0 Site Operations

symbolics™

Revised for Genera 7.2

Cambridge, Massachusetts

Site Operations
 # 999050

February 1988

This document corresponds to Genera 7.2 and later releases.

Symbolics, Inc.
11 Cambridge Center
Cambridge, MA 02142

Text written and produced on Symbolics 3600™-family computers by the Documentation Group of Symbolics, Inc.
Text masters produced on Symbolics 3600™-family computers using Symbolics Concordia™, a system for supporting document development, and printed on Symbolics LGP2 Laser Graphics Printers.

Printed in the United States of America.

Printing year and number: 90 89 88 5 4 3 2 1

Table of Contents

List of Figures

1. Introduction to Site Operations

This document provides information for various operations that are necessary to maintain Symbolics machines at your site. These include the following:

- performing backups and dumps
- using the distribution subsystem
- using the carry tape system
- using the FEP-Tape system
- adding LMFS partitions
- maintaining the file system
- saving new versions of world loads
- distributing worlds to your users

2. Booting, Netbooting, and Autobooting

Once you receive a distribution world from Symbolics you have many choices about modifying and customizing that world to suit your needs.

Here is a typical sequence of events:

1. Boot the original distribution world on one machine.
2. Make changes to the world to customize it for your site.
3. Repeat the process for as many customized worlds as your site needs.
4. Save the site-customized worlds to use over again, making either complete new worlds or using Incremental Disk Save (IDS).
5. Distribute the site-customized worlds to other machines at the site.
6. Boot the new worlds on the user machines.

In addition, individual users at your site may wish to add further customizations for the worlds they will be using.

Your choices for distributing the new worlds to individual systems are to copy the world to the local disk of each machine and boot it from there, or to use netbooting to boot a world from a remote machine.

2.1. Booting a World

Booting a world is handled by the Front-end Processor, or FEP. You give commands to the FEP to boot the world. A world to boot can be found either on the local disk of the machine you wish to boot or on a remote netboot server.

If you wish to load a world from the local disk, issue the Load World command to the FEP. If you wish to load a world from a remote netboot server, issue the Netboot command to the FEP. These actions can be combined. See the section "Netbooting IDS Worlds", page 12.

Netbooting does not necessarily save a lot of disk space. In netbooting, the disk space not used for a world-load file is used for another paging file. The extra paging file enables the netbooted world operate as efficiently as a world loaded from the local disk. Users who are netbooting do not, however, have to have enough room for two worlds on their local disk, so to that degree netbooting saves disk space. See the section "Netbooting", page 8.

Here is the sequence of commands for a boot file that boots a world from the local disk.

```
Clear Machine
Declare Paging-files files-names
Load Microcode microcode-file-name
Load World distribution-world-file-name
Enable IDS
Set Chaos-Address this-machine's-chaos-address
Start
```

The sequence of commands for netbooting is quite similar. The Load World command with the pathname of a world is replaced by a Netboot command with the description of a world.

```
Clear Machine
Declare Paging-files files-names
Load Microcode microcode-file-name
Netboot world-description
Enable IDS
Set Chaos-Address this-machine's-chaos-address
Start
```

If you encounter any errors with a particular command, try that command again.

Declare Paging-files declares *file-names* to be the paging files for all subsequent Load World commands until a new Declare Paging-files command overrides it. This command replaces the Add Paging-files command in boot files. For more information about the Declare Paging-files command: See the section "FEP System Commands: General Usage". *microcode-file-name* refers to the microcode needed for a specific system and hardware configuration. For example, if you are installing Genera 7.0 on a 3640, then the appropriate microcode file name is 3640-MIC.MIC.*version-number*. *Version-number* is the released microcode version, which changes from release to release. Note that if you have an IFU, you need a different microcode. For a list of Genera 7.2 microcode types: See the section "Genera 7.2 Microcode Types".

In the Load World command *distribution-world-file-name* refers to the name of the world distributed on the disk (for example, >release-7-2-dist.load).

In the Netboot command *world-description* refers to a world kept on a netboot server. See the section "World Description for Netbooting".
This-machine's-chaos-address refers to the chaosnet address of the Symbolics computer. You must select a chaos address for each machine at your site. For more information about chaosnet addresses: See the section "Choosing Chaosnet Addresses". It is uneccessary to type all these commands every time you boot a world. Put the correct series of commands in a boot file (default name `boot.boot`) and put that file on your local disk. You can edit a boot file at any time, and you can also have more than one boot file.

The FEP Boot command takes the name of a boot file and executes all the commands in the file.

You can have more than one boot file. This is useful if your site uses various world loads, some of which may have special programs loaded into their environment. You can also have boot files to netboot and boot files to boot a world from the local disk.

2.2. Transferring Worlds to Other Machines

Once you have loaded a distribution world on one machine, you may want to transfer it to another machine. There are two methods of doing this, copying the world to the local disk of the other machine or netbooting. This section describes how to copy the world to another machine. For more information on netbooting: See the section "Netbooting", page 8.

Use the Copy World command to copy worlds. When a machine receives a new distribution world, it may also need new microcode. See the section "Genera 7.2 Software Installation Guide".

Before transferring a world to another machine, it is necessary to look at the contents of FEP directory on that machine, to see what world loads and microcode files are currently on the machine. To look at the FEP directory, use the Show FEP Directory command.

Show FEP Directory Command

Show FEP Directory *keywords*

Displays a description of the FEP files on the local host. The :Host keyword allows you to specify another host.

keywords :Format, :Highlight Files In Use, :Highlighting Mode, :Host :Output Destination, :Type, :Unit

:Format {Normal, Detailed} How much information to include in the display. The default is Normal, meaning file name, length in blocks, and file comment (if any) are displayed for each file. Detailed means that all information, including creation date and author, is displayed.

:Host A host on the network. The default is local.

:Highlight Files In Use
{Yes, No} Whether to indicate files that are currently in use. The default is Yes.

