

HP AdvanceNet

NS

Cross-System NFT Reference Manual

HP AdvanceNet

NS
Cross-System NFT

Reference Manual



11000 WOLFE ROAD, CUPERTINO, CA. 95014

Part No. 5958-8563
E0891

Printed in U.S.A. AUG 1991

NOTICE

The information contained in this document is subject to change without notice.

HEWLETT-PACKARD MAKES NO WARRANTY OF ANY KIND WITH REGARD TO THIS MATERIAL, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance or use of this material.

Hewlett-Packard assumes no responsibility for the use or reliability of its software on equipment that is not furnished by Hewlett-Packard.

This document contains proprietary information which is protected by copyright. All rights are reserved. No part of this document may be photocopied, reproduced or translated to another language without the prior written consent of Hewlett-Packard Company.

PRINTING HISTORY

New editions are complete revisions of the manual. Update packages, which are issued between editions, contain additional and replacement pages to be merged into the manual by the customer. The dates on the title page change only when a new edition or a new update is published. No information is incorporated into a reprinting unless it appears as a prior update; the edition does not change when an update is incorporated.

The software code printed alongside the date indicates the version level of the software product at the time the manual or update was issued. Many product updates and fixes do not require manual changes and, conversely, manual corrections may be done without accompanying product changes. Therefore, do not expect a one to one correspondence between product updates and manual updates.

First Edition	OCT 1986
Update 1.	DEC 1987
Second Edition	JAN 1989
Third Edition.	AUG 1991



LIST OF EFFECTIVE PAGES

The List of Effective Pages gives the date of the most recent version of each page in the manual. To verify that your manual contains the most current information, check the dates printed at the bottom of each page with those listed below. The date on the bottom of each page reflects the edition or subsequent update in which that page was printed.

Effective Pages	Date
All.	AUG 1991

PREFACE

AdvanceNet is the collective name for Hewlett-Packard's data communications and data management products. Hewlett-Packard's Network Services (NS) product is the networking software that allows HP computer systems to communicate with each other. Network File Transfer (NFT) is one of the user services provided by NS.

PURPOSE

The *NS Cross-System NFT Reference Manual* is provided as a supplement to the NS manual sets for the following products:

- NS-ARPA/1000
- NS3000/V
- NS3000/XL
- NS/9000
- NS for the DEC VAX computer
- NS on the PC.

This manual is the primary reference source for programmers and users who will be using NFT to copy files between *systems of different types*.

The *NS Cross-System NFT Reference Manual* does not explain how to copy files between *systems of the same type*. For this information, refer to your NS user manual.

NOTE

Throughout this manual, the name HP 3000 is used for computers running either the MPE V or MPE XL operating systems; and the name NS/3000 refers to NS3000/V and NS3000/XL.

The MS-DOS* and PC-DOS operating systems are jointly referred to in this manual as DOS. Throughout this manual, the name PC is used for computers running either of these operating systems.

*MS-DOS is a U.S. registered trademark of Microsoft Corporation.

PREFACE (continued)

ASSUMPTIONS

It is assumed that readers of this manual have read the NFT section of the user manual provided with their particular NS implementation. Readers of this manual should understand the three-node model and Interchange Format concepts and be familiar with the syntax of the DSCOPY command as it is implemented on their local NS system.

SUPPORTED NS VERSIONS

Cross-system NFT as documented in this manual is supported for the following version numbers on the different computer systems:

For NS/1000: Version 5.1 or later.

For NS3000/V: Version V-Delta-4 or later.

For NS3000/XL: Version 1.2 or later.

For NS/9000: Version 7.0 or later.

For NS for the VAX Computer: Version 2.1 or later.
(VMS Operating System: Version 4.4 or greater)

For NS on the HP OfficeShare Network (PC): Version B.00.00 or later.
(MS-DOS Operating System: Versions 3.1 to 3.3)

NOTE

If you are running an NS version different than the one listed above for your system, you may encounter NFT options and functionalities that are not covered in this manual. This manual is updated periodically to include new NFT enhancements; however, because of the large number of systems and software releases involved, the manual sometimes may not contain the latest software updates for *all* systems.

CONVENTIONS USED IN THIS MANUAL

NOTATION **DESCRIPTION**

nonitalics Words in syntax statements which are not in italics must be entered exactly as shown. Punctuation characters other than brackets, braces and ellipses must also be entered exactly as shown. For example:

EXIT;

italics Words in syntax statements which are in italics denote a parameter which must be replaced by a user-supplied variable. For example:

CLOSE *filename*

[] An element inside brackets in a syntax statement is optional. Several elements stacked inside brackets means the user may select any one or none of these elements. For example:

$\left[\begin{array}{l} A \\ B \end{array} \right]$ User *may* select A or B or neither.

{ } When several elements are stacked within braces in a syntax statement, the user must select one of those elements. For example:

$\left\{ \begin{array}{l} A \\ B \\ C \end{array} \right\}$ User *must* select A or B or C.

... A horizontal ellipsis in a syntax statement indicates that a previous element may be repeated. For example:

[,*itemname*]...;

In addition, vertical and horizontal ellipses may be used in examples to indicate that portions of the example have been omitted.

⋮ A shaded delimiter preceding a parameter in a syntax statement indicates that the delimiter *must* be supplied whenever (a) that parameter is included or (b) that parameter is omitted and any *other* parameter which follows is included. For example:

itema[⋮*itemb*][,*itemc*]

means that the following are allowed:

itema
itema,itemb
itema,itemb,itemc
itema,,itemc

CONVENTIONS (continued)

Δ When necessary for clarity, the symbol Δ may be used in a syntax statement to indicate a required blank or an exact number of blanks. For example:

```
SET[(modifier)] $\Delta$ (variable);
```

underlining When necessary for clarity in an example, user input may be underlined. For example:

```
NEW NAME? ALPHA
```

In addition, brackets, braces or ellipses appearing in syntax or format statements which must be entered as shown will be underlined. For example:

```
LET var[[subscript]] = value
```

shading

Shading represents inverse video on the terminal's screen. In addition, it is used to emphasize key portions of an example.

The symbol may be used to indicate a key on the terminal's keyboard. For example, RETURN indicates the carriage return key.

CONTROLchar

Control characters are indicated by CONTROL followed by the character. For example, CONTROLY means the user presses the control key and the character Y simultaneously.

CONTENTS

Section 1

INTRODUCTION

Overview	1-1
Three-Node Model	1-1
File Copying Formats	1-2
Using This Manual	1-3
How This Manual Is Organized	1-3
How to Use This Manual	1-4
Step One	1-4
Step Two	1-4
Step Three	1-4

Section 2

DSCOPY INITIATOR COMMAND SYNTAX AND EXAMPLES

NS-ARPA/1000 DSCOPY Command Syntax	2-2
NS-ARPA/1000 DSCOPY Commands	2-4
NS-ARPA/1000 DSCOPY Examples	2-5
NS/3000 DSCOPY Command Syntax	2-6
NS/3000 DSCOPY Examples	2-7
NS/9000 DSCOPY Command Syntax	2-9
NS/9000 DSCOPY Examples	2-10
DSCOPY for the VAX Computer	2-12
DSCOPY Examples for VAX Computers	2-13
PC DSCOPY Command Syntax	2-14
DSCOPY Examples for the PC	2-16
Interchange Format Options	2-17
Common Options	2-18
Special Characters	2-19

Section 3

HP 1000 SOURCE FILE

RTE-A File System	3-1
HP 1000 Logon Syntax	3-1
HP 1000 File Name Syntax	3-1
Supported Target Nodes	3-2
HP 3000 Target Node	3-3
HP 3000 Logon Syntax	3-3
HP 3000 File Name Syntax	3-3
Wild Card Characters in NS-ARPA/1000 to NS/3000 NFT	3-4
HP 1000 to HP 3000 Interchange Format Defaults	3-4
Special Cases	3-5
Interchange Format Options	3-5
Supported Interchange Format Options	3-6
Interchange Format Option Results	3-6
File Mappings	3-8

CONTENTS (continued)

HP 9000 Target Node	3-10
NS/9000 Logon Syntax	3-10
HP 9000 File Name Syntax	3-10
HP 1000 to HP 9000 Interchange Format Defaults	3-11
Interchange Format Options	3-12
Supported Interchange Format Options	3-12
Unsupported Options	3-12
Interchange Format Option Results	3-13
File Mappings	3-14
DEC VAX Target Node	3-16
VAX System Logon Syntax	3-16
VAX System File Name Syntax	3-16
DECnet VAX System File Name Syntax	3-18
File Access Modes	3-19
HP 1000 to VAX Computer Interchange Format Defaults	3-20
Interchange Format Options	3-20
Supported Interchange Format Options	3-21
Unsupported Options	3-21
Interchange Format Option Results	3-21
File Mappings	3-24
PC Target Node	3-26
DOS Stream Files	3-26
PC Logon Syntax	3-26
PC File Name Syntax	3-26
HP 1000 to PC Interchange Format Defaults	3-27
Interchange Format Options	3-28
Supported Interchange Format Options	3-28
Unsupported Options	3-29
Interchange Format Option Results	3-29
File Mappings	3-30

Section 4

HP 3000 SOURCE FILE

MPE File System	4-1
HP 3000 Logon Syntax	4-1
HP 3000 File Name Syntax	4-2
Supported Target Nodes	4-2
HP 1000 Target Node	4-3
HP 1000 Logon Syntax	4-3
HP 1000 File Name Syntax	4-3
HP 1000 File Masks and Wild Card Characters	4-4
HP 3000 Wild Card Characters	4-4
HP 3000 to HP 1000 Interchange Format Defaults	4-5
Interchange Format Options	4-6
Supported Interchange Format Options	4-6
Unsupported Options	4-6
Interchange Format Option Results	4-7
File Mappings	4-8

CONTENTS (continued)

HP 9000 Target Node	4-10
NS/9000 Logon Syntax	4-10
HP 9000 File Name Syntax	4-10
HP 3000 to HP 9000 Interchange Format Defaults	4-11
Interchange Format Options	4-12
Supported Interchange Format Options	4-12
Unsupported Options	4-12
Interchange Format Option Results	4-13
File Mappings	4-14
DEC VAX Target Node	4-16
VAX System Logon Syntax	4-16
VAX System File Name Syntax	4-16
DECnet VAX System File Name Syntax	4-18
Using the FILEPASSWORD Option with DECnet	4-19
File Access Modes	4-19
HP 3000 to VAX/VMS Interchange Format Defaults	4-20
Interchange Format Options	4-21
Supported Options	4-21
Unsupported Options	4-21
Interchange Format Option Results	4-22
File Mappings	4-24
PC Target Node	4-26
DOS Stream Files	4-26
PC Logon Syntax	4-26
PC File Name Syntax	4-26
HP 3000 to PC Interchange Format Defaults	4-27
Interchange Format Options	4-28
Supported Interchange Format Options	4-28
Unsupported Options	4-28
Interchange Format Option Results	4-29
File Mappings	4-30

Section 5

HP 9000 SOURCE FILE

HP-UX File System	5-1
HP 9000 Logon Syntax	5-1
HP 9000 File Name Syntax	5-2
Supported Target Nodes	5-2
HP 1000 Target Node	5-3
HP 1000 Logon Syntax	5-3
HP 1000 File Name Syntax	5-3
HP 9000 to HP 1000 Interchange Format Defaults	5-4
Interchange Format Options	5-5
Supported Interchange Format Options	5-5
Unsupported Options	5-6
Interchange Format Option Results	5-6
File Mappings	5-8

CONTENTS (continued)

HP 3000 Target Node.	5-9
HP 3000 Logon Syntax	5-9
HP 3000 File Name Syntax	5-9
HP 9000 to HP 3000 Interchange Format Defaults	5-10
Special Cases.	5-10
Interchange Format Options.	5-11
Supported Interchange Format Options	5-11
Unsupported Options	5-12
Interchange Format Option Results	5-12
File Mappings	5-13
DEC VAX Target Node.	5-15
VAX System Logon Syntax.	5-15
VAX System File Name Syntax	5-15
DECnet VAX System File Name Syntax	5-17
File Access Modes.	5-18
HP 9000 to VAX/VMS Interchange Format Defaults	5-19
Interchange Format Options.	5-19
Supported Options	5-19
Unsupported Options	5-20
Interchange Format Option Results	5-20
File Mappings	5-22
PC Target Node	5-25
DOS Stream Files.	5-25
PC Logon Syntax	5-25
PC File Name Syntax	5-25
HP 9000 to PC Interchange Format Defaults	5-26
Interchange Format Options.	5-27
Supported Interchange Format Options	5-27
Unsupported Options	5-27
Interchange Format Option Results	5-28
File Mappings	5-29

Section 6

VAX/VMS SOURCE FILE

RMS File System.	6-1
VAX System Logon Syntax	6-1
VAX System File Name Syntax.	6-2
DECnet VAX File Name Syntax	6-3
Using the FILEPASSWORD Option with DECnet	6-4
File Access Modes	6-4

CONTENTS (continued)

Supported Target Nodes	6-5
HP 1000 Target Node.	6-5
HP 1000 Logon Syntax	6-6
HP 1000 File Name Syntax	6-6
VAX/VMS to NS/1000 Interchange Format Defaults.	6-7
Interchange Format Options.	6-8
Supported Interchange Format Options	6-8
Unsupported Options	6-8
Interchange Format Option Results	6-9
File Mappings	6-11
HP 3000 Target Node.	6-12
HP 3000 Logon Syntax	6-12
HP 3000 File Name Syntax	6-12
VAX/VMS to HP 3000 Interchange Format Defaults.	6-13
Special Cases.	6-13
Interchange Format Options.	6-14
Supported Interchange Format Option	6-14
Unsupported Options	6-14
Interchange Format Option Results	6-15
File Mappings	6-16
HP 9000 Target Node.	6-18
NS/9000 Logon Syntax.	6-18
HP 9000 File Name Syntax	6-18
VAX/VMS to HP 9000 Interchange Format Defaults.	6-19
Interchange Format Options.	6-19
Supported Options	6-20
Unsupported Options	6-20
Interchange Format Option Results	6-21
File Mappings	6-22
PC Target Node	6-23
DOS Stream Files.	6-23
PC Logon Syntax	6-23
PC File Name Syntax	6-23
DEC VAX/VMS to PC Interchange Format Defaults	6-24
Interchange Format Options.	6-25
Supported Interchange Format Options	6-25
Unsupported Options	6-25
Interchange Format Option Results	6-26
File Mappings	6-27

Section 7

PC SOURCE FILE

DOS Stream Files	7-1
PC Logon Syntax.	7-1
PC File Name Syntax	7-2

CONTENTS (continued)

Supported Target Nodes	7-2
HP 1000 Target Node.	7-3
HP 1000 Logon Syntax	7-3
HP 1000 File Name Syntax	7-3
PC to HP 1000 Interchange Format Defaults	7-4
Interchange Format Options.	7-5
Supported Interchange Format Options	7-5
Unsupported Options	7-5
Interchange Format Option Results	7-6
File Mappings	7-7
HP 3000 Target Node.	7-8
HP 3000 Logon Syntax	7-8
HP 3000 File Name Syntax	7-8
PC to HP 3000 Interchange Format Defaults	7-9
Special Cases.	7-9
Interchange Format Options.	7-10
Supported Interchange Format Options	7-10
Unsupported Options	7-11
Interchange Format Option Results	7-11
File Mappings	7-12
HP 9000 Target Node.	7-14
HP 9000 Logon Syntax	7-14
HP 9000 File Name Syntax	7-14
PC to HP 9000 Interchange Format Defaults	7-14
Interchange Format Options.	7-15
Supported Interchange Format Options	7-15
Unsupported Options	7-16
Interchange Format Option Results	7-16
File Mappings	7-18
DEC VAX/VMS Target Node.	7-19
VAX/VMS Logon Syntax.	7-19
VAX/VMS File Name Syntax	7-19
DECnet VAX File Name Syntax	7-21
File Access Modes.	7-21
PC to VAX/VMS Interchange Format Defaults.	7-22
Interchange Format Options.	7-22
Supported Interchange Format Options	7-23
Unsupported Options	7-23
Interchange Format Option Results	7-23
File Mappings	7-25

OVERVIEW

Network File Transfer (NFT) is one of the services provided by Hewlett-Packard's Network Services (NS) product. NS is implemented on the HP 1000, HP 3000, HP 9000, DEC VAX computer systems, and PC computers. NFT enables you to copy files between NS systems in your network. The following subsection describes the general concept of Network File Transfer and Interchange Format.

NOTE

Throughout this manual, the name HP 3000 is used for computers running either the MPE V or MPE XL operating systems; and the name NS/3000 refers to NS3000/V or NS3000/XL.

The name PC is used for computers running either MS-DOS or PC-DOS operating systems.

NS on the DEC VAX computer is in support life.

THREE-NODE MODEL

NFT utilizes a **three-node** model to copy files between systems. Under the three-node model there are three logical participants in the file copy process:

- **The Initiator.** Located on the system where the copy request originates, the Initiator receives the user request and *initiates* the copy process. The initiator system determines the NFT command syntax you use.
- **The Producer.** Located on the same node as the source file, the Producer accesses that file and *produces* the data which is to be copied.
- **The Consumer.** Located on the same node as the target file, the Consumer *consumes* the data and writes it into the target file.

All three participants are logically distinct. They may be three separate processes on three separate nodes, or any two, or all three, may reside on the same node. This is because the copy request does not have to originate from either the source or the target node. Figure 1-1 is a conceptual view of the three-node model in NFT.

THREE-NODE MODEL

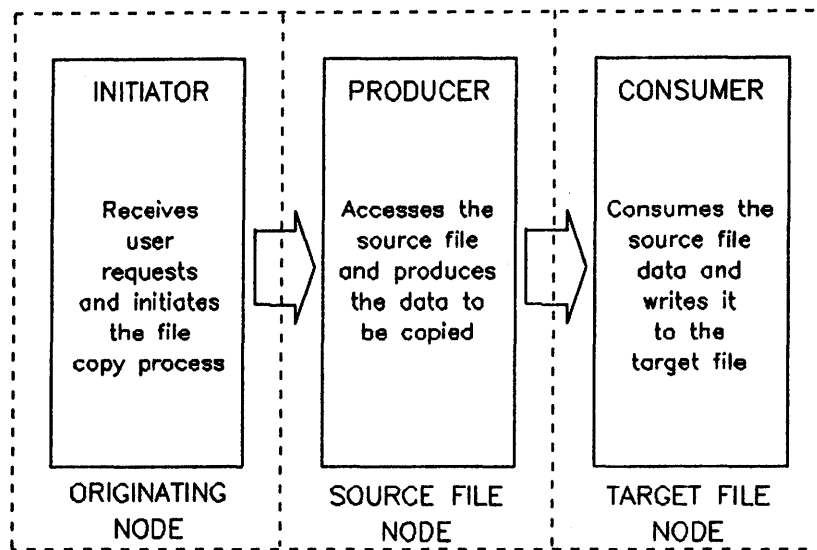


Figure 1-1. Three-Node Model.

NOTE

When a PC takes part in a file transfer using NFT, it *must* be the initiator. To transfer files to or from a PC, you must initiate the transfer from the PC.

HP 9000 workstations are *not* the same as PC's. In an NFT transfer, the workstation does not need to be the initiator.

FILE COPYING FORMATS

NFT uses two file copying formats: **Transparent Format** and **Interchange Format**. Transparent Format is used by default when the source and the target system have the *same type of operating system*.

Interchange Format is invoked by default for cross-system NFT.

When a file is copied using Interchange Format, it is translated into Interchange Format at the source system before being copied to the target system. At the target system, the file is mapped from Interchange Format into the target system's file format. Interchange Format's standard file attributes enable the target computer to map the source file into a target file that has attributes that match the source file's as closely as possible. You can also use the Interchange Format options to give a target file a different set of attributes from those that characterized the source file.

Figure 1-2 is a conceptual view of Interchange Format.

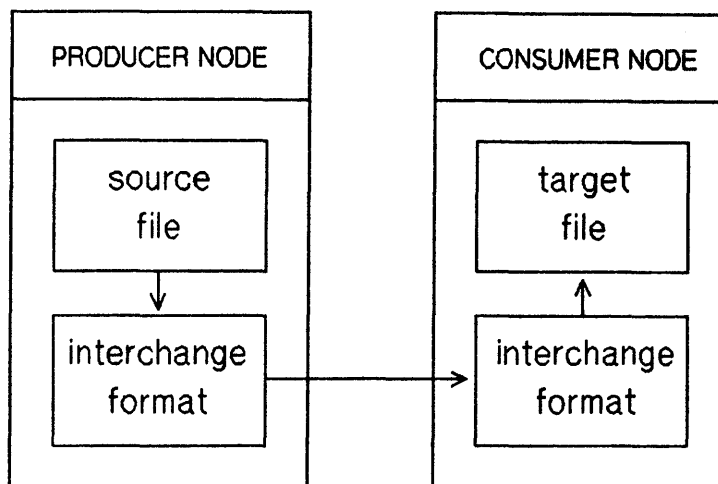


Figure 1-2. Interchange Format.

USING THIS MANUAL

The *NS Cross-System NFT Reference Manual* is a generic reference manual; it may be used by *any* NS user, regardless of whether the user's local computer is an HP 1000, HP 3000, HP 9000, VAX system, or PC computer.

This manual explains how to copy files between *different types* of NS systems. It does *not* explain how to copy a file between systems of the same type. For this information, refer to the appropriate NS user manual.

You should refer to the *NS Cross-System NFT Reference Manual* when you want to:

- Copy a file from your local system to a remote system of a different type.
- Copy a file to your local system from a remote system of a different type.
- Copy a file between remote systems of different types.

It is assumed that users of this manual are familiar with NFT on their local system.

How This Manual Is Organized

The first section provides this **Introduction**.

Section 2 provides cross-system DSCOPY examples, basic syntax syntax of the DSCOPY command for each initiator system, an Interchange Format Options Table, and a Special Character Table showing the special characters that need to be escaped in the DSCOPY command for each initiator node.

Sections 3 through 6 are divided into sections according to the **location of the source file that you want to copy**. Within each of these sections, the subsections are then subdivided into supported target systems.

Introduction

Each section is preceded by a tabbed divider that identifies the location of the source file. The sections are organized as follows:

- **Section 3, "HP 1000 Source File,"** explains how to copy an NS/1000 file to an NS/3000, NS/9000, NS for the VAX node, or PC node.
- **Section 4, "HP 3000 Source File,"** explains how to copy an NS/3000 file to an NS/1000, NS/3000, NS/9000, NS for the VAX node, or PC node.
- **Section 5, "HP 9000 Source File,"** explains how to copy an NS/9000 file to an NS/1000, NS/3000, NS for the VAX node, or PC node.
- **Section 6, "DEC VAX/VMS Source File,"** explains how to copy a file from an NS for the VAX node to an NS/1000, NS/3000, NS/9000, or PC node.
- **Section 7, "PC Source File,"** explains how to copy a file from a PC node to an NS/1000, NS/3000, NS/9000, or NS for the VAX node.
- **Appendix A, "DSCOPY Examples,"** shows DSCOPY command examples for NS/1000, NS/3000, NS/9000, NS for the VAX node, or PC.

How to Use This Manual

Section 2 provides the basic DSCOPY syntax for each HP system and a list of cross-system DSCOPY examples. This may be all you need.

Section 3 through Section 7 provides detailed information on file transfer between different source and target systems. Refer to these sections if you need more information about NFT transfer between specific source and target systems. Section 3 (HP 1000 Source File) provides information about NFT from an HP 1000 source node to the supported target nodes; Section 4 (HP 3000 Source File) provides information about NFT from an HP 3000 source node, and so on.

DSCOPY INITIATOR COMMAND SYNTAX AND EXAMPLES

SECTION

2

This section lists the basic syntax of the DSCOPY command for each initiator system:

- NS-ARPA/1000
- NS/3000
- NS/9000
- NS on the VAX computer
- NS on the PC.

This section also provides the following tables:

- Interchange Format Options Table - which shows the availability and syntax of the Interchange Format options for each initiator system,
- Common Options Table - which lists the common DSCOPY options available for cross-system NFT, and
- Special Character Table - which shows the special characters that need to be escaped in the DSCOPY command for each initiator node.

At the end of the section, a list of cross-system DSCOPY examples provides a quick reference of DSCOPY syntax.

NS-ARPA/1000 DSCOPY COMMAND SYNTAX

This subsection contains the basic syntax for NS-ARPA/1000 DSCOPY command for interactive file transfer.

Syntax

```
DSCOPY sourcefile[slogon]>snode TO targetfile[tlogon]>tnode, [options]
```

Parameters

sourcefile

The name of the source file.

If the file is not in the current working directory (or group), you must specify the full pathname for the file. The format to specify a full pathname is shown below. Use the appropriate format for the system on which the file resides.

HP 1000: */directory/subdir/.../file.ext*

HP 3000: *file.group.account*

HP 9000: */directory/subdir/.../file.ext*

VAX computer: *device:[directory]file.type;version*

PC: *device:\directory\subdir\...\file.ext*

[*slogon*]

The logon and password, if any, at the source node. Brackets are required if logon is used. The logon is not necessary if the source file resides on the local (initiator) node.

The format to specify a logon is shown below. Use the appropriate logon format for the system on which the file resides.

HP 1000: *account/password*

HP 3000: *user/userpass.account/acctpass*

HP 9000: *user:password*

VAX computer: *user:password*

PC is a single-user system; there is no logon for PC.

>*snode*

The name of the source node. Must be preceded by ">" or comma. >*snode* is not required if the source node is the local (initiator) node.

- targetfile* The name of the target file.
- If the file is not in the current working directory (or group), you must specify the full pathname for the file. The format to specify a full pathname is shown above under *sourcefile*. Use the appropriate file pathname format for the system on which the target file resides.
- [*tlogon*] The logon and password, if any, at the source node. Brackets are required if logon is used. The logon is not necessary if the source file resides on the local (initiator) node.
- The format to specify a logon is shown under *slogon*. Use the appropriate logon format for the system on which the target file resides.
- >*tnode* The name of the target node. Must be preceded by ">." >*tnode* is not required if the target node is the local (initiator) node.
- options* One or more of the options in DSCOPY. There is no limit to the number of options you can specify. Each option must be separated by a comma, semicolon or space. If conflicting options are given (for example, ASCII and BINARY), DSCOPY will issue a warning and the *last option given will take precedence*.
- The Interchange Format Options Table provided later in this section (Table 2-1) lists the Interchange Format options available for this initiator system. The effects of these options for file transfer between two systems are described in detail in the section that explains that particular Producer/Consumer configuration.
- For detailed information about the Transparent Format options, refer to the *NS-ARPA/1000 User/Programmer Reference Manual*.

NOTE

HP 1000 commands entered at the RTE Command Interpreter (CI) prompt are automatically upshifted. This might cause problems for HP 9000 producer or target nodes since HP-UX commands are case sensitive. To avoid this problem, enclose the part of the HP 1000 DSCOPY command string that references an HP 9000 node and must not be upshifted with grave accents (`). The grave accent is *not* the same as the single quotation mark (').

In interactive mode, commands typed after the DSCOPY> prompt are not automatically upshifted; you need not use the grave accents.

CONTROL D can be used to terminate the interactive mode.

NS-ARPA/1000 DSCOPY Commands

NS-ARPA/1000 additionally provides ten commands that can be used with the DSCOPY program:

- **+CLEAR.** Clears all the copy descriptor defaults previously set with the +DEFAULT command.
- **+DEFAULT.** Sets defaults for selected portions of subsequently issued copy descriptors.
- **+ECHO.** Causes commands to be echoed, or not echoed, to the list file or device.
- **+EX.** Exits DSCOPY.
- **+LL.** Changes the list file or device.
- **+RU.** Runs a program from within DSCOPY.
- **+SHOW.** Shows all currently active copy descriptor defaults set with the +DEFAULT command.
- **+TRANSFER.** Transfers control to a command file or device.
- **+WD.** Displays or changes the current working directory.
- **?** Requests help information for any DSCOPY command or copy descriptor option. Can also be used to provide a general help summary.

Each command, with the exception of ?, must begin with a plus (+) so that DSCOPY can distinguish it from a copy descriptor. Only one command can be issued per line. For complete syntax of each command, refer to the *NS-ARPA/1000 User/Programmer Reference Manual*.

NS-ARPA/1000 DSCOPY Examples

Example 1. Initiator: NS/1000. Producer: NS/9000. Consumer: NS/1000.

```
dscopy `file9k[user1:passwd]>uxnode` to testfile
```

where:

source file =	file9k	target file =	testfile
source logon =	user1:passwd	target logon =	none, default logon
source node =	uxnode	target node =	none, local node

Example 2. Initiator: NS/1000. Producer: NS/1000. Consumer: NS/3000.

```
dscopy file1k to file3k.testgrp[lab/labpass.test/testpass]>node3k
```

where:

source file =	file1k	target file =	file3k.testgrp
source logon =	none, default logon	target logon =	lab/labpass.test/testpass
source node =	none, local node	target node =	node3k

Example 3. Initiator: NS/1000. Producer: NS/9000. Consumer: VAX/VMS.

```
dscopy `/nftlabs/tfile[acct1]>node9k` to tfile.dat[user1:userpw]>vxnode
```

where:

source file =	/nftlabs/tfile	target file =	tfile.dat
source logon =	acct1	target logon =	user1:userpw
source node =	node9k	target node =	vxnode

NS/3000 DSCOPY COMMAND SYNTAX

This subsection contains the basic syntax for NS/3000 DSCOPY for interactive file transfer.

Syntax

```
DSCOPY sourcefile:snode[slogon] TO targetfile:tnode[tlogon][options]
```

Parameters

sourcefile

The name of the source file.

If the file is not in the current working directory (or group), you must specify the full pathname for the file. The format to specify a full pathname is shown below. Use the appropriate format for the system on which the file resides.

HP 1000: */directory/subdir/.../file.ext*

HP 3000: *file.group.account*

HP 9000: */directory/subdir/.../file.ext*

VAX computer: *device:[directory]file.type;version*

PC: *device:\directory\subdir\...\file.ext*

:snode

The name of the source node. Must be preceded by ":" or comma. *:snode* is not required if the source node is the local (initiator) node.

[slogon]

The logon and password, if any, at the source node. Brackets are required if logon is used. The logon is not necessary if the source file resides on the local (initiator) node.

The format to specify a logon is shown below. Use the appropriate logon format for the system on which the file resides.

HP 1000: *account/password*

HP 3000: *user/userpass.account/acctpass*

HP 9000: *user:password*

VAX computer: *user:password*

PC is a single-user system; there is no logon for PC.

- targetfile* The name of the target file.
- If the file is not in the current working directory (or group), you must specify the full pathname for the file. The format to specify a full pathname is shown under *sourcefile*. Use the appropriate format for the system on which the file resides.
- :tnode* The name of the target node. Must be preceded by ":" or comma. *tnode* is not required if the target node is the local (initiator) node.
- tlogon* The logon and password, if any, at the target node. Brackets are required if logon is used. The logon is not necessary if the target file resides on the local (initiator) node.
- options* One or more of the options in DSCOPY. There is no limit to the number of options that you can specify. The Interchange Format Options Table provided later in this section (Table 2-1) lists the Interchange Format options available for this initiator system. The effects of these options for file transfer between two systems are described in detail in the section that explains that particular Producer/Consumer configuration.
- For detailed information about the Transparent Format options, refer to the NS user manual that covers NFT on your HP 3000.

NS/3000 DSCOPY Examples

Example 1. Initiator: HP 3000. Producer: HP 1000. Consumer: HP 9000.

```
dscopy test.dat:rtenode[acct1/passwd] to /examples/test2:uxnode[usr2:pwd]
```

where:

source file =	test.dat	target file =	/examples/test2
source logon =	acct1/passwd	target logon =	usr2:pwd
source node =	rtenode	target node =	uxnode

Example 2. Initiator: HP 3000. Producer: HP 3000. Consumer: HP 9000.

```
dscopy file1:node3[lp/rosebud.test/passwd] to /examples/file2:node6[usr1:secret]
```

where:

source file =	file1	target file =	/examples/file2
source logon =	lp/rosebud.test/passwd	target logon =	usr1:secret
source node =	node3	target node =	node6

DSCOPY Initiator Command Syntax and Examples

Example 3. Initiator: HP 3000. Producer: VAX. Consumer: HP 1000.

```
dscopy "[.subdir]test5.dat":node4[usr:usrpw] to /test/file2.dat:node7[acct1/pw1]
```

where:

source file =	[.subdir]test5.dat	target file =	/test/file2.dat
source logon =	usr:usrpw	target logon =	acct1/pw1
source node =	node4	target node =	node7

NS/9000 DSCOPY COMMAND SYNTAX

This subsection contains the basic syntax for NS/9000 DSCOPY for interactive file transfer.

Syntax

```
DSCOPY [options] snode#slogon#sourcefile tnode#tlogon#targetfile
```

Parameters

options Can be one or more DSCOPY options. If multiple options are specified, they must be separated by a space. Options can be specified in any order. The Interchange Format Options Table provided later in this section (Table 2-1) lists the Interchange Format options available for this initiator system. The effects of these options for file transfer between two systems are described in detail in the section that covers that particular Producer/Consumer configuration.

snode The name of the source node. *snode* is not required if the source node is the local (initiator) node.

slogon The logon and password, if any, at the source node. The logon is not necessary if the source file resides on the local (initiator) node.

The format to specify a logon is shown below. Use the appropriate logon format for the system on which the file resides.

