server.

Using “Logical” Paths

- A user on WS1 creates a document on file server FS2 and defines its location as ADM:WP\CLF\MEMO.
- A user on WS2 retrieves the MEMO from FS2.
- MEMO is located by WS2 because the document's full path is the same as for WS1. Users do not need to worry about the F:\ and G:\ drive letter designations, which connections are transparent due to mapping.
- The document can be retrieved by any workstation in the network by specifying the “logical” path and filename.

Mapping Utilities

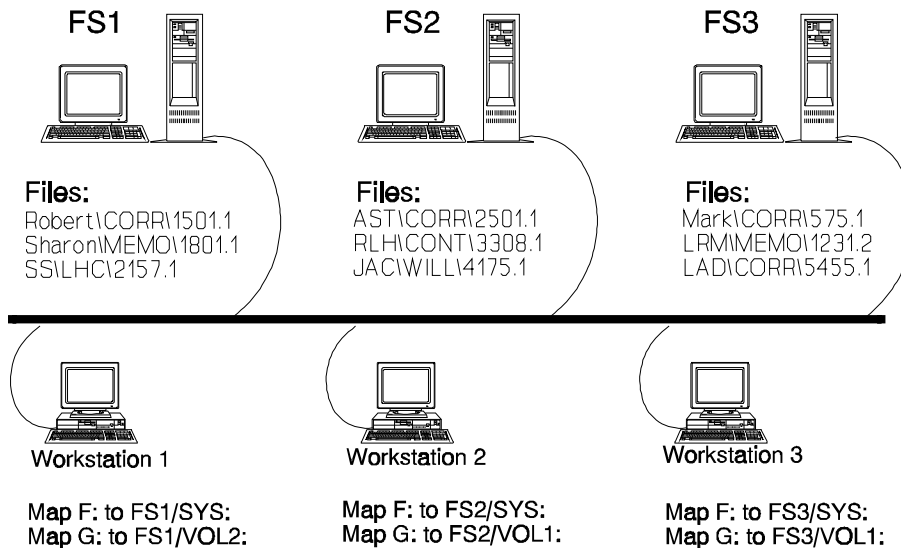
Many applications require a document's path to begin with a drive letter. “Logical” paths or paths including *[File Server]/[Volume]:\[DIR1]\[DIR]* must be changed to include a drive letter (e.g., “F:”). This can be performed automatically by SoftSolutions so that the location in the Profile may still contain a “logical” path or a path beginning with file server and volume names for the advantages previously discussed.

Under NetWare networks, dynamic file server attachments are supported so that users are not required to attach to all necessary file servers before accessing SoftSolutions. This is done using the /NOVMAP switch (see “OAM Switches” on page 68). When users attempt to access documents located on a server to which they are not attached, SoftSolutions can automatically establish a temporary file server attachment if the file server and volume name are provided.

SoftSolutions may obtain the file server and volume name directly from the **Location** field in the Profile; or, the “logical” path contained in the **Location** field may be converted to contain a file server and volume by a *.MAP file for the purpose of establishing an attachment. Once the file server and volume name are obtained, the next available drive letter is automatically mapped to this location and the path in the **Location** field in the Profile is converted to contain this drive letter for the purpose of accessing the document.

For example, the location is ADM:JAC\CORR. The *.MAP file indicates that “ADM:” should be replaced with “FS1/VOL1:.” The location is automatically converted to FS1/VOL1:JAC\CORR. SoftSolutions then maps the next available drive letter to FS1/VOL1. When exiting the document, the location in the Profile is still ADM:JAC\CORR.

NetWare Dynamic Drive Attachments



Scenarios:

1. A user on Workstation 1 creates a Profile with the Location "FS1/SYS:\ROBERT\CORR\1501.1." When the Location is passed to the application, Dynamic Drive Mapping checks to see if FS1/SYS: is mapped to a drive letter. Since F: is currently mapped to FS1/SYS:, "F:\ROBERT\CORR\1501.1" is passed to the application.
2. A user on Workstation 1 retrieves the document "FS2/VOL1:\WILL\4175.1." Dynamic Drive Mapping recognized that a drive letter is not mapped to FS2/VOL1:, so the next available drive letter (H:) is temporarily mapped. Thus, "H:\JAC\WILL\4175.1" is passed to the application. When the application is exited, the H: mapping is deleted.

NOTE: Most applications do not understand file servers/volumes, and therefore a Location with a drive letter is required for NetWare. SoftSolutions' Dynamic Drive Mapping automatically converts file servers/volumes to drive letters.

Figure 5

For other networks, dynamic file server attachments are not currently supported. Therefore, users need to attach to all file servers containing the needed documents before invoking SoftSolutions. However, each user may map their drive letters differently (*e.g.*, one user may have file server “1” mapped as “F:” and another have file server “1” mapped as “H:”). *.MAP files are used to simply convert the path to include a specified drive letter for the purpose of accessing the document.

If the path cannot be converted to include drive letters for applications, the **Document Size** field in the DOCUMENT PROFILE window will display zero, and the **Open**, **Launch**, and **Preview** options in the **File** menu *cannot* be used for the document.

***.MAP Files**

*.MAP files are simple text files which define a “search-and-replace” method for accessing drives through SoftSolutions. This allows **Location** in the DOCUMENT PROFILE window to point to the document location without using an actual drive letter.

There are two types of *.MAP files: the DEFAULT.MAP and personalized *.MAP files. DEFAULT.MAP can supply the mapping data for all SoftSolutions users, but there may likely be some who desire different mappings. Therefore, SoftSolutions also handles personalized *.MAP files.

NOTE

If you are running SoftSolutions on NetWare 4.0, Bindery Emulation must be enabled in order for dynamic drive mapping to be performed. By default, Bindery Emulation is enabled. For additional information, contact Novell GroupWare Customer Support Services (*see Appendix E*).

DEFAULT.MAP File

Use any text editor and create a file called “DEFAULT.MAP” in the SoftSolutions data directory. This file can contain comment lines, blank lines, and text lines in any order. The format of a *text* line is as follows:

```
[Search String], [Replace String], [Case], [Separator]
```