:Highlighting Mode
{Bold, Arrow} How to indicate that a file is in use. Bold means display the filename in boldface, Arrow means prefix the filename with an arrow. The arrow is useful from a remote terminal. The default is Bold.

:Output Destination
: {Buffer, File, Printer, Stream, Window} Where to redirect the typeout done by this command. The default is the stream ***standard-output***.

:Type
: {All, Boot, Fep, LMFS, Microcode, Paging, World} The type of file(s) to display information about. The default is All. More than one type can be given, separated by commas.

:Unit
: {*disk-unit-number* All} The default is All. *disk-unit-number* is an integer, interpreted as a disk unit number on the specified host.

Show FEP Directory first displays an estimate of the number of free blocks and the proportion of blocks used on each disk unit. It then displays a summary of the files on each unit grouped by file type.

The following function copies a world from another machine:

Copy World Command

Copy World *file destination keywords*

Makes a copy of a world load. This includes the original *parent world* as well as all modifications made in the form of an Incremental Disk Save (IDS) World. See the section "Incremental Disk Save (IDS)", page 27. Copy World also also copies a netboot cores. See the section "Netbooting", page 8.

file
: FEP file spec. The world to copy. The default is constructed from the version of the world that you have booted.

destination
: FEP file spec. The file pathname of the new world. The default is a wildcard pathname assuring the correct hierarchical pathname relationship for the *parent* world and an *IDS* world.

After you issue the Copy World Command, Genera puts up a menu allowing you to specify the actions you want it to take in the course of the copy:

```
There are 40450 blocks total in the files to be transferred.
  There are 809 blocks free on FEP1.
  You need 39641 more blocks.
  Possible actions to make space right now: Run dired  Expunge directory  Selectively delete

Parent IDS files to transfer: All parents  Missing parents  Just requested files  Selective
Update boot file to load FEP1:>Base-System-372-0.load.1: Yes  No
  Boot file to update: FEP0:>boot.boot.newest
  Update FEP0:>boot.boot.newest to load microcode 410: Yes  No
Transfer mode: Transfer-And-Checksum  Transfer-Only  Checksum-Only
Attempt automatic error recovery: Yes  No
<ABORT> aborts, <END> uses these values
```

Figure 1. Copy World

keywords :Automatic, :End Block, :File Set, :Output Destination, :Query, :Start Block, :Transfer Mode, :Update Boot File.

:Automatic {Yes, No} Whether or not to attempt automatic error recovery. The default is yes.

:End Block {*integer*} The number of the last block to copy from source. The default is the last block, meaning copy until the end.

:File Set {All parents, Missing Parents, Just Requested Files, Selective} Which parent IDS files to transfer. The default is Missing Parents.

:Output Destination {Buffer, File, Printer, Stream, Window} Where to redirect the typeout done by this command. The default is the stream ***standard-output***.

:Query {Yes, No} Whether or not to present a menu of transfer parameters. The default is yes.

:Start Block {*integer*} The number of the block to start copying from the source. The default is 0, meaning begin at the beginning.

:Transfer Mode {Transfer-and-Checksum, Transfer-Only, Checksum-Only} Whether to verify the integrity of the copied world. The default is Transfer-and-Checksum. You can use Checksum-Only to checksum a band that you copied previously but were unable to checksum due to network problems.

:Update Boot file {*FEP-file-spec*, none}. Boot file to update to load the new world. The default for IDS or complete worlds is boot.boot. The default for netboot cores is none.

You do not have to have a world load resident on your FEP to boot a world. You can boot a world from a remote world server with only a netboot core on your FEP. See the section "Netbooting", page 8.

You install microcode on a specific machine using the Copy Microcode command. You can do this after installing the new microcode at the site, with the distribution loader.

Copy Microcode Command

Copy Microcode *{version or pathname} destination keywords*

Installs a version of microcode.

version or pathname
Microcode version number or pathname to copy. *version* is a microcode version number (in decimal). *pathname* rarely needs to be supplied. It defaults to a file on FEP*n*:> (where *n* is unit number of the boot disk) whose name is based on the

microcode name and version. (The file resides in the logical directory sys:l-ucode;.) The *version* actually stands for the file *appropriate-hardware*-MIC.MIC.*version* on FEP*n*:>. (See the Section "Genera 7.2 Microcode Types" in *Software Installation Guide*)

destination — FEP file specification. The pathname on your FEP*n*:> directory. The default is created from the microcode version.

keywords — :update boot file

:update boot file — {*FEP-file-spec*, None, Query}. The pathname of the boot file you want it to update. The default is the current default boot file name.

The logical directory *sys:l-ucode;* includes multiple types of microcode for each version number. The correct microcode to install depends upon the particular hardware configuration of your machine. When your machine is shipped, the default microcode filename is correct, but if your machine is upgraded (for example, an FPA board is installed) you might need to override the default used by the Copy Microcode command to get the correct microcode type for your configuration. Below is an example of how you would get the microcode type for your configuration.

```
Copy Microcode 3640-fpa-mic.mic.389
```

If you use the wrong microcode for your configuration, your machine will not boot, except in the case where your system has an FPA and you use a non-FPA microcode. In this case, the machine functions normally, but does not make use of the FPA at all.

If you are copying copying microcode from a machine running Genera 7.0 to a machine running Release 6.1, **you must specify the microcode file name**, rather than just specifying the microcode version. The name of the microcode files changed from Release 6.1 to Genera 7.0, and the Copy Microcode command does not understand the name change. It tries to find a microcode version with the same name components as the old microcode.

For example, to distribute a new microcode using the Copy Microcode command on the machine that will receive the new microcode, type:

```
Copy Microcode   SYS:L-UCODE; 3640-mic.mic.389
```

For a list of microcode types: See the section "Genera 7.2 Microcode Types".

When the Copy Microcode command asks if you want to update your boot file, answer Yes. The file is now updated.

2.3. Netbooting

Netbooting is the ability to boot a world from a remote machine and run it without having the booted world-load file on the local disk.