HP 1000: *account/password*

HP 3000: *user/userpass.account/acctpass*

HP 9000: *user:password*

VAX computer: *user:password*

PC is a single-user system; there is no logon for PC.

sourcefile The name of the source file.

If the file is not in the current working directory (or group), you must specify the full pathname for the file. The format to specify a full pathname is shown below. Use the appropriate format for the system on which the file resides.

HP 1000: */directory/subdir/.../file.ext*

HP 3000: *file.group.account*

DSCOPY Initiator Command Syntax and Examples

HP 9000: */directory/subdir/.../file.ext*
VAX computer: *device:[directory]file.type;version*
PC: *device:\directory\subdir\...\file.ext*

tnode The name of the target node. *tnode* is not required if the target node is the local (initiator) node.

tlogon The logon and password, if any, at the target node. *tlogon* is not required if the target node is the local (initiator) node.

The format to specify a logon is shown under *slogon*. Use the appropriate logon format for the system on which the target file resides.

targetfile The name of the target file.

If the file is not in the current working directory (or group), you must specify the full pathname for the file. The format to specify a full pathname is shown under *sourcefile*. Use the appropriate format for the system on which the file resides.

NOTE

You may specify `dscopy -i` to invoke Interactive Mode. When NFT is used in Interactive Mode, `dscopy` prompts you for copy descriptors until you enter **CONTROL D** to indicate End-Of-File (EOF). (EOF is configurable and may be different for your system.)

NS/9000 DSCOPY Examples

Example 1. Initiator: HP 9000. Producer: HP 9000. Consumer: HP 1000.

```
dscopy uxnode#user1:pass#/examples/sfile node1k#acct1/pw#file2
```

where:

source file =	<i>/examples/sfile</i>	target file =	<i>file2</i>
source logon =	<i>user1:pass</i>	target logon =	<i>acct1/pw</i>
source node =	<i>uxnode</i>	target node =	<i>node1k</i>

Example 2. Initiator: HP 9000. Producer: HP 3000. Consumer: VAX.

```
dscopy mpenode#lp/rosebud.test/testpass#file1 vxnode#user1:pw1#testfile
```

where:

source file =	file1	target file =	testfile
source logon =	lp/rosebud.test/testpass	target logon =	user1:pw1
source node =	mpenode	target node =	vxnode

Example 3. Initiator: HP 9000. Producer: HP 1000. Consumer: HP 3000.

```
dscopy node1k#usr/xxx222#/test/file1 node3k#lp/rosebud.test/testpass#file3.tgroup
```

where:

source file =	/test/file1	target file =	file3.tgroup
source logon =	usr/xxx222	target logon =	lp/rosebud.test/testpass
source node =	node1k	target node =	node3k

DSCOPY FOR THE DEC VAX COMPUTER

This subsection contains the basic syntax for the DSCOPY command on NS for the VAX computer.

Syntax

```
DSCOPY [options] snode#slogon#sourcefile tnode#tlogon#targetfile
```

Parameters

options Can be one or more DSCOPY options. DSCOPY options can be specified in any order and can be concatenated. Each option must be preceded by a slash (/). The options are described fully in the *Network Services for the DEC VAX Computer Manual*.

snode The name of the source node. *snode* is not required if the source node is the local (initiator) node.

slogon The logon and password, if any, at the source node. The logon is not necessary if the source file resides on the local (initiator) node.

The format to specify a logon is shown below. Use the appropriate logon format for the system on which the file resides.

HP 1000: *account/password*

HP 3000: *user/userpass.account/acctpass*

HP 9000: *user:password*

VAX computer: *user:password*

PC is a single-user system; there is no logon for PC.

sourcefile The name of the source file.

If the file is not in the current working directory (or group), you must specify the full pathname for the file. The format to specify a full pathname is shown below. Use the appropriate format for the system on which the file resides.

HP 1000: */directory/subdir/.../file.ext*

HP 3000: *file.group.account*

HP 9000: */directory/subdir/.../file.ext*

VAX computer: *device:[directory]file.type;version*

PC: *device:\directory\subdir\...\file.ext*

tnode The name of the target node. *tnode* is not required if the target node is the local (initiator) node.

tlogon The logon and password, if any, at the target node. *tlogon* is not required if the target node is the local (initiator) node. The format to specify the logon is shown under *slogon*. Use the appropriate logon format for the system on which the target file resides.

targetfile The name of the target file.

If the file is not in the current working directory (or group), you must specify the full pathname for the file. The format to specify a full pathname is shown under *sourcefile*. Use the appropriate format for the system on which the file resides.

NOTE

You may use DSCOPY /INTERACTIVE to invoke the interactive mode, which allows you to transfer several files between the same two computers, within the same DSCOPY command. DSCOPY/INTERACTIVE will generate the DSCOPY> prompt, which lets you enter any command options, source and target file arguments. To exit from the interactive mode, type **CONTROL**Z.

NS for the DEC VAX Computer is in support life.

DSCOPY Examples for VAX Computers

Example 1. Initiator: VAX/VMS. Producer: HP 1000. Consumer: HP 9000.

```
dscopy node1k#user6#test.dat node9k#acct1:pw1234#vx1.dat
```

where:

source file =	test.dat	target file =	vx1.dat
source logon =	user6	target logon =	acct1:pw1234
source node =	node1k	target node =	node9k

PC DSCOPY COMMAND SYNTAX

This subsection contains the basic syntax for PC DSCOPY for interactive file transfer.

NOTE

The MS-DOS and PC-DOS operating systems are jointly referred to in this manual as DOS. Throughout this manual, the name PC is used for computers running either of these operating systems.

When a PC takes part in a file transfer using NFT, it *must* be the initiator. To transfer files to or from a PC, you must initiate the transfer from the PC.

Syntax

```
DSCOPY [options] snode#slogon#sourcefile tnode#tlogon#targetfile
```

Parameters

options Can be one or more DSCOPY options. If multiple options are specified, they must be separated by a space. Options can be specified in any order. The Interchange Format Options Table provided later in this section (Table 2-1) lists the Interchange Format options available for this initiator system. The effects of these options for file transfer between two systems are described in detail in the section that covers that particular Producer/Consumer configuration.

snode The name of the source node. *snode* is not required if the source node is the local (initiator) node.

slogon The logon and password, if any, at the source node. *slogon* is not necessary if the source file resides on the local (initiator) node.

The format to specify a logon is shown below. Use the appropriate logon format for the system on which the file resides.

HP 1000: *account/password*

HP 3000: *user/userpass.account/acctpass*

HP 9000: *user:password*

VAX computer: *user:password*

PC is a single-user system; there is no logon for PC.

sourcefile

The name of the source file.

If the file is not in the current working directory (or group), you must specify the full pathname for the file. The format to specify a full pathname is shown below. Use the appropriate format for the system on which the file resides.

HP 1000: /*directory/subdir/.../file.ext*

HP 3000: *file.group.account*

HP 9000: /*directory/subdir/.../file.ext*

VAX computer: *device:[directory]file.type;version*

PC: *device:\directory\subdir\...\file.ext*

tnode

The name of the target node. *tnode* is not required if the target node is the local (initiator) node.

tlogon

The logon and password, if any, at the target node. *tlogon* is not required if the target node is the local (initiator) node.

The format to specify a logon is shown under *slogon*. Use the appropriate logon format for the system on which the target file resides.

targetfile

The name of the target file.

If the file is not in the current working directory (or group), you must specify the full pathname for the file. The format to specify a full pathname is shown under *sourcefile*. Use the appropriate format for the system on which the file resides.

NOTE

You may specify DSCOPY -i to invoke an interactive session with NFT. In interactive mode, DSCOPY prompts you for copy descriptors until you enter **CONTROL C** to terminate the interactive session.

DSCOPY Examples for the PC

Example 1. Initiator: PC. Producer: PC. Consumer: HP 9000.

```
dscopy A:\status\may uxnode#acct1:passwd#/examples/mayfile
```

where:

source file =	A:\status\may	target file =	/examples/mayfile
source logon =	no logon	target logon =	acct1:passwd
source node =	none, local node	target node =	uxnode

Example 2. Initiator: PC. Producer: HP 3000. Consumer: PC.

```
dscopy mpenode#lp/rosebud.test/testpass#feb.status testfile
```

where:

source file =	feb.status	target file =	testfile
source logon =	lp/rosebud.test/testpass	target logon =	no logon
source node =	mpenode	target node =	none, local node

INTERCHANGE FORMAT OPTIONS

This subsection provides a table of Interchange Format options that are available at each initiator system and contains the syntax for each option. The actual effects of each Interchange Format option for file transfers between every Producer/Consumer combination are discussed in the particular section that covers that particular source and target node. For example, the effect of the BINARY option for file transfers between an HP 1000 and HP 9000 would be covered in Section 3, "HP 1000 Source File," under the subsection of "HP 9000 Target Node."

TABLE 2-1. INTERCHANGE FORMAT OPTIONS.

Option Mnemonic	NS/1000	NS/3000	NS/9000	VAX/VMS	PC
Binary	BI[NARY]	BIN	-B	/BINARY	-B
ASCII	AS[CII]	ASC	-A ***	*	-A
Fixed	FI[XED]	FIX	-F	/FIXED	-F
File Size	FS[IZE]= <i>filesize</i>	F S IZE= <i>filesize</i>	**	**	**
Record Size	RS[IZE]= <i>recsize</i>	R S IZE= <i>recsize</i>	-L [<i>recsize</i>]	/RECL E N= <i>recsize</i>	-L
Append	AP[PEND]	APP	**	**	**
Direct	**	DIR	**	**	**
Insert Character	IC[HAR]= <i>char</i>	I C HAR= <i>char</i>	- <i>dchar</i>	/DELIMITE R = <i>char</i>	-d
Search Character	SC[HAR]= <i>char</i>	S C HAR= <i>char</i>	- <i>schar</i>	/SEARCH= <i>char</i>	-s
Interchange	*	*	*	*	*
Variable	VA[RIBLE]	VAR	**	**	-V
Strip	ST[RIP]	STRIP	**	**	**
Sequential	**	SEQ	*	*	**

* indicates that the option is the default at that system and has no effect.

** indicates that the option is not available at that system.

*** indicates that the option is only meaningful for cross-system NFT. The default setting is off.

COMMON OPTIONS

Besides the Interchange Format options (shown in Table 2-1), the DSCOPY command for each initiator system provides common options to use for file transfer between similar systems. These options can also be used for cross-system NFT. Table 2-2 lists the common options available for each initiator that can be used for cross-system NFT.

TABLE 2-2. COMMON NFT OPTIONS.

Option Explanation	NS on 1000	NS on 3000	NS on 9000	VAX/VMS	PC
Replaces existing target file with source file.	REPLACE	REP	-r	/REPLACE	-r
Deletes source file after NFT.	MOVE	MOVE	**	**	**
Suppresses all output except error messages.	QUIET	QUIET	*	*	*
Prints results of file transfer.	*	*	-p	/PRINT	-p
Invokes interactive DSCOPY session.	RETURN	RETURN	-i	/INTERACTIVE	-i
Overwrites existing target file with source file.	OVER	OVER	**	**	**
Prompts for lockword.	**	**	-P	**	-P

* indicates that the option is the default at that system.

** indicates that the option is not available at that system.

For more information on these options, refer to the user/programmer manual of your initiator system.

To invoke interactive DSCOPY on the NS/1000 or NS/3000 initiator, enter a **RETURN** after the DSCOPY keyword. To terminate an interactive DSCOPY session, enter the following after the DSCOPY prompt: an +EX command on the NS/1000, two slashes (//) on the NS/3000, **CONTROL** D on the NS/9000, **CONTROL** Z on the VAX node, or **CONTROL** C on the PC.

The following results occur when you use the -r option from NS/9000.

- If the producer is an HP-UX computer and the consumer is a non-HP-UX computer, the target file receives the default file protection and/or ownership of the target login.
- If the producer node is a non-HP-UX computer and the consumer is an HP-UX computer, the target file acquires the file mode 666 (rw-rw-rw-).

- If both the producer and consumer are non-HP-UX computers, the target file receives the default file protection of the target login.

For more information about this option, refer to the *Using Network Services* manual.

For explanation of the difference between REP and OVER, refer to the user/programmer manual for NS/1000 or NS/3000.

SPECIAL CHARACTERS

Every system has a set of characters that have special meaning to the system. If used as part of the command string in DSCOPY, these characters must be escaped to ensure that they are interpreted literally by DSCOPY. To escape a special character, you can do one of the following:

- Enclose the command string with single quotation marks in the HP 1000 and HP 9000 initiator.

The following example shows quotation marks enclosing the target filename specification on a VAX filename specification to escape the square brackets around the directory.

```
dscopy source to '[quickdir]tfile'[acct1:pw]>vnode1
```

- Enclose the command string with double quotation marks in the HP 3000 and VAX initiator node.

```
dscopy source9k to "[quickdir]tfile2"[acct1:pw]>vnode1
```

- Precede the special character with a back slash (\) in the HP 9000 initiator node.

```
dscopy source800 vnode1#acct1:pw#sys\$manager:testfile
```

The following table lists the special characters that must be escaped at each initiator node.

TABLE 2-3. SPECIAL CHARACTERS.

HP 1000	HP 3000	HP 9000	VAX/VMS
/ @ .	[] #	? & *	/ , Δ
- >]	: = ;	; ()	
[: Δ	; + "	\ \$	
	' Δ	" \ , `	
		Δ	

CROSS-SYSTEM DSCOPY EXAMPLES

Below are some cross-system DSCOPY examples.

Example 1. Initiator: NS/1000. Producer: NS/9000. Consumer: NS/3000.

The grave accent (`) is used to enclose the command line so as to escape some characters that have special meaning to the HP 1000.

```
dscopy `/nftlabs/tfile[acct1]>node9k to tfile.nftgrp[qa/passwd.test/tpass]>3knode`
```

Example 2. Initiator: NS/1000. Producer: VAX/VMS. Consumer: NS/3000.

```
dscopy tfile.dat[user1:userpw]>vxnode tfile[acct1.isys,testgrp]>node3k
```

Example 3. Initiator: NS/3000. Producer: NS/3000. Consumer: NS/9000.

```
dscopy file1:node3k[lp/rosebud.test/tpass] to /examples/file2:node6[user:secret]
```

Example 4. Initiator: NS/3000. Producer: VAX/VMS. Consumer: NS/1000.

```
dscopy "[.subdir]tfile.dat":vxnode[user1:userpw] to test.dat:node1k[acct1]
```

Example 5. Initiator: NS/3000. Producer: NS/1000. Consumer: NS/9000

```
dscopy tfile.dat:node1k[acct1] to test.dat:node9k[acct1]
```

Example 6. Initiator: NS/9000. Producer: NS/3000. Consumer: NS/9000.

```
dscopy 3knode#lp/rosebud.test/tpass#file3.nftgrp 9knode#user1:psswd#/test/tfile
```

Example 7. Initiator: NS/9000. Producer: NS/3000. Consumer: VAX/VMS.

```
dscopy node3k#"acct1.isys,testgrp"#tfile1 vxnode#user1:pw1#tfile1
```

Example 8. Initiator: NS/9000. Producer: NS/9000. Consumer: NS/1000.

```
dscopy tfile node1k#acct1#tfile
```

Example 9. Initiator: VAX/VMS. Producer: NS/1000. Consumer: NS/9000.

```
dscopy node1k#user6#test.dat node9k#acct1#vx1.dat
```

Example 10. Initiator: VAX/VMS. Producer: NS/9000. Consumer: NS/3000.

```
dscopy node9k#acct1#vx1.dat node3k#"acct1.isys,testgrp"#tfile2
```

Example 11. Initiator: NS/9000. Producer: NS/9000. Consumer: NS/1000.

In the following example, the `-s` option is used to define a search character that will be used to divide the source file into logical records at the target node. This command will create a new logical record in the HP 1000 target file `/bedtime/story` every time a FORMFEED (decimal ASCII value 012) is encountered.

```
dscopy -s012 storyfile remnode3#don/secret#/bedtime/story
```

Example 12. Initiator: PC. Producer: PC. Consumer: NS/1000.

```
DSCOPY MYLIST1.TXT DataCen#bob/locked#products/partlist
```

Example 13. Initiator: PC. Producer: NS/3000. Consumer: PC.

```
DSCOPY admin#sue/nolimit.sales/nopass,pubgrp#memo memo3.txt
```

Example 14. Initiator: PC. Producer: PC. Consumer: VAX.

In the following example, the target file resides on a VAX node in a DECnet network. File transfer is routed through a VAX node that runs both the NS network and the DECnet network. This shared node is called `share` in our example. The VAX target node containing the target file is called `far` in the example.

```
DSCOPY VX.TXT "share#pub1:pswd#far\"jim sec\"::sys$mfg[med]samp.edt"
```

Example 15. Initiator: NS/3000. Producer: NS/3000. Consumer: NS/9000.

In this example, you are logged to an HP 3000, and need to push a file "sourcefile" on the local HP 3000 to a "targetfile" on an HP 9000.

```
DSCOPY sourcefile TO targetfile:9knodename[user:password]
```

Example 16. Initiator: NS/3000. Producer: NS/9000. Consumer: NS/3000.

In this example, you are logged to an HP 3000 and need to pull a file "sourcefile" from an HP 9000 to the "targetfile" on the local HP 3000.

```
DSCOPY sourcefile:9knodename[user:password] TO targetfile
```


RTE-A FILE SYSTEM

HP 1000 computers that run NS-ARPA/1000 have RTE-A operating systems. The RTE-A operating system includes a record structure and categorizes files into certain "types." These file types are defined as follows:

- **Type 1.** These files have fixed-length records of 256 bytes and NFT assumes these files contain binary data.
- **Type 2.** These files also have fixed-length records, but the record length is defined by the user at file creation. NFT assumes these files contain binary data.
- **Type 3.** These files have variable-length records and are assumed to contain binary data. Type 3 files transferred to any system will produce a binary target file unless the ASCII option is specified in the file transfer.
- **Type 4.** These files also have variable-length records and contain ASCII data. Type 4 files are text files and may be altered with the RTE-A text editor, EDIT/1000.
- **Greater than Type 4.** File types 5, 6 and 7 contain binary data in different forms and have variable-length records. File types greater than 7 are user-defined. NOTE: Type 6 files are RTE-A program files and may not be moved in Interchange mode.

HP 1000 Logon Syntax

Logons for NS-ARPA/1000 systems take the form:

```
accountname [/password]
```

The *accountname* parameter is the logon name for the computer. The *password* parameter is optional if the logon has not password. If it is used, a slash (/) must separate the *accountname* from the *password* parameter. If the logon has no password, the slash is omitted.

HP 1000 File Name Syntax

The RTE-A file system has a hierarchical file structure. Files are catalogued in directories. Directories can also contain similar information about other directories, called **subdirectories**. Subdirectories have the same characteristics as directories; the term subdirectory means only that the directory is catalogued in the next higher level directory or subdirectory.

HP 1000 Source File

Each account or logon on an RTE-A system has a "home" or default logon working directory. This directory is automatically made available to the user when he or she logs on.

If the HP 1000 source file resides in the hierarchical file system, the file name syntax is as follows:

```
[/][directory/][directory/]... filename
```

If the *directory* parameter is omitted, the default logon working directory for the logon specified is used. If the initial slash (/) is omitted and the *directory* parameter is specified, the directory is assumed to be in the default logon working directory for the logon provided. The maximum file path name, including the file name, is 63 characters.

If the HP 1000 source file resides on a FMGR cartridge, the file name must be specified as follows:

```
filename::directory
```

The *filename* parameter may be a maximum of six characters long. The *directory* parameter may be up to two ASCII characters or a positive or negative integer.

For an extensive discussion of the HP 1000 file system, refer to the *RTE-A User's Manual*.

NOTE

NFT cannot be used to copy files to or from non-disc devices in the HP 1000.

SUPPORTED TARGET NODES

HP 1000 source files can be copied to HP 3000, HP 9000, VAX/VMS systems, and PC target nodes.

This section provides all the information you will need to know when copying an NS-ARPA/1000 source file to each supported target node. This information includes:

- Target file system information.
- File name and logon syntax.
- Interchange Format defaults.
- The effect of each Interchange Format option on the file translation process.

HP 3000 Target Node

HP 3000 computers utilize the Multiprogramming Executive (MPE) operating system. MPE files may contain ASCII or binary data, fixed or variable length records, and may be Relative I/O (RIO) or non-RIO files. RIO files are random access; non-RIO files are sequential access.

Table 3-1 introduces four abbreviations for MPE files with particular attributes:

TABLE 3-1. ABBREVIATIONS OF MPE FILE ATTRIBUTES.

Abbreviation	MPE File Attributes
VA	Variable-length records of ASCII data.
VB	Variable-length records of Binary data.
FA	Fixed-length records of ASCII data.
FB	Fixed-length records of Binary data.

HP 3000 Logon Syntax

Each field in the HP 3000 logon sequence must begin with a letter, contain only alphanumeric characters, and be at most eight characters long. The logon syntax is as follows:

```
username[/userpass].account[/acctpass][,groupname[/grouppass]]
```

HP 3000 File Name Syntax

Each field in an HP 3000 file name must begin with a letter, contain only alphanumeric characters, and be at most eight characters long. The group and account parameters default to the logon group and account. The file name syntax is as follows:

```
filename[/lockword][.group[.account]]
```

Wild Card Characters in NS-ARPA/1000 to NS/3000 NFT

The "at" (@) character can be used as a target file mask in the HP 3000 target file name parameter. The "at" character matches zero or more characters in the filename.

Because of differences between the HP 1000 and HP 3000 file systems and naming conventions, certain changes will take place when HP 1000 files are copied to HP 3000 target nodes using file masks:

- If an HP 1000 source file mask references one or more directories of files and the target node is an HP 3000, the files within those directories will be copied to the HP 3000, but directory information will *not* be included and the hierarchical structure of the files will not be preserved at the target node.
- If a target file mask is used when an HP 1000 source file name containing non-alphanumeric characters is copied to an HP 3000 target node, NFT will remove the non-alphanumeric characters from the target file name and NFT will issue a warning.
- If a target file mask is used when an HP 1000 source file name containing more than eight characters is copied to an HP 3000 target node, the file name will be truncated to eight characters at the target node and NFT will issue a warning. (HP 3000 file names cannot be longer than eight characters.)

All of the file mask features defined by the RTE-A file system can be used as source file masks in the source file parameter. NFT appends the "d" qualifier to the source file name whenever a file is copied using a wild card mask in the source file parameter. (If any directory matches the mask, the "d" qualifier causes all of the files in that directory to also match the mask. This feature can be overridden by using the "n" qualifier.)

HP 1000 to HP 3000 Interchange Format Defaults

Figure 3-1 illustrates the file translation process that occurs when an HP 1000 source file is copied to an HP 3000 target node. However, the following exceptions should be noted:

- You must override the Interchange Format default if the HP 1000 source file contains records longer than 4400 bytes. Files copied from an HP 1000 in Interchange Format cannot have records longer than 4400 bytes. Records can be truncated with the Record Size option. (The Record Size option is explained later in this subsection.)
- When a sparse (i.e., Type 1 or 2 with missing extents) HP 1000 source file is copied in Interchange Format, it does not remain sparse at the target node. The missing records are deleted from the target file. This cannot be overridden.
- When an HP 1000 source file is copied to an HP 3000 target node, an NFT warning message may be printed stating that source and target file attributes differ. This message does *not* indicate that the file copy process was unsuccessful. It is printed as a result of the record size negotiation process that occurs between the HP 1000 and HP 3000.
- If an HP 1000 source file containing variable length records is copied to an HP 3000, the record size of the file will be 4 bytes less than the record size of the source file rounded up to the nearest multiple of 256 bytes.
- On an HP 3000, direct access of variable length records is not allowed. As a result, NFT will return an error if you specify the Direct and Variable options in the same command or if the

source file has variable length records and the Direct option is specified but the Fixed option is not in effect.

- If you want to transfer a type 3 file from the HP 1000 to an ASCII file in the HP 3000, you must specify the ASCII option if the option is available on your initiator node. If the option is not available on your initiator node, you must convert the type 3 file to a type 4 file before transferring the file to the HP 3000.

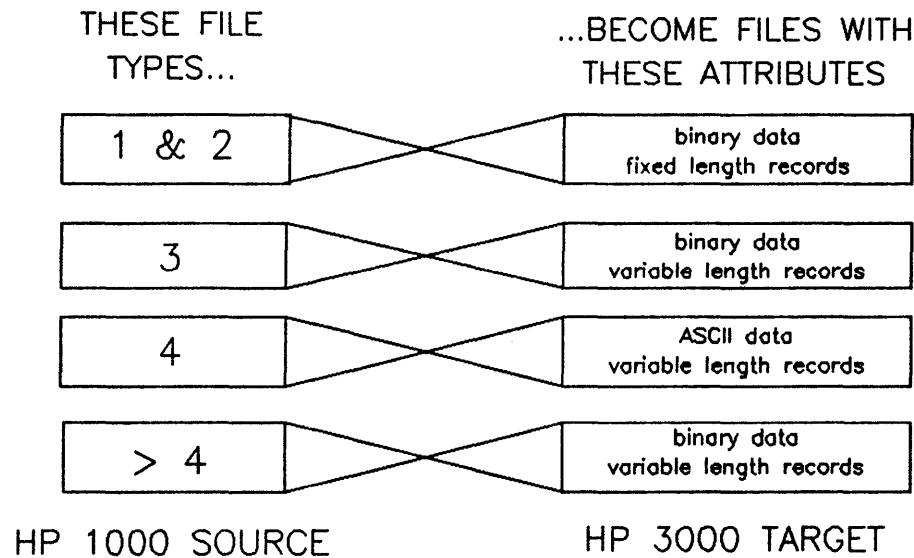


Figure 3-1. Interchange Format Defaults: HP 1000 to HP 3000.

Special Cases

When transferring records from HP 1000 to HP 3000, it is possible that the target file will be truncated if the source file's average record size is less than 15 bytes. This is due to the algorithm used to approximate the target file size.

To circumvent this problem, use the Record Size option to specify the length of the records in the target file. (However, this will not work if you also specify the Fixed option). If you need to get fixed records in the target file, first transfer the file without the Fixed option, then do a local HP 3000 DSCOPY using the Fixed option.

Interchange Format Options

Most of the Interchange Format options described in your NS user's manual may be used when copying HP 1000 source files to HP 3000 target nodes. However, your initiator system will determine which options are actually available for use. Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available at each initiator system.

Supported Interchange Format Options

The following Interchange Format options can be used when copying an HP 1000 source file to an HP 3000 target node. (Consult Table 2-1 to see if the options are available for your initiator system before using.)

- Append
- ASCII
- Binary
- Direct
- Fixed
- File Size
- Interchange
- Record Size
- Sequential
- Strip
- Variable

The following options cannot be used when copying an HP 1000 source file to an HP 3000 target node:

- Insert Character
- Search Character

Interchange Format Option Results

Table 3-2 describes the results of using these options when copying a file from an HP 1000 to an HP 3000.

TABLE 3-2. INTERCHANGE FORMAT OPTIONS: HP 1000 TO HP 3000.

Option Mnemonic	Description
Append	Appends the source file onto an existing target file. The attributes of the source file must match those of the target file. If the target file does not exist, NFT will return an error message.
ASCII	Specifies that target file records contain printable ASCII characters and that ASCII SPACE characters should be used as padding when creating fixed-length records. This option can be used in conjunction with the Strip option to indicate that spaces should be stripped from the ends of records.
Binary	Specifies that target file records contain binary information and that NULL characters be used as padding when creating fixed-length records. This option may be used in conjunction with the Strip option to indicate that NULL characters should be stripped from the ends of records.
Direct	Specifies that the target file will be organized to allow direct access. Do not use the Direct option if the target file is to be variable (either because the Variable option is specified or because the source file has variable-length records and the Fixed option is not in effect); this will cause an error to occur.
Fixed	Specifies that target file records should be created as fixed-length. Record size can be specified by using the Record Size option. The type of padding used can be specified using the ASCII or Binary options. If the Record Size option is <i>not</i> used, the record size at the target will be set to the length of the longest record in the source file. Shorter records may be padded (with ASCII SPACE characters if the file contains ASCII data, with NULL characters if it contains binary data).
File Size	Specifies how much space to allocate for the target file. The File Size option argument is the number of maximum size records in the target file. This option is specifically provided for use with target nodes whose files are not dynamically extendable (i. e., HP 3000).
Interchange	Causes the file to be copied using Interchange Format. This option has no effect in cross-system NFT since Interchange mode is the default.
Record Size	Specifies the record size of the target file in bytes. If fixed-length records are being produced, the Record Size option argument is interpreted as the size of each record. If variable-length records are being produced, the Record Size option argument limits the size of the largest record. Record size cannot be set larger than 4400 bytes. NFT will issue a warning if it must truncate records in order to execute this option.

TABLE 3-2. INTERCHANGE FORMAT OPTIONS: HP 1000 TO HP 3000 (cont).

Option Mnemonic	Description
Sequential	Causes the target file to be organized to allow sequential access. Records in the source file will be sent to the target node contiguously.
Strip	Strips any record padding from the ends of target file records. You can use this option to create variable-length records from fixed-length records. (Also see the Variable option.) The type of padding to strip is based on the type of the source file. For HP 1000 Type 4 files, SPACE characters are stripped. In other HP 1000 file types, NULL characters are stripped. You can use this option in conjunction with the Record Size option to truncate records. Records will be truncated before padding is stripped.
Variable	Specifies that target file records should be created as variable-length. The maximum size of a variable-length record may be given using the Record Size option.

File Mappings

Table 3-3 lists the options required to map a particular RTE-A file type into a particular kind of MPE file. The "Comments" column of the table describes important details about each transfer. Other options besides the required options can also be used.

Table 3-3. FILE MAPPING FROM HP 1000 SOURCE FILE TO HP 3000 TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
Type 1, Type 2	VA	ASCII, Variable	The default record length is determined by the MPE system; it is never less than the record length of the source file, and is commonly slightly larger. Any existing padding is transferred with the data.
Type 1, Type 2	VB	Variable	The default record length is determined by the MPE system; it is never less than the record length of the source file, and is commonly slightly larger. Any existing padding is transferred with the data.
Type 1, Type 2	FA	ASCII	The default record length for the target file is the same as the record length in the source file. Any existing padding is transferred with the data.

Table 3-3. FILE MAPPING FROM HP 1000 SOURCE TO HP 3000 TARGET (cont).

Source File Type	Target File Type	Options Required	Comments
Type 1, Type 2	FB	none	The default record length for the target file is the same as the record length in the source file. Any existing padding is transferred with the data. If you use the Record Size option, any necessary padding uses the ASCII NULL character.
Type 3, Type >4	VA	ASCII	The default record length is determined by the MPE system; it is never less than the record length of the source file, and is commonly slightly larger.
Type 3, Type >4	VB	none	The default record length is determined by the MPE system; it is never less than the record length of the source file, and is commonly slightly larger.
Type 3, Type >4	FA	ASCII, Fixed	The default record length for the target file is the same as the record length in the source file.
Type 3, Type >4	FB	Fixed	The default record length for the target file is the same as the record length in the source file. If you use the Record Size option, any necessary padding uses the ASCII NULL character).
Type 4	VA	none	The default record length is determined by the MPE system; it is never less than the record length of the source file, and is commonly slightly larger.
Type 4	VB	Binary	The default record length is determined by the MPE system; it is never less than the record length of the source file, and is commonly slightly larger.
Type 4	FA	Fixed	The default record length for the target file is the same as the record length in the source file. If you use the Record Size option, any necessary padding uses the ASCII SPACE character.
Type 4	FB	Binary, Fixed	The default record length for the target file is the same as the record length in the source file.

HP 9000 Target Node

HP 9000 computers have HP-UX operating systems. There is only one type of file on HP-UX systems: a stream of bytes. The bytes may represent ASCII characters or some type of binary data.

The HP-UX file system has no concept of records. However, many HP-UX programs interpret LINEFEED characters as delimiters for streams of bytes that are treated like "records". NFT uses this interpretation when the ASCII option of Interchange Format is used.

HP 1000 commands typed at the RTE Command Interpreter (CI) prompt are automatically upshifted. This might cause problems for HP 9000 producer or target nodes since HP-UX commands are case sensitive. To avoid this problem, enclose with grave accent (`) the part of the NS-ARPA/1000 DSCOPY command that must not be upshifted.

Refer to Appendix A for examples of DSCOPY.

NS/9000 Logon Syntax

Each parameter in an HP 9000 logon may contain any ASCII character, with the exception of the colon (:), and must not exceed eight characters. The following syntax is an NFT convention and may *not* be used to log on to an HP-UX operating system.

```
username[:password]
```

HP 9000 File Name Syntax

The syntax of an HP 9000 file name is as follows:

```
[/][dir1/dir2/.../dirn/]filename
```

The */dir1* and */dir2* parameters denote a directory within a path name. If the first character is a slash (/), the search starts from the root, otherwise the search starts from the user's default logon working directory. Directory and file names may contain any of the ASCII characters, except for null and slash (/), and can each be up to 255 characters long, with a maximum path length of 1023 characters.

NOTE

You cannot transfer directories to NS/9000 system using NFT.

Access Control Lists (ACLs) cannot be set using NFT.

There is no security code for HP 9000 source files; protection bits must control access. When the target node is an HP-UX system and the source node is an HP 1000 system, the target HP-UX file will be created with a default mode of 0666 octal (rw-rw-rw-). The owner will be determined by the login name you use. If *umask* is set to a value other than zero at the target node, the default file mode may be altered.