Netbooting should be considered a site-management tool, rather than a disk-conservation tool. While you do not have to have a world-load file on your local disk, you should add an extra paging file of the same size (approximately 35,000 blocks) if you want equivalent performance. On the other hand, you never need to have two worlds on your local disk.

Each user needs only a netboot core (and a backup netboot core) on the local disk to take advantage of netbooting. Netboot cores have the `.load` file type, but use only approximately 100 blocks each. Of course, you continue to have other files on the local disk, such as microcode files, FEP overlays (flods), paging files, boot files, and the like.

There are no clear rules on when netbooting is the correct choice. Symbolics is interested in hearing from users about how they use netbooting.

In the meantime, some general principles can be suggested.

The choice of booting methods is a site-management issue. The parameters are the number of machines at the site, the number of worlds commonly used, and the number of users.

Netbooting a world is slower than booting a world from the local disk. On the other hand, netbooting is much faster than copying a world to the local disk and then booting it. This means that users who reboot the same world frequently may prefer booting from the local disk.

Users who run many different worlds on the same machine may want to pay the price of slower boot time in exchange for eliminating the need to constantly juggle paging and world space on the local disk.

If several worlds are in common use at your site, netbooting may be a better choice because you'll only have to maintain two copies of each world (one of them a backup) per site, instead of one copy for each user. Netbooting is particularly useful if systems loaded in those worlds are patched frequently.

Here is the setup for a site using netbooting.

- One or more netboot servers with world-load files for all worlds used at the site.

- User machines, each with a netboot core. This is a special kind of world-load file used to accomplish netbooting. Each user machine must also have a boot file including the Netboot command at the point that the Load World command usually appears.

Here are the limitations on netbooting:

- Netboot servers must run Genera 7.2 or a later release.

- Only worlds based on Genera 7.2 or a later release can be netbooted.
- Netbooting works only on the local Ethernet. You can netboot only when there is a netboot server on the same subnet.

2.3.1. Netboot Servers

Here is how you set up a system as a netboot server:

- Use the Copy Flod Files command to copy the new NFEP-Overlays to the FEP file system on the server running the Genera 7.2 world.
- Add the following services to the server's Host Object in your local namespace:

```
Service: NETBOOT SLAP NETBOOT
  . . .
Server Machine: Yes
```

  The effect of `Server Machine: Yes` is to cause the server to boot with services disabled. This prevents any user system from netbooting before the server itself is completely booted. This means you must enable services after the server system is booted.

  See the section "Enable Services Command" in *Genera Handbook*. See the function **sys:enable-services** in *Networks*.

- Copy all world-load files that users will be netbooting to a top-level FEP directory. This can be any top-level directory on the netboot server.

When you enable services on the netboot server, netboot service is enabled by default. You can also enable netboot service by name.

Netbooting includes a queuing mechanism. Once a user machine requests a world to netboot, no further user attention is necessary.

Netboot service is provided serially, on a first-come-first-serve basis. This minimizes boot time. Sites with many machines may find this a problem when they all try to boot at once, such as after a power failure.

Therefore you will probably want to have a backup netboot server, but it need not have netboot service enabled. If the primary netboot server goes down, you must enable netboot service on the backup so users can access the worlds there.

A server can be a user machine, but the system may be slow while it is serving up a world. A user machine is fine as a backup server. Symbolics does not recommend using a file server as a netboot server because file service may be severely degraded when the server is also serving a netbooting client.

You can have two netboot servers without having identical sets of worlds on them. The queuing mechanism can find a world on any netboot server.

If you have the same world on several servers, take care that the worlds are really the same, that is, at the same patch level. Otherwise, there is some possibility of booting the older version. If you decide to stop supporting a world on a particular server, you should delete all copies of the world. They need not be expunged; marking them for deletion is good enough.

2.3.2. Netbooting User Systems

Here is how you set up a user machine for using netbooting:

- Use the Copy Flod Files command to copy the Genera 7.2 FEP overlay files (flods) to the FEP file system. See the section "Copy Flod Files Command", page 81. Copy the netboot core to the FEP file system of the local machine using the Copy World command. This file has the same name as the distribution world, prefixed by `Netboot-core-from-`, such as `Netboot-Core-from-7-2.load`, and should be found on the FEP file system of the same machine as the distribution world.

  If you have multiple disks, it is best to put the core on FEP0. (But if you explicitly mount another disk in your boot file, before the Netboot command, you can safely put the netboot core on that disk.)

- Add extra paging space, equivalent to the size of the world you will netboot. This paging space is in addition to the paging space you normally use. You must add this extra paging space if you want the netbooted world to have equivalent performance to a world booted from the local disk.

Then, to netboot, use the FEP Netboot command in your boot file where you would normally use the Load World command.

2.3.2.1. Netboot

Netboot *world-description*

Netboots the world described by *world-description*. Polls the netboot servers on the local subnet for worlds that match *world-description* and netboots the most recent of those worlds.

Here is an example:

```
Netboot chip-poker
```

In effect, the Netboot command replaces Load World if you are netbooting.

See the section "Netbooting", page 8.

See the section "World Description for Netbooting".

The Load World command can initiate netbooting in certain circumstances. See the section "Netbooting IDS Worlds", page 12.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

(In some cases, the Load World command can also initiate netbooting. See the section "Netbooting IDS Worlds", page 12.)

For information on the world description: See the section "World Description for Netbooting". Here is an sample boot file with a recommended sequence of commands for using Netboot:

```
Clear Machine
Declare Paging-files  fep0:>page.page aux.page betty.page satchel.page
Load Microcode FEP0:>3640-mic.mic.418
Netboot chip-poker
Enable IDS
Set Chaos-Address 14159
Start
```

(If your system runs DNA, the boot file should also include a Set Ethernet-Address command.)

A world matching the world description is netbooted.

User machines that have been netbooted do not need any world in their FEP file system except the netboot core (and a backup copy of the core). If you have a Genera 7.2 world-load file on your FEP, you can use the world-load file as a netboot core. That is, any Genera 7.2 world load can be used as a netboot core, but the netboot core files are much smaller.