HP 1000 to HP 9000 Interchange Format Defaults

When an HP 1000 ASCII file is copied to an HP 9000 system, the HP 9000 appends a LINEFEED character (or the character specified by the Insert Character option in the Interchange Format) to the end of each incoming record. When an HP 1000 binary file is copied to an HP 9000 system, the LINEFEED character is *not* appended; instead, incoming records are placed directly in the target file. The LINEFEED character is the only difference between ASCII and binary files copied to an HP 9000; whether or not the incoming records are fixed or variable-length is of no consequence.

Figure 3-2 relates the HP 1000 file types to the attributes these files acquire when they are copied to an HP 9000 node. (The term "records" is used in the following figure to describe data located between delimiters in the HP-UX file.)

Files copied from an HP 1000 using Interchange Format are limited to a record size of 4400 bytes. Records larger than 4400 bytes must be truncated using the Record Size option. (The Record Size option is explained later in this subsection.)

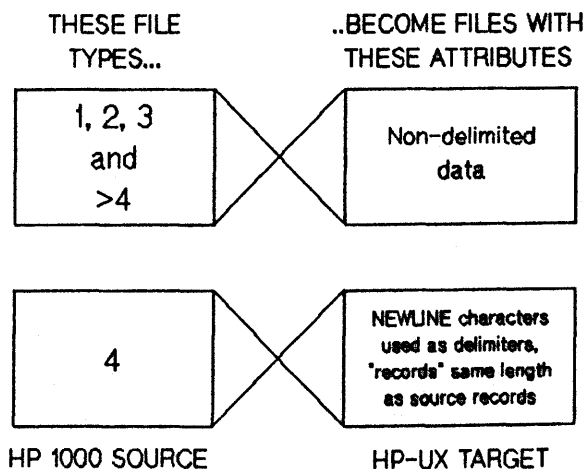


Figure 3-2. Interchange Format Defaults: HP 1000 to HP 9000.

Interchange Format Options

Most of the Interchange Format options described in your NS user's manual may be used when copying HP 1000 source files to HP 9000 target nodes. However, your initiator system will determine which options are actually available for use. Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available at each initiator system.

Supported Interchange Format Options

The following Interchange Format options can be used when copying an HP 1000 source file to an HP 9000 target node. (Consult Table 2-1 to see if the options are available for your initiator system before using.)

- Binary
- ASCII
- Fixed
- Record Size
- Insert Character
- Interchange
- Variable
- Strip
- Sequential

Unsupported Options

The following options cannot be used when copying an HP 1000 source file to an HP 9000 target node:

- Append
- Direct
- Search Character

Interchange Format Option Results

Table 3-4 describes the results of using these options when copying a file from an HP 1000 to an HP 9000.

TABLE 3-4. INTERCHANGE FORMAT OPTIONS: HP 1000 to HP 9000.

Option Mnemonic	Description
ASCII	Specifies that target file records contain printable ASCII characters and that ASCII SPACE characters should be used as padding when creating fixed-length records. This option can be used in conjunction with the Strip option to indicate that spaces should be stripped from the ends of records.
Binary	Source file records are placed directly in the target file. LINEFEED characters are not appended to the ends of records.
Fixed	Records are padded or truncated in order to make them all the same length. The length is determined by the length of the longest record in the source file, unless the Record Size option is used to specify another length. Padding, if required, will be with ASCII SPACE for Type 4 source files and with ASCII NULL characters for all other types of source files. This may be overridden by the use of the ASCII or Binary option.
Insert Character	Defines a record delimiter character to be inserted after each logical record in the source file before it is placed in the target file. The delimiter may be either be a single ASCII character or a decimal number from 0 to 255. If a decimal number is used, it must be preceded by a zero.
Interchange	Causes the file to be copied using Interchange Format. This is the default for cross-system NFT.
Record Size	<p>Specifies the maximum length of any record. Records longer than the specified record size will be truncated. For fixed length files (Type 1 or 2), or when the Fixed option is used, all records will be truncated or padded to this length. Padding, if required, will be with ASCII SPACE for Type 4 source files and with ASCII NULL characters for all other types of source files.</p> <p>If the Binary option is used in conjunction with the Record Size option, records longer than the Record Size argument are truncated and the resulting data is placed directly in the target file; LINEFEED characters are not appended to the ends of the record.</p> <p>Record size cannot be set longer than 4400 bytes.</p> <p>A warning message notifies you if any records are truncated during the transfer.</p>

TABLE 3-4. INTERCHANGE FORMAT OPTIONS: HP 1000 to HP 9000 (cont).

Option Mnemonic	Description
Sequential	Causes the target file to be organized to allow sequential access. Sparse source files of Type 1 or 2 will always be sent as a set of sequential records; however, a warning message will be generated if the Sequential option is not included.
Strip	Strips any record padding from the ends of records. You can use this option to create variable length records from fixed length records. To strip from Type 1 or 2 files, you must also specify the Variable option. The type of padding to strip is based on the type of source file. For HP 1000 type 4 files, ASCII SPACE characters are stripped. In other HP 1000 file types, ASCII NULL characters are stripped. You can use this option in conjunction with the Record Size option to truncate records. Records will be truncated before padding is stripped.
Variable	Specifies that target file records should be created as variable-length. This option has no effect on the target file unless it is used with the Strip option to strip ASCII blanks or NULL characters from the end of fixed-length records.

File Mappings

Table 3-5 lists the options required to map a particular RTE-A file type into an HP-UX file. The "Comments" column of the table describes important details about each transfer. Other options besides the required options can also be used.

TABLE 3-5. FILE MAPPING FROM HP 1000 SOURCE FILE TO HP 9000 TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
Type 1, Type 2, Type 3, Type >4	HP-UX	none	No record delimiters are appended. No padding or truncation occurs.
Type 4	HP-UX	none	Record delimiters are appended. No padding or truncation occurs.
Type 1, Type 2, Type 3, Type >4	HP-UX	Fix	No record delimiters are appended. The default record length for the target file is the same as the record length in the source file. If you use the Record Size option, any necessary padding uses the ASCII NULL character.

TABLE 3-5. FILE MAPPING FROM HP 1000 SOURCE TO HP 9000 TARGET (cont).

Source File Type	Target File Type	Options Required	Comments
Type 4	HP-UX	Fix	Record delimiters are appended. The default record length for the target file is the same as the record length in the source file. If you use the Record Size option, any necessary padding uses the ASCII SPACE character.
Type 4	HP-UX	Binary, Fix	No record delimiters are appended. The default record length for the target file is the same as the record length in the source file. If you use the Record Size option, any necessary padding uses the ASCII NULL character.
Type 3, Type >4	HP-UX	ASCII	Record delimiters are appended. No padding or truncation occurs.

DEC VAX Target Node

DEC VAX computers use the Record Management Services (RMS) file system. RMS provides several different types of files with various record formats. NS for the DEC VAX computer supports most RMS sequential file formats for transfers to VAX nodes. This manual uses the following abbreviations for RMS record formats:

TABLE 3-6. ABBREVIATIONS OF DEC VAX TARGET FILE TYPES.

Abbreviation	File Type
Stream_LF	Stream file with LINEFEED record terminators.
Stream_CR	Stream file with carriage return record terminators.
Stream	Stream file with record terminators other than carriage return. These terminators could be FF, VT, LF, or CRLF.
Variable	Variable record format file.
Fix	Fixed record format file.

VAX System Logon Syntax

NFT logons for the VAX nodes take the form:

```
username[:password]
```

The *username* parameter is a logon name for the computer. The *password* is optional if the logon has no password; if you use it, a colon (:) must separate the *username* from the *password* parameter.

VAX System File Name Syntax

Under VAX/VMS operating system, files reside in a hierarchy of **devices**, **directories** and **subdirectories**. The device is the physical unit (e.g., the disc) where the file exists. Files can be catalogued in any directory or subdirectory on the device. A subdirectory is a directory that is catalogued in the next higher level directory or subdirectory. Each account or logon on a VAX system has a default logon device and working directory, which is immediately available to the user when he or she logs on.

A file specification consists of a device name, followed by a directory path, followed by the file name. The general syntax of a VAX/VMS file specification is as follows:

```
device: [[directory]]file_name.type;version
```

- The *device* is a device name, which if omitted (along with the colon), defaults to the current device, commonly the logon device.
- The *directory* is a directory name, which defaults to the current working directory. The *directory* can be an absolute directory path, or it can be relative to the current working directory. To specify a subdirectory, include a period and a subdirectory name before the closing bracket. You can use a subdirectory several layers deep by chaining subdirectory names together with periods between.

NOTE

Square brackets - "[" and "]" - are required by the VAX/VMS file name syntax whenever a directory or subdirectory is specified. However, you can omit the brackets and the directory name when using the default directory. Also note that if a *device* is specified, there is no default *directory*. Instead of square brackets, you can also use angle brackets - "<" and ">."

For initiator nodes other than the VAX system, if the square brackets are used to specify the directory, the whole filename path should be enclosed in quotation marks to ensure that the DSCOPY command interprets the square brackets as part of the filename path. Refer to Appendix A for examples of DSCOPY.

- Either the *file_name* or the *type* (or both) must be present, and the period separating them is always present.
- The *version* is an integer, which may be omitted if you want to use the file with the highest version number. A file of the same name (and version, if specified) in the target directory will only be replaced if you use the Replace option. If you do not use the Replace option and a file of the same name exists in the target directory, a file with an incremented version number is created.

Each field except the version number in a file name must begin with a letter and contain only alphanumeric characters without spaces. The maximum length allowed for the entire VAX/VMS file specification is 39 characters.

NOTE

DSCOPY can be made to automatically execute the user's LOGIN.COM. This command file is used to define logical names that NFT can use, or it can set the default directory. To have DSCOPY execute the LOGIN.COM command file, you must set the DTRACE bit to 283. The manual, *Network Services for the DEC VAX Computer*, describes how to do this, and the details of how it affects your NFT transfers.

DECnet VAX System File Name Syntax

If your AdvanceNet for the DEC VAX Computer is also connected to DECnet, you can copy files to any DECnet node from an NS-ARPA/1000 node. This "network-to-network" transfer requires that you explicitly route the copy through the DEC VAX node that is running both AdvanceNet and DECnet. Therefore, in the target file *node_name* and *login* fields of the DSCOPY command you specify the VAX node that is running both AdvanceNet and DECnet. You then use the *file_path_name* field to specify the DECnet node, along with the appropriate device, directory, etc. as shown. The file path name syntax to specify a file on a DECnet node is as follows:

```
dec_node"username password"::device:[[directory]]file_name.type;version
```

- The *dec_node* field is the name of the DECnet node.
- There is a space separating *username* and *password*.
- The rest of the parameters are described in "VAX System File Name Syntax" above.

The string containing the user name and password is preceded and followed by quotation marks, and always contains a blank between the user name and password. Quotation marks and blanks have special meaning to the initiator; therefore, if your transfer requires the user name and password, you must use escape characters in the dscopy command. Table 3-7 below shows the required syntax at each initiator system.

For example, the basic syntax for HP 1000 DSCOPY is:

```
DSCOPY sourcefile[slogon]>snode TO targetfile[tlogon]>tnode
```

If you are copying a local HP 1000 file to a DecNet node through a VAX node that runs both DecNet and AdvanceNet, the proper syntax would be:

```
DSCOPY file1k to ``DecNetNode"user pw"::vaxfile``[user:pw]>VAXNode
```

where: ``DecNetNode"*user pw*"::*vaxfile*`` is the *targetfile* field of the DSCOPY command, and it specifies the logon and node for the DecNet as well as the target file residing on that node.

TABLE 3-7. SYNTAX AT EACH INITIATOR.

Initiator System	DECnet VAX File Name Syntax
HP 1000	<code>`dec_node"username password"::VAX/VMS_file_name`</code>
HP 3000	<code>`dec_node"username password"::VAX/VMS_file_name`</code>
HP 9000	<code>`dec_node"username password"::VAX/VMS_file_name`</code>
VAX/VMS	<code>"dec_node""username password""::VAX/VMS_file_name"</code>

For HP 1000, the apostrophes (') enclosing the grave accents (`) are needed around the above string to keep HP`1000 from putting a comma between the user name and password.

For additional information about the use of escape characters, see the Network Services manual for your initiator system.

File Access Modes

RMS on the VAX computer does not differentiate between ASCII and binary data types, but NFT assumes that all RTE-A files, except Type 4 files, contain binary data. Binary data are handled differently for VAX consumers. NFT on a VAX computer operates as follows:

- ASCII data:

When a VAX consumer receives ASCII data records (either from a Type 4 source file or whenever you use the ASCII option), it adds headers or terminating characters to each record depending on the target file type. This is called **Record Mode I/O**. This record formatting information is inserted for the exclusive use of RMS and is generally not available (or useful) to the user.

- BINARY data:

When a VAX consumer receives BINARY data records (either from a source file *other than* type 4, or whenever you use the Binary option), it concatenates them to form blocks of data. *No record formatting occurs*; the records from the source file are written to the target file without attention to record structure. This is called **Block Mode I/O**. Unless the target file is Fix, the absence of record formatting means you will not be able to access individual records in the target file.

When you transfer binary files (or use the Binary option), be sure you understand the type of file you are creating. In particular, when transferring binary data to an RMS Stream_LF, Stream_CR, Stream or Variable target file, note that the lack of record formatting in the target file may restrict the ability of RMS to access the records. You can use the ASCII option to force the VAX/VMS operating system to perform record formatting; then RMS can access each record in the target file. This will not alter the contents of the records (see ASCII option description in Table 3-8).

HP 1000 to VAX Computer Interchange Format Defaults

Files of Type 1 and Type 2 become Fix files on the DEC VAX target system. Files of Type 3, Type 4 and greater than Type 4 become Variable files on the VAX/VMS target system. Figure 3-3 relates the HP 1000 file types to the VAX/VMS file types.

Files copied from an HP 1000 using Interchange Format are limited to a record size of 4400 bytes. Records larger than 4400 bytes must be truncated using the Record Size option. (The Record Size option is explained later in this subsection.)

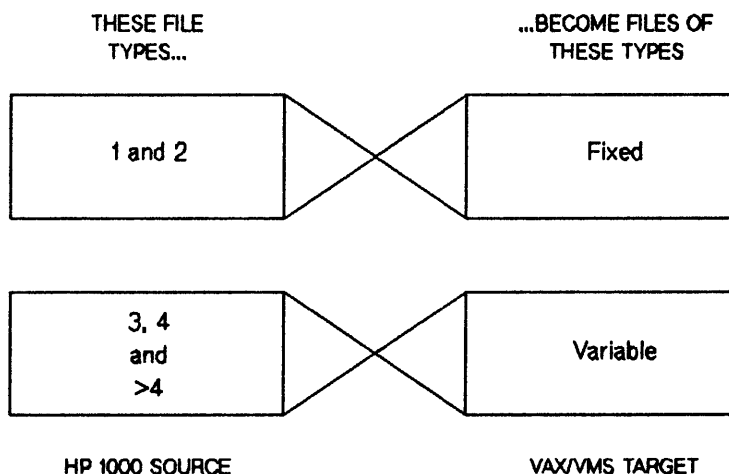


Figure 3-3. Interchange Format Defaults: HP 1000 to VAX/VMS System.

NOTE

When you transfer a Type 3 or greater than Type 4 file to a VAX/VMS target node, an RMS Variable target file is created using Block Mode I/O by default. This target file consists of data without record formatting information. You can use the ASCII option to force the VAX/VMS consumer to use Record Mode I/O; then record formatting information will occur when the file is created. The data in the records is not affected.

Interchange Format Options

Most of the Interchange Format options described in your NS user's manual may be used when copying HP 1000 source files to VAX/VMS target nodes. However, your initiator system will determine which options are actually available for use. Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available at each initiator system.

Supported Interchange Format Options

The following Interchange Format Options can be used when copying an HP 1000 source file to a VAX/VMS target node. (Consult Table 2-1 to see if the options are available for your initiator system before using.)

- ASCII
- Binary
- Fixed
- Insert Character
- Interchange
- Record Size
- Sequential
- Strip
- Variable

Unsupported Options

The following options cannot be used when copying an HP 1000 source file to a VAX/VMS target node:

- Append
- Direct
- File Size
- Search Character

Interchange Format Option Results

Table 3-8 describes the results of using these options when copying a file from an HP 1000 to a VAX/VMS target node.

TABLE 3-8. INTERCHANGE FORMAT OPTIONS: HP 1000 to VAX/VMS Computer.

Option Mnemonic	Description
ASCII	<p>Specifies that records are padded, if necessary, with the ASCII SPACE character. Also specifies that the file is transferred using Record Mode I/O (instead of Block Mode I/O). Record Mode I/O causes RMS to insert record headers or delimiters depending on the target file type. Using the ASCII option does not alter the contents of the records; however, because it forces the VAX/VMS operating system to use Record Mode I/O, RMS can access each record in the target file.</p>
Binary	<p>Specifies that any padding is done with the ASCII NULL character. Also specifies that the file is transferred using Block Mode I/O (instead of Record Mode I/O). With Block Mode I/O, RMS puts records directly into the target file without inserting headers or delimiters.</p> <p>NOTE: When you use the Binary option (or transfer binary files), be sure you understand the type of file you are creating. In particular, when transferring binary data to an RMS Stream_LF, Stream_CR, Stream or Variable target file, note that the lack of record formatting in the target file may restrict the ability of RMS to access the records (see ASCII).</p>
Fixed	<p>Creates a Fix type target file on the VAX/VMS node. You specify the length of the records with the Record Size qualifier, which defaults to the length of the longest record in the source file. Any record shorter than the specified or default length will be padded, and if you override the default length, truncation may occur.</p> <p>If the source file is Type 4 (or if you use the ASCII option), padding is done with the ASCII SPACE character.</p> <p>If the source file is <i>other</i> than Type 4 (or if you use the Binary option), padding is done with the ASCII NULL character.</p>
Record Size	<p>This option specifies the length of fixed-length records. Longer records will be truncated to the specified length, and shorter records will be padded to the specified length. The padding character is ASCII NULL if the source file is other than Type 4 or if you use Binary option, and ASCII SPACE if the source file is Type 4 or if you use the ASCII option.</p> <p>This option also specifies the maximum length of variable length records. Records longer than the specified length are truncated; shorter records are unaffected (also see Fixed).</p> <p>Record size cannot be set longer than 4400 bytes.</p> <p>A warning message notifies you if any records are truncated during the transfer.</p>

TABLE 3-8. INTERCHANGE FORMAT OPTIONS: HP 1000 to VAX Computer (cont).

Option Mnemonic	Description
Insert Character	<p>This option forces the VAX/VMS target file to be one of the following types:</p> <p>If the parameter is ASCII LF, a Stream__LF file is created.</p> <p>If the parameter is ASCII CR, a Stream__CR file is created.</p> <p>If the parameter is anything else, a Stream file is created, with CRLF appended to all records.</p> <p>If the source file has fixed-length records (i.e., Type 1 or Type 2), the Variable option <i>must</i> accompany the Insert Character option (see Variable).</p> <p>Note that the file will have the specified type regardless of what other options, if any, you use. However, if the source file is other than Type 4 or if you use the Binary option, the Insert Character option only determines the file type of the target file. No record formatting occurs.</p>
Interchange	Causes the file to be copied using Interchange Format. This option is the default for cross-system NFT.
Sequential	Causes the target file to be organized to allow sequential access. Sparse source files of Type 1 or 2 will always be sent as a set of sequential records.
Strip	<p>Strips any record padding from the ends of records. You can use this option to create variable length records from fixed length records. To strip from Type 1 or 2 files, you must also specify the Variable option.</p> <p>The type of padding to strip is based on the type of source file. For HP 1000 Type 4 files, ASCII SPACE characters are stripped. In other HP 1000 file types, ASCII NULL characters are stripped. You can use this option in conjunction with the Record Size option to truncate records. Records will be truncated before padding is stripped.</p>
Variable	Specifies that the target file be an RMS Variable type file (but see Insert Character). Note that unless the source file is Type 4 or you use the ASCII option, Block Mode I/O writes will be made to the file. In this case, the absence of record formatting may restrict the ability of RMS to access the file.

File Mappings

Table 3-9 lists the options required to map a particular RTE-A file type into a particular RMS sequential file. The "Comments" column of the table describes important details about each transfer. Other options besides the required options can also be used.

TABLE 3-9. FILE MAPPING FROM HP 1000 SOURCE FILE TO VAX/VMS TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
Type 1, Type 2,	Variable	Variable	Block mode writes are made to the target file, so no record formatting occurs (see note on next page).
Type 1, Type 2	Fix	none	The default record length for the target file is the same as the record length in the source file. If you use the Record Size option, any necessary padding uses the ASCII NULL character.
Type 1, Type 2	Stream__LF, Stream__CR, Stream	Variable, Insert Character	If the parameter to the Insert Character option is LF, the target file is Stream__LF. If the option parameter is CR, the target file is Stream__CR. If it is anything else, the target file is Stream. Block mode writes are made to the target file. The target file has the specified type but no record formatting occurs (see note on next page).
Type 3, Type >4	Stream__LF, Stream__CR, Stream	Insert Character	(Same as last comment.)
Type 3, Type >4	Variable	none	Block mode writes are made to the target file, so no record formatting occurs (see note on next page).
Type 3, Type 4, Type >4	Fix	Fixed	The default record length for the target file is the same as the record length in the source file. If you use the Record Size option, the character used for padding is determined as follows: If the source file is Type 4 or if you use the ASCII option, padding uses the ASCII SPACE character. Otherwise, padding is with the ASCII NULL character.

TABLE 3-9. FILE MAPPING FROM HP 1000 SOURCE TO VAX/VMS TARGET (cont).

Source File Type	Target File Type	Options Required	Comments
Type 4	Stream_LF, Stream_CR, Stream	Insert Character	If the parameter to the Insert Character option is LF, the target file is Stream_LF. If the option parameter is CR, the target file is Stream_CR. If it is anything else, the target file is Stream with CRLF record terminators.
Type 4	Variable	none	The default record length for the target file is the same as the record length in the source file.

NOTE

The absence of record formatting may (in the cases noted above) restrict the ability of RMS to access the file. Using the ASCII option forces the VAX/VMS system to include record formatting, and permits full RMS accessibility. This option does not affect the data in the records.

PC Target Node

NOTE

Whenever a PC takes part in a file transfer using NFT, it *must* be the initiator. To transfer a file to or from a PC, you must initiate the transfer from the PC. If you are unfamiliar with the DOS operating system, refer to Section 2 of this manual or the *NFT User's Guide, HP OfficeShare Network* manual to learn more about the syntax and usage of NFT on a PC.

DOS Stream Files

The DOS operating system maintains a single type of file: a stream of bytes, usually called a **stream file**. The bytes may represent any kind of data, including ASCII characters. DOS imposes no structure on the file; under this scheme, a file is a place to read or write bytes. Higher level concepts, like "records," are not present in DOS.

However, many PC applications use a convention for finding the beginning and end of logical records of data. Under this convention, the next record in the file consists of all the data up to *but not including* the next occurrence of an ASCII CR character followed by an ASCII LF character. This sequence of characters, ASCII CR-LF, is called the record delimiter or terminator.

Most applications use the record delimiter only as a marker, and discard it as soon as they read it. NFT uses this interpretation unless you employ the Binary option; then NFT does not use record delimiters at all. Details about the Interchange Format options are later in this subsection.

PC Logon Syntax

Because the PC is always local,* you do not need to give a node name; and because it is a single-user computer, there is no logon.

Therefore, when you specify the PC file, leave the *node_name* and *logon* fields blank. Supply only the file name, as described below.

PC File Name Syntax

<code>[device:][\][dir1\dir2\...\dirN\]filename.ext</code>
--

*Whenever a PC takes part in a file transfer using NFT, it *must* be the initiator. See the NOTE above.

The *device:* is used to specify a disc drive other than the active drive (or specify a shared disc directory on the server). Prefix the pathname with the device letter (for example, "A:").

The *dir1* through *dirN* parameters denote directories within a pathname. If the first character of the pathname is a backslash (\), the search starts from the root of the active or specified drive; otherwise the search starts from the current directory on the active or specified drive.

Directory and file names can always contain letters and numbers; check your DOS reference manual for restrictions if you want to use punctuation characters.

Each filename can be up to eight characters long; the optional extension has a three character maximum. The maximum pathname length is 63 characters.

NOTE

You cannot use NFT to transfer directories to a PC.

HP 1000 to PC Interchange Format Defaults

NFT assumes that:

- HP 1000 Type 4 files contain ASCII data
- HP 1000 files of all other Types contain binary data.

You can override these assumptions using the Binary and ASCII Interchange Format options.

So far as NFT is concerned, the only difference between an ASCII file and a binary file is whether or not record terminators are appended to incoming records:

ASCII	When an HP 1000 ASCII file is copied to a PC, the PC appends (by default) a record terminator to the end of each incoming record. The default terminator is an ASCII CR-LF sequence (see the Insert Character option).
Binary	When an HP 1000 Binary file is copied to a PC, the PC does <i>not</i> (by default) append record terminators. Instead, incoming records are copied directly into the target file without delimiters.

Records copied from an HP 1000 are limited to a maximum length of 4400 bytes. Records larger than this must be truncated with the Record Size option, as explained later in Table 3-10.

Figure 3-4 relates the HP 1000 file types to the attributes of the associated target file on a PC when no Interchange Format options are applied. The term "records" as used below means the data located between record delimiters in the PC file.

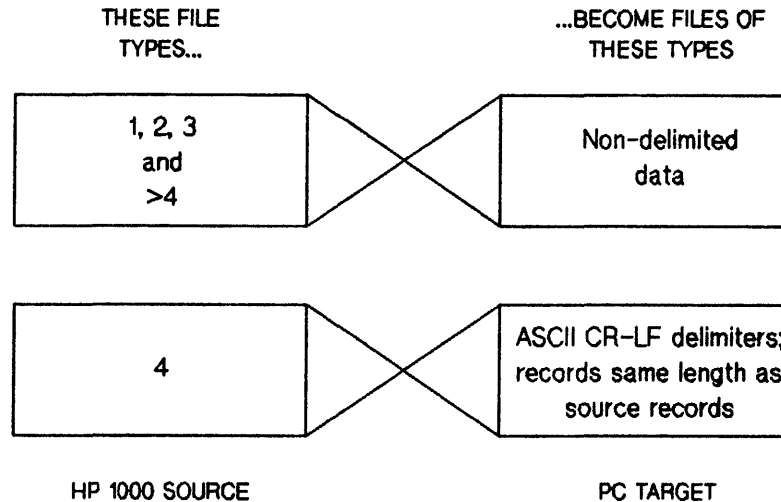


Figure 3-4. Interchange Format Defaults: HP 1000 to PC.

Interchange Format Options

Table 2-1 in Section 2 provides a quick reference of the Interchange Format syntax available at each initiator system.

Because the PC must always be the initiator,* the following supported PC options are available for NFT between a PC and another system.

Supported Interchange Format Options

The following Interchange Format options can be used when copying an HP 1000 source file to a PC:

- ASCII
- Binary
- Fixed
- Record Size
- Insert Character
- Search Character**
- Variable**

The options marked with two asterisks (**) can be used without causing an error, but they have no effect in transfers between these two machines. They are not included in the following tables.

*Whenever a PC takes part in a file transfer using NFT, it *must* be the initiator. See the NOTE at the beginning of this subsection.

Unsupported Options

The following options cannot be used when copying an HP 1000 file to a PC:

- Append
- Direct
- File Size
- Interchange
- Sequential
- Strip

Interchange Format Option Results

Table 3-10 describes the results of using these options when copying a file from an HP 1000 to a PC.

TABLE 3-10. INTERCHANGE FORMAT OPTIONS: HP 1000 to PC.

Option Mnemonic	Description
ASCII	Specifies that the ASCII SPACE character should be used as padding when creating fixed length records. Also specifies that record terminators should be appended in the target file. You cannot use the ASCII and Binary options together.
Binary	Source file records are placed directly in the target file without record delimiters. Use the Record Size option to change the record size; longer records are truncated, shorter records are unaffected. If you use this option, the Insert Character option is ignored. You cannot use the ASCII and Binary options together.
Fixed	<p>Source file records are truncated or padded, as necessary, so they are all the same length. The length defaults to the length specified in the source file, or you can specify it with the Record Size option.</p> <p>Records from a Type 4 source file (or any source file if you use the ASCII option) are padded, if necessary, with ASCII SPACE characters.</p> <p>If the file contains binary data (either it is not Type 4 or you are using the Binary option), records are padded, if necessary, with ASCII NULL characters.</p>

TABLE 3-10. INTERCHANGE FORMAT OPTIONS: HP 1000 to PC (cont).

Option Mnemonic	Description
Insert Character	<p>The argument to the Insert Character option specifies a single character that defines the record delimiter to use in the target file. The default record delimiter is the ASCII CR-LF pair; the Insert Character option defines only single-character delimiters. You can give the parameter either as a literal, or as the decimal ASCII equivalent, preceded by zero.</p>
Record Size	<p>The target file is formed with a maximum record length as defined by the Record Size argument. If the source file contains a record that is longer than the Record Size argument, the record is truncated and NFT prints a warning message. There is no limit to record size beyond the physical limitations of the system.</p> <p>If you also use the Fixed option, or if the source file is Type 1 or Type 2, shorter records are padded to the specified length. If the source file is Type 4, padding is with ASCII SPACE characters; if the source file is any other Type, padding is with ASCII NULL characters (see Fixed).</p> <p>If you use both the Binary and Fixed options along with the Record Size option, padding is always with ASCII NULL characters (see Binary and Fixed).</p>

File Mappings

Table 3-11 lists the options required to map a particular RTE-A file **Type** into a PC file with particular attributes. *Note that each pair occurs more than once in the table.* The "Comments" column of the table describes important details about each transfer. Other options (Insert Character, for example) can also be used.

TABLE 3-11. FILE MAPPING FROM HP 1000 SOURCE FILE TO PC TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
Type 1, Type 2, Type 3, Type >4	DOS	none	No record delimiters are appended. No padding or truncation occurs.
Type 1, Type 2, Type 3, Type >4	DOS	ASCII	Record delimiters are appended. No padding or truncation occurs.

TABLE 3-11. FILE MAPPING FROM HP 1000 SOURCE TO PC TARGET FILE (cont).

Source File Type	Target File Type	Options Required	Comments
Type 1, Type 2, Type 3, Type >4	DOS	Fixed	No record delimiters are appended. The default record length for the target file is the same as the record length in the source file. Any necessary padding uses the ASCII NULL
Type 1, Type 2, Type 3, Type >4	DOS	ASCII, Fixed	Record delimiters are appended. The default record length for the target file is the same as the record length in the source file. Any necessary padding uses the ASCII SPACE character.
Type 4	DOS	none	Record delimiters are appended. No padding or truncation occurs.
Type 4	DOS	Binary	No record delimiters are appended. No padding or truncation occurs.
Type 4	DOS	Fixed	Record delimiters are appended. The default record length for the target file is the same as the record length in the source file. Any necessary padding uses the ASCII SPACE character.
Type 4	DOS	Binary, Fixed	No record delimiters are appended. The default record length for the target file is 256. Any necessary padding uses the ASCII NULL character.

NOTE

The default delimiter character is an ASCII CR-LF pair.

MPE FILE SYSTEM

HP 3000 computers utilize the Multiprogramming Executive (MPE) operating system. MPE files may contain ASCII or binary data, fixed or variable length records, and may be Relative I/O (RIO) files or non-RIO files. (RIO files are random access, non-RIO files are sequential access.)

Table 4-1 introduces four abbreviations for MPE files with particular attributes:

TABLE 4-1. ABBREVIATIONS OF MPE FILE ATTRIBUTES.

Abbreviation	MPE File Attributes
VA	Variable-length records of ASCII data.
VB	Variable-length records of Binary data.
FA	Fixed-length records of ASCII data.
FB	Fixed-length records of Binary data.

HP 3000 Logon Syntax

Each field in the HP 3000 logon sequence must begin with a letter, contain only alphanumeric characters, and be at most eight characters long. The logon syntax is as follows:

```
username[/userpass].account[/acypass][,groupname[/grouppass]]
```

The parameter *userpass* is the user's password; *acypass* is the account password; and *grouppass* is the group password.

HP 3000 File Name Syntax

Each field in an HP 3000 file name must begin with a letter, contain only alphanumeric characters, and be at most eight characters long. The group and account parameters default to the logon group and account. The file name syntax is as follows:

```
filename[/lockword][.group[.account]]
```

SUPPORTED TARGET NODES

HP 3000 source files can be copied to HP 1000, HP 9000, DEC VAX systems, and PC target nodes.

This section provides all the information you will need to know when copying an HP 3000 source file to each supported target node. This information includes:

- Target file system information.
- File name and logon syntax.
- Interchange Format defaults.
- The effect of each Interchange Format option in the file translation process.

HP 1000 Target Node

HP 1000 computers that run NS/1000 have RTE-A operating systems. The RTE-A file system includes a record structure and categorizes files into certain "types". These file types are defined as follows:

- **Type 1.** These files have fixed-length records of 256 bytes and contain binary data.
- **Type 2.** These files also have fixed-length records, but the record length is defined by the user at file creation. They contain binary data.
- **Type 3.** These files have variable-length records and are assumed to contain binary data. Type 3 files transferred to any system will produce a binary target file unless the ASCII option is specified in the file transfer.
- **Type 4.** These files also have variable-length records and contain ASCII data. Type 4 files are text files and may be altered with the RTE-A text editor, EDIT/1000.
- **Greater than Type 4.** File types 5, 6 and 7 contain binary data in different forms and have variable-length records. File types greater than 7 are user-defined.