When you netboot, the console screen goes blank and then provides an attenuated narrative of what is happening. In most cases, you won't need this information.

You can see how far along the netbooting has proceeded by observing the progress bar on the lower right of the screen. The leading dotted line indicates how many pages of the world have been requested to be transferred and the trailing solid line indicates how much of the world has actually been transferred. The progress-bar label describes what phase the netboot process is in. In addition, there are short bars indicating, from left to right, Disk-wait, CPU-run, and Net-wait.

You can halt netbooting at any time by pressing `h-c FUNCTION`, as indicated in the upper left of the screen.

If, for any reason, your netboot core does not work, you may need to use the FEP command Set World-to-netboot.

2.3.2.2. Set World-to-netboot

Set World-to-netboot *world-description*

Replacement for the Netboot command in certain rare circumstances. Usually, the Netboot command selects the "best" world to use as a netboot core. If the "best" core does not, for some reason, seem to work, you can use Load World to load another world to use as a netboot core and then use this

command to netboot using that world rather than selecting the "best" core. In other words, this command uses the currently loaded world to netboot, rather than searching for a "better" one. See the section "World Description for Netbooting".

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

2.3.2.3. Netbooting IDS Worlds

There is no limit on using netbooting with IDS worlds. If the Netboot command in your boot file identifies an IDS world on a netboot server, all the parents of that world are also netbooted, provided they are on the same server.

You can also combine loading an IDS world from your local disk with netbooting. If you are accustomed to using a standard world for your site with your own custom IDS built on top of it, you can continue to do so. In this case, however, use the Load World command and specify the pathname of the IDS world on your local disk in your boot file. Do not use the Netboot command in this case.

Note: Not having all parents of an IDS world on the local disk was once an error, but is now interpreted as a request to netboot the missing parents.

Two conditions must be met to use this technique:

- You have an IDS world on your local disk, but not all its parents.
- All the remaining parent world loads are available on the same enabled netboot server.

In this case, the Load World command in the boot file loads the IDS world normally and searches the local disk for its parents. When some parent is not found, netbooting is initiated to boot that and all remaining parent worlds. When any parent world on the netboot server is deleted, you must, of course, create a new IDS on your local disk. You can netboot the new parent worlds and then create the new IDS.

Note: If there is no enabled netboot server on the local subnet, the Start command will start Lisp and Lisp will wait forever for a netboot server. No error is returned. If your site does not support netbooting, you will have to use `h-c FUNCTION` to halt the netboot process and get back to the FEP and boot from the local disk.

2.4. Autobooting

Systems with G208 FEP EPROMS support autobooting. Autobooting is automatic, unattended booting of the Lisp machine after any FEP reset, such as a power-up event.

The autoboot sequence is specified by an `autoboot.boot` file.

The presence of autoboot.boot enables autobooting.

This single file must include all of the FEP commands which are normally specified by both the hello.boot and boot.boot files, in the order in which these commands would normally be executed.

Specifically, autoboot.boot must contain commands to:

- Scan the required flods.
- Initialize the hardware tables.
- Clear the machine.
- Load the microcode and world.
- Declare the paging files.
- Set the Chaos address.
- Do whatever else both boot files normally do.
- Start the machine.

The autoboot process is enabled by the the autoboot.boot file. If this file does not exist, autobooting is not activiated. autoboot.boot must reside on the lowest numbered disk unit implemented in that specific machine. In most cases this will be unit 0.

During disk drive spin-up, the software will wait for up to three minutes for unit 0 to respond. If unit 0 has not responded by the end of that time period, the system checks for autoboot.boot file existence will be checked for on the lowest numbered disk unit that has responded.

You can abort autobooting by pressing any key during the disk spin-up and autoboot.boot file search activities. If you abort autobooting, you must then boot using the standard hello.boot and boot.boot approach.

In a system with disks already spun-up, the above time delay will approach zero. A ten-second-long fixed abort window is provided for such instances. The operator abort window can be extended by inclusion of the Autoboot Delay command as the first command in the autoboot.boot file. This optional command permits specifying a length of time to wait before actually beginning the autoboot sequence. A keystroke during this period of time will abort autobooting. This command is provided for those instances where the ten-second fixed delay may be inadequate to permit operator intervention. The argument to the Autoboot Delay command is decimal seconds. The default value is 10 seconds.

See the section "Autoboot Delay FEP Command".

2.5. Customizing and Saving Worlds

Symbolics distributes new software on distribution tapes in the form of world loads. These distribution worlds, once they are installed on your machines, may be altered in many ways. The sequence of steps you can take with the initial distribution world load to achieve a desired outcome is varied. This section outlines the procedures for accomplishing certain customizations for your site. All of the commands mentioned in the procedures are documented in this section.

If you want to simply load a new world and save an incremental version of it, here is a basic sequence of steps you might take:

- Boot the new world. For information on booting: See the section "Booting a World", page 2.
- Use the Set Site command to make the world site-specific
- Use the Optimize World command to optimize the world
- Use the Save World Incremental command to save an incremental world on your machine

If you want to add some special programs, or systems, to the initial distribution world load, and then create an incremental world, this is the sequence of steps you might take:

- Boot the new world. For information on booting: See the section "Booting a World", page 2.
- Use the Set Site command to make the world site-specific
- Use the Load System command to load special software into your world
- Use the Optimize World command to optimize the world
- Use the Save World Incremental command to save an incremental world on your machine

If you want to add patches (updates to the software distributed by Symbolics) to the initial distribution world load, and then create an incremental world, this is the sequence of steps you might take:

- Boot the new world. For information on booting: See the section "Booting a World", page 2.
- Use the Set Site command to make the world site-specific
- Use the Load Patches command to add updated software to your world

- Use the Optimize World command to optimize the world
- Use the Save World Incremental command to save an incremental world on your machine

For information about Incremental Disk Save (IDS): See the section "Incremental Disk Save (IDS)", page 27.

For information about the Optimize World command: See the section "The Optimize World Command", page 22.

Note: If you want to supply a world to a machine which does not have FEP eproms of version 127 or greater, you should follow one of the procedures listed above, but use the Save World Complete command, instead of the Save World Incremental command.

If your site needs to build a distribution world, follow these steps:

- Boot the new world. For information on booting: See the section "Booting a World", page 2.
- Use the Set Site command to make the world site-specific
- Use the Load System command to load any additional software
- Use the Set Site command with the site name as Distribution
- Use the function **(si:full-gc)** to garbage-collect the world
- Use the function **(si:reorder-memory)** to optimize paging performance
- Use the Save World Complete command to save the world

2.5.1. Commands Used to Customize and Save Worlds

2.5.1.1. The Load World Command

Load World *file-name*
Restores enough of the world on the local disk that you can start up the machine. It prints both the desired microcode version for this world and the currently loaded microcode version. The default value of *file-name* is the last file name given to the Load World command. Its initial default is >Genera-World.load.

```
FEP Command:  Load World FEP1:>Genera-7-2.load
```

Netbooting is an alternate to loading a world off the local disk. You can combine loading an IDS world off the local disk with netbooting. See the section "Netbooting", page 8.See the section "Netbooting IDS Worlds", page 12.

2.5.1.2. Setting and Defining Sites

The location of your machines is called a site. When you boot new software (distributed as a world load), your site is not set. The herald gives the machine name as *DIS-LOCAL-HOST*. In order to begin operating your computer at a site, you must configure the new software at a site by using the Set Site or the Define Site command.

If you are configuring software at a site that already exists, you must use the Set Site comand. This gives your machine access to the network capabilities of an already existing site.

If you are configuring software at a site that does not yet exist, you must use the Define Site command. This creates a new site for your machine. Later, you can expand the site to include other users, hosts, printers, and networks. For more information: See the section "Customizing and Saving Worlds", page 14.

Set Site Command

Set Site *site-name*

Starts a dialogue to set the current site to be *site-name*. It should be the first thing you type to your machine after booting a new distribution world. The command makes the identity of all objects included in the site's namespace known to your machine.

In order for the Set Site command to work:

1. Your machine must be declared as a host in the site's namespace. See the section "Registering Hosts", page 79.

2. The namespace database for each site is stored on a machine called the namespace server. If the site's namespace server is not the local host, you must know the name and chaosnet address of the namespace server. For more information: See the section "Choosing Chaosnet Addresses".

Set Site Dialogue

Here is an explanation of each line of the Set Site dialogue for a site. Following each line of the Set Site dialogue is the explanatory text.

```
Command: Set Site (site name [default Get from network]) downunder
```

First you must issue the Set Site command. If you accept the default (called "Get from network"), your machine automatically queries other network machines to learn their site name, and uses that site name.

If you wish to change your site, or if your site information is unavailable, you must enter a site name manually. In this example *downunder* is the site name.

When you enter the Set Site command, the system prompts:

```
What host is a namespace server for DOWNUNDER
(default: Local):
```

This question prompts for the name of the namespace server for your site.

1. If the site's namespace server is the local host (the machine on which you are typing), you must press RETURN to accept the default. **Do not type the name of the local host.** The machine then asks for the location of the *descriptor file.* The *descriptor file* holds namespace data files. The default value is the same one offered when defining a site:

   ```
   Where is the descriptor file for DOWNUNDER
   (default local:>sys>site>downunder-namespace.text):RETURN
   ```

   The machine is now on the network at the site *downunder.*

2. If the site's namespace server is not the local host, then you must provide the name of the namespace server for the site. The machine then asks for the server's chaosnet address, and asks you to confirm the server's name. In this example, *sydney* is the name of the server.

   ```
   What host is a namespace server for DOWNUNDER
   (default: Local): sydney
   Chaosnet address for sydney: xxxxx
   Host responds as SYDNEY, ok? (Y or N) Yes.
   ```

   The site's namespace determines all other relevant information, and the machine is now on the network at the site *downunder.*

Your machine now has access to all objects located at the site, and you are ready to login and use the new world load. You may want to save the new, site-specific version of your world. For more information:See the section "Customizing and Saving Worlds", page 14.

Define Site Command

Define Site *site-name*

Starts a dialogue to create a new site called *site-name.*

The default values for the site's namespace server, SYS host, computer storing the namespace data files, and host for bug reports, are all the local host. If you provide a non-local host for any of these, you must know the non-local Chaosnet address(es) and operating system(s). For more information:See the section "Choosing Chaosnet Addresses". Here are the default values for the Define Site command:

System Type*: LISPM
Service: CHAOS-STATUS CHAOS-SIMPLE CHAOS-STATUS
Service: SHOW-USERS CHAOS NAME
Service: TIME CHAOS-SIMPLE TIME-SIMPLE
Service: UPTIME CHAOS-SIMPLE UPTIME-SIMPLE
Service: LOGIN CHAOS TELNET
Service: LOGIN CHAOS SUPDUP
Service: LOGIN CHAOS 3600-LOGIN
Service: SEND CHAOS CONVERSE
Service: SEND CHAOS SEND
Service: NAMESPACE CHAOS NAMESPACE
Service: NAMESPACE-TIMESTAMP CHAOS-SIMPLE NAMESPACE-TIMESTAMP
Service: LISPM-FINGER CHAOS-SIMPLE LISPM-FINGER
Service: FILE CHAOS NFILE
Service: FILE CHAOS QFILE
Service: CONFIGURATION CHAOS CONFIGURATION
Address: CHAOS *nnnnn*

In the Address attribute line, *nnnnn* represents a valid 5-digit octal Chaos address.