HP 1000 Logon Syntax

Logons for HP 1000 systems take the form:

accountname[/*password*]

The *accountname* parameter is the logon name for the computer. The *password* parameter is optional if the logon has no password. If it is used, a slash (/) must separate the *accountname* from the *password* parameter. If the logon has no password, the slash is omitted.

HP 1000 File Name Syntax

The RTE-A file system has a hierarchical file structure. Files are catalogued in directories. Directories can also contain similar information about other directories, called **subdirectories**. Subdirectories have the same characteristics as directories; the term subdirectory means only that the directory is catalogued in the next higher level directory or subdirectory. Each account or logon on an RTE-A system has a "home" or default logon working directory. This directory is automatically made available to the user when he or she logs on.

For more information on the RTE-A file system, refer to the *RTE-A User's Manual*.

HP 3000 Source File

If the HP 1000 file resides in the hierarchical file system, the file name syntax is as follows:

```
[/][directory/][directory/]...filename
```

If the *directory* parameter is omitted, the default logon working directory for the logon specified is used. If the slash (/) is omitted and the *directory* parameter is specified, the directory is assumed to be in the default logon working directory for the logon provided. The maximum file path name, including the file name, is 63 characters.

If the HP 1000 file you wish to access resides on an FMGR cartridge, the file name must be specified as follows:

```
filename::directory
```

The *filename* parameter may be a maximum of six characters. The *directory* parameter may contain up to two ASCII characters or a positive or negative integer.

HP 1000 File Masks and Wild Card Characters

The wild card character "at" (@) can be used in the target file parameter to form target file names; you cannot use any of the other RTE-A file mask features. NFT uses this wild card character as a target file mask in the same way as the RTE-A Command Interpreter.

If a target file mask is used when an HP 3000 source file is copied to an HP 1000 target node, the group and account parts of the file name will be truncated. For example, if the file MEMO.PUB.SYS is copied from an HP 3000 to an HP 1000, the file will become MEMO at the target node.

HP 3000 Wild Card Characters

The following wildcard characters can be used in the source file name parameter for source files that reside on HP 3000 nodes.

- @ Matches zero or more characters in any of the fields except the lockword.
- ? Matches exactly one alphanumeric character.
- # Matches exactly one numeric character.

NOTE

NFT cannot copy files to a non-disc device at the HP 1000.

HP 3000 to HP 1000 Interchange Format Defaults

Figure 4-1 illustrates the translation process that occurs when an HP 3000 source file is copied to an HP 1000 target node. However, the following exceptions should be noted:

- When an HP 3000 RIO file is copied in Interchange Format, the random access attribute of the file is not retained in the target file. This cannot be overridden.
- Image root and set files can be copied into an HP 1000 target node, but only if the correct file system code (-400 and -401, respectively) is provided and the HP 3000 logon used has system manager, privileged mode capability, or is the creator of the file. File system codes can be specified by using the NS/3000 File Code option. (The File Code option is explained in the *NS/3000 User/Programmer Reference Manual*.)
- HP 3000 message files cannot be copied into an HP 1000 target node.
- You must override the default record size if the HP 3000 source file contains records longer than 4400 bytes. Files copied to an HP 1000 in Interchange Format cannot have records longer than 4400 bytes. Records can be truncated with the Record Size option. (The Record Size option is explained later in this subsection.)

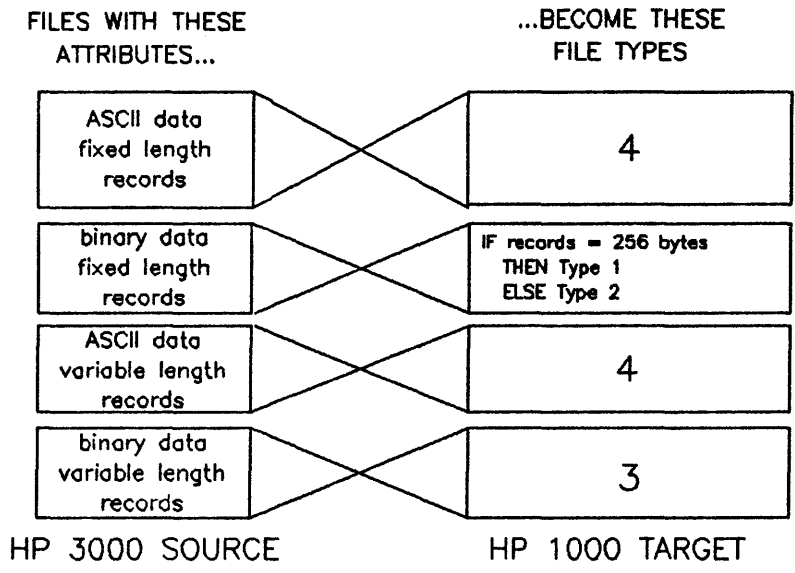


Figure 4-1. Interchange Format Defaults: HP 3000 to HP 1000.

Interchange Format Options

Most of the Interchange Format options described in your NS user's manual may be used when copying HP 3000 source files to HP 1000 target nodes. However, your initiator system will determine which options are actually available for use. Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available at each initiator system.

Supported Interchange Format Options

The following Interchange Format options can be used when copying an HP 3000 source file to an HP 1000 target node. (Consult Table 2-1 to see if the options are available for your initiator system before using.)

- ASCII
- Binary
- Direct
- Fixed
- File Size
- Interchange
- Record Size
- Sequential
- Strip
- Variable

Unsupported Options

The following options cannot be used when copying an HP 3000 source file to an HP 1000 target node:

- Append
- Insert Character
- Search Character

Interchange Format Option Results

Table 4-2 describes the results of using these options when copying a file from an HP 3000 to an HP 1000.

TABLE 4-2. INTERCHANGE FORMAT OPTIONS: HP 3000 to HP 1000.

Option Mnemonic	Description
ASCII	Specifies that records contain printable ASCII characters and that spaces should be used as padding when creating fixed length records. This option may be used in conjunction with the Strip option to indicate that spaces should be stripped from the ends of records.
Binary	Specifies that records contain binary information and that ASCII NULL characters should be used as padding when creating fixed length records. This option may be used in conjunction with the Strip option to indicate that ASCII NULL characters should be stripped from the ends of records.
Direct	Causes the HP 1000 target file to be created as a Type 1 or Type 2 file regardless of the data type. This option may only be used for files with fixed length records or when the Fixed option is also used.
Fixed	Specifies that target file records should be formed into fixed length records. Record size can be specified using the Record Size option and the type of padding used can be specified using the ASCII or Binary options. If the Record Size option is not used, the record size created in the target file is determined by the type of source file being copied. Shorter records will be padded (with ASCII SPACE characters if the file contains ASCII data, ASCII NULL characters if it contains binary data).
File Size	Specifies how much space to allocate for the target file. If the target file has fixed length records, the File Size argument is in records. If the target file has variable length records, the File Size argument indicates number of maximum size records. This option can be used instead of the HP 1000 file descriptor size parameter to specify the size of an HP 1000 target file.
Interchange	Causes the file or files to be copied using Interchange Format. This option is always used for cross-system file transfers.
Record Size	Specifies the record size of the target file in bytes. If fixed length records are being produced, this option specifies the size of each record. If variable length records are being produced, this option limits the size of the largest record and records may be padded or truncated. Record size cannot be set to longer than 4400 bytes. NFT will issue a warning if it must truncate records in order to execute this option.
Sequential	Causes RIO files to be transferred sequentially. "Missing" records will not be included in the target files. This will have the effect of compressing the file.

TABLE 4-2. INTERCHANGE FORMAT OPTIONS: HP 3000 to HP 1000 (cont).

Option Mnemonic	Description
Strip	Strips any record padding from the ends of records. This option only applies when the source file contains variable length records. You must use this option with the Variable option to strip files containing fixed length records, and create variable length records from fixed length records. The type of padding stripped depends on whether the file is ASCII or binary. You can use this option in conjunction with the Record Size option to truncate records. Records will be truncated before padding is stripped.
Variable	Specifies that source file records should be formed into variable length records. The maximum size of a variable length record may be given using the Record Size option.

File Mappings

Table 4-3 lists the options required to map a particular kind of MPE file into a particular RTE-A file type. The "Comments" column of the table describes important details about each transfer. Other options besides the required options can also be used. Consult Table 4-1 for the meaning of the source file abbreviations.

TABLE 4-3. FILE MAPPING FROM HP 3000 SOURCE FILE TO HP 1000 TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
VA	Type 1, Type 2	Binary, Fixed	The default record length for the target file is the same as the record length in the source file. If the record length is 256, the target file is Type 1; if the record length is not 256, the target file is Type 2. If you use the Record Size option, any necessary padding uses the ASCII NULL character.
VB	Type 1, Type 2	Fixed	The default record length for the target file is the same as the record length in the source file. If the record length is 256, the target file is Type 1; if the record length is not 256, the target file is Type 2. If you use the Record Size option, any necessary padding uses the ASCII NULL character.
FA	Type 1, Type 2	Binary	The default record length for the target file is the same as the record length in the source file. If the record length is 256, the target file is Type 1; if the record length is not 256, the target file is Type 2. If you use the Record Size option, any necessary padding uses the ASCII NULL character.

TABLE 4-3. FILE MAPPING FROM HP 3000 SOURCE TO HP 1000 TARGET (cont).

Source File Type	Target File Type	Options Required	Comments
FB	Type 1, Type 2	none	The default record length for the target file is the same as the record length in the source file. If the record length is 256, the target file is Type 1; if the record length is not 256, the target file is Type 2. If you use the Record Size option, any necessary padding uses the ASCII NULL character.
VA	Type 3	Binary	The default record length for the target file is the same as the record length in the source file.
VB	Type 3	none	The default record length for the target file is the same as the record length in the source file.
FA	Type 3	Binary, Variable	The default record length for the target file is the same as the record length in the source file.
FB	Type 3	Variable	The default record length for the target file is the same as the record length in the source file.
VA, FA	Type 4	none	The default record length for the target file is the same as the record length in the source file. If you use the Fixed option, any necessary padding uses the ASCII SPACE character.
VB, FB	Type 4	ASCII	The default record length for the target file is the same as the record length in the source file. If you use the Fixed option, any necessary padding uses the ASCII SPACE character.

HP 9000 Target Node

HP 9000 computers have HP-UX operating systems. There is only one type of file on HP-UX systems: a stream of bytes. The bytes may represent ASCII characters or some type of binary data.

The HP-UX file system has no concept of records. However, many HP-UX programs interpret LINEFEED characters as delimiters for streams of bytes that are treated like "records." NFT uses this interpretation when the ASCII option of Interchange Format is used.

HP 1000 commands typed at the RTE Command Interpreter (CI) prompt are automatically upshifted. This might cause problems for NS/9000 producer or target nodes since HP-UX commands are case sensitive. To avoid this problem, enclose with grave accents (`) the part of the NS/1000 DSCOPY command string that must not be upshifted.

Refer to Appendix A for examples of DSCOPY.

NS/9000 Logon Syntax

Each parameter in an HP 9000 logon may contain any ASCII character, with the exception of the colon (:) and must not exceed eight characters. The following syntax is an NFT convention and should *not* be used to log on to an HP-UX operating system except as part of a DSCOPY command.

```
username[:password]
```

HP 9000 File Name Syntax

The syntax of an HP 9000 file name is as follows:

```
[/][dir1/dir2/.../dirn/]filename
```

The *dir1* and *dir2* parameters denote a directory within a path name. If the first character is a slash (/), the search starts from the root; otherwise, the search starts from the user's default logon working directory. Directory and file names can contain any ASCII characters, except for null and slash (/), and can each be up to 255 characters long, with a maximum path length of 1023 characters.

NOTE

HP 9000 target file names cannot contain wild cards.

Access Control Lists (ACLs) cannot be set using NFT.

There are no lockwords or passwords for HP 9000 source files; protection bits control access. When the target node is an HP-UX system and the source node is an HP 3000 system, the target HP-UX file will be created with a default mode of 0666 octal (rw-rw-rw-). The owner will be the login name you use. (If *umask* is set to a value other than zero at the target node, the default file mode may be altered.)

HP 3000 to HP 9000 Interchange Format Defaults

Figure 4-2 relates HP 3000 file attributes to the HP 9000 file attributes. The term "records" is used in the following figure to describe data located between delimiters in the HP-UX file.

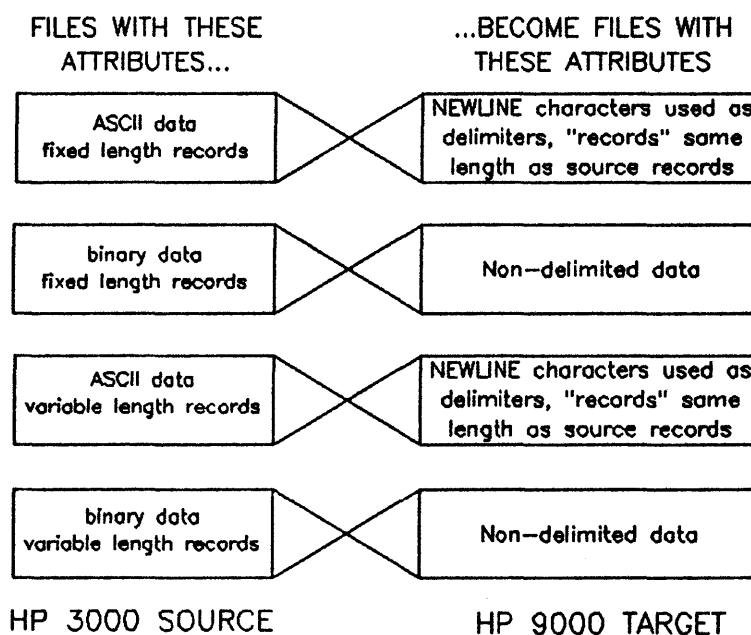


Figure 4-2. Interchange Format Defaults: HP 3000 to HP 9000.

Interchange Format Options

Most of the Interchange Format options described in your NS user's manual may be used when copying HP 3000 source files to HP 9000 target nodes. However, your initiator system will determine which options are actually available for use. Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available at each initiator system.

Supported Interchange Format Options

The following Interchange Format options can be used when copying an HP 3000 source file to an HP 9000 target node. (Consult Table 2-1 to see if the options are available for your initiator system before using.)

- ASCII
- Binary
- File Size
- Fixed
- Insert Character
- Interchange
- Record Size
- Sequential
- Strip
- Variable

Unsupported Options

The following options cannot be used when copying an HP 3000 source file to an HP 9000 target node:

- Append
- Direct
- Search Character

Interchange Format Option Results

Table 4-4 describes the results of using these options when copying a file from an HP 3000 to an HP 9000.

TABLE 4-4. INTERCHANGE FORMAT OPTIONS: HP 3000 to HP 9000.

Option Mnemonic	Description
ASCII	Specifies that records contain ASCII characters and that ASCII SPACE should be used as padding when creating fixed length record.
Binary	Source file records are placed directly in the target file. LINEFEED characters are not appended to the ends of records. No padding or truncation occurs.
File Size	Specifies how much space to allocate for the target file. The File Size argument specifies the number of records to allocate for the target file.
Fixed	<p>Source file records are truncated or padded (with ASCII SPACE characters), if necessary, so that they are the length of the source file. If the Record Size option is used, the records are truncated or padded to the size specified by the Record Size option. LINEFEED characters are appended to the ends of records.</p> <p>If the Binary option is used in conjunction with the Fixed option, records are padded with ASCII NULL characters, if necessary, so that they are the length of the source file. If the Record Size option is used, records are truncated or padded to the size specified by the Record Size option. LINEFEED characters are not appended to the ends of records.</p>
Insert Character	Defines a record delimiter character to be inserted after each logical record in the source file before it is placed in the target file. The delimiter character can be entered in literal form or in its equivalent decimal ASCII value. The ASCII value must be prefixed with a zero.
Interchange	Causes the file to be copied using Interchange Format. This option is the default for cross-system NFT.
Sequential	Causes RIO files to be transferred sequentially. "Missing" records will not be included in the target files. This will have the effect of compressing the file.
Strip	Strips any record padding from the ends of records. You can use this option to create variable length records from fixed length records. (Also see the Variable option.) The type of padding to strip is based on the type of source file. For HP 3000 ASCII files, ASCII SPACE characters are stripped. In HP 3000 binary files, ASCII NULL characters are stripped. You can use this option in conjunction with the Record Size option to truncate records. Records will be truncated before padding is stripped.

TABLE 4-4. INTERCHANGE FORMAT OPTIONS: HP 3000 to HP 9000 (cont).

Option Mnemonic	Description
Record Size	<p>Specifies the record size to be used in the target file. Source file records are truncated, if necessary. LINEFEED characters are appended to the ends of the resulting records. If the Record Size argument is omitted, it is set to 160 bytes for ASCII and 256 bytes for binary transfers. If it is zero (0) bytes, the record size is set to the length of the longest record in the source file.</p> <p>If the Binary option is used in conjunction with the Record Size option, records longer than the Record Size argument are truncated and the resulting data is placed directly in the target file; LINEFEED characters are not appended to the ends of the record.</p> <p>If the Fixed option is used in conjunction with the Record Size option, source file records are truncated or padded (with ASCII SPACE characters), if necessary, so that they are the Record Size argument bytes long. LINEFEED characters are appended to the ends of records.</p> <p>The maximum record size allowed at HP 9000 is $(2^{**}31) - 1$.</p> <p>A warning message notifies you if any records are truncated during the transfer.</p>
Variable	<p>Specifies that source file records should be formed into variable length records. The maximum size of a variable record may be given using the Record Size option.</p>

File Mappings

Table 4-5 lists the options required to map a particular MPE file into an HP-UX file with certain characteristics. The "Comments" column of the table describes important details about each transfer. Other options besides the required options can also be used.

TABLE 4-5. FILE MAPPING FROM HP 3000 SOURCE FILE TO HP 9000 TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
VA	HP-UX	none	Record delimiters are appended. If you use the Record Size option, records longer than the specified length are truncated, shorter records are unaffected.
VA	HP-UX	Binary	No record delimiters are appended. If you use the Record Size option, records longer than the specified length are truncated, shorter records are unaffected.

TABLE 4-5. FILE MAPPING FROM HP 3000 SOURCE TO HP 9000 TARGET (cont).

Source File Type	Target File Type	Options Required	Comments
VA	HP-UX	Fixed	Record delimiters are appended. The default record length for the target file is the same as the record length in the source file. If you use the Record Size option, any necessary padding uses the ASCII SPACE character.
VA	HP-UX	Fixed, Binary	No record delimiters are appended. The default record length for the target file is the same as the record length in the source file.
VB	HP-UX	none	No record delimiters are appended. If you use the Record Size option, records longer than the specified length are truncated; shorter records are unaffected.
VB	HP-UX	Fixed	No record delimiters are appended. The default record length for the target file is the same as the record length in the source file. If you use the Record Size option, any necessary padding uses the ASCII NULL character.
FA	HP-UX	none	Record delimiters are appended. The default record length for the target file is the same as the record length in the source file. If you use the Record Size option, any necessary padding uses the ASCII SPACE character.
FA	HP-UX	Binary	No record delimiters are appended. The default record length for the target file is the same as the record length in the source file.
FB	HP-UX	none	No record delimiters are appended. The default record length for the target file is the same as the record length in the source file. If you use the Record Size option, any necessary padding uses the ASCII NULL character.

DEC VAX Target Node

VAX/VMS computers use the Record Management Services (RMS) file system. RMS provides several different types of files with various record formats. NS for the DEC VAX computer supports most RMS sequential file formats for transfers to VAX/VMS nodes. This manual uses the following abbreviations for RMS record formats:

TABLE 4-6. ABBREVIATIONS OF VAX/VMS TARGET FILE TYPES.

Abbreviation	File Type
Stream_LF	Stream file with LINEFEED record terminators.
Stream_CR	Stream file with carriage return record terminators.
Stream	Stream file with record terminators other than carriage return. These terminators could be LF, VT, CRLF, or FF.
Variable	Variable record format file.
Fix	Fixed record format file.

VAX System Logon Syntax

Logons for the VAX nodes take the form:

```
username[:password]
```

The *username* parameter is a logon name for the computer. The *password* is optional if the logon has no password; if you use it, a colon (:) must separate the *username* from the *password* parameter. For password security, see the Network Services manual for the initiator system you are using.

VAX System File Name Syntax

Under the VAX/VMS system, files reside in a hierarchy of **devices**, **directories** and **subdirectories**. The device is the physical unit (e.g., the disc) where the file exists. Files can be catalogued in any directory or subdirectory on the device. A subdirectory is a directory that is catalogued in the next higher level directory or subdirectory. Each account or logon on a VAX system has a default logon device and working directory, which is immediately available to the user when he or she logs on.

A file specification consists of a device name, followed by a directory path, followed by the file name. The general syntax of a VAX/VMS file specification is as follows:

```
device: [[directory]]file_name.type;version
```

- The *device* is a device name which, if omitted (along with the colon), defaults to the current device, commonly the logon device.
- The *directory* is a directory name, which defaults to the current working directory. The *directory* can be an absolute directory path, or it can be relative to the current working directory. To specify a subdirectory, include a period and a subdirectory name before the closing bracket. You can use a subdirectory several layers deep by chaining subdirectory names together with periods between.

NOTE

Square brackets - "[" and "]" - are required by the VAX/VMS file name syntax whenever a directory or subdirectory is specified. However, you can omit the brackets and the directory name when using the default directory. Also note that if a device is specified, there is no default *directory*. Instead of square brackets, you can also use angle brackets - "<" and ">."

For all initiator nodes other than the VAX system, if the square brackets are used to specify the directory, the whole filename path specification should be enclosed in quotation marks to ensure that the DSCOPY command interprets the square brackets as part of the filename path. Refer to Appendix A for examples of DSCOPY.

- Either the *file_name* or the *type* (or both) must be present, and the period separating them is always present.
- The *version* is an integer, which may be omitted if you want to use the file with the highest version number. A file of the same name (and version, if specified) in the target directory will only be replaced if the Replace option is also used. If Replace is not used and a file of the same name exists in the target directory, a file with an incremented version number is created.

Each field, except the version number, file name must begin with a letter and contain only alphanumeric characters (including hyphen and underscore) without spaces.

NOTE

DSCOPY can be made to automatically execute the user's LOGIN.COM. This command file is used to define logical names that NFT can use, or it can set the default directory. To have DSCOPY execute the LOGIN.COM command file, you must set the DTRACE bit to 283. The manual, *Network Services for the DEC VAX Computer*, describes how to do this, and the details of how it affects your NFT transfers.

DECnet VAX System File Name Syntax

If your AdvanceNet VAX system is also connected to DECnet, you can copy files to any DECnet node from an HP 3000 node. This "network-to-network" transfer requires that you explicitly route the copy through the VAX node that is running both AdvanceNet and DECnet. Therefore, in the target file *node_name* and *login* fields of the *dscopy* command you specify the VAX node that is running both AdvanceNet and DECnet. You then use the *file_path_name* field to specify the DECnet node, along with the appropriate device, directory, etc. as shown. The file path name syntax to specify a file on a DECnet node is as follows:

<pre><i>dec_node</i>"<i>username password</i>"::<i>device</i>: [[<i>directory</i>]]<i>file_name.type;version</i></pre>
--

- The *dec_node* field is the name of the DECnet node.
- There is a space separating *username* and *password*.
- The rest of the parameters are described in "VAX System File Name Syntax" above.

The string containing the user name and password is preceded and followed by quotation marks, and always contains a blank. These characters have special meaning to the initiator; therefore, if your transfer requires an user name and password, you must use escape characters in the *dscopy* command. The following table 4-7 shows the required syntax at each initiator system.

For example, the basic syntax for HP 3000 DSCOPY is:

```
DSCOPY sourcefile:snode[slogon] TO targetfile:tnode[tlogon]
```

If you are copying a local HP 3000 file to a DecNet node through a VAX node that runs both DecNet and AdvanceNet, the proper syntax would be:

```
DSCOPY file3k to `DecNetNode"user pw"::vaxfile`:vaxnode[user:pw]
```

where: `DecNetNode"user pw"::vaxfile` is the *targetfile* field of the DSCOPY command, and it specifies the logon and node for the DecNet node as well as the target file name residing on that node.

TABLE 4-7. DECNET VAX SYSTEM FILE NAME SYNTAX.

Initiator System	DECnet VAX File Name Syntax
HP 1000	'`dec_node"username password"::VAX/VMS_file_name`'
HP 3000	'dec_node"username password"::VAX/VMS_file_name'
HP 9000	'dec_node"username password"::VAX/VMS_file_name'
VAX/VMS	"dec_node""username password""::VAX/VMS_file_name"

For HP 1000, apostrophes (') enclosing the grave accents (`) are needed around the above string to keep the HP 1000 from putting a comma between the user name and the password.

For additional information about the use of escape characters, see the *Network Services* manual for your initiator system.

Using the FILEPASSWORD Option with DECnet

When you use the **Filepassword** option for a transfer between a DECnet VAX computer and an HP 3000 system, you will receive a spurious prompt for a file lock word caused by the two colons (::) in the DECnet VAX file specification. In response to the spurious prompt, enter two colons (::). The transfer will then proceed normally.

File Access Modes

RMS on the VAX/VMS computer does not differentiate between ASCII and binary data types. Binary data are handled differently by VAX/VMS consumers. NFT on a VAX computer operates as follows:

- ASCII data:

When a VAX consumer receives ASCII data records (either from an ASCII source file, or whenever you use the ASCII option), it adds headers or terminating characters to each record depending on the target file type. This is called Record Mode I/O. This record formatting information is inserted for the exclusive use of RMS and is generally not available (or useful) to the user.

- BINARY data:

When a VAX consumer receives BINARY data records (either from a binary-data source file, or whenever you use the Binary option), it concatenates them to form blocks of data. *No record formatting occurs*; the records from the source file are written to the target file without attention to record structure. This is called Block Mode I/O. Unless the target is an RMS Fix file, the absence of record formatting means you will not be able to access individual records in the target file.

When you transfer binary files (or use the Binary option), be sure you understand the type of file you are creating. In particular, when transferring binary data to an RMS Stream_LF, Stream_CR, Stream or Variable target file, note that the lack of record formatting in the target file may restrict the ability of RMS to access the records. The ASCII option forces the VAX/VMS system to perform record formatting; then RMS can access each record in the target file. This does not alter the contents of the records. Note that the Binary and ASCII options are mutually exclusive, and that using either may change the pad character (see Table 4-9).

HP 3000 to VAX/VMS Interchange Format Defaults

Below are four abbreviations for MPE files with particular attributes. An MPE file of fixed-length records (FA or FB) copied to a VAX/VMS node becomes an RMS Fix file. An MPE file of variable-length records (VA or VB) becomes an RMS Variable file.

TABLE 4-8. ABBREVIATIONS OF MPE FILE ATTRIBUTES.

Abbreviation	MPE File Attributes
VA	Variable-length records of ASCII data.
VB	Variable-length records of Binary data.
FA	Fixed-length records of ASCII data.
FB	Fixed-length records of Binary data.

Figure 4-3 relates the VAX/VMS file types to the HP 3000 file attributes.

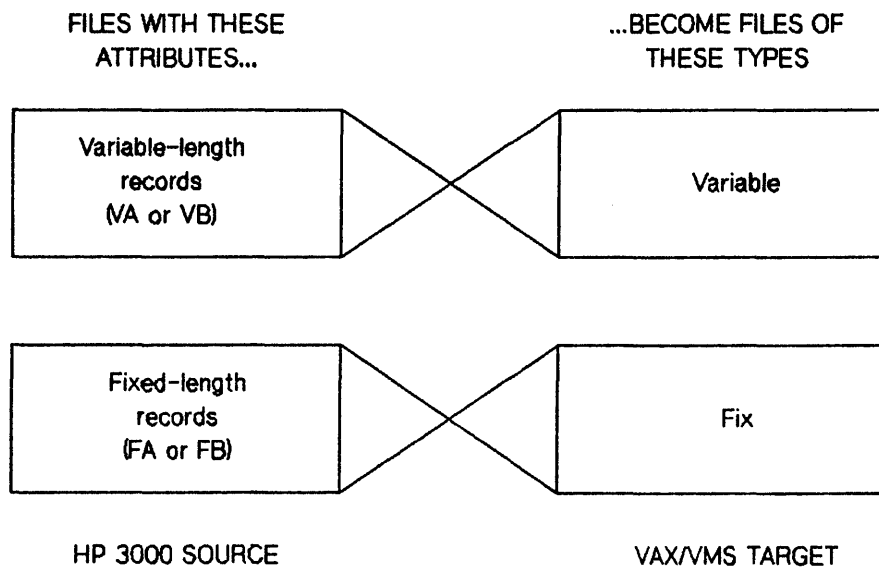


Figure 4-3. Interchange Format Defaults: HP 3000 to VAX/VMS.

NOTE

When you transfer a source file of binary data with variable length records (VB) to a VAX/VMS target node, by default an RMS Variable target file is created using Block Mode I/O. This target file does not contain RMS record formatting information, so the individual records are not accessible through Record Mode I/O. Using the ASCII option forces the VAX consumer to use Record Mode I/O, and it therefore inserts RMS record formatting information when the file is created. The data in the records is unchanged.

Interchange Format Options

Most of the Interchange Format options described in your NS user's manual may be used when copying HP 3000 source files to VAX/VMS target nodes. However, your initiator system will determine which options are actually available for use. Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available at each initiator system.

Supported Options

The following Interchange Format options can be used when copying an HP 3000 source file to a VAX/VMS target node. (Consult Table 2-1 to see if the options are available for your initiator system before using.)

- ASCII
- Binary
- File Size
- Fixed
- Insert Character
- Interchange
- Record Size
- Variable

Unsupported Options

The following Interchange Format options cannot be used when copying an HP 3000 source file to a VAX/VMS target node:

- Append
- Direct

- Search Character
- Sequential
- Strip

Interchange Format Option Results

Table 4-9 describes the results of using these options when copying a file from an HP 3000 to a VAX/VMS target node.

TABLE 4-9. INTERCHANGE FORMAT OPTIONS: HP 3000 to VAX/VMS.

Option Mnemonic	Description
ASCII	<p>Specifies that whenever padding is necessary, the ASCII SPACE character is used. This option also specifies that the file is written using Record Mode I/O (instead of Block Mode I/O). Record Mode I/O causes RMS to insert record headers or delimiters depending on the target file type. The ASCII and Binary options are mutually exclusive.</p>
Binary	<p>Specifies that records are padded, if necessary, with the ASCII NULL character. This option also specifies that the file is written using Block Mode I/O (instead of Record Mode I/O). With Block Mode I/O RMS puts records directly into the target file without inserting record headers or delimiters. The ASCII and Binary options are mutually exclusive.</p> <p>When you use the Binary option (or transfer binary files), be sure you understand the type of file you are creating. In particular, when transferring binary data to an RMS Stream__LF, Stream__CR, Stream or Variable target file, note that the lack of record formatting in the target file may restrict the ability of RMS to access the records (see ASCII).</p>
File Size	<p>Specifies how much space to allocate for the target file. The File Size argument specifies the number of records to allocate for the target file.</p>
Fixed	<p>Creates an RMS Fix type target file on the VAX/VMS node. The record length defaults to the length of the longest record in the source file. If you use the Record Size option source file records longer than the specified length may be truncated. Records shorter than the default or specified record length will be padded. The Fixed and Variable options are mutually exclusive.</p> <p>If the source file contains binary data (or if you use the Binary option), the padding character is ASCII NULL (see Binary). If the source file contains ASCII data (or if you use the ASCII option), the padding character is ASCII SPACE.</p>

TABLE 4-9. INTERCHANGE FORMAT OPTIONS: HP 3000 to VAX/VMS (cont).

Option Mnemonic	Description
Record Size	<p>This option specifies the length of fixed-length records. Source records longer than the specified length will be truncated; shorter records will be padded to the specified length (see Fixed).</p> <p>This option also specifies the maximum length of variable length records. Source records longer than the specified length will be truncated; shorter records will not be padded.</p> <p>A warning message notifies you if any records are truncated during a transfer.</p> <p>The maximum record size at the VAX/VMS system is 32,765 bytes.</p>
Insert Character	<p>This option forces the VAX/VMS target file to be one of the following types:</p> <p>If the parameter is ASCII LF, a Stream__LF file is created.</p> <p>If the parameter is ASCII CR, a Stream__CR file is created.</p> <p>If the parameter is anything else, a Stream file is created, with CRLF appended to all records.</p> <p>If the source file has fixed-length records, the Variable option must accompany the Insert Character option (see Variable).</p> <p>Note that the file will have the specified type regardless of what other options, if any, you use. However, if the data is binary or if you use the Binary option, the Insert Character option only determines the file type of the target file. No record formatting occurs (see Binary).</p>
Interchange	<p>Causes the file to be copied using Interchange Format. This option has no effect in cross-system NFT since Interchange mode is the default.</p>
Variable	<p>Specifies that the target file be an RMS Variable type file (but see Insert Character). Note that if the source file contains binary data or if you use the Binary option, Block Mode I/O writes will be made to the file. In this case, the absence of record formatting may restrict the ability of RMS to access the file. The Fixed and Variable options are mutually exclusive.</p>

File Mappings

Table 4-10 lists the options required to map a particular MPE file into a particular RMS sequential file type. The "Comments" column of the table describes important details about each transfer. Other options besides the required options can also be used.