Define Site Dialogue

Here is an explanation of each line of the Define Site dialogue for a site. Following each line of the Define Site dialogue is the explanatory text.

```
Command: Define Site (site name) downunder
```

First you issue the command and give the site name. This might be the name of your company, or, if you are more whimsical-minded, it might be related to the machine names you have chosen. In this example, *downunder* is the name of an example site.

```
Define a new site named DOWNUNDER (as opposed to looking for an
existing definition of DOWNUNDER on disk)? (Yes or No) Y
```

verifies that you want to create a new site rather than setting your site to an existing one.

```
What host is to be a namespace server for DOWNUNDER
(default LOCAL): RETURN
```

This question prompts for the name of the machine where the namespace database is stored. If the namespace server is not the local host, then you can provide the name of the namespace server for the site. **If the namespace server is the local host, do not type the host's name, instead use the default by pressing** RETURN.

```
What host is to be the SYS host for DOWNUNDER
(default LOCAL): RETURN
```

This question prompts for the name of the machine where Symbolics-supplied source and documentation files are to be stored by default. If the SYS host is not

the local host, then you can provide the name of the SYS host for the site. For more information: See the section "Logical Pathnames and the SYS Host".

```
What Symbolics computer will store the namespace data
files for DOWNUNDER (default LOCAL): RETURN
```

This question prompts for the name of the machine that stores the namespace data files. Although this machine must be a Symbolics computer, it does not have to be the same machine that is the namespace server, but we strongly suggest that it be the same machine.

```
What host is to be used for bug reports for DOWNUNDER
(default LOCAL): RETURN
```

This question prompts for the name of the machine that receives bug reports. This will most likely be the machine that you use for your mailer. For information about the mailer: See the section "Installing and Configuring the Mailer", page 203.

```
What is the real name of the local host: koala
```

This question prompts for your machine's name at the site. Machine names can be different at different sites. In this example, *koala* is the name of the local host. For more information: See the section "Why do you name machines and printers?" in *Genera User's Guide*. The default values for the namespace server, the SYS host, the computer storing the namespace data files, and the host for bug reports, are all the local host. If you provide a non-local host for any of these, the computer prompts for the non-local chaosnet address(es) and operating system(s)

```
What directory on KOALA will hold the namespace data files
(default >sys>site>): KOALA:>sys>site>
```

This question prompts for the location of your namespace data files. These files are accessed each time the site is set. We strongly suggest that you use the default value to avoid possible confusion.

```
What directory on KOALA corresponds to SYS: SITE;
(default >sys>site>): KOALA:>sys>site>
```

This question prompts for the directory that holds your systems translations files. These are necessary for logical pathname translations. We strongly suggest that you use the default value to avoid possible confusion.

```
What account should be used for the system to login to KOALA
(default: LISPM): wombat
```

This question prompts for a name the system uses automatically when it needs to access files. In this example we are using the name wombat. Do not supply your own name or that of another real user. When the system needs to log in automatically (for example, to read the namespace data files for the first time) the namespace server logs in with this user name.

```
What is the local timezone [default EST]: EST
```

This question prompts for the timezone. The default is Eastern Standard Time (EST) or Eastern Daylight Time (EDT) depending on whether Daylight Savings Time is in effect. For more timezone information: See the section "Specifying a Time Zone for Your Site".

```
Is DOWNUNDER a standalone site (there are no servers to respond
to a Who-am-I request)? (Y or N) Yes.
```

This question prompts for whether your site is a standalone site consisting of one Symbolics machine, or if it has many Symbolics machines.

Several notification messages follow and confirm the definition of your site.

You are now ready to login and use the newly configured software. The site of your local host is automatically set at the newly defined site. You may want to save the new, site-specific version of your world. For more information: See the section "Customizing and Saving Worlds", page 14.

Load System Command

Load System *system keywords*

Loads a system into the current world.

system — Name of the system to load. The default is the last system loaded.

keywords — :Branch, :Component Version, :Condition, :Include Components :Load Patches, :Query,:Redefinitions Ok :Silent, :Simulate, :Version

:Branch — Reserved for future use.

:Component Version — {Released, Latest, Newest, *version-designator*} The version of any component systems to load. Released means the version designated as released in the journal file. Latest means the most recent version recorded in the journal file. Newest means to ignore the versions in the journal file and just find the newest files. The default is the version with which the system was compiled.

:Condition — {Always, Never, Newly-Compiled} Under what conditions to load each file in the system. Always means load each file. Newly-compiled means load a file only if it has been compiled since the last load. The default is newly-compiled.

:Include Components — {Yes, No} Whether to load component systems. The default is Yes.

:Load Patches — {Yes, No} Whether to load patches after loading the system. The default is Yes.

:Query — {Everything, Confirm-only, No} Whether to query before loading. Everything means query before loading each file. Confirm-only means create a list of all the files to be loaded and then ask for confirmation before proceeding. No means just go ahead and load the system without asking any

questions. The default is No. The mentioned default is Everything.

:Redefinitions Ok {Yes, No} Controls what happens if the system asks for confirmation of any redefinition warnings during the loading process. Yes means assume that all requests for confirmation are answered yes and proceed. No means pause at each redefinition and await confirmation. The default is No. The mentioned default is Yes. This allows you to start loading a system that you know will take a long time to load and leave it to finish by itself without interruption for questions such as "Warning: *function-name* being redefined, ok? (Y or N)".

:Silent {Yes, No} Whether to turn off output to the console while the system is loading. The default is No. The mentioned default is Yes.

:Simulate {Yes, No} Print a simulation of what compiling and loading would do. The default is No. The mentioned default is Yes.

:Version {Released, Latest, Newest, *version-designator*} Which version number to load. Released means the version designated as released in the journal file. Latest means the most recent version recorded in the journal file. Newest means to ignore the versions in the journal file and just find the newest files. The default is Released.

Note: This command only loads a system. If you want to compile and load a system: See the section "Compile System Command" in *Genera Handbook*.

Load Patches Command

Load Patches *system keywords*

Loads patches into the current world for all systems, locally maintained systems, or the indicated systems.

system {All Local *system-name1, system-name2* ... } The system(s) for which to load patches. The default is All.

keywords :Dangerous Patch Action, :Include Components, :Output Destination, :Query, :Save, :Show

:Dangerous Patch Action
{Skip, Query, Load} Whether to skip loading *dangerous* patches, that is patches that might make data structures in your world inconsistent, causing unexpected behavior. The default is Skip.