TABLE 4-10. FILE MAPPING FROM HP 3000 SOURCE FILE TO VAX/VMS TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
VA, VB	Variable	none	If the source file is VB, Block Mode writes are made to the target file (see note below).
VA, VB	Stream_LF, Stream_CR, Stream	Insert Character	If the option parameter is LF, the target file is Stream_LF. If the parameter is CR, the target file is Stream_CR. If it is anything else, the target file is Stream. If the source file is VA, records in a Stream target file have CRLF record terminators. If the source file is VB, Block Mode writes are made to the target file; the target file has the specified type but no record formatting occurs (see note below).
VA, VB	Fix	Fixed	The default record length for the target file is the same as the record length in the source file. If you use the Record Size option to specify another length, the character used for padding is determined as follows: If the source file is VA or you use the ASCII option, padding is with the ASCII SPACE character. If the source file is VB or if you use the Binary option, any padding uses the ASCII NULL character.
FA, FB	Fix	none	The default record length for the target file is the same as the record length in the source file. If you use the Record Size option to specify another length, the character used for padding is determined as follows: If the source file is FA or you use the ASCII option, padding is with the ASCII SPACE character. If the source file is FB or if you use the Binary option, any padding uses the ASCII NULL character.

TABLE 4-10. FILE MAPPING FROM HP 3000 SOURCE TO VAX/VMS TARGET (cont).

Source File Type	Target File Type	Options Required	Comments
FA, FB	Stream_LF, Stream_CR, Stream	Variable, Insert Character	If the Insert Character option parameter is LF, the target file is Stream_LF. If it is CR, the target file is Stream_CR. If it is anything else, the target file is Stream. If the source file is FA, records in a Stream target file have CRLF record terminators. If the source file is FB, Block Mode writes are made to the target file; the target file has the specified type but no record formatting occurs (see note below).
FA, FB	Variable	Variable	If the source file is FB, Block Mode writes are made to the target file (see note below).

NOTE

The absence of record formatting may (in the cases noted above) restrict the ability of RMS to access the file. Using the ASCII option forces the VAX/VMS system to include record formatting, and permits full RMS accessibility. This option does not affect the data in the records.

PC Target Node

NOTE

Whenever a PC takes part in a file transfer using NFT, it *must* be the initiator. To transfer a file to or from a PC, you must initiate the transfer from the PC. If you are unfamiliar with the DOS operating system, refer to Section 2 of this manual or the *NFT User's Guide, HP OfficeShare Network* manual to learn more about the syntax and usage of NFT on a PC.

DOS Stream Files

The DOS operating system maintains a single type of file: a stream of bytes, usually called a **stream file**. The bytes may represent any kind of data, including ASCII characters. DOS imposes no structure on the file; under this scheme, a file is a place to read or write bytes. Higher level concepts, like "records," are not present in DOS.

However, many PC applications use a convention for finding the beginning and end of logical records of data. Under this convention, the next record in the file consists of all the data up to *but not including* the next occurrence of an ASCII CR character followed by an ASCII LF character. This sequence of characters, ASCII CR-LF, is called the **record delimiter** or **terminator**.

Most applications use the record delimiter only as a marker, and discard it as soon as they read it. NFT uses this interpretation unless you employ the Binary Interchange Format option; then NFT does not use record delimiters at all. Details about the Interchange Format options are discussed later in this subsection.

PC Logon Syntax

Because the PC is always local,* you do not need to give a node name; and because it is a single-user computer, there is no logon.

Therefore, when you specify the PC file, leave the *node_name* and *Logon* fields blank. Supply only the file name, as described below.

PC File Name Syntax

<code>[device:][\][dir1\dir2\...\dirN]filename.ext</code>

*Whenever a PC takes part in a file transfer using NFT, it *must* be the initiator. See the NOTE above.

The *device*: is used to specify a disc drive other than the active drive (or specify a shared disc directory on the server). Prefix the pathname with the device letter (for example, "A:").

The *dir1* through *dirN* parameters denote directories within a pathname. If the first character of the pathname is a backslash (\), the search starts from the root of the active or specified drive; otherwise the search starts from the current directory on the active or specified drive.

Directory and file names can always contain letters and numbers; check your DOS reference manual for restrictions if you want to use punctuation characters. Each filename can be up to eight characters long; the optional extension has a three character maximum. The maximum pathname length is 63 characters.

NOTE

You cannot use NFT to transfer directories to a PC.

HP 3000 to PC Interchange Format Defaults

So far as NFT is concerned, the only difference between an ASCII file and a binary file is whether or not record terminators are appended to incoming records:

When an HP 3000 ASCII file is copied to a PC, the PC appends (by default) a record terminator to the end of each incoming record. The default terminator is ASCII CR-LF (see the Insert Character option).

When an HP 3000 binary file is copied to a PC, the PC does *not* (by default) append record terminators. Instead, incoming records are copied directly into the target file without delimiters.

Figure 4-4 relates the HP 3000 file types to the attributes of the associated target file on a PC when no Interchange Format options are applied. The term "records" as used below means the data located between record delimiters in the PC file.

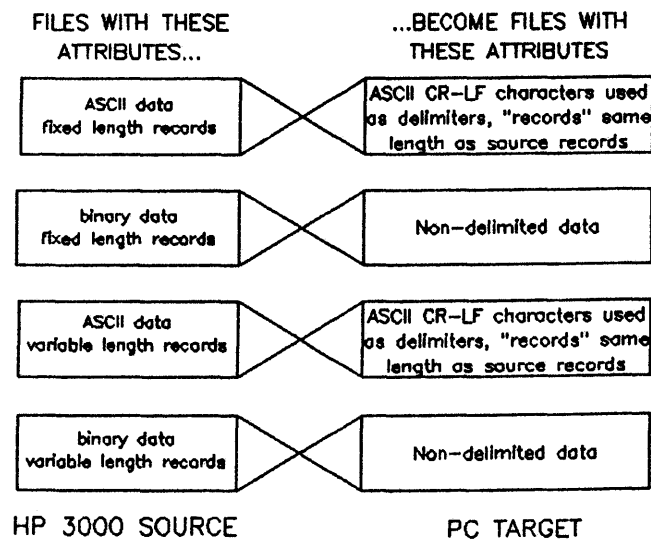


Figure 4-4. Interchange Format Defaults: HP 3000 to PC.

Interchange Format Options

Table 2-1 in Section 2 provides a quick reference of the Interchange Format options available for the PC.

Because the PC must always be the initiator,* the following options are available.

Supported Interchange Format Options

The following Interchange Format options can be used when copying an HP 3000 source file to a PC:

- ASCII
- Binary
- Fixed
- Record Size
- Insert Character
- Search Character**
- Variable**

The options marked with two asterisks (**) can be used without causing an error, but they have no effect in transfers between these two machines. They are not included in the following tables.

Unsupported Options

The following options cannot be used when copying an HP 3000 file to a PC:

- Append
- Direct
- File Size
- Interchange
- Sequential
- Strip

*Whenever a PC takes part in a file transfer using NFT, it *must* be the initiator. See the NOTE at the beginning of this subsection.

Interchange Format Option Results

Table 4-11 describes the results of using these options when copying a file from an HP 3000 to a PC.

TABLE 4-11. INTERCHANGE FORMAT OPTIONS: VAX/VMS to HP 1000.

Option Mnemonic	Description
ASCII	Specifies that the ASCII SPACE character should be used as padding when creating fixed length records. Also specifies that record terminators should be appended in the target file. You cannot use the ASCII and Binary options together.
Binary	Source file records are placed directly in the target file without record delimiters. Use the Record Size option to change the record size; longer records will be truncated, shorter records will be unaffected. If you use this option, the Insert Character option is ignored. You cannot use the ASCII and Binary options together.
Fixed	<p>Source file records are truncated or padded, as necessary, so they are all the same length. The length defaults to the record length of the source file; you can specify the length using the Record Size option.</p> <p>If the source is an ASCII file (or any file if you use the ASCII option), padding is with ASCII SPACE characters; if the source is a binary file (or any file if you use the Binary option), padding is with ASCII NULL characters.</p> <p>You cannot use the Fixed and Variable options together.</p>
Insert Character	The argument to the Insert Character option specifies a single character that defines the record delimiter to use in the target file. The default record delimiter is the ASCII CR-LF pair; the Insert Character option defines only single-character delimiters.
Record Size	<p>The target file is formed with a maximum record length as defined by the Record Size argument. If the source file contains a record that is longer than the Record Size argument, the record is truncated and NFT prints a warning message. There is no limit to record size beyond the physical limitations of the system.</p> <p>If you use the Binary option along with the Record Size option, no record terminators are inserted. (See Binary.)</p> <p>If you use the Fixed option along with the Record Size option, records in the source file are truncated or padded so that all records are the same length. If the source is an ASCII file (or any file if you use the ASCII option), padding is with ASCII SPACE characters; if the source is a binary file (or any file if you use the Binary option), padding is with ASCII NULL characters. (See Fixed.)</p>

File Mappings

Table 4-12 lists the options required to map a particular MPE file into a PC file with certain characteristics. The "Comments" column of the table describes important details about each transfer. Other options (Insert Character, for example) can also be used. This table is not exhaustive; other mappings are possible by using the Interchange Format options.

TABLE 4-12. FILE MAPPING FROM HP 3000 SOURCE FILE TO PC TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
VA	DOS	none	Record delimiters are appended. If you use the Record Size option, records longer than the specified length are truncated; shorter records are unaffected.
VA	DOS	Binary	No record delimiters are appended. If you use the Record Size option, records longer than the specified length are truncated, shorter records are unaffected.
VA	DOS	Fixed	Record delimiters are appended. The default record length for the target file is the same as the record length in the source file. Any necessary padding uses the ASCII SPACE character.
VA	DOS	Fixed, Binary	No record delimiters are appended. The default record length for the target file is the same as the record length in the source file. Any necessary padding uses the ASCII NULL character.
VB	DOS	none	No record delimiters are appended. If you use the Record Size option, records longer than the specified length are truncated; shorter records are unaffected.
VB	DOS	ASCII	Record delimiters are appended. If you use the Record Size option, records longer than the specified length are truncated, shorter records are unaffected.
VB	DOS	Fixed	No record delimiters are appended. The default record length for the target file is the same as the record length in the source file. Any necessary padding uses the ASCII NULL character.

TABLE 4-12. FILE MAPPING FROM HP 3000 SOURCE TO PC TARGET FILE (cont).

Source File Type	Target File Type	Options Required	Comments
FA	DOS	none	Record delimiters are appended. The default record length for the target file is the same as the record length in the source file. If you use the Record Size option, any necessary padding uses the ASCII SPACE character.
FA	DOS	Binary	No record delimiters are appended. The default record length for the target file is the same as the record length in the source file. If you use the Record Size option, any necessary padding uses the ASCII NULL character.
FB	DOS	none	No record delimiters are appended. The default record length for the target file is the same as the record length in the source file. If you use the Record Size option, any necessary padding uses the ASCII NULL character.
FB	DOS	ASCII	Record delimiters are appended. The default record length for the target file is the same as the record length in the source file. If you use the Record Size option, any necessary padding uses the ASCII SPACE character.

NOTE

The default delimiter character is an ASCII CR-LF pair.

HP-UX FILE SYSTEM

HP 9000 computers have HP-UX operating systems. There is only one type of file on HP-UX systems: a stream of bytes. The bytes may represent ASCII characters or some type of binary data.

The HP-UX file system has no concept of records. In fact, no structure is imposed on a file by the system, and no meaning is attached to its contents. However, many HP-UX programs interpret LINEFEED characters as delimiters for streams of bytes that are treated like "records". NFT uses this interpretation when the ASCII option of Interchange Format is used. HP 1000 commands typed at the RTE Command Interpreter (CI) prompt are automatically upshifted. This might cause problems for NS/9000 producer or target nodes since HP-UX commands are case sensitive. To avoid this problem, enclose with grave accents (`) the part of the NS/1000 DSCOPY command string that must not be upshifted.

Refer to Appendix A for examples of DSCOPY.

HP 9000 Logon Syntax

Each parameter in an HP 9000 logon may contain any ASCII character, with the exception of the colon (:), and must not exceed eight characters. The following syntax is an NFT convention and should not be used to log on to an HP-UX operating system except as part of a DSCOPY command.

```
username[:password]
```

HP 9000 File Name Syntax

The syntax of an HP 9000 file name is as follows:

```
[/][dir1/dir2/.../dirn/]filename
```

The *dir1* and *dir2* parameters denote a directory within a path name. If the first character is a slash (/), the search starts from the root; otherwise, the search starts from the default logon working directory. Directory and file names can contain any ASCII characters, except for null and slash (/), and can each be up to 255 characters long, with a maximum path length of 1023 characters.

NOTE

HP 9000 source file names cannot contain wild cards.

Access Control Lists (ACLs) cannot be set using NFT.

There are no lockwords or passwords for HP 9000 source files; protection bits control access.

SUPPORTED TARGET NODES

HP 9000 source files can be copied to HP 1000, HP 3000, VAX/VMS systems, and PC target nodes.

This section provides all the information you will need to know when copying an NS/9000 source file to each supported target node. This information includes:

- Target file system information.
- File name and logon syntax.
- Interchange Format defaults.
- The effect of each Interchange Format option on the file translation process.

HP 1000 Target Node

HP 1000 computers that run NS/1000 have RTE-A operating systems. The RTE-A file system includes a record structure and categorizes files into certain "types". These file types are defined as follows:

- **Type 1.** These files have fixed-length records of 256 bytes and contain binary data.
- **Type 2.** These files also have fixed-length records, but the record length is defined by the user at file creation. They contain binary data.
- **Type 3.** These files have variable-length records and are assumed to contain binary data. Type 3 files transferred to any system will produce a binary target file unless the ASCII option is specified in the file transfer. binary data.
- **Type 4.** These files also have variable-length records and contain ASCII data. Type 4 files are text files and may be altered with the RTE-A text editor, EDIT/1000.
- **Greater than Type 4.** File types 5 and 7 contain binary data in different forms and have variable-length records. File types greater than 7 are user-defined. Note that type 6 files are RTE-A program files and cannot be moved in Interchange mode.

HP 1000 Logon Syntax

Logons for NS/1000 systems take the form:

accountname[/password]

The *accountname* parameter is the logon name for the computer. The *password* parameter is optional if the logon has no password. If it is used, a slash (/) must separate the *accountname* from the *password* parameter. If the logon has no password, the slash is omitted.

HP 1000 File Name Syntax

The RTE-A file system found on NS/1000 systems has a hierarchical file structure similar to the HP-UX file system. Files are catalogued in directories. Directories can also contain similar information about other directories, called **subdirectories**. Subdirectories have the same characteristics as directories; the term subdirectories means only that the directory is catalogued in the next higher level directory or subdirectory. Each account or logon on an RTE-A system has a "home" or default working directory. This directory is automatically made available to the user when he or she logs on.

If the HP 1000 file resides in the hierarchical file system, the file name syntax is as follows:

```
[/][directory/][directory/]...filename
```

If the *directory* parameter is omitted, the default working directory for the logon specified is used. If the slash (/) is omitted and the *directory* parameter is specified, the directory is assumed to be in the default working directory for the logon provided. The maximum file path name, including the file name, is 63 characters.

If the HP 1000 file you wish to access resides on a FMGR cartridge, the file name must be specified as follows:

```
filename::directory
```

The *filename* parameter may be a maximum of six characters. The *directory* parameter may be up to two ASCII characters or a positive or negative integer.

For more information on the RTE-A file system, refer to the *RTE-A User's Manual*.

NOTE

NFT cannot be used to copy files to a non-disc device at the HP 1000.

HP 9000 to HP 1000 Interchange Format Defaults

Figure 5-1 relates the HP-UX file attributes to the HP 1000 file types. As shown in Figure 5-1, all files copied from an HP-UX system to an NS/1000 system become ASCII files with variable-length records (Type 4) at the target system by default. LINEFEED characters found in the source file are used to form records in the target file, but are not transmitted to the target node.

If you want to create a binary file (Type 1, 2, 3 or greater than Type 4) at the target node, you must override the Interchange Format default attributes by explicitly specifying the Binary option. The Binary option also forces the files to be sent in fixed length records. The default length is 256 bytes. Similarly, to create a target file with fixed-length (rather than variable-length) records at the target node, you must override the Interchange Format default by explicitly specifying the Fixed option. If Binary is used, delimiter characters are transferred.

In addition, because NFT uses LINEFEED characters found in HP-UX source files as delimiters when creating logical records in HP 1000 target files, you must use the Search Character option to define a different delimiter character if your source file does not use LINEFEED characters.

The Binary, Fixed and Search Character options are explained later in this subsection.

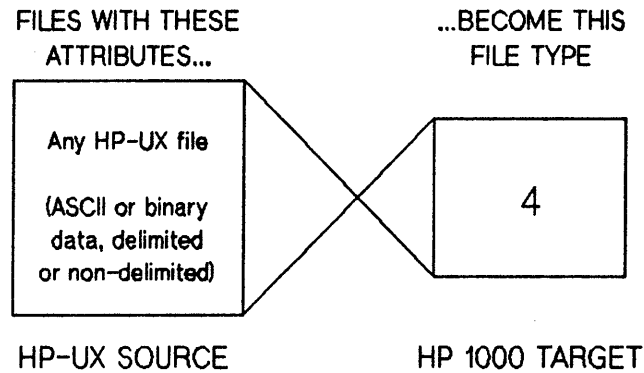


Figure 5-1. Interchange Format Defaults: HP 9000 to HP 1000.

NOTE

You must override the Interchange Format default attributes if the HP 9000 source file contains records longer than 4400 bytes. Files copied to an HP 1000 cannot have records longer than 4400 bytes. Records can be truncated with the Record Size option.

Interchange Format Options

Most of the Interchange Format options described in your NS user's manual may be used when copying HP 9000 source files to HP 1000 target nodes. However, your initiator system will determine which options are actually available for use. Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available at each initiator system.

Supported Interchange Format Options

The following Interchange Format options can be used when copying an HP 9000 source file to an HP 1000 target node. (Consult Table 2-1 to see if the options are available for your initiator system before using.)

- ASCII
- Binary
- File Size
- Fixed
- Interchange
- Record Size
- Search Character

HP 9000 Source File

- Sequential
- Variable

Unsupported Options

The following options cannot be used when copying an HP 9000 file to an HP 1000 target node:

- Append
- Direct
- Insert Character
- Strip

Interchange Format Option Results

Table 5-1 describes the results of using these options when copying a file from an HP 9000 to an HP 1000.

TABLE 5-1. INTERCHANGE FORMAT OPTIONS: HP 9000 to HP 1000.

Option Mnemonic	Description
ASCII	Specifies that target file records contain printable ASCII characters and that ASCII SPACE characters should be used as padding when creating fixed length records.
Binary	Source file data is divided into records every 256 bytes before being placed in the target file. LINEFEED characters found in the source file are considered part of the data. You must use this option if you want to create a binary target file. The resulting HP 1000 file type is Type 3. The record size can be change from 256 bytes by using the Record Size option.
File Size	Specifies how much space to allocate for the target file. If the target file will have fixed length records, the File Size argument is in records. If the target file will have variable length records, the File Size argument indicates number of maximum size records.

TABLE 5-1. INTERCHANGE FORMAT OPTIONS: HP 9000 to HP 1000 (cont).

Option Mnemonic	Description
Fixed	<p>Data between LINEFEED delimiters in the source file is truncated or padded (with ASCII SPACE characters), if necessary, so that the data is 160 bytes long. The resulting data is then formed into records and placed in the target file. The resulting HP 1000 file type is Type 4.</p> <p>The Binary option can be used in conjunction with the Fixed option to create a binary target file. The conversion process is the same as when the Binary option is used separately except that the last record in the target file is padded with ASCII NULL characters if it is less than 256 bytes long. The resulting HP 1000 file type is Type 1. The record size can be changed from 256 bytes by using the Record Size option.</p>
Interchange	<p>Causes the file to be copied using Interchange Format. This option is the default for cross-system file transfer.</p>
Record Size	<p>The target file is formed with a maximum record length as defined by the Record Size argument. If the source file contains a record (i.e., data between LINEFEED delimiters) that is longer than the Record Size argument, the record will be truncated and NFT will print a warning message. The resulting HP 1000 file type is Type 4. Record size cannot be set to greater than 4400 bytes.</p> <p>If the Binary option is specified in conjunction with the Record Size option, LINEFEED characters found in the source file are considered part of the data, and no truncation occurs. The resulting HP 1000 file type is Type 3.</p> <p>If the Fixed option is used in conjunction with the Record Size option, data between LINEFEED characters is truncated or padded if necessary, so that it is the Record Size argument number of bytes long. Padding is with ASCII SPACE characters unless the Binary option is used, in which case padding is with ASCII NULL characters.</p> <p>If the Binary and Fixed options are both used in conjunction with the Record Size option, LINEFEED characters found in the source file are considered part of the data, and the last record is padded with ASCII NULL characters if it is less than the Record Size argument. The resulting HP 1000 file type is Type 2.</p>
Search Character	<p>The Search Character option specifies a search character that is used to divide source file data into logical records in the target file.</p>
Sequential	<p>Causes the target file to be organized to allow sequential access. Records in the source file will be sent to the target node contiguously.</p>
Variable	<p>Specifies that target file records should be created as variable length. The maximum size of a variable length record may be given using the Record Size option.</p>

File Mappings

Table 5-2 lists the options required to map an HP-UX file with particular attributes into a particular Type of RTE-A file. The "Comments" column of the table describes important details about each transfer. Note that some pairs occur more than once in the table. Other options besides the required options can also be used.

TABLE 5-2. FILE MAPPING FROM HP 9000 SOURCE FILE TO HP 1000 TARGET NODE.

Source File Type	Target File Type	Options Required	Comments
HP-UX	Type 4	none	Use the Search Character option to extract logical records that are delimited by a character other than ASCII LF.
HP-UX	Type 4	Fixed	The default record length is 160 bytes. Truncation or padding (with the ASCII SPACE character) may occur.
HP-UX	Type 4	Record Size	Truncation may occur if a source record is longer than the specified length. No padding occurs.
HP-UX	Type 3	Binary	All bytes (including any record delimiters) are considered data and transferred. The default record length is 256 bytes. No padding or truncation occurs.
HP-UX	Type 1, Type 2	Binary, Fixed	All bytes (including any record delimiters) are considered data and transferred. The default record length is 256 bytes. No truncation occurs. The last record may be padded if the number of bytes in the file is not evenly divisible by the record length. The default target file is Type 1; if you use the Record Size option and specify a length other than 256, the target file is Type 2.

HP 3000 Target Node

HP 3000 computers utilize the Multiprogramming Executive (MPE) operating system. MPE files may contain ASCII or binary data, fixed or variable length records, and may be Relative I/O (RIO) files or non-RIO files. (RIO files are random access, non-RIO files are sequential access.)

Table 5-3 introduces four abbreviations for MPE files with particular attributes:

TABLE 5-3. ABBREVIATIONS OF MPE FILE ATTRIBUTES.

Abbreviation	MPE File Attributes
VA	Variable-length records of ASCII data.
VB	Variable-length records of Binary data.
FA	Fixed-length records of ASCII data.
FB	Fixed-length records of Binary data.

HP 3000 Logon Syntax

Each field in the HP 3000 logon sequence must begin with a letter, contain only alphanumeric characters, and be at most eight characters long. The logon syntax is as follows:

```
username[/userpass].account[/acctpass][,groupname[/grouppass]]
```

HP 3000 File Name Syntax

Each field in an HP 3000 file name must begin with a letter, contain only alphanumeric characters, and be at most eight characters long. The group and account parameters default to the logon group and account. The file name syntax is as follows:

```
filename[/lockword][.group[.account]]
```

The "at" character (@) can be used as a target file wild card in the HP 3000 target file name parameter. The "at" character (@) can only be used stand alone; for example @ is allowed, but not T@. The "at" character matches zero or more characters in the file name. In this case, the group and account will be

those from the logon. This will result in a target file name with the same name as the source file. However, the resulting HP 3000 target file name will be truncated to eight characters, and non-alphanumeric characters will be stripped from it. NFT will issue a warning indicating this name change.

HP 9000 to HP 3000 Interchange Format Defaults

Any file copied from an HP 9000 defaults to an ASCII file with variable-length records at the target node. LINEFEED characters found in the source file are used to form records in the target file, but are not transmitted to the target node. An ASCII data, variable-length record file is created at the target node regardless of whether the HP-UX source file contains ASCII or binary data. Figure 5-2 relates the HP-UX file attributes to the HP 3000 file attributes.

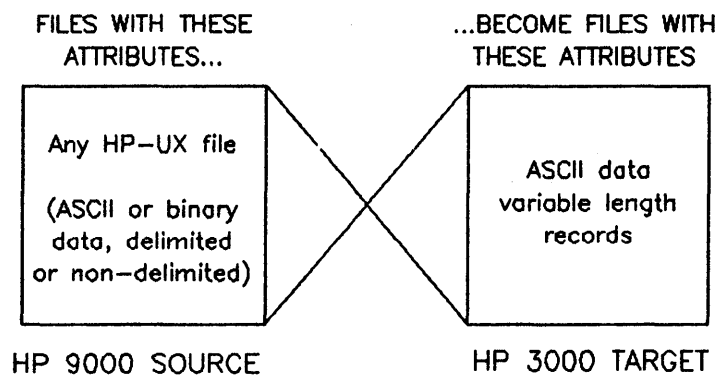


Figure 5-2. Interchange Format Defaults: HP 9000 to HP 3000.

Special Cases

In certain unusual situations, file transfers between HP 9000s and HP 3000s may produce unexpected results. These situations and their causes are described below.

Before an HP 9000 HP-UX file is copied to an HP 3000 system, the HP 3000 allocates space for the target file. To determine how much space to preallocate, the HP 3000 must know the maximum file size (in number of logical records). This information is provided by the HP 9000. Unless the length of the target file's records are explicitly specified by the use of the Record Size option, the HP 9000 HP-UX system approximates the file size (in number of logical records) by dividing the number of bytes in the file by 15. If the target file's record length is specified with the Record Size option, the file size will be the total number of bytes in the source file divided by the specific record length.

If the file size approximation is used, one of the following conditions might occur:

- A transfer of an HP-UX ASCII file to an HP 3000 may result in inadequate preallocated space if the file's average record size is less than 15 bytes. If not enough space has been allocated, the transfer will be unsuccessful, and NFT will print an error message. You can circumvent this failure by using the Record Size option to specify the length of the records in the target file. (However, this will not work if you also specify the Fixed option.)

If you need to get fixed records in the target file, you should transfer the file from the HP 9000 to the HP 3000 without the Fixed option, then execute a local DSCOPY on the HP 3000 using the Fixed option.

- A transfer of an HP-UX ASCII file may result in records being truncated if the source file's records are too long for the preallocated target file record size. The transfer does not fail, but a warning is displayed stating that records have been truncated to fit the target file record size. You can circumvent this failure by specifying a Record Size option value that is larger than the largest record in the source file.
- Another special case occurs if you use the Record Size option to specify an odd record length for the target file when copying a file from an HP 9000 to an HP 3000. If you specify an odd record length, the HP 3000 assigns one more byte to the record size. If record truncation is required, it occurs based on the odd record length specified but target file records may have an extra byte appended to them. If the file is later copied back to an HP-UX system, this extra byte is stripped off, and a warning is displayed stating that the source and target file attributes differ.

NOTE

When an HP 9000 file is copied to an HP 3000, an NFT warning message may be printed stating that the source and target file attributes differ. This message does not indicate that the file copy process was unsuccessful. It is printed as a result of the record size negotiation process that occurs between the HP 9000 and HP 3000.

Interchange Format Options

Most of the Interchange Format options described in your NS user's manual may be used when copying HP 9000 source files to HP 3000 target nodes. However, your initiator system will determine which options are actually available for use. Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available at each initiator system.

Supported Interchange Format Options

The following Interchange Format options can be used when copying an HP 9000 file to an HP 3000 target node. (Consult Table 2-1 to see if the options are available for your initiator system before using.)

- Append
- ASCII
- Binary
- File Size
- Fixed
- Interchange

HP 9000 Source File

- Record Size
- Search Character
- Sequential
- Variable

Unsupported Options

The following options cannot be used when copying an HP 9000 file to an HP 3000 target node:

- Direct
- Insert Character
- Strip

Interchange Format Option Results

Table 5-4 describes the results of using these options when copying a file from an HP 9000 to an HP 3000.

TABLE 5-4. INTERCHANGE FORMAT OPTIONS: HP 9000 to HP 3000.

Option Mnemonic	Description
Append	Appends the source file onto an existing target file. The attributes of the source file must match those of the target file. If the target file does not exist, NFT will return an error message.
ASCII	Specifies that target file records contain printable ASCII characters and that ASCII SPACE characters should be used as padding when creating fixed length records. This option can be used in conjunction with the Strip option to indicate that spaces should be stripped from the ends of records.
Binary	Specifies that target file records contain binary information and that ASCII NULL characters be used as padding when creating fixed length records. This option may be used in conjunction with the Strip option to indicate that ASCII NULL characters should be stripped from the ends of records.

TABLE 5-4. INTERCHANGE FORMAT OPTIONS: HP 9000 to HP 3000 (cont).

Option Mnemonic	Description
Fixed	Specifies that target file records should be created as fixed length. Record size can be specified by using the Record Size option. The type of padding used can be specified using the ASCII or Binary options. If the Record Size option is not used, the record size created in the target file is determined by the type of source file being copied. The record size will default to 160 bytes if the source file contains ASCII data, 256 bytes if it contains binary data.
File Size	Specifies how much space to allow for the target file. If the target file will be created with fixed length records, the File Size option argument is interpreted as the number of records in the target file. If the target file will be created with variable length records, the File Size option argument is interpreted as the number of maximum size records in the target file. This option is specifically provided for use with target nodes whose files are not dynamically extendable (i. e., HP 3000).
Interchange	Causes the file to be copied using Interchange Format. This option is the default for cross-system NFT.
Record Size	Specifies the record size of the target file in bytes. If fixed length records are being produced, the Record Size option argument is interpreted as the size of each record. If variable length records are being produced, the Record Size option argument limits the size of the largest record and records may be padded or truncated. The maximum record size at the HP 3000 is $(2^{15})-1$. NFT will issue a warning if it must truncate records in order to execute this option.
Search Character	The Search Character option specifies a search character that is used to divide source file data into logical records into the target file.
Sequential	Causes the target file to be organized to allow sequential access. Records in the source file will be sent to the target node contiguously.
Variable	Specifies that target file records should be created as variable length. The maximum size of a variable length record may be given using the Record Size option.

File Mappings

Table 5-5 lists the options required to map an HP-UX file into an MPE file with certain attributes. The "Comments" column of the table describes important details about each transfer. Other options besides the required options can also be used.

TABLE 5-5. FILE MAPPING FROM HP 9000 SOURCE FILE TO HP 3000 TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
HP-UX	VA	none	The default record length is 252 bytes. Use the Search Character option to extract logical records that are delimited by a character other than ASCII LF; the search character is not transferred with the data.
HP-UX	VB	Binary	The default record length is 508 bytes. You cannot use the Search Character option.
HP-UX	FA	Fixed	<p>The default record length is 160 bytes. If you use the Record Size option, any necessary padding uses the ASCII SPACE character. Use the Search Character option to extract logical records that are delimited by a character other than ASCII LF; the search character is not transferred with the data.</p> <p>There is a special case for files with a short average record length. Because of the way the HP-UX system approximates the target file size, if the average record size is less than 15 bytes, the MPE system does not preallocate enough space for the file; the transfer will fail and you will receive an error message. Use the Record Size option to specify the length of the records in the target file; or add one or more long dummy records to the source file to increase the average record length to more than 15 bytes; or pad existing source records to increase the average record length to more than 15 bytes.</p>
HP-UX	FB	Fixed, Binary	The default record length is 256 bytes. You can specify another record length by using the Record Size option; the last record may be padded with the ASCII NULL character if the number of bytes in the source file (including the record delimiters) is not evenly divisible by the record length. You cannot use the Search Character option.

DEC VAX Target Node

VAX/VMS computers use the Record Management Services (RMS) file system. RMS provides several different types of files with various record format. NS for the DEC VAX computer supports most RMS sequential file formats for transfers to VAX/VMS nodes. This manual uses the following abbreviations for RMS record formats:

TABLE 5-6. ABBREVIATIONS OF VAX/VMS TARGET FILE TYPES.

Abbreviation	File Type
Stream_LF	Stream file with linefeed record terminators.
Stream_CR	Stream file with carriage return record terminators.
Stream	Stream file with record terminators other than carriage return. These terminators could be LF, VT, CRLF, or FF.
Variable	Variable record format file.
Fix	Fixed record format file.

VAX System Logon Syntax

Logons for the VAX nodes take the form:

```
username[:password]
```

The *username* parameter is a logon name for the computer. The *password* is optional if the logon has no password. If it is used, a colon (:) must separate the *username* from the *password* parameter.

VAX System File Name Syntax

Under VAX/VMS system, files reside in a hierarchy of **devices**, **directories**, and **subdirectories**. The device is the physical unit (e.g., the disc) where the file exists. Files can be catalogued in any directory or subdirectory on the device. A subdirectory is a directory that is catalogued in the next higher level directory or subdirectory. Each account or logon on a VAX system has a default logon device and working directory which is immediately available to the user when he or she logs on.

A file specification consists of a device name, followed by a directory, followed by the file name. The general syntax of a VAX/VMS file specification is as follows:

```
device: [[directory]]file_name.type;version
```

- The *device* is a device name which, if omitted (along with the colon) defaults to the current device, commonly the logon device.
- The *directory* is a directory name, which defaults to the current working directory. The *directory* can be an absolute directory path or it can be relative to the current working directory. To specify a subdirectory, include a period and a subdirectory name before the closing bracket. You can use a subdirectory several layers deep by chaining subdirectory names together with periods between.

NOTE

Square brackets - "[" and "]" - are required by the VAX/VMS file name syntax whenever a directory or subdirectory is specified. However, you can omit the brackets and the directory name when using the default directory. Also note that if a *device* is specified, there is no default *directory*. Instead of square brackets, you can also use angle brackets - "<" and ">."

For all initiator nodes other than the VAX system, if the square brackets are used to specify the directory, the whole filename path specification should be enclosed in quotation marks to ensure that the DSCOPY command interprets the square brackets as part of the filename path. Refer to Appendix A for examples of DSCOPY.

- Either the *file_name* or the *type* (or both) must be present, and period separating them is always present.
- The *version* is an integer, which may be omitted if you want to use the file with the highest version number. A file of the same name (and version, if specified) in the target directory will only be replaced if the Replace option is also used. If Replace is not used and a file of the same name exists in the target directory, a file with an incremented version number is created.

Each field except the version number in a VAX file name must begin with a letter and contain only alphanumeric characters without spaces.

NOTE

DSCOPY can be made to automatically execute the user's LOGIN.COM. This command file is used to define logical names that NFT can use, or it can set the default directory. To have DSCOPY execute the LOGIN.COM command file, you must set the DTRACE bit to 283. The manual, *Network Services for the DEC VAX Computer*, describes how to do this, and the details of how it affects your NFT transfers.

DECnet VAX System File Name Syntax

If your AdvanceNet VAX computer is also connected to DECnet, you can copy files to a DECnet node from an NS/9000 node. This "network-to-network" transfer requires that you explicitly route the copy through the VAX node that is running AdvanceNet and DECnet. Therefore, in the target file *node_name* and *login* fields of the *dscopy* command you specify the VAX node that is running both AdvanceNet and DECnet. You then use the *file_pathname* field to specify the DECnet node, along with the appropriate device, directory, etc. as shown. The file path name syntax to specify a file on DECnet node is as follows:

```
dec_node"username password"::device:[[directory]]file_name.type;version
```

- The *dec_node* field is the name of the DECnet node.
- A space separates *username* and *password*.
- The rest of the parameters are described in VAX/VMS File Name Syntax above.

The string that contains the user name and password is preceded and followed by quotation marks, and always contains a blank. These characters have special meaning to the initiator; therefore, if your transfer requires a user name string, you must use escape characters in the *dscopy* command. The following table shows the required syntax at each initiator system.

TABLE 5-7. DECNET VAX SYSTEM FILE NAME SYNTAX.

Initiator System	DECnet VAX File Name Syntax
HP 1000	<code>` `dec_node"username password"::VAX/VMS_file_name` `</code>
HP 3000	<code>'dec_node"username password"::VAX/VMS_file_name'</code>
HP 9000	<code>'dec_node"username password"::VAX/VMS_file_name'</code>
VAX/VMS	<code>"dec_node""username password""::VAX/VMS_file_name"</code>

For HP 1000, apostrophes (') enclosing the grave accents (`) are needed around the above string to keep the HP`1000 from putting a comma between the user name and the password.

If the DECnet node contains an user name and password string, the string must be placed in two pairs of quotation marks and the entire file name must be quoted as shown. This is because VAX/VMS system NFT interprets quotation marks as escape characters. Since quotation marks are required in the user name and password string, they must be escaped in this way.

For additional information about the use of escape characters, see the *Network Services* manual for your initiator system.

NOTE

When you use the -P option and one of the computers involved is a remote DECnet VAX system, you receive a valid prompt for the file password, and a spurious prompt for a file password.

The spurious prompt is caused by the two colons in the DECnet file specification. The valid and spurious prompts appear in the same order as the two-colon keys in the command.

In response to the spurious password prompt, enter two colons (:). The transfer will then proceed normally.

File Access Modes

RMS does not differentiate between ASCII and binary data types. Therefore, the Binary option is handled differently for VAX/VMS consumers. NFT on a VAX computer operates as follows:

- **WITHOUT THE BINARY OPTION:**

Without the Binary option, when a VAX consumer receives data records from a producer, it adds headers or terminating characters to each record depending on the target file type. This is called Record Mode I/O. This record formatting information is inserted for the exclusive use of RMS.

- **WITH THE BINARY OPTION:**

If the Binary option is used, the VAX/VMS consumer writes blocks of data without attention to record structure. *No record formatting occurs*; the contents of the source file are copied directly. This is called Block Mode I/O. When you use the Binary option, be sure that you understand the type of file you are creating. The lack of record formatting can have important implications depending on the type of the target file. For example, if an HP-UX file is transferred to a VAX/VMS target node using the Binary option alone, an RMS Variable type file will be created. But this target file will not contain the RMS record formatting information required to use Record Mode I/O to access it.

HP 9000 to VAX/VMS Interchange Format Defaults

Any file copied from an HP/9000 node to a VAX/VMS target node defaults to a Stream_LF file. Because an HP-UX file and an RMS Stream_LF file are structurally equivalent, NFT's transparent mode is used to speed up the transfer. Figure 5-3 relates the HP-UX file attributes to the VAX/VMS file types.

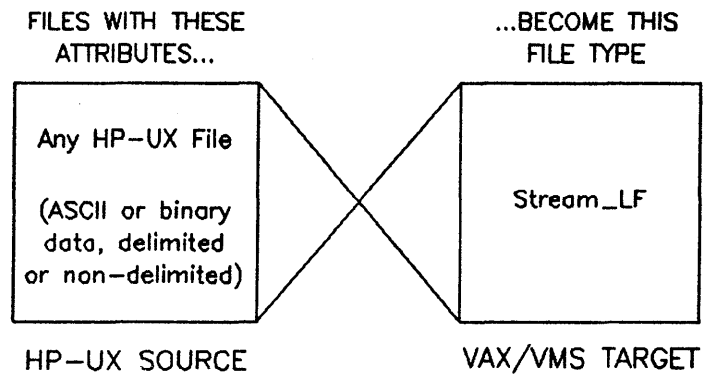


Figure 5-3. Interchange Format Defaults: HP 9000 to VAX/VMS Computer.

Interchange Format Options

Most of the Interchange Format options described in your NS user's manual may be used when copying HP 9000 source files to HP 3000 target nodes. However, your initiator system will determine which options are actually available for use. Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available at each initiator system.

Supported Options

The following Interchange Format options can be used when copying an HP 9000 file to a VAX/VMS target node. (Consult Table 2-1 to see if the options are available for your initiator system before using.)

- Binary
- Fixed
- Insert Character
- Interchange
- Record Size
- Search Character
- Sequential

Unsupported Options

The following Interchange Format options cannot be used when copying an HP/9000 source file to a VAX/VMS target node:

- ASCII
- File Size
- Append
- Direct
- Variable
- Strip

Interchange Format Option Results

Table 5-8 describes the results of using these options when copying a file from an HP 9000 to a VAX/VMS target node.

TABLE 5-8. INTERCHANGE FORMAT OPTIONS: HP 9000 to VAX Computer.

Option Mnemonic	Description
Binary	<p>Creates a Variable type target file on the VAX/VMS node. Specifies that the transfer uses Block Mode I/O (instead of Record Mode I/O). Block mode I/O causes NFT to forgo internal record formatting and transfer everything "verbatim" in a block of data.</p> <p>If the Fixed option is used with the Binary option, a Fixed type target file is created. Padding occurs at the end of the file if the record length used does not divide evenly into the total number of bytes in the source file. Any required padding of the last record is accomplished using the ASCII NULL character. No truncation occurs.</p>
Fixed	<p>Creates a Fixed type target file on the VAX/VMS node. The length of the records is specified with the Record Size qualifier (note defaults below). Truncation or padding may occur. The padding character depends on the use of the Binary option.</p> <p>If the Binary option is used and the Record Size qualifier is not specified, the size defaults to 256 bytes (see Binary).</p> <p>If the Binary option is not used and the Record Size qualifier is not specified, the size defaults to 160 bytes. The ASCII SPACE character is used for padding. The transfer uses Record Mode I/O.</p>
Record Size	<p>Used alone, this creates a Variable type target file on the VAX/VMS node. Often used with the Fixed option (see Fixed).</p> <p>In a Variable target file, no record will be longer than the length given in the option parameter. Any or all records may be shorter.</p> <p>In a Fixed target file, all records are padded or truncated (if necessary) to the length specified in the option parameter.</p> <p>Record size cannot be set to greater than 32,765 bytes.</p> <p>A warning message notifies you if any records are truncated during the transfer.</p>
Interchange	<p>Causes the file to be copied using Interchange Format. This is the default for cross-system NFT.</p>

TABLE 5-8. INTERCHANGE FORMAT OPTIONS: HP 9000 to VAX Computer (cont).

Option Mnemonic	Description
Insert Character	<p>Forces the VAX/VMS target file to be one of the following types:</p> <p>If no parameter is used, or if the parameter is 010, a Stream__LF file is created. (010=ASCII code for LF)</p> <p>If the parameter is 013, a Stream__CR file is created. (013= ASCII code for CR)</p> <p>If the parameter is anything else, a Stream file is created.</p> <p>Note that the file will have the specified type regardless of what other options, if any, are used. If the Binary option is used, however, this option only determines the file type of the target file. No record delimiters are appended.</p>
Search Character	<p>The Search Character option specifies a search character that is used to divide source file data into logical records in the target file. The character can be entered in literal form or in its equivalent decimal ASCII value. The ASCII value must be prefixed with a zero. If this option is not specified, NFT uses LINEFEED characters by default. The search character itself is not transferred with the data.</p>
Sequential	<p>Causes the target file to be organized to allow sequential access. Records in the source file will be sent to the target node contiguously.</p>

File Mappings

Table 5-9 lists the options required to map an HP-UX source file into a particular RMS sequential file type. Note that pairs occur more than once in the table. The "Comments" column of the table describes important details about each transfer. Other options besides the required options can also be used.

TABLE 5-9. FILE MAPPING FROM HP 9000 SOURCE FILE TO VAX/VMS TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
HP-UX	Stream_LF	none	Transparent mode transfer. The target file is structurally identical to source file.
HP-UX	Stream_LF, Stream_CR, Stream	Insert Character	If parameter is LF, target file is Stream_LF. If parameter is CR, target file is Stream_CR. If parameter is anything else, target file is Stream with CRLF delimiters.
HP-UX	Fix	Fixed	The default record length is 160 bytes. The Search Character option can be used to extract logical records from the source file. Padding (if necessary) is with the ASCII SPACE character. Truncation may occur.
HP-UX	Variable	Record Size	Source file records that exceed the length specified in the parameter are truncated; shorter records are not affected.
HP-UX	Stream_LF, Stream_CR, Stream	Binary, Insert Character	Block mode writes are used. Any record delimiters in source file are transferred with the data. The default record length is 256 bytes. If the parameter of the Insert Character option is LF, the target file is Stream_LF. If parameter is CR, target file is Stream_CR. If parameter is anything else, target file is Stream. Note that when you use the Insert Character option with the Binary option, the target file has the specified type, but because Block Mode is used to write the file, NFT does not append any record terminators.

TABLE 5-9. FILE MAPPING FROM HP 9000 SOURCE TO VAX TARGET (cont).

Source File Type	Target File Type	Options Required	Comments
HP-UX	Fix	Binary, Fixed	Block mode writes are used. Any record delimiters in source file are transferred with data. The default record length is 256 bytes. The last record may be padded with ASCII NULL characters if the number of bytes in the file (including any record delimiters) is not evenly divisible by the record length.
HP-UX	Variable	Binary	Block mode writes are used. Any record delimiters in source file are transferred with data. The default record length is 256 bytes.

PC Target Node

NOTE

Whenever a PC takes part in a file transfer using NFT, it *must* be the initiator. To transfer a file to or from a PC, you must initiate the transfer from the PC. If you are unfamiliar with the DOS operating system, refer to Section 2 of this manual or the *NFT User's Guide, HP OfficeShare Network* manual to learn more about the syntax and usage of NFT on a PC.

DOS Stream Files

The DOS operating system maintains a single type of file: a stream of bytes, usually called a **stream file**. The bytes may represent any kind of data, including ASCII characters. DOS imposes no structure on the file; under this scheme, a file is a place to read or write bytes. Higher level concepts, like "records," are not present in DOS.

However, many PC applications use a convention for finding the beginning and end of logical records of data. Under this convention, the next record in the file consists of all the data up to *but not including* the next occurrence of an ASCII CR character followed by an ASCII LF character. This sequence of characters - ASCII CR-LF - is called the record **delimiter** or **terminator**.

Most applications use the record delimiter only as a marker, and discard it as soon as they read it. NFT uses this interpretation unless you employ the Binary Interchange Format option; then NFT does not use record delimiters at all. Details about the Interchange Format options are later in this subsection.

PC Logon Syntax

Because the PC is always local,* you do not need to give a node name; and because it is a single-user computer, there is no logon.

Therefore, when you specify the PC file, leave the *node_name* and *logon* fields blank. Supply only the file name, as described below.

PC File Name Syntax

<code>[device:] [\] [dir1\dir2\...\dirN\]filename.ext</code>

*Whenever a PC takes part in a file transfer using NFT, it *must* be the initiator. See the NOTE above.

The *device*: is used to specify a disc drive other than the active drive (or specify a shared disc directory on the server). Prefix the pathname with the device letter (for example, "A:").

The *dir1* through *dirN* parameters denote directories within a pathname. If the first character of the pathname is a backslash (\), the search starts from the root of the active or specified drive; otherwise the search starts from the current directory on the active or specified drive.

Directory and file names can always contain letters and numbers; check your DOS reference manual for restrictions if you want to use punctuation characters. Each filename can be up to eight characters long; the optional extension has a three character maximum. The maximum pathname length is 63 characters.

NOTE

You cannot use NFT to transfer directories to a PC.

HP 9000 to PC Interchange Format Defaults

When an HP 9000 file is copied to a PC,* by default:

- the HP-UX system uses an ASCII LF as the record delimiter character in the source file, and
- the PC appends an ASCII CR-LF pair to the end of each incoming record as the record terminator in the target file.

Figure 5-4 show how HP 9000 stream files map to DOS stream files when no Interchange Format options are applied. The term "records" as used below means the data located between record delimiters in the HP-UX or DOS stream files.

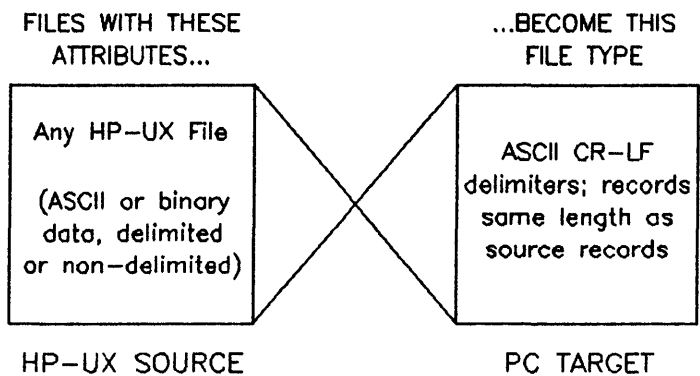


Figure 5-4. Interchange Format Defaults: HP 9000 to PC.

*See the Insert Character and Search Character options.

Interchange Format Options

Table 2-1 in Section 2 provides a quick reference of the PC Interchange Format options.

Because the PC must always be the* initiator, the following options are available.

Supported Interchange Format Options

The following Interchange Format options can be used when copying an HP 9000 source file to a PC:

- ASCII**
- Binary
- Fixed
- Record Size
- Insert Character
- Search Character
- Variable**

The options marked with two asterisks (**) can be used without causing an error, but they have no effect in transfers between these two machines. They are not included in the following tables.

Unsupported Options

The following options cannot be used when copying an HP 9000 file to a PC:

- Append
- Direct
- File Size
- Interchange
- Sequential
- Strip

*Whenever a PC takes part in a file transfer using NFT, it *must* be the initiator. See the NOTE at the beginning of this subsection.

Interchange Format Option Results

Table 5-10 describes the results of using these options when copying a file from an HP 9000 to a PC.

TABLE 5-10. INTERCHANGE FORMAT OPTIONS: HP 9000 to PC.

Option Mnemonic	Description
Binary	<p>Source file records are placed directly in the target file without record delimiters. If you use this option, the Search Character and Insert Character options are ignored. You cannot use the ASCII and Binary options together.</p>
Fixed	<p>Source file records are truncated or padded as necessary so they are all the same length. The length defaults to 160 bytes, or to 256 bytes if the Binary option is also in effect. You can specify the length using the Record Size option. Unless the Binary option is also in effect, records are padded as necessary with ASCII SPACE characters.</p> <p>You can use the Binary and Fixed options together. Data is extracted from the source file and written to the target file in blocks of 256 bytes (or the length you specify with the Record Size option). If the number of bytes in the file is not evenly divisible by 256 (or the length you specify with the Record Size option), the last record is padded with ASCII NULL characters.</p>
Insert Character	<p>The argument to the Insert Character option specifies a single character that defines the record delimiter to use in the target file. The character can be entered in literal form or in its equivalent decimal ASCII value; the ASCII value must be prefixed with a zero. The ASCII CR-LF pair is the default target file record delimiter in transfers from an HP-UX system to a PC; the Insert Character option defines only single-character delimiters. If you use the Binary option, this option is ignored.</p>
Record Size	<p>The target file is formed with a maximum record length as defined by the Record Size argument. If the source file contains a record that is longer than the Record Size argument, the record is truncated and NFT prints a warning message. There is no limit to record size beyond the physical limitations of the system.</p> <p>If you use the Binary option along with the Record Size option records are never truncated. In effect, the Record Size option is ignored.</p> <p>If you use the Fixed option along with the Record Size option, records in the source file are truncated or padded so that all records are the specified length. Padding is with ASCII SPACE characters.</p> <p>If you use both the Binary and Fixed options along with the Record Size option, the last record may be padded with ASCII NULL characters if the number of bytes in the file (including any record delimiters) is not evenly divisible by the specified record size.</p>

TABLE 5-10. INTERCHANGE FORMAT OPTIONS: HP 9000 TO PC (cont).

Option Mnemonic	Description
Search Character	The argument to the Insert Character option specifies a single character that defines the record delimiter to use in the source file. The character can be entered in literal form or in its equivalent decimal ASCII value; the ASCII value must be prefixed with a zero. The search character itself is not transferred with the data. The ASCII LF character is the default source file record delimiter in transfers from an HP-UX system to a PC; the Search Character option defines only single-character delimiters. If you use the Binary option, this option is ignored.

File Mappings

Table 5-11 lists the options required to map an HP-UX file into a PC file with certain attributes. The "Comments" column of the table describes important details about each transfer. Details about each option are in Table 5-10. Other options (Record Size, for example) can also be used.

TABLE 5-11. FILE MAPPING FROM HP 9000 SOURCE FILE TO PC TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
HP-UX	DOS	none	Record delimiters are appended to each record. The ASCII CR-LF pair is the default record delimiter.
HP-UX	DOS	Fixed	Fixed length records. Record delimiters are appended to each record. The record length defaults to 160 bytes. Longer records may be truncated. Any necessary padding uses the ASCII SPACE character.
HP-UX	DOS	Binary	Record delimiter characters in the source file (if any) are transferred with the data. NFT inserts no record delimiters in the target file. The result is that the target file is structurally identical to the source file.
HP-UX	DOS	Binary, Fixed	Fixed length records. Record delimiter characters (if any) are transferred with the data. NFT inserts no record delimiters. The default record length is 256 bytes. The last record may be padded (with ASCII NULL characters) if the record length does not divide evenly into the total number of bytes in the source file (including any record delimiter characters).

NOTE

The default record delimiter in the PC target file is the ASCII CR-LF pair.
The default record delimiter in the HP-UX source file is the ASCII LF character.

RMS FILE SYSTEM

VAX/VMS computers use the Record Management Services (RMS) file system. RMS provides several different types of files with various record formats. NS for the DEC VAX computer supports all RMS sequential file formats for transfers from VAX/VMS nodes. This manual uses the following abbreviations for VAX/VMS record formats:

TABLE 6-1. ABBREVIATIONS OF VAX/VMS SOURCE FILE TYPES.

Abbreviation	File Type
Stream_LF	Stream file with LINEFEED record terminators.
Stream_CR	Stream file with carriage return record terminators.
Stream	Stream file with record terminators other than carriage return. These terminators could be LF, VT, CRLF, or FF.
Variable	Variable record format file.
Fixed	Fixed record format file.
VFC- (Variable Fixed Control Format) UDF- (User Defined Control Format)	VFC and UDF files are supported for VAX source files only. Use of Binary option is required.

If you need to transfer any other file type, such as relative files or indexed files, you must first translate the file into one of the sequential file types using RMS (VMS Record Management Services) utilities.

VAX System Logon Syntax

Logons for the VAX nodes take the form:

```
username[:password]
```

The *username* parameter is a logon name for the computer. The *password* is optional if the logon has no password. If it is used, a colon (:) must separate the *username* from the *password* parameter. For password security, see the Network Services manual for the initiator system you are using.

VAX System File Name Syntax

Under VAX/VMS system, files reside in a hierarchy of devices, directories, and subdirectories. The device is the physical unit (e.g., the disc) where the file exists. Files can be catalogued in any directory or subdirectory on the device. A subdirectory is a directory that is catalogued in the next higher level directory or subdirectory. Each account or logon on a VAX system has a default logon device and working directory, which is immediately available to the user when he or she logs on. A file specification consists of a device name, followed by a directory path, followed by the file name. The general syntax of a VAX/VMS file specification is as follows:

device: [[*directory*]]*file_name.type;version*

- The *device* is a device name which, if omitted (along with the colon), defaults to the current device, commonly the logon device.
- The *directory* is a directory name which defaults to the current working directory. The *directory* can be an absolute directory path, or it can be relative to the current working directory. To specify a subdirectory, include a period and a subdirectory name before the closing bracket. You can use a subdirectory several layers deep by chaining subdirectory names together with periods between.

NOTE

Square brackets - "[" and "]" - are required by the VAX/VMS file name syntax whenever a directory or subdirectory is specified. However, you can omit the brackets and directory name when using the default directory. Also note that if a *device* is specified, there is no default directory. Instead of square brackets, you can also use angle brackets - "<" and ">."

For all initiator nodes other than the VAX system, if the square brackets are used to specify the directory, the whole filename path specification should be enclosed in quotation marks to ensure that the DSCOPY command interprets the square brackets as part of the filename path. Refer to Appendix A for examples of DSCOPY.

- Either the *file_name* or the *type* (or both) must be present, and the period separating them is always present.

- The *version* is an integer. If version number is omitted, the highest numbered version of the file is used as the source.

NFT does not currently support directory or "wildcard" transfers involving a VAX/VMS producer. Each field (except the version number) in a VAX/VMS file name must begin with a letter and contain only alphanumeric characters without spaces.

NOTE

DSCOPY can be made to automatically execute the user's LOGIN.COM. This command file is used to define logical names that NFT can use, or it can set the default directory. To have DSCOPY execute the LOGIN.COM command file, you must set the DTRACE bit to 283. The manual, *Network Services for the DEC VAX Computer*, describes how to do this, and the details of how it affects your NFT transfers.

DECnet VAX File Name Syntax

If your AdvanceNet VAX system is also connected to DECnet, you can copy files from any DECnet node. This "network-to-network" transfer requires that you explicitly route the copy through the VAX node that is running both AdvanceNet and DECnet. Therefore, in the source file *node_name* and *login* fields of the *dscopy* command you specify the VAX node that is running both AdvanceNet and DECnet. You then use the *file_path_name* field to specify the DECnet node, along with the appropriate device, directory, etc. as shown. The file path name syntax to specify a file on a DECnet node is as follows:

<pre><i>dec_node</i>"<i>username password</i>"::<i>device</i>:<i>[[directory]]file_name.type;version</i></pre>
--

- The *dec_node* field is the name of the DECnet node.
- A space separates *username* and *password*.
- The rest of the parameters are described in "VAX System File Name Syntax" above.

The string that contains the user name and password is preceded and followed by quotation marks, and always contains a blank. These characters have special meaning to the initiator; therefore, if your transfer requires an user name and password string, you must use escape characters in the DSCOPY command. The following table shows the required syntax at each initiator system.

TABLE 6-2. DECNET VAX SYSTEM FILE NAME SYNTAX.

Initiator System	DECnet VAX File Name Syntax
HP 1000	<code>`dec_node"username password":VAX/VMS_file_name`</code>
HP 3000	<code>'dec_node"username password":VAX/VMS_file_name'</code>
HP 9000	<code>'dec_node"username password":VAX/VMS_file_name'</code>
VAX/VMS	<code>"dec_node""username password""":VAX/VMS_file_name"</code>

For HP 1000, apostrophes (') enclosing the grave accents (`) are needed around the above string to keep the HP 1000 from putting a comma between the user name and the password.

If the DECnet node contains an user name and password string, the string must be placed in two pairs of quotation marks and the entire file name must be quoted as shown. This is because VAX/VMS system NFT interprets quotation marks as escape characters. Since quotation marks are required in an user name and password string, they must be escaped in this way.

For additional information about the use of escape characters, see the *Network Services* manual for your initiator system.

Using the FILEPASSWORD Option with DECnet

When you use the `filepassword` option for a transfer between a DECnet VAX computer and an HP 3000 system, you will receive a spurious prompt for a file lock word caused by the two colons (::) in the DECnet VAX file specification. In response to the spurious prompt, enter two colons (::). The transfer will then proceed normally.

File Access Modes

RMS does not differentiate between ASCII and binary data types. Therefore, the Binary option is handled differently for VAX/VMS producers. NFT on a VAX computer operates as follows:

- **WITHOUT THE BINARY OPTION:**

Without the Binary option, when a VAX producer reads data records it extracts each logical record by removing RMS record formatting information from each record depending on the source file type. This is called **Record Mode I/O**.

- **WITH THE BINARY OPTION:**

If the Binary option is used, the VAX/VMS producer read blocks of data without attention to record structure. *No logical record extraction occurs*; the contents of the source file are copied directly. This is called **Block Mode I/O**. When you use the Binary option, be sure that you understand the type of file you are creating. In particular, a file created with Record Mode I/O

and transferred using Block Mode I/O will contain the RMS record formatting information along with the data. Because RMS Fix files do not contain record formatting information, Block Mode reads are of no special concern when the source file is Fix.

SUPPORTED TARGET NODES

VAX/VMS source files can be copied to NS/1000, NS/3000, NS/9000, or PC target nodes.

This section provides all the information you will need to know when copying an VAX/VMS source file to each supported target node. This information includes:

- Target file system information
- File name and logon syntax
- Interchange Format defaults
- The effect of each Interchange Format option on the file translation process.

HP 1000 Target Node

HP 1000 computers that run NS/1000 have RTE-A operating systems. The RTE-A file system includes a record structure and categorizes files into certain types. These file types are defined as follows:

- **Type 1.** These files have fixed-length records of 256 bytes and contain binary data.
- **Type 2.** These files also have fixed-length records, but the record length is defined by the user at file creation. They contain binary data.
- **Type 3.** These files have variable-length records and are assumed to contain binary data. Type 3 transferred to any system will produce a binary target file unless the ASCII option is specified in the file transfer.
- **Type 4.** These files also have variable-length records, and NFT assumes they contain ASCII data. Type 4 files are text files and may be altered with the RTE-A text editor, EDIT/1000.
- **Greater than Type 4.** File types 5, 6 and 7 have variable-length records. File types greater than 7 are user-defined. NFT assumes files of Type >4 contain binary data.

HP 1000 Logon Syntax

Logons for NS/1000 systems take the form:

```
accountname[/password]
```

The *accountname* parameter is the logon name for the computer. The password parameter is optional if the logon has no password; if you use it, a slash (/) must separate the *accountname* from the *password* parameter. If the logon has no password, the slash is omitted.

For NFT between a VAX/VMS initiator/producer and an HP 1000 target node, the following syntax should be used for the HP 1000 logon sequence in order to invoke password prompting.

```
node1000#logon: :#file1000
```

The two colons will result in the following DSCOPY prompt on the VAX initiator/producer node:

```
TARGET PASSWORD>
```

You should respond with a slash (/) followed immediately with the proper password. The slash is necessary so that the HP 1000 target node will parse the string correctly.

HP 1000 File Name Syntax

The RTE-A file system has a hierarchical file structure. Files are catalogued in directories. Directories can also contain similar information about other directories, called subdirectories. Subdirectories have the same characteristics as directories; the term subdirectory means only that the directory is catalogued in the next higher level directory or subdirectory. Each account or logon on an RTE-A system has a home or default logon working directory. This directory is automatically made available to the user when he or she logs on.

If the HP 1000 file resides in the hierarchical file system, the file name syntax is as follows:

```
[/][directory/][directory/]:filename
```

If the *directory* parameter is omitted, the default logon working directory for the logon specified is used. If the slash (/) is omitted and the *directory* parameter is specified, the directory is assumed to be in the default logon working directory for the logon provided. The maximum file path name, including the file name, is 63 characters.

If the HP 1000 file you wish to access resides on an FMGR cartridge, the file name must be specified as follows:

filename::directory

The *filename* parameter may be a maximum of six characters. The *directory* parameter may be up to two ASCII characters or a positive or negative integer.

VAX/VMS to NS/1000 Interchange Format Defaults

By default, any file copied from a VAX/VMS source node to an HP 1000 node creates an RTE-A Type 4 file. You can override the default with *dscopy* options. Figure 6-1 relates the RTE-A file attributes to the VAX/VMS file types.

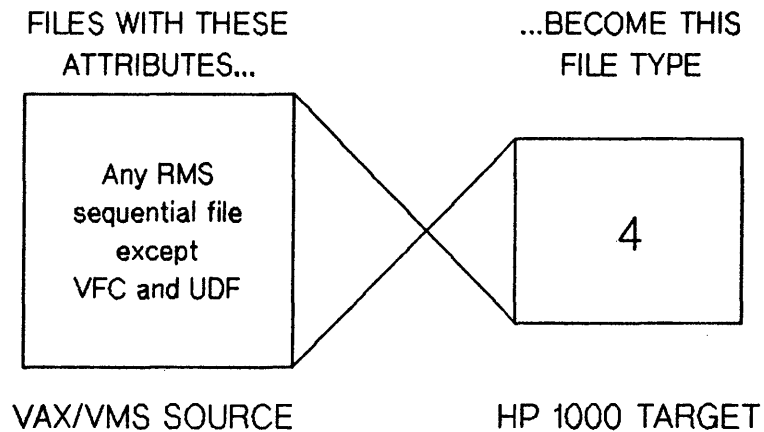


Figure 6-1. Interchange Format Defaults: VAX/VMS to NS/1000.

NOTE

Two RMS sequential file types - VFC and UDF - can only be transferred using the Binary option. No Interchange Format defaults exist for these source files.

Files copied to an NS/1000 system cannot have records longer than 4400 bytes (2200 words). Records can be truncated using the Record Size option.

Interchange Format Options

Most of the Interchange Format options described in your NS user's manual may be used when copying VAX/VMS source files to HP 1000 target nodes. However, your initiator system will determine which options are actually available for use. Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available at each initiator system.

Supported Interchange Format Options

The following Interchange Format options can be used when copying a VAX/VMS source file to an HP 1000 target node. (Consult Table 2-1 to see if the options are available for your initiator system before using.)

- ASCII
- Binary
- File Size
- Fixed
- Interchange
- Record Size
- Sequential
- Variable

Unsupported Options

The following options cannot be used when copying a VAX/VMS source file to an HP 1000 target node:

- Append
- Direct
- Insert Character
- Strip
- Search Character

Interchange Format Option Results

Table 6-3 describes the results of using these options when copying a file from a VAX/VMS source node to an HP 1000 target node.

TABLE 6-3. INTERCHANGE FORMAT OPTIONS: VAX/VMS to HP 1000.

Option Mnemonic	Description
ASCII	Specifies that the file is transferred using Record Mode I/O (instead of Block Mode I/O). Record Mode I/O causes RMS to transfer the logical records only; RMS record formatting information is not transferred. This option also specifies that records are padded, if necessary, with the ASCII SPACE character.
Interchange	Causes the file or files to be copied using Interchange Format. This option is the default for cross-system file transfer.
Fixed	<p>Used alone, this option creates a Type 4 target file in which all records are of equal length. The record length is specified by the Record Size option, and defaults to the length of the longest record in the file if that is known; otherwise the default record length is 160 bytes (see Record Size). If necessary, source file records are truncated or padded with ASCII SPACE characters.</p> <p>This option is sometimes used with the Binary option (see Binary).</p>
Binary	<p>If the source is an RMS Fix file, the target will be a Type 1 or Type 2, depending on the source file record length (see Record Size). If the source is not an RMS Fix file, the target is a Type 3 file in which all records are 512 bytes long, except the last record which may be less than 512 bytes. Block Mode reads are used, so the target file records are composed of blocks of adjacent bytes from the source file, including any RMS record formatting information. This option is required to transfer UDF and VFC files.</p> <p>If the Fixed option is also used the target file type depends on the length of the records (see Record Size). The last record may be padded with ASCII NULL characters if the number of bytes in the file (including any RMS record formatting information) is not evenly divisible by the specified or default record size (see Fixed).</p> <p>NOTE: Whenever you use the Binary option, be sure that you understand the kind of file you are creating. In particular, a file created with Record Mode I/O and transferred using Block Mode I/O may contain RMS record formatting information along with the data (see ASCII).</p>

TABLE 6-3. INTERCHANGE FORMAT OPTIONS: VAX/VMS to HP 1000 (cont).