:Include Components
{Yes, No} Whether to load patches for any component systems. The default is No. The mentioned default is Yes.

:Output Destination
: {Buffer, File, Printer, Stream, Window} Where to redirect the typeout done by this command. The default is the stream ***standard-output***.

:Query
: {Yes, No, Ask} Yes asks for confirmation before beginning the load patches process and again before loading each patch. Ask asks whether or not it should query before each patch, and then for confirmation before beginning the load patches process. The default is No. The mentioned default is Yes.

:Save
: {*pathname*, Prompt, No-Save} The file in which to save the world with all patches loaded. Omitting this keyword means do not save the world. The mentioned default is Prompt, which means save the world and then prompt for a pathname.

:Show
: {Yes, No, Ask} Whether to print the patch comments as each patch is loaded. The default is Yes.

See the function **load-patches** in *Program Development Utilities.*

2.5.1.3. Optimizing a World

If you load special software or programs into your world you may want to use the Optimize World command to improve paging performance.

The Optimize World Command

The Optimize World command prepares a world to be saved after you have loaded your own special programs or layered software into the original distribution world. Optimize World improves paging performance by reorganizing compiled functions and certain data in virtual memory, so related things are on the same page. Paging performance is the speed at which virtual memory is swapped in order to access data. Once you have used Optimize World, your world load will page less than if it had been saved without running this program beforehand. Optimize World does not move objects that were already optimized in the distribution world load; it only moves things that you have loaded into that world.

Why Use the Optimize World Command

You should use Optimize World if you load your own programs or layered software into the distribution world and want to save this world for later use. Once you have processed a world with the Optimize World command you can boot and reuse this world over-and-over again, with the continued benefit of improved paging performance. The Optimize World command is one of the SmartStore storage management facilities.

When to Use the Optimize World Command

It is recommended that you use the Optimize World command shortly before you plan to disk-save a world load.

When Is It Unnecessary to Use the Optimize World Command

It is not necessary to use Optimize World if you have simply loaded the distribution world and customized it for your site without loading any additional programs.

Optimizing IDS Worlds

You should use the Optimize World command on IDS worlds only if the parent world has itself been optimized or if the parent world does not need optimizing. If you violate this rule, your IDS worlds will be considerably larger than they need to be. If you follow the rule, Optimize World does not significantly increase the size of the IDS world. The actual size difference is usually less than 10 per cent.

Distribution worlds are optimized. Site-configured worlds which have only had Set Site done on them (and no systems loaded) do not need to be optimized.

What Happens When the Command Is Used

After you enter the Optimize World command it asks you if you are ready to begin an optimization. If you respond yes, the command begins optimization. When the program finishes, another message appears informing you that the process is complete. However, if you would like to see status reports of the reordering process, you can specify the keyword **:show** when you enter the command.

Optimize World typically takes about one-half hour to execute, but this time period varies according to the size of the world load and the total amount of main memory. **NOTE: During the time that Optimize World is running your machine does not respond to the network nor to the terminal. You are not able to use your machine.**

Optimize World Command

Optimize World *keywords*

Optimizes the world that is currently loaded into your enviroment by reorganizing the world to improve paging performance. Use this command if you load special programs or systems in addition to the distribution.

keywords :Show

:Show Displays the progress of the optimization process on the screen.

si:reorder-memory &key (*incremental* **t**) (*run-without-interrupts* **t**) *Function*

This function can be run after **si:full-gc**. It moves objects around in memory in a manner that optimizes paging performance. It takes between

40 and 60 minutes. While it is running, the machine cannot be used for any other purpose. Usually this function is invoked with the intention of saving the world immediately thereafter.

Note: The Command Processor command Optimize World is the preferred high-level interface to the functions **si:full-gc**, **si:reorder-memory**, and **si:optimize-compiled-functions**. See the section "Optimize World Command", page 23.

Using the Optimize World Command

Here is the recommended procedure for making a customized world with your own programs loaded into the world:

1. Boot a suitable base world, such as the site-customized version of the distribution world load.

2. Then type the following sequence of commands at the Command Processor prompt. (Note that the default Command Processor prompt in an uncustomized system is `Command:`. The prompt for your site may be different, however, due to customization.)

```
Command: Start GC :Ephemeral
Command: Disable Services
Command: Login a suitable user :Init File none
Command: Load System name of system
Command: Optimize World
```

   To see the progress of the optimization display on the screen type:

```
Command: Optimize World :show
```

   And then use the Save World command:

```
Command: Save World Incremental filename
```

   - For most applications it is better not to do **(si:full-gc)**. This is because using Incremental Disk Save after **(si:full-gc)** renders your world the same size as the world from which you started. **(si:full-gc)** generally does not gain anything if done in a freshly booted world that has only had a program loaded with the EGC turned on, assuming that the freshly booted world is a distribution world (which is a garbage-free world). If you decide to use **(si:full-gc)**, do it immediately before Optimize World, and use Complete rather than Incremental in the Save World.

   - At the point in the sequence above when you use the Load System command, you can type Load System and/or Load File as many times as you need to in order to load all of your special systems and additional programs.

   - If you choose to save your optimized world with the Save World Incremental command, then you must keep the parent world from which

you made the optimized world. The parent world must be available to the machine in order for the incremental world to boot. For more information about incremental worlds, see the section *Incremental Disk Save* in *Installation and Site Operations*.

- The Optimize World command improves paging performance by reorganizing the pages in the world. This increases the actual size of the world by a small amount.

2.5.1.4. Saving a World

To save a world, use the Save World command or **zl:disk-save**.