Option Mnemonic	Description
File Size	Specifies how much space to allocate for the target file. If the target file will have fixed length records, the File Size argument is in records. If the target file will have variable length records, the File Size argument indicates number of maximum size records.
Record Size	<p>Determines the length of fixed-length records. A length of zero (0) bytes is treated as unspecified, and the above default applies (see Fixed).</p> <p>Also determines the type of a target file of binary data with fixed-length records. If the Record Size is 256 the target file is Type 1. If the Record Size is unspecified and the source file is Fix with a record length of 256, the target file is Type 1. Otherwise the target file is Type 2.</p> <p>Record size cannot be set to greater than 4400 bytes.</p> <p>A warning message notifies you if any records are truncated during NFT.</p>
Sequential	Causes the target file to be organized to allow sequential access. Records in the source file will be sent to the target node contiguously.
Variable	Specifies that target file records be variable length. The maximum size of a variable record may be given using the Record Size option.

NOTE

It is not possible to map into an RTE-A file of Type >4.

File Mappings

Table 6-4 lists the options required to map a particular RMS sequential file type into a particular RTE-A file type. The "Comments" column of the table describes important details about each transfer. Other options besides the required options can also be used.

TABLE 6-4. FILE MAPPING FROM VAX/VMS SOURCE FILE TO HP 1000 TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
Stream_LF Stream_CR, Stream, Variable, Fix	Type 4	none	Record mode reads are used; only the records are transferred. If the VAX/VMS system can provide the length of the records in source file, that is the default record length for target file; otherwise the record length defaults to 160 bytes. You can use the Record Size option to specify a different record length. Records longer than the default or specified record length may be truncated.
Stream_LF, Stream_CR, Stream, Variable	Type 3	Binary	Block mode reads are used; RMS record formatting information is transferred with the data. The default record length is 512 bytes.
Fix	Type 3	Binary, Variable	Any existing padding is transferred with the records.
Stream_LF, Stream_CR, Stream, Variable	Type 1, Type 2	Binary, Fixed	The default record length is 512 bytes. The target file is Type 1 only if you use the the Record Size option and specify a record length of 256. Otherwise the target file is Type 2. Block mode reads are used; RMS record formatting information is transferred with the data.
Fix	Type 1, Type 2	Binary	The default record length for the target file is the same as the record length in the source file. The target file is Type 1 if the source file records are 256 bytes long, or if you use the Record Size option to specify a record length of 256. Otherwise the target file is Type 2. No truncation or padding occurs.

HP 3000 Target Node

HP 3000 computers utilize the Multiprogramming Executive (MPE) operating system. MPE files may contain ASCII or binary data, fixed or variable length records, and may be Relative I/O (RIO) or non-RIO files. (RIO files are random access, non-RIO files are sequential access.) A VAX/VMS producer can only create non-RIO files on an HP 3000 target node.

This table introduces four abbreviations for MPE files with particular attributes:

TABLE 6-5. ABBREVIATIONS OF MPE FILE ATTRIBUTES.

Abbreviation	MPE File Attributes
VA	Variable-length records of ASCII data.
VB	Variable-length records of Binary data.
FA	Fixed-length records of ASCII data.
FB	Fixed-length records of Binary data.

HP 3000 Logon Syntax

Each field in the HP 3000 logon sequence must begin with a letter, contain only alphanumeric characters, and must not exceed eight characters. The logon syntax is as follows:

```
username[/userpass].account[/acctpass][,groupname[/grouppass]]
```

For NFT between a VAX/VMS initiator/producer and an HP 3000 target node, the following syntax should be used for the HP 3000 logon sequence to invoke password prompting:

```
username/::account/::,groupname/::
```

HP 3000 File Name Syntax

Each field in an HP 3000 file name must begin with a letter, contain only alphanumeric characters, and must not exceed eight characters. The *group* and *account* parameters default to the logon group and account. The syntax of an HP 3000 file name is as follows:

```
filename[/lockword][.group[.account]]
```

VAX/VMS to HP 3000 Interchange Format Defaults

An RMS Fix file copied to an NS/3000 system becomes an ASCII file with fixed-length records (FA). Any other supported RMS file type becomes an ASCII file with variable-length records (VA). In each case, the RMS file is read using Record Mode I/O.

Figure 6-2 relates the MPE file attributes to the VAX/VMS file types.

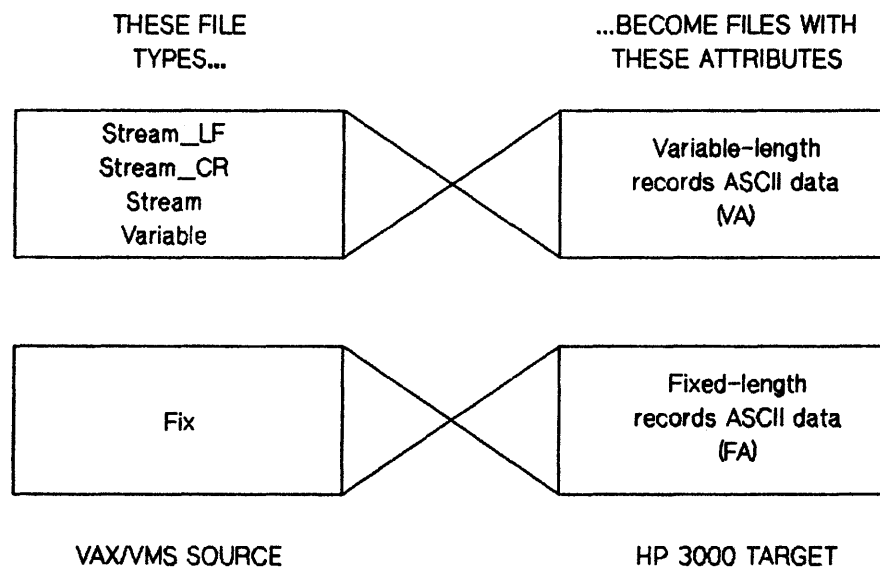


Figure 6-2. Interchange Format Defaults: VAX/VMS to HP 3000.

NOTE

Two RMS sequential file types - VFC and UDF - can only be transferred using the Binary option. No Interchange Format defaults exist for these source files.

Special Cases

A transfer of a VAX source file to the HP 3000 may result in inadequate allocation of space for the target file on the HP 3000 if the file's average record size is less than 15 bytes. If not enough space has been allocated, the transfer will be unsuccessful, and NFT will print an error message. You can circumvent this failure by using the Record Size option to specify the length of the record in the target file. (However, this will not work if you also specify the Fixed option.) If you need to get fixed records

in the target file, you should transfer the file from the VAX node to the HP 3000 without the Fixed option, then execute a local DSCOPY within the HP 3000, using the Fixed option.

Interchange Format Options

Most of the Interchange Format options described in your NS user's manual may be used when copying VAX/VMS source files to HP 3000 target nodes. However, your initiator system will determine which options are actually available for use. Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available at each initiator system.

Supported Interchange Format Option

The following Interchange Format options can be used when copying a VAX/VMS source file to an HP 3000 target node. (Consult Table 2-1 to see if the options are available for your initiator system before using.)

- Append
- ASCII
- Binary
- File Size
- Fixed
- Interchange
- Record Size
- Sequential
- Variable

Unsupported Options

The following options cannot be used when copying a VAX/VMS source file to an HP 3000 target node:

- Direct
- Insert Character
- Search Character
- Strip

Interchange Format Option Results

Table 6-6 describes the results of using these options when copying a file from a VAX/VMS source node to an HP 3000 target node.

TABLE 6-6. INTERCHANGE FORMAT OPTIONS: VAX/VMS to HP 3000.

Option Mnemonic	Description
Append	Appends the source file onto an existing target file. The attributes of the source file must match those of the target file. If the target file does not exist, NFT will return an error message.
ASCII	Specifies that whenever padding is necessary, the ASCII SPACE character is used. Also specifies that the transfer uses Record Mode I/O to read the source. Record Mode causes NFT to transfer only the logical records; RMS record formatting information is not transferred. The ASCII and Binary options are mutually exclusive.
Binary	<p>Specifies that if padding is necessary, the ASCII NULL character is used. In addition, Block Mode reads are used, so RMS record formatting information (if any) is transferred with the data. If the source is an RMS Fix file, a binary file of fixed-length records (FB) is created; the record length is the same as in the source file. If the source is any other RMS file type, a binary file of variable-length records (VB) is created; by default these records are 764 bytes long. This option is required to transfer RMS UDF and VFC files. The ASCII and Binary options are mutually exclusive.</p> <p>If you use both the Fixed option and the Binary option, a binary file of fixed-length records (FB) is created regardless of the type of the source file. If the source is an RMS Fix file, the record length defaults to be the same as in the source. If the source is any other RMS file type, the default record length is 512 bytes (see Fixed).</p> <p>NOTE: Whenever you use the Binary option, be sure that you understand the type of file you are creating. In particular, a file created with Record Mode I/O and transferred using Block Mode I/O may contain RMS record formatting information along with the data.</p>
File Size	Specifies how much space to allow for the target file. If the target file will be created with fixed length records, the File Size option argument is interpreted as the number of records in the target file. If the target file will be created with variable length records, the File Size option argument is interpreted as the number of maximum size records in the target file.

TABLE 6-6. INTERCHANGE FORMAT OPTIONS: VAX/VMS to HP 3000 (cont).

Option Mnemonic	Description
Fixed	<p>Specifies that the target file is to be an ASCII file of fixed-length records (FA). You can specify the record length using the Record Size option, but note the defaults below. Source file records may be truncated or padded (with ASCII SPACE characters) as necessary. The Fixed and Variable options are mutually exclusive.</p> <p>If the source is an RMS Fix file, the default record length is the same as the source file record length. If the source file is any other RMS file type the default record length is 160 bytes; unless you also use the Binary option, then the default record length is 512 bytes (see Binary).</p> <p>If you use both the Fixed option and the Binary option, a binary file of fixed-length records (FB) is created. The last record alone may be padded if the number of bytes in the source file (including any RMS record formatting information) is not evenly divisible by the record length. In this case, the padding uses ASCII NULL characters (see Binary).</p>
Interchange	<p>Causes the file or files to be copied using Interchange Format. This option is the default for cross-system file transfers.</p>
Record Size	<p>Determines the length of fixed-length records, and sets a maximum length for variable-length records. Source file records longer than the specified length are truncated in the target file. If the target file has fixed-length records, shorter records are padded; if the target file has variable-length records, shorter records are not affected. A warning message notifies you if any records are truncated during a transfer.</p>
Sequential	<p>Causes the target file to be organized to allow sequential access. Records in the source file will be sent to the target node contiguously.</p>
Variable	<p>Specifies that the target file records should be variable length. This option is only useful if the source file is an RMS Fix file. The maximum size of a variable length record can be specified using the Record Size option. The Fixed and Variable options are mutually exclusive.</p>

File Mappings

Table 6-7 lists the options required to map a particular RMS sequential file type into a particular MPE file type. The "Comments" column of the table describes important details about each transfer. Other options besides the required options can also be used.

TABLE 6-7. FILE MAPPING FROM VAX/VMS SOURCE FILE TO HP 3000 TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
Stream_LF, Stream_CR, Stream, Variable	VA	none	Record mode reads are used. The default record length is 252 bytes. If you use the Record Size option, source records longer than the specified length are truncated.
Stream_LF, Stream_CR, Stream, Variable	FA	Fixed	If the VAX/VMS system can provide the length of the records in source file, that is the default record length for target file; otherwise the record length defaults to 160 bytes. If you use the Record Size option, any necessary padding uses the ASCII SPACE character.
Stream_LF, Stream_CR, Stream, Variable	VB	Binary	Block reads are used; RMS record formatting information is transferred with the data. The default record length is 764 bytes.
Stream_LF, Stream_CR, Stream, Variable	FB	Binary, Fixed	Block reads are used; RMS record formatting information is transferred with the data. The default record length is 512 bytes.
Fix	VA, VB	Variable, [Binary]	The default record length is determined by the MPE system; it is never less than the length of the source file records, and is commonly slightly larger. Use the Binary option only if you want a VB target file.
Fix	FA, FB	none, [Binary]	The default record length is determined by the MPE system; it is never less than the length of the source file records, and is commonly slightly larger. If you use the Record Size option, any necessary padding uses the ASCII SPACE character. Use the Binary option only if you want an FB target file; padding will then use the ASCII NULL character.

HP 9000 Target Node

HP/9000 computers have HP-UX operating systems. There is only one type of file on HP-UX systems: a stream of bytes. The bytes can represent ASCII characters or some type of binary data. There is no concept of records in the HP-UX file system. In fact, no structure is imposed on a file by the system, and no meaning is attached to its contents. The meaning of the bytes in a file depends solely on the programs that interpret that file by looking for LINEFEED characters. The LINEFEED characters are treated as record delimiters by these programs.

HP 1000 commands typed at the RTE Command Interpreter (CI) prompt are automatically upshifted. This might cause problems for NS/9000 producer or target nodes since HP-UX commands are case sensitive. To avoid this problem, enclose with grave accents (`) the part of the NS/1000 DSCOPY command string that must not be upshifted.

Refer to Appendix A for examples of DSCOPY.

NS/9000 Logon Syntax

Each parameter in an HP/9000 logon may contain any ASCII character, with the exception of the colon (:), and must not exceed eight characters. The following syntax is an NFT convention and should *not* be used to log on to an HP-UX operating system.

```
username[:password]
```

HP 9000 File Name Syntax

The syntax of an HP/9000 file name is as follows:

```
[/][dir1/dir2/.../dirn/]filename
```

The */dir1* and */dir2* parameters denote a directory within a path name. If the first character is a slash (/), the search starts from the root, otherwise the search starts from the user's default logon working directory. Directory and file names may contain any of the ASCII characters, except for null and slash (/), and can be up to 255 characters long, with a maximum path length of 1023 characters.

The target file on the HP-UX system will be created with a default mode of 0666 octal (rw-rw-rw-).

Access Control Lists (ACLs) cannot be set using NFT.

VAX/VMS to HP 9000 Interchange Format Defaults

Any file copied from a VAX/VMS source node to an HP/9000 node creates an HP-UX stream file. Because an RMS Stream_LF file and an HP-UX file are functionally equivalent, NFT's transparent mode is used to speed up this type of transfer. All other transfers occur in Interchange mode. Records are delimited by ASCII LF characters in the target file.

Figure 6-3 relates the HP-UX file attributes to the VAX/VMS file types.

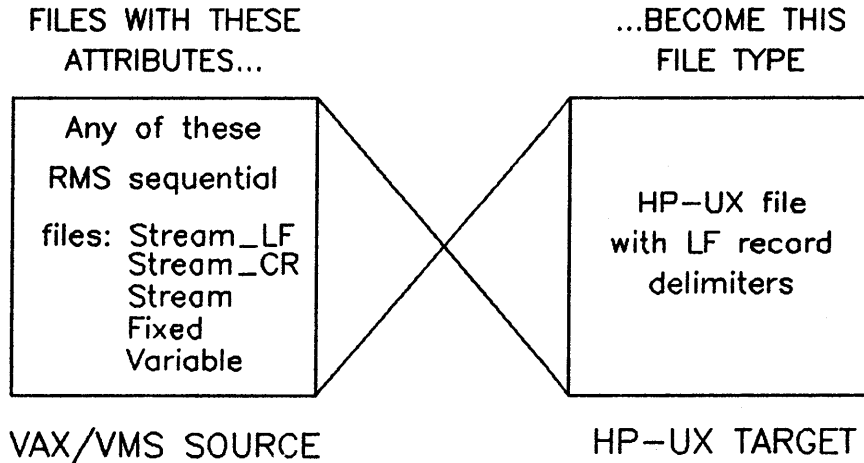


Figure 6-3. Interchange Format Defaults: VAX/VMS to HP 9000.

NOTE

Two RMS sequential file types - VFC and UDF - can only be transferred using the Binary option. Therefore, there are no Interchange Format defaults for these source files.

Interchange Format Options

Most of the Interchange Format options described in your NS user's manual may be used when copying VAX/VMS source files to HP 9000 target nodes. However, your initiator system will determine which options are actually available for use. Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available at each initiator system.

Supported Options

The following Interchange Format options can be used when copying a VAX/VMS source file to an HP 9000 target file. (Consult Table 2-1 to see if the options are available for your initiator system before using.)

- Binary
- Fixed
- Record Size
- Insert Character

Unsupported Options

The following options cannot be used when copying a VAX/VMS source file to an HP/9000 target node:

- ASCII
- File Size
- Append
- Direct
- Interchange
- Variable
- Strip
- Search
- Sequential

Interchange Format Option Results

Table 6-8 describes the results of using these options when copying a file from a VAX/VMS source node to an HP/9000 target node.

TABLE 6-8. INTERCHANGE FORMAT OPTIONS: VAX/VMS to HP 9000.

Option Mnemonic	Description
Binary	<p>Source file records are placed directly in the target file. Block Mode reads are used. NFT appends no record delimiters. This option is required to transfer UDF and VFC files. Padding of the last record may occur if the Fixed option is also used (see Fixed).</p> <p>Whenever you use the Binary option, be sure that you understand the type of file you are creating. In particular, a file created with Record Mode I/O and transferred using Block Mode I/O may contain RMS record formatting information along with the data.</p>
Fixed	<p>If Record Size is not specified, target file record length defaults to the source file's assigned maximum record size if it exists, or to 512 bytes if the Binary option is in effect, or to 160 bytes otherwise. If Insert Character is not specified, LF record delimiters are inserted. Source file records are truncated or padded as necessary. If the Binary option is in effect, padding of the last record (if necessary) is with ASCII NULL characters; otherwise each record is padded (if necessary) with ASCII SPACE characters.</p>
Record Size	<p>Source file records longer than the specified length are truncated to the specified length. A length of zero (0) bytes is treated as unspecified, and the above defaults apply (see Fixed).</p> <p>A warning message notifies you if any records are truncated during the transfer.</p>
Insert Character	<p>Defines a record delimiter character to be inserted after each logical record in the source file before it is placed in the target file. The delimiter character can be entered in literal form, or in its equivalent decimal ASCII value which must be prefixed with a zero (0).</p>

File Mappings

Table 6-9 lists the options required to map a particular RMS sequential file type into an HP-UX file with certain attributes. The "Comments" column of the table describes important details about each transfer. Other options besides the required options can also be used.

TABLE 6-9. FILE MAPPING FROM VAX/VMS SOURCE FILE TO HP 9000 TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
Stream_LF	HP-UX	none	Transparent mode transfer. The target file is structurally identical to the source file.
Stream_CR, Stream Fix	HP-UX	none	The ASCII LF character is the default record delimiter, and it is appended to each record.
Stream_LF, Stream_CR, Stream, Variable, Fix	HP-UX	Fixed	Fixed length records. The ASCII LF character is the default record delimiter. If the VAX/VMS system can provide the length of the records in source file, that is the default record length for target file; otherwise the record length defaults to 160 bytes. Longer records may be truncated. If you use the Record Size option, any necessary padding uses the ASCII SPACE character.
Stream_LF, Stream_CR, Stream, Variable, Fix	HP-UX	Binary	Block mode reads are used; RMS record formatting information (if any) is transferred with the data. The target file is structurally identical to the source file. NFT inserts no record delimiters.
Stream_LF, Stream_CR, Stream, Variable, Fix	HP-UX	Binary, Fixed	Fixed length records. Block mode reads are used; RMS record formatting information (if any) is transferred with the data. NFT inserts no record delimiters. The default record length is 512 bytes. The last record may be padded (with ASCII NULL characters) if the record length does not divide evenly into the total number of bytes in the source file (including any record formatting information).

PC Target Node

NOTE

Whenever a PC takes part in a file transfer using NFT, it *must* be the initiator. To transfer a file to or from a PC, you must initiate the transfer from the PC. If you are unfamiliar with the DOS operating system, refer to Section 2 of this manual or the *NFT User's Guide, HP OfficeShare Network* manual to learn more about the syntax and usage of NFT on a PC.

DOS Stream Files

The DOS operating system maintains a single type of file: a stream of bytes, usually called a **stream** file. The bytes may represent any kind of data, including ASCII characters. DOS imposes no structure on the file; under this scheme, a file is a place to read or write bytes. Higher level concepts, like "records," are not present in DOS.

However, many PC applications use a convention for finding the beginning and end of logical records of data. Under this convention, the next record in the file consists of all the data up to *but not including* the next occurrence of an ASCII CR character followed by an ASCII LF character. This sequence of characters - ASCII CR-LF - is called the record **delimiter** or **terminator**.

Most applications use the record delimiter only as a marker, and discard it as soon as they read it. NFT uses this interpretation unless you employ the Binary Interchange Format option; then NFT does not use record delimiters at all. Details about the Interchange Format options are discussed later in this subsection.

PC Logon Syntax

Because the PC is always local,* you do not need to give a node name; and because it is a single-user computer, there is no logon.

Therefore, when you specify the PC file, leave the *node_name* and *logon* fields blank. Supply only the file name, as described below.

PC File Name Syntax

<code>[device:][\][dir1\dir2\...\dirN\]filename.ext</code>
--

*Whenever a PC takes part in a file transfer using NFT, it *must* be the initiator. See the NOTE above.

The *device:* is used to specify a disc drive other than the active drive (or specify a shared disc directory on the server). Prefix the pathname with the device letter (for example, "A:").

The *dir1* through *dirN* parameters denote directories within a pathname. If the first character of the pathname is a backslash (\), the search starts from the root of the active or specified drive; otherwise the search starts from the current directory on the active or specified drive.

Directory and file names can always contain letters and numbers; check your DOS reference manual for restrictions if you want to use punctuation characters.

Each filename can be up to eight characters long; the optional extension has a three character maximum. The maximum pathname length is 63 characters.

NOTE

You cannot use NFT to transfer directories to a PC.

DEC VAX/VMS to PC Interchange Format Defaults

Any file copied from a DEC VAX/VMS computer to a PC creates a DOS stream file. All transfers occur in Interchange mode. The ASCII CR-LF pair is the default record delimiter in the target file. (See the Insert Character option).

Figure 6-4 show how DEC VAX/VMS stream files map to DOS stream files when no Interchange Format options are applied. The term "records" as used below means the data located between record delimiters in the DOS stream file.

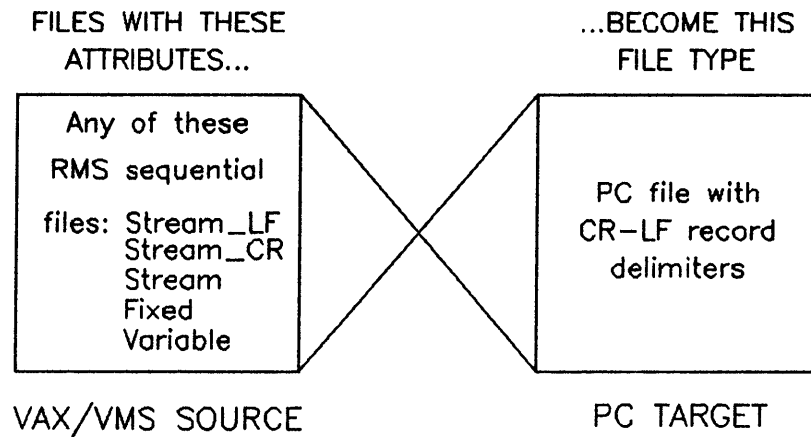


Figure 6-4. Interchange Format Defaults: VAX/VMS system to PC.

Interchange Format Options

Table 2-1 in Section 2 provides a quick reference of the syntax for using the options.

Because the PC must always be the initiator,* the following options are available.

Supported Interchange Format Options

The following Interchange Format options can be used when copying a DEC VAX/VMS computer source file to a PC:

- ASCII**
- Binary
- Fixed
- Record Size
- Insert Character
- Search Character**
- Variable**

The options marked with two asterisks (**) can be used without causing an error, but they have no effect in transfers between these two machines. They are not included in the following tables.

Unsupported Options

The following options cannot be used when copying an DEC VAX/VMS computer file to a PC:

- Append
- Direct
- File Size
- Interchange
- Sequential
- Strip

*Whenever a PC takes part in a file transfer using NFT, it *must* be the initiator. See the NOTE at the beginning of this subsection.

Interchange Format Option Results

Table 6-10 describes the results of using these options when copying a file from an DEC VAX/VMS system to a PC.

TABLE 6-10. INTERCHANGE FORMAT OPTIONS: VAX/VMS to PC.

Option Mnemonic	Description
Binary	<p>Source file records are placed directly in the target file without record delimiters. Block Mode I/O is used to read the source file. NFT does not append record delimiters in the target file. This option is required to transfer UDF and VFC files. Use the Record Size option to change the record size; longer records are truncated, shorter records are unaffected. Padding (with ASCII NULL characters) may occur at the end of the file if you also use the Fixed option (see Fixed). If you use this option, the Insert Character option is ignored. You cannot use the ASCII and Binary options together.</p> <p>Whenever you use the Binary option, be sure that you understand the type of file you are creating. In particular, a file created with Record Mode I/O and transferred using Block Mode I/O may contain RMS record formatting information along with the data.</p>
Fixed	<p>Source file records are truncated or padded as necessary so they are all the same length. The length defaults to the source file's assigned maximum record size if it exists, or to 512 bytes if the Binary option is in effect, or to 160 bytes otherwise; you can specify a length using the Record Size option. Records are padded as necessary with ASCII SPACE characters. The default terminator is the ASCII CR-LF pair.</p> <p>You can use the Binary and Fixed options together. If the number of bytes in the file is not evenly divisible by 512 (or the length you specify with the Record Size option), the last record is padded with ASCII NULL characters. When using the Binary option, record terminators are not appended.</p> <p>You cannot use the Fixed and Variable options together.</p>
Insert Character	<p>The argument to the Insert Character option specifies a single character that defines the record delimiter to use in the target file. The default record delimiter is an ASCII CR-LF pair; the Insert Character option defines only single-character delimiters. A request for an ASCII LF delimiter is not honored; the default is used instead.</p>

TABLE 6-10. INTERCHANGE FORMAT OPTIONS: VAX/VMS to PC (cont).

Option Mnemonic	Description
Record Size	<p>The target file is formed with a maximum record length as defined by the Record Size argument. If the source file contains a record that is longer than the Record Size argument, the record is truncated and NFT prints a warning message. The maximum record length is controlled by the VAX/VMS computer, and cannot exceed 32767</p> <p>If you use the Binary option along with the Record Size option, records are never truncated. In effect, the Record Size option is ignored.</p> <p>If you use the Fixed option along with the Record Size option, records in the source file are truncated or padded so that all records are the same length. Padding is with ASCII SPACE characters.</p> <p>If you use both the Binary and Fixed options along with the Record Size option, the last record may be padded with ASCII NULL characters if the number of bytes in the file is not evenly divisible by the length you specify.</p>

File Mappings

Table 6-11 lists the options required to map a particular RMS sequential file type into a PC file with certain attributes. The "Comments" column describes important details about each transfer. Details about each option are in Table 6-10. Other options (Record Size, for example) can also be used.

TABLE 6-11. FILE MAPPING FROM VAX/VMS SOURCE FILE TO PC TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
Stream_LF, Stream_CR, Stream, Variable, Fix	DOS	none	The ASCII CR-LF pair is the default record delimiter, and it is appended to each record.
Stream_LF, Stream_CR, Stream, Variable, Fix	DOS	Fixed	Fixed length records. The ASCII CR-LF pair is the default record delimiter. If the VAX/VMS system can provide the length of the records in source file, that is the default record length for target file; otherwise the record length defaults to 160 bytes. Longer records may be truncated. If you use the Record Size option, any necessary padding uses the ASCII SPACE character.

TABLE 6-11. FILE MAPPING FROM VAX/VMS SOURCE TO PC TARGET FILE (cont).

Source File Type	Target File Type	Options Required	Comments
Stream_LF, Stream_CR, Stream, Variable, Fix	DOS	Binary	Block mode reads are used; RMS record formatting information (if any) is transferred with the data. The target file is structurally identical to the source file. NFT inserts no record delimiters.
Stream_LF, Stream_CR, Stream, Variable, Fix	DOS	Binary, Fixed	Fixed length records. Block mode reads are used; RMS record formatting information (if any) is transferred with the data. NFT inserts no record delimiters. The default record length is 512 bytes. The last record may be padded (with ASCII NULL characters) if the record length does not divide evenly into the total number of bytes in the source file (including any record formatting information).

NOTE

The default record delimiter character is an ASCII CR-LF pair.

NOTE

Whenever a PC takes part in a file transfer using NFT, it *must* be the initiator. To transfer a file to or from a PC, you must initiate the transfer from the PC. If you are unfamiliar with the DOS operating system, refer to Section 2 of this manual or the *NFT User's Guide, HP OfficeShare Network* manual to learn more about the syntax and usage of NFT on a PC.

DOS STREAM FILES

The DOS operating system maintains a single type of file: a stream of bytes, usually called a *stream* file. The bytes may represent any kind of data, including ASCII characters. DOS imposes no structure on the file; under this scheme, a file is a place to read or write bytes. Higher level concepts, like "records," are not present in DOS.

However, many PC applications use a convention for finding the beginning and end of logical records of data. Under this convention, the next record in the file consists of all the data up to *but not including* the next occurrence of an ASCII CR character followed by an ASCII LF character. This sequence of characters, ASCII CR-LF, is called the record **delimiter** or **terminator**.

Most applications use the record delimiter only as a marker, and discard it as soon as they read it. NFT uses this interpretation unless you employ the Binary Interchange Format option; then NFT does not use record delimiters at all. Details about the Interchange Format options are discussed in each subsection.

PC Logon Syntax

The PC is always local.* When you specify a PC file, you do not need to supply a node name; and because it is a single-user computer, there is no logon.

Therefore, when you specify the PC file, leave the `node_name` and `logon` fields blank. Supply only the file name, as described next.

*Whenever a PC takes part in a file transfer using NFT, it *must* be the initiator. See the NOTE above.

PC File Name Syntax

```
[device:][\][dir1\dir2\...\dirN\]filename.ext
```

The *device:* is used to specify a disc drive other than the active drive (or specify a shared disc directory on the server). Prefix the pathname with the device letter (for example, A:).

The *dir1* through *dirN* parameters denote directories within a pathname. If the first character of the pathname is a backslash (\), the search starts from the root of the active or specified drive; otherwise the search starts from the current directory on the active or specified drive.

Directory and file names can always contain letters and numbers; check your DOS reference manual for restrictions if you want to use punctuation characters.

Each filename can be up to eight characters long; the optional extension has a three character maximum. The maximum pathname length is 63 characters.

NOTE

NFT does not support the use of any wildcard characters, and you cannot use NFT to transfer directories from a PC.

SUPPORTED TARGET NODES

PC source files can be copied to HP 1000, HP 3000, HP 9000, and DEC VAX/VMS computers that run HP AdvanceNet software.

The subsections of this section provide all the information you need to know when copying a PC source file to each supported target. This information includes:

- Target file system information.
- File name and logon syntax.
- Interchange Format defaults.
- The effect of each Interchange Format option on the file translation process, and how to create a specific target file type.

HP 1000 Target Node

HP 1000 computers that run NS/1000 have RTE-A operating systems. The RTE-A file system includes a record structure and categorizes files into certain **Types**. These file types are defined as follows:

- **Type 1.** These files have fixed-length records of 256 bytes and contain binary data.
- **Type 2.** These files also have fixed-length records, but the record length is defined by the user at file creation. They contain binary data.
- **Type 3.** These files have variable-length records, and NFT assumes they contain binary data.
- **Type 4.** These files also have variable-length records, and NFT assumes they contain ASCII data. Type 4 files are text files and may be altered with the RTE-A text editor, EDIT/1000.
- **Greater than Type 4.** File types 5, 6 and 7 have variable-length records. File types greater than 7 are user-defined. NFT assumes files of Type >4 contain binary data.

HP 1000 Logon Syntax

Logons for HP 1000 systems take the form:

accountname [/password]

The *accountname* parameter is the logon name for the computer. The password parameter is optional; if you use it, you must include a slash between the *accountname* and the *password* parameter. If no password is required, omit the slash. For password security, see the NFT manual for your PC.

HP 1000 File Name Syntax

The RTE-A file system has a hierarchical file structure. Files are catalogued in directories. Directories can also contain similar information about other directories, called **subdirectories**. Subdirectories have the same characteristics as directories; the term subdirectory means only that the directory is catalogued in the next higher level directory or subdirectory. Each account or logon on an RTE-A system has a *home* or default logon working directory. This directory is automatically available to the user when he or she logs on.

If the HP 1000 file resides in the hierarchical file system, the file name syntax is as follows:

[/[*directory*/]][*directory*/]...*filename*

PC Source File

If you omit the *directory* parameter, the default logon working directory for the logon specified is used. If you omit the slash (/) and specify a *directory*, the directory is assumed to be in the default logon working directory for the logon provided. The maximum file path name, including the file name, is 63 characters.

If the HP 1000 file you wish to access resides on an FMGR cartridge, specify the file name as follows:

filename::directory

The *filename* parameter may be a maximum of six characters. The *directory* parameter can be two ASCII characters or a positive or negative integer.

PC to HP 1000 Interchange Format Defaults

By default, any file copied from a PC source file to an HP 1000 creates an RTE-A Type 4 file. You can override the default with Interchange Format options. Figure 7-1 relates the RTE-A file attributes to the PC file type.

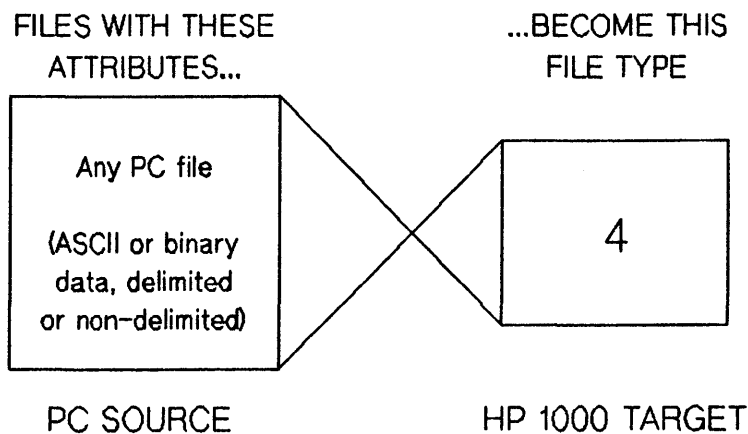


Figure 7-1. Interchange Format Defaults: PC to HP 1000

NOTE

Files copied to an HP 1000 system cannot have records longer than 4400 bytes (2200 words). You can truncate records by using the Record Size option.

Interchange Format Options

Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available at each initiator system.

Because the PC must always be the initiator,* the following options are available:

Supported Interchange Format Options

The following Interchange Format options can be used when copying a PC source file to an HP 1000 system.

- ASCII**
- Binary
- Fixed
- Insert Character**
- Record Size
- Search Character
- Variable**

The options marked with two asterisks (**) can be used without causing an error, but they have no effect in transfers between these two machines. They are not included in the following tables.

Unsupported Options

The following options cannot be used when copying a PC source file to an HP 1000:

- Append
- Direct
- File Size
- Interchange
- Sequential
- Strip

*Whenever a PC takes part in a file transfer using NFT, it *must* be the initiator. See the NOTE at the beginning of this section.

Interchange Format Option Results

Table 7-1 describes the results of using these options when copying a file from a PC to an HP 1000.

TABLE 7-1. INTERCHANGE FORMAT OPTIONS: PC to HP 1000.

Option Mnemonic	Description
Binary	<p>Source file data is divided into records every 256 bytes before being placed in the target file. Any record delimiters in the source file are transferred along with the data. You must use this option to create a binary target file. Unless you also use the Fixed option, the resulting HP 1000 file is Type 3 (see Fixed). Use the Record Size option to change the record size. You cannot use the ASCII and Binary options together.</p>
Fixed	<p>Data between record delimiters in the source file is truncated or padded as necessary (with ASCII SPACE characters) to the default length of 160 bytes, or the length specified with the Record Size option. The resulting records are copied to the target file. The resulting HP 1000 file is Type 4.</p> <p>You can use the Binary option in conjunction with the Fixed option to create a binary target file. The conversion process is the same as when the Binary option is used alone, except that the last record in the target file is padded with ASCII NULL characters if it is less than 256 bytes long. The resulting HP 1000 file is Type 1. The record size can be changed using the Record Size option; then the resulting HP 1000 file is Type 2.</p>
Record Size	<p>The target file is formed with a maximum record length as defined by the Record Size argument. If the source file contains a record (i.e., data between delimiters) that is longer than the Record Size argument, the record is truncated and NFT prints a warning message. The resulting HP 1000 file is Type 4. Record size cannot exceed 4400 bytes.</p> <p>If you use the Binary option with the Record Size option, delimiter characters in the source file are transferred along with the data, and no truncation occurs. The resulting HP 1000 file is Type 3.</p> <p>If you use the Fixed option with the Record Size option, data between delimiter characters in the source file is truncated or padded so the record length matches the specified length. Padding is with ASCII SPACE unless the Binary option is used, which changes padding to ASCII NULL.</p> <p>If you use both the Binary and Fixed options along with the Record Size option, delimiter characters in the source file are transferred along with the data, and the last record may be padded with ASCII NULL if it is shorter than the specified length. The resulting HP 1000 file is Type 2.</p>
Search Character	<p>The argument to the Search Character option specifies a single character that defines the record delimiter in the source file. This character is not transferred to the target file with the data. The default record delimiter is the ASCII CR-LF pair; the Search Character option allows only single-character delimiters to be defined.</p>

File Mappings

Table 7-2 lists the options required to map a PC file with particular attributes into a particular Type of RTE-A file. The "Comments" column of the table describes important details about each transfer. *Note that some pairs occur more than once in the table.* Other options (Record Size, for example) can also be used.

TABLE 7-2. FILE MAPPING FROM PC SOURCE FILE TO HP 1000 TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
DOS	Type 4	none	Use the Search Character option to extract logical records that are delimited by a character other than ASCII CR-LF pair.
DOS	Type 4	Fixed	The default record length is 160 bytes. Truncation or padding (with the ASCII SPACE character) may occur.
DOS	Type 4	Record Size	Truncation may occur if a source record is longer than the specified length. No padding occurs.
DOS	Type 3	Binary	All bytes (including any record delimiters) are considered data and transferred. The default record length is 256 bytes. No padding or truncation occurs.
DOS	Type 1, Type 2	Binary, Fixed	All bytes (including any record delimiters) are considered data and transferred. The default record length is 256 bytes. No truncation occurs. The last record may be padded if the number of bytes in the file is not evenly divisible by the record length. The default target file is Type 1 ; if you use the Record Size option and specify a length <i>other than 256</i> , the target file is Type 2 .

HP 3000 Target Node

HP 3000 computers utilize the Multiprogramming Executive (MPE) operating system. MPE files may contain ASCII or binary data, fixed or variable length records, and may be Relative I/O (RIO) or non-RIO files. (RIO files are random access, non-RIO files are sequential access.) A PC producer can only create non-RIO files on an HP 3000 target node.

Table 7-3 introduces four abbreviations for MPE files with particular attributes:

TABLE 7-3. ABBREVIATION OF MPE FILE ATTRIBUTES.

Abbreviation	MPE File Attributes
VA	Variable-length records of ASCII data.
VB	Variable-length records of Binary data.
FA	Fixed-length records of ASCII data.
FB	Fixed-length records of Binary data.

HP 3000 Logon Syntax

Each field in the HP 3000 logon sequence must begin with a letter, contain only alphanumeric characters, and must not exceed eight characters. The logon syntax is as follows:

```
username [/userpass].account [/acctpass][,groupname [/grouppass]]
```

HP 3000 File Name Syntax

Each field in an HP 3000 file name must begin with a letter, contain only alphanumeric characters, and must not exceed eight characters. The *group* and *account* parameters default to the logon group and account. The syntax of an HP 3000 file name is as follows:

```
filename [/lockword][.group [.account]]
```

PC to HP 3000 Interchange Format Defaults

Any file copied from a PC system defaults to an ASCII file with variable-length records (VA), regardless of what type of data the file actually contains. Record delimiters are not transferred with the data. The default delimiter is an ASCII CR-LF character pair. Figure 7-2 relates the PC file to the HP 3000 file attributes.




Figure 7-2. Interchange Format Defaults: PC to HP 3000.

Special Cases

In certain unusual situations, file transfers between PC systems and HP 3000s may produce unexpected results. These situations and their causes are described below.

Before a PC file is copied to an HP 3000 system, the HP 3000 allocates space for the target file. To determine how much space to preallocate, the HP 3000 must know the maximum file size (in number of logical records). This information is provided by the PC system. Unless the length of the target file's records are explicitly specified by the use of the Record Size option, the PC system approximates the file size (in number of logical records) by dividing the number of bytes in the file by 15. If the target file's record length is specified with the Record Size option, the file size will be the total number of bytes in the source file divided by the specific record length.

If the file size approximation is used, one of the following conditions might occur:

- A transfer of a PC ASCII file to an HP 3000 may result in inadequate preallocated space if the file's average record size is less than 15 bytes. If not enough space has been allocated, the transfer will be unsuccessful, and NFT will print an error message. You can circumvent this failure by using the Record Size option to specify the length of the records in the target file. (However, this will not work if you also specify the Fixed option.)

If you need to get fixed records in the target file, you should transfer the file from the PC to the HP 3000 without the Fixed option, then execute a local DSCOPY on the HP 3000, using the Fixed option.

PC Source File

- A transfer of a PC ASCII file may result in records being truncated if the source file's records are too long for the preallocated target file record size. The transfer does not fail, but a warning is displayed stating that records have been truncated to fit the target file record size. You can circumvent this failure by specifying a Record Size option value that is larger than the largest record in the source file.
- Another special case occurs if you use the Record Size option to specify an odd record length for the target file. If you specify an odd record length, the HP 3000 assigns one more byte to the record size. NFT prints a warning that the source and target file attributes differ. If record truncation is required, it occurs based on the odd record length specified, but target file records will have an extra byte appended to them. If the file is later copied back to a PC system, this extra byte is stripped off.

NOTE

When a PC file is copied to an HP 3000, an NFT warning message may be printed stating that the source and file attributes differ. This message does not indicate that the file copy was unsuccessful. It is printed as a result of the record size negotiation process that occurs between the PC system and HP 3000.

Interchange Format Options

Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available for each initiator system.

Because the PC system must always be the initiator,* the following options are available:

Supported Interchange Format Options

The following Interchange Format options can be used when copying a PC source file to an HP 3000 system.

- ASCII**
- Binary
- Fixed
- Record Size
- Insert Character**
- Search Character

*Whenever a PC system takes part in a file transfer using NFT, it *must* be the initiator. See the NOTE at the beginning of this section.

- Variable**

The options marked with two asterisks (**) can be used without causing an error, but they have no effect in transfers between these two machines. They are not included in the following tables.

Unsupported Options

The following options cannot be used when copying a PC source file to an HP 3000 target node:

- Append
- Direct
- File Size
- Interchange
- Sequential
- Strip

Interchange Format Option Results

Table 7-4 describes the results of using these options when copying a file from a PC to an HP 3000.

TABLE 7-4. INTERCHANGE FORMAT OPTIONS: PC to HP 3000.

Option Mnemonic	Description
Binary	Source file data is divided into records every 508 bytes before being placed in the target file. Any record delimiters in the source file are transferred along with the data. You must use this option to create a binary target file. Use the Record Size option to change the record size. You cannot use the ASCII and Binary options together.
Fixed	<p>Specifies that the target file records should be of fixed length. Data between record delimiters in the source file is truncated or padded (with ASCII SPACE characters), if necessary, to the default length of 160 bytes, or the length specified with the Record Size option. The resulting records are copied to the target file.</p> <p>You can use the Binary option in conjunction with the Fixed option to create a binary target file. The last record in the target file may be padded with ASCII NULL characters if it is less than 256 bytes long (or the length you specify with the Record Size option).</p> <p>You cannot use the Fixed and Variable options together.</p>

TABLE 7-4. INTERCHANGE FORMAT OPTIONS: PC to HP 3000 (cont).

Option Mnemonic	Description
Search Character	<p>The argument to the Search Character option specifies a single character that defines the record delimiter in the source file. This character is not transferred to the target file with the data. The default record delimiter is the ASCII CR-LF pair; the Search Character option allows only single-character delimiters to be defined.</p>
Record Size	<p>The target file is formed with a maximum record length as defined by the Record Size argument. If the source file contains a record (i.e., data between delimiters) that is longer than the Record Size argument, the record will be truncated and NFT will print a warning message. Record size cannot exceed (2**15)-1 bytes.</p> <p>If you use the Binary option along with the Record Size option, delimiter characters found in the source file are transferred along with the data, and no truncation occurs.</p> <p>If you use the Fixed option along with the Record Size option, data between delimiter characters in the source file is truncated or padded so that the record length matches the specified length. Padding is with ASCII SPACE characters.</p> <p>If you use both the Binary and Fixed options along with the Record Size option, delimiter characters in the source file are transferred along with the data, and the last record alone may be padded with ASCII NULL characters if it is shorter than the specified length.</p>

File Mappings

Table 7-5 lists the options required to map a PC file into an MPE file with certain attributes. The "Comments" column of the table describes important details about each transfer. Other options (Record Size, for example) can also be used.

TABLE 7-5. FILE MAPPING FROM PC SOURCE FILE TO HP 3000 TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
DOS	VA	none	<p>The default record length is 252 bytes. Use the Search Character option to extract logical records that are delimited by a single character rather than an ASCII CR-LF pair; the search character is not transferred with the data.</p>

TABLE 7-5. FILE MAPPING FROM PC SOURCE FILE TO HP 3000 TARGET FILE (cont)

Source File Type	Target File Type	Options Required	Comments
DOS	VB	Binary	The default record length is 508 bytes. You cannot use the Search Character option.
DOS	FA	Fixed	<p>The default record length is 160 bytes. If you use the Record Size option, any necessary padding uses the ASCII SPACE character. Use the Search Character option to extract logical records that are delimited by a single character rather than an ASCII CR-LF pair; the search character is not transferred with the data.</p> <p>There is a special case for files with a short average record length. Because of the way the PC system approximates the target file size, if the average record size is less than 15 bytes, the MPE system does not preallocate enough space for the file; the transfer will fail and you will receive an error message. Use the Record Size option to specify the length of the records in the target file; or add one or more long dummy records to the source file to increase the average record length to more than 15 bytes; or pad existing source records to increase the average record length to more than 15 bytes.</p>
DOS	FB	Fixed, Binary	The default record length is 256 bytes. You can specify another record length by using the Record Size option; the last record may be padded with the ASCII NULL character if the number of bytes in the source file (including the record delimiters) is not evenly divisible by the record length. You cannot use the Search Character option.

NOTE

The default search character is an ASCII CR-LF pair.

HP 9000 Target Node

HP 9000 computers have HP-UX operating systems. There is only one type of file on HP-UX systems: a stream of bytes. The bytes can represent ASCII characters or some type of binary data. There is no concept of records in the HP-UX file system. In fact, no structure is imposed on a file by the system, and no meaning is attached to its contents. The meaning of the bytes in a file depends solely on the programs that interpret that file by looking for ASCII LF characters. The ASCII LF characters are treated as record delimiters by these programs.

HP 9000 Logon Syntax

Each parameter in an HP 9000 logon may contain any ASCII character, with the exception of the colon (:), and must not exceed eight characters. The following syntax is an NFT convention and should not be used to log on to an HP-UX operating system except as part of a DSCOPY command.

```
username[:password]
```

HP 9000 File Name Syntax

The syntax of an HP 9000 file name is as follows:

```
[[/]dir1/dir2/.../dirn/]filename
```

The */dir1* and */dir2* parameters denote a directory within a path name. If the first character is a slash (/), the search starts from the root, otherwise the search starts from the default logon working directory. Directory and file names may contain any of the ASCII characters, except for NULL and slash (/), and can be up to fourteen characters long.

PC to HP 9000 Interchange Format Defaults

Any file copied from a PC to an HP 9000 creates an HP-UX stream file. All transfers occur in Interchange mode. The default record delimiter in the source file is the ASCII CR-LF pair; the default record delimiter in the target file is the ASCII LF character.

Figure 7-3 relates the HP-UX file attributes to the PC file type.

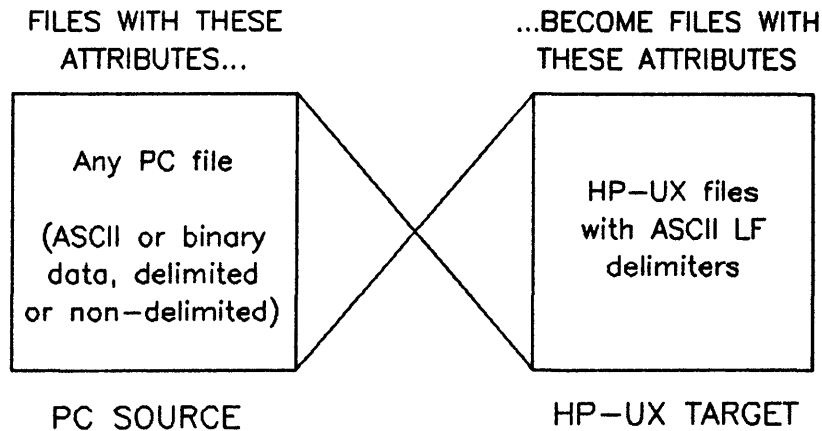


Figure 7-3. Interchange Format Defaults: PC to HP 9000.

Interchange Format Options

Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available at each initiator system.

Because the PC system must always be the initiator,* the following options are available:

Supported Interchange Format Options

The following Interchange Format options can be used when copying a PC source file to an HP 9000 system.

- ASCII**
- Binary
- Fixed
- Record Size
- Insert Character
- Search Character
- Variable**

The options marked with two asterisks (**) can be used without causing an error, but they have no effect in transfers between these two machines. They are not included in the following tables.

*Whenever a PC system takes part in a file transfer using NFT, it *must* be the initiator. See the NOTE at the beginning of this section.

Unsupported Options

The following options cannot be used when copying a PC source file to an HP 9000 target node:

- Append
- Direct
- File Size
- Interchange
- Sequential
- Strip

Interchange Format Option Results

Table 7-6 describes the results of using these options when copying a file from a PC system to an HP 9000.

TABLE 7-6. INTERCHANGE FORMAT OPTIONS: PC to HP 9000.

Option Mnemonic	Description
Binary	<p>Every byte in the source file is considered data and is eligible for transfer; this includes any record delimiters that may be present. When the Binary option is used alone, the target file is identical to the source file. The Search Character and Insert Character options are ignored.</p> <p>When the Fixed option is used together with the Binary option, padding with ASCII NULL characters may occur at the end of the file if the total number of bytes in the file cannot be evenly divided by the record size. The default record size is 256 bytes; use the Record Size option to specify a different size.</p>
Fixed	<p>Data between record delimiters in the source file may be truncated or padded to the length of the records specified with the Record Size qualifier (note defaults below). The resulting records are copied to the target file.</p> <p>If you do not specify a record size, the size defaults to 160 bytes. Any required padding is accomplished with ASCII SPACE characters.</p> <p>If you use the Binary option together with the Fixed option, see Binary description above.</p>

TABLE 7-6. INTERCHANGE FORMAT OPTIONS: PC to HP 9000 (cont).

Option Mnemonic	Description
Insert Character	<p>The argument to the Insert Character option specifies a single character that defines the record delimiter to use in the target file. The default record delimiter is the ASCII LF character; the Insert Character option defines only single-character delimiters. You can give the parameter either as a literal, or as the decimal ASCII equivalent, preceded by a zero.</p>
Search Character	<p>The argument to the Search Character option specifies a single character that defines the record delimiter in the source file. Specify the Search Character as a literal, or in its decimal ASCII value prefixed with a zero. This character is not transferred to the target file with the data. The default record delimiter is the ASCII CR-LF pair; the Search Character option allows only single-character delimiters to be defined.</p>
Record Size	<p>No record in the target file will be longer than the length given in the option parameter. Any or all records may be shorter. If any record is truncated, NFT will print a warning message. There is no effective limit to record size.</p> <p>If you use the Fixed option along with the Record Size option, records are padded or truncated (as necessary) to the length specified in the option parameter.</p> <p>If you use the Binary option along with the Record Size option, delimiter characters in the source file are transferred along with the data, and no truncation occurs. In effect, the Record Size option is ignored.</p> <p>If you use both the Binary and Fixed options along with the Record Size option, delimiter characters in the source file are transferred along with the data, and the last record may be padded with ASCII NULL characters if it is shorter than the specified length.</p>

File Mappings

Table 7-7 lists the options required to map a PC file into an HP-UX file with certain attributes. The "Comments" column of the table describes important details about each transfer. Details about each option are in Table 7-6. Other options (Record Size, for example) can also be used.

TABLE 7-7. FILE MAPPING FROM PC SOURCE FILE TO HP 9000 TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
DOS	HP-UX	none	Record delimiters in the source file are not transferred with the data; the default record delimiter for the target file is ASCII LF character, and it is appended to each record.
DOS	HP-UX	Fixed	Fixed length records. Record delimiters in the source file are not transferred with the data; the default record delimiter for the target file is an ASCII LF character, and it is appended to each record. The record length defaults to 160 bytes. Longer records may be truncated. Any necessary padding uses the ASCII SPACE character.
DOS	HP-UX	Binary	Record delimiter characters (if any) are transferred with the data. The target file is structurally identical to the source file. NFT does not insert record delimiters.
DOS	HP-UX	Binary, Fixed	Fixed length records. Record delimiter characters (if any) are transferred with the data. NFT inserts no record delimiters. The default record length is 512 bytes. The last record may be padded (with ASCII NULL characters) if the record length does not divide evenly into the total number of bytes in the source file (including any record delimiter characters). Use the Record Size option to change the record length.

NOTE

The default record delimiter in the source file is the ASCII CR-LF pair.

DEC VAX/VMS Target Node

You can exchange files with a DEC VAX/VMS computer that is running "Network Services for the DEC VAX Computer" software.

DEC VAX/VMS computers use the Record Management Services (RMS) file system. RMS provides several different types of files with various record formats. NFT on a DEC VAX/VMS computer supports most RMS sequential file formats for transfers to VAX/VMS nodes. This manual uses the following abbreviations for VAX/VMS record formats:

TABLE 7-8. ABBREVIATION OF VAX/VMS TARGET FILE TYPES.

Abbreviation	File Type
Stream_LF	Stream file with linefeed record terminators.
Stream_CR	Stream file with carriage return record terminators.
Stream	Stream file with record terminators other than carriage return. These terminators could be FF, VT, LF, or CR-LF.
Variable	Variable record format file.
Fix	Fixed record format file.

VAX/VMS Logon Syntax

Logons for the VAX/VMS take the form:

```
username[:password]
```

The *username* parameter is a logon name for the computer. The *password* is optional; a colon (:) must separate the *username* from the *password* parameter. For password security, see the *NFT User's Guide*, *HP OfficeShare Network* manual.

VAX/VMS File Name Syntax

Under VAX/VMS, files reside in a hierarchy of devices, directories and subdirectories. The device is the physical unit (e.g., the disc) where the file exists. Files can be catalogued in any directory or subdirectory on the device. A subdirectory is a directory that is catalogued in the next higher level directory or subdirectory. Each account or logon on a VAX/VMS system has a default logon device and working directory, which is immediately available to the user when he or she logs on.

A file specification consists of a device name, followed by a directory path, followed by the file name. The general syntax of a VAX/VMS file specification is as follows:

device: [[*directory*]]*file_name.type;version*

- The *device* is a device name which, if omitted (along with the colon), defaults to the current device, commonly the logon device.
- The *directory* is a directory name, which defaults to the current working directory. The *directory* can be an absolute directory path, or it can be relative to the current working directory. To specify a subdirectory, include a period and a subdirectory name before the closing bracket. You can use a subdirectory several layers deep by chaining subdirectory names together with periods between.

NOTE

Square brackets - "[" and "]" - are required by the VAX/VMS file name syntax whenever a directory or subdirectory is specified. Omit the brackets when you use the default directory. Note that if you specify a *device*, there is no default *directory*.

- Either the *file_name* or the *type* (or both) must be present, and the period between them is *always* present.
- The *version* is an integer, which you can omit. A file of the same name (and version, if specified) in the target directory will only be replaced if you use the Replace option. If you do not use the Replace option and a file of the same name exists in the target directory, a file with an incremented version number is created.

Each field except the version number in a VAX/VMS file name must begin with a letter and contain only alphanumeric characters (including hyphen and underscore) without spaces.

NOTE

DSCOPY can be made to automatically execute the user's LOGIN.COM. This command file is used to define logical names that NFT can use, or it can set the default directory. To have DSCOPY execute the LOGIN.COM command file, you must set the DTRACE bit to 283. The manual, *Network Services for the DEC VAX Computer*, describes how to do this, and the details of how it affects your NFT transfers.

DECnet VAX File Name Syntax

If your AdvanceNet VAX is also connected to DECnet, you can copy files to any DECnet node from a PC. This "network-to-network" transfer requires that you explicitly route the copy through the VAX/VMS node that is running both AdvanceNet and DECnet. Therefore, in the target file *node_name* and *login* fields of the *dscopy* command you specify the VAX/VMS node that is running *both* AdvanceNet and DECnet. You then use the *file_path_name* field to specify the DECnet node, along with the appropriate device, directory, etc. as shown. The file path name syntax to specify a file on a DECnet node is as follows:

```
dec_node"username password":VAX/VMS_file_name
```

- The *dec_node* field is the name of the DECnet node.
- A space separates *username* and *password*.
- The *VAX/VMS_file_name* is described in "VAX/VMS File Name Syntax" above.

The string that contains the user name and password is preceded and followed by quotation marks, and always contains a blank as noted above. These characters have special meaning to the initiator; therefore, if your transfer requires an access control string, you must escape the special meaning by using backslashes and quotes in the file specification. The following shows the required syntax. Note the backslash (\) preceding the embedded quotes, and the enclosing quotes:

```
dec_node\"username password\"":VAX/VMS_file_name
```

For additional information about using escape characters, see the *NFT User's Guide, HP OfficeShare Network* manual.

File Access Modes

RMS does not differentiate between ASCII and binary data types. VAX/VMS consumers handle the Binary option differently. NFT on a VAX/VMS computer operates as follows:

- WITHOUT THE BINARY OPTION:

Without the Binary option, when a VAX/VMS consumer receives data records from a PC producer, it adds headers or terminating characters to each record depending on the target file type. This is called **Record Mode I/O**. This record formatting information is inserted for the exclusive use of RMS, and is generally not available (or useful) to the user.

- WITH THE BINARY OPTION:

If you use the Binary option, the VAX/VMS consumer writes blocks of data without attention to record structure. *No record formatting occurs*; the contents of the source file are copied directly. This is called **Block Mode I/O**. Unless the target file is of type **Fix**, the absence of record formatting means you may not be able to access individual records in the target file.

When you use the Binary option, be sure that you understand the type of file you are creating. The lack of record formatting can have important implications depending on the type of the target file. In particular, when transferring binary data to an RMS **Stream_LF**, **Stream_CR**, **Stream** or **Variable** type file, the lack of record formatting information in the target file may restrict the ability of RMS to access the records.

PC to VAX/VMS Interchange Format Defaults

Any file copied from a PC to a VAX/VMS target node defaults to a **Variable** file. Figure 7-4 relates the PC file attributes to the VAX/VMS file types.

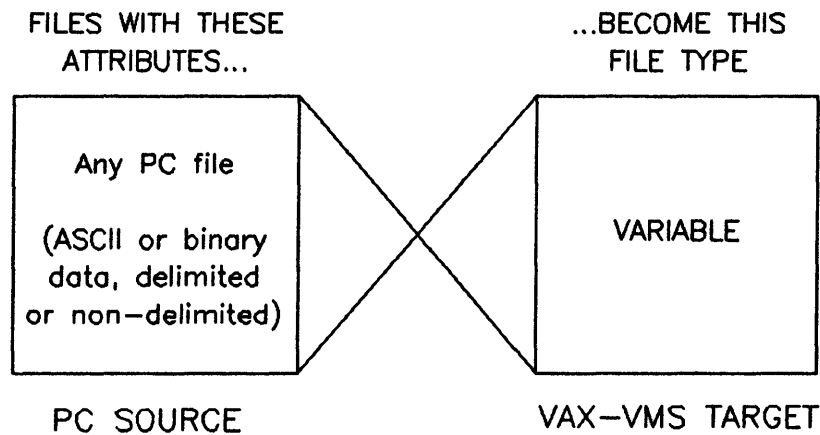


Figure 7-4. Interchange Format Defaults: PC to VAX/VMS.

Interchange Format Options

Table 2-1 in Section 2 provides a quick reference of the Interchange Format option syntax available at each initiator system.

Because the PC system must always be the initiator,* the following options are available:

*Whenever a PC system takes part in a file transfer using NFT, it *must* be the initiator. See the NOTE at the beginning of this section.

Supported Interchange Format Options

The following Interchange Format options can be used when copying a PC source file to a DEC VAX/VMS system.

- ASCII**
- Binary
- Fixed
- Record Size
- Insert Character
- Search Character
- Variable**

The options marked with two asterisks (**) can be used without causing an error, but they have no effect in transfers between these two machines. They are not included in the following tables.

Unsupported Options

The following options cannot be used when copying a PC source file to a DEC VAX/VMS target node:

- Append
- Direct
- File Size
- Interchange
- Sequential
- Strip

Interchange Format Option Results

Table 7-9 describes the results of using these options when copying a file from a PC system to a VAX/VMS target node.

TABLE 7-9. INTERCHANGE FORMAT OPTIONS: PC to DEC VAX/VMS.

Option Mnemonic	Description
Binary	<p>Creates a Variable type target file on the VAX/VMS node. Specifies that the transfer uses Block Mode I/O (instead of Record Mode I/O). Block Mode I/O causes NFT on the VAX/VMS computer to forgo internal record formatting and accept everything "verbatim" in a block of data. Thus, any record delimiters in the source file are transferred along with the data. If you also use the Fixed option, a Fix type target file is created (see Fixed).</p>
Fixed	<p>Creates a Fix type target file on the VAX/VMS computer. Data between record delimiters in the source file may be truncated or padded to the length of the records specified with the Record Size qualifier (note defaults below). The resulting records are copied to the target file.</p> <p>If you do not specify a record size, the size defaults to 160 bytes. Any required padding is accomplished with ASCII SPACE characters.</p> <p>If you use the Binary option in conjunction with the Fixed option, and do not specify a record size, the size defaults to 256 bytes. No truncation occurs. The last record in the target file may be padded with ASCII NULL characters if it is less than 256 bytes long (or the length you specify with the Record Size option).</p>
Insert Character	<p>Specifies the record delimiter character in the target file, and forces the target file on the VAX/VMS computer to be as follows:</p> <p>If no parameter is used, or if the parameter is 010, a Stream__LF file is created. (010 = decimal ASCII code for LF).</p> <p>If the parameter is 013, a Stream__CR file is created. (013 = decimal ASCII code for CR).</p> <p>If the parameter is anything else, a Stream file is created.</p> <p>Note that the file will have the specified type regardless of what other options, if any, are used. If you use the Binary option, however, this option only determines the file type of the target file. No record delimiters are appended.</p>
Search Character	<p>The argument to the Search Character option specifies a single character that defines the record delimiter in the source file. Specify the Search Character as a literal, or in its decimal ASCII value prefixed with a zero. This character is not transferred to the target file with the data. The default record delimiter is the ASCII CR-LF pair; the Search Character option allows only single-character delimiters to be defined.</p>

TABLE 7-9. INTERCHANGE FORMAT OPTIONS: PC to DEC VAX/VMS (cont).

Option Mnemonic	Description
Record Size	<p>No record in the target file will be longer than the length given in the option parameter. Any or all records may be shorter. If any record is truncated, NFT prints a warning message. Record size cannot exceed 32765 bytes for Variable, 32766 for Stream, and 32767 for Fixed target files.</p> <p>If you use the Fixed option along with the Record Size option, records are padded or truncated (as necessary) to the length specified in the option parameter.</p> <p>If you use the Binary option along with the Record Size option, delimiter characters in the source file are transferred along with the data, and no truncation occurs.</p> <p>If you use both the Binary and Fixed options along with the Record Size option, delimiter characters in the source file are transferred along with the data, and the last record may be padded with ASCII NULL characters if it is shorter than the specified length.</p>

File Mappings

Table 7-10 lists the options required to map a PC source file into a particular RMS sequential file type. *Note that pairs occur more than once in the table.* The "Comments" column of the table describes important details about each transfer. Other options (Record Size, for example) can also be used.

TABLE 7-10. FILE MAPPING FROM PC SOURCE FILE TO VAX/VMS TARGET FILE.

Source File Type	Target File Type	Options Required	Comments
DOS	Variable	none	Record delimiters in the source file are not transferred, only the data between delimiters. If you use the Record Size option, source file records that exceed the length specified in the parameter are truncated; shorter records are not affected.
DOS	Stream_LF, Stream_CR, Stream	Insert Character	If parameter is LF, target file is Stream_LF. If parameter is CR, target file is Stream_CR. If parameter is anything else, target file is Stream with CRLF delimiters.

TABLE 7-10. FILE MAPPING FROM PC SOURCE TO VAX/VMS TARGET FILE (cont).

Source File Type	Target File Type	Options Required	Comments
DOS	Fix	Fixed	The default record length is 160 bytes. You can use the Search Character option to extract logical records from the source file. Padding (if necessary) is with the ASCII SPACE character. Truncation may occur.
DOS	Variable	Binary	Block mode writes are used. Any record delimiters in source file are transferred with the data. The default record length is 256 bytes.
DOS	Stream_LF, Stream_CR, Stream	Binary, Insert Character	Block mode writes are used. Any record delimiters in source file are transferred with the data. The default record length is 256 bytes. If the parameter of the Insert Character option is LF, the target file is Stream_LF. If parameter is CR, target file is Stream_CR. If parameter is anything else, target file is Stream. Note that when you use the Insert Character option with the Binary option, the target file has the specified type, but because Block Mode is used to write the file, <i>NFT does not append any record delimiters.</i>
DOS	Fix	Binary, Fixed	Block mode writes are used. Any record delimiters in source file are transferred with data. The default record length is 256 bytes. The last record may be padded with ASCII NULL characters if the number of bytes in the file (including any record delimiters) is not evenly divisible by the record length.

NOTE

The default search character is the ASCII CR-LF pair.

Customer Order Number:

5958-8563 E0891

Printed in U.S.A. 08/91

**** For HP Internal Reference Only ****

Manufacturing Part Number:

5960-1634