Save World Command

Save World (Complete or Incremental) *pathname*

Saves the current world. The system prompts (Complete or Incremental) if Incremental Disk Save is enabled. You specify Complete to do a save of the entire world, or Incremental to do an Incremental Disk Save. The default is Complete. If Incremental Disk Save is not enabled, the prompt is (Complete).

pathname — The pathname to use for the saved world. The default is the FEP file specification for the local machine, based on the version number of the current system and on whether the save is to be complete or incremental.

zl:disk-save &optional *destination-file* (*incremental* **':ask**) *Function*

Saves out the current Lisp environment into a FEP world-load file.

destination-file is the pathname of the FEP world-load file in which the Lisp environment is to be saved. It defaults to a file on FEP: (the boot unit) with a file type of .load and a file name of Release-*major-minor* (for released systems) or System-*major-minor* (for unreleased systems). You are queried for the pathname if you do not supply this argument.

incremental specifies whether an incremental image or full image of the world should be saved. This argument has meaning only if IDS is enabled. A full image is always saved if IDS is not enabled. If IDS is enabled, specifying **t** causes an incremental image of the world to be saved. Specifying **nil** causes a full image to be saved. If IDS is enabled, and if **:ask** is specified or if the *incremental* argument is not supplied, **zl:disk-save** asks whether an incremental or full image should be saved.

The following examples show the actions of **zl:disk-save** if IDS is enabled.

The following queries for the pathname and saves a full world.

```
(disk-save nil nil)
```

The following queries for the IDS pathname and saves an incremental world.

```
(disk-save nil t)
```

The following queries whether it should save a full or incremental world and for the pathname. It then saves the appropriate world.

```
(disk-save) or
(disk-save nil) or
(disk-save nil :ask)
```

The following saves a full world under the name "new-world.load".

```
(disk-save "fep0:>new-world.load" nil)
```

The following saves an incremental world under the name "inc-new-world-from-new-world.load".

```
(disk-save "fep1:>inc-new-world-from-new-world.load" t)
```

The following queries whether it should save a full or incremental world and saves the world under the name "test-world.load"

```
(disk-save "test-world.load") or
(disk-save "test-world.load" :ask)
```

zl:disk-save asks for firmation only if the file to be written is one of the same file from which the currently running world was booted.

When the FEP file system does not have enough room to save the world-load file, **zl:disk-save** tells you how many more blocks you need and offers a choice of actions that you can take on the root directory of the FEP file system of the unit on which you tried to save:

```
Estimated size is M blocks.
Not enough room in FEP filesystem, need at least N more blocks.
Run Dired, Expunge directory, List directory, or Selectively
delete load files (D, E, L, or S)?
```

D — Runs Dired on the directory. If you try to run Dired, **zl:disk-save** warns:

```
Running Dired will substantially increase the size of
the saved world. You should, therefore, boot again
before saving if you run it. Run Dired anyway? (Y or
N)
```

E — Expunges the directory.

L — Lists all files in the directory.

S — Selectively deletes load files. It goes into a loop asking about each ".load" file in directory (except the files you are currently running). It asks whether to delete each one, expunging after each successful deletion until you have enough room. This is probably the only option you need to use.

If a world-load file to be written already exists, **zl:disk-save** offers a choice of superseding the file, overwriting the file, or providing a new pathname.

zl:disk-save displays the herald for the world to be saved. It asks whether you want to add an additional comment to the herald or change the comment if already present. The value of the variable **si:system-additional-info** is set to the string you type, and the Show Herald command displays this string in parentheses at the end of the first line of the herald.

zl:disk-save constructs a title for the world load based on the comment and versions of the systems loaded in the environment; it defaults to the value returned by **(si:system-version-info t)**, but you can supply an alternate title. The title is a property of the world-load file; the Show Disk Label command displays this title.

zl:disk-save then displays a few messages telling you what it is doing: It runs the initializations in the **si:before-cold-initialization-list**, logs out, and runs the initializations in the **si:system-shutdown-initialization-list** (see below). Then the machine seems to act as it does when it cold boots, although actually it is copying from the paging file to a FEP file rather than the other way around. When it finishes it displays the herald message as if cold booting had just completed.

(Note that the entries on the **zl-user:system-shutdown** list should all be for subsystems that are required for almost anything else to run. Currently there are entries for the network, the TIME system, and the Lisp Machine file system. User programs should add themselves to the **before-cold** list rather than to **zl-user:system-shutdown**.)

zl:disk-save offers to update the FEP boot file to load the world that has just been saved.

2.5.1.5. Incremental Disk Save (IDS)

Incremental Disk Save (IDS) is a facility that allows you to save modified worlds using much less disk space than if you saved the entire world. IDS saves a world by saving only the pages that have been changed; it does not save a copy of the entire world. This copy is called the *incremental world* and contains only the pages that have been modified. The original world, before it is modified, is the *ancestor* of the modified world and is referred to as the *parent world.*

IDS has two major advantages over a full world save. One is that you can save an incremental world using much less disk space than if you saved the entire world. The second is that you can easily keep multiple incremental versions of a world on the same disk because the incremental disk saves use so much less disk space. In other words, you can take a parent world, make multiple different incremental worlds, and save all the incremental worlds using a minimum of disk space.

Performance is another consideration when using IDS. IDS uses extra wired memory which slightly impairs page faulting. Thus, an IDS-world will be slower than a non-IDS world.

Netbooting and IDS interact. See the section "Netbooting IDS Worlds", page 12.

The following diagram shows how you can have multiple versions of an incremental world. You might start with a distribution world (world 1) and modify it to reflect your site-specific information. You would then have one parent world and one incremental world (incremental world 2). You could then load another layered product, such as Pascal, into the incremental world. After using IDS to save the incremental changes, you would have one parent world and two incremental worlds (incremental worlds 2 and 3). You could continue creating more incremental worlds from the site-specific incremental world. For example, you might boot your machine with the site-specific world and then load your own application into it. After incrementally disk-saving this, you would now have one parent world and three incremental worlds (incremental worlds 2, 3, and 4). You might continue to modify one of these incremental worlds, such as by loading patches into the incremental world with Pascal. This would give you incremental world 5 and so on.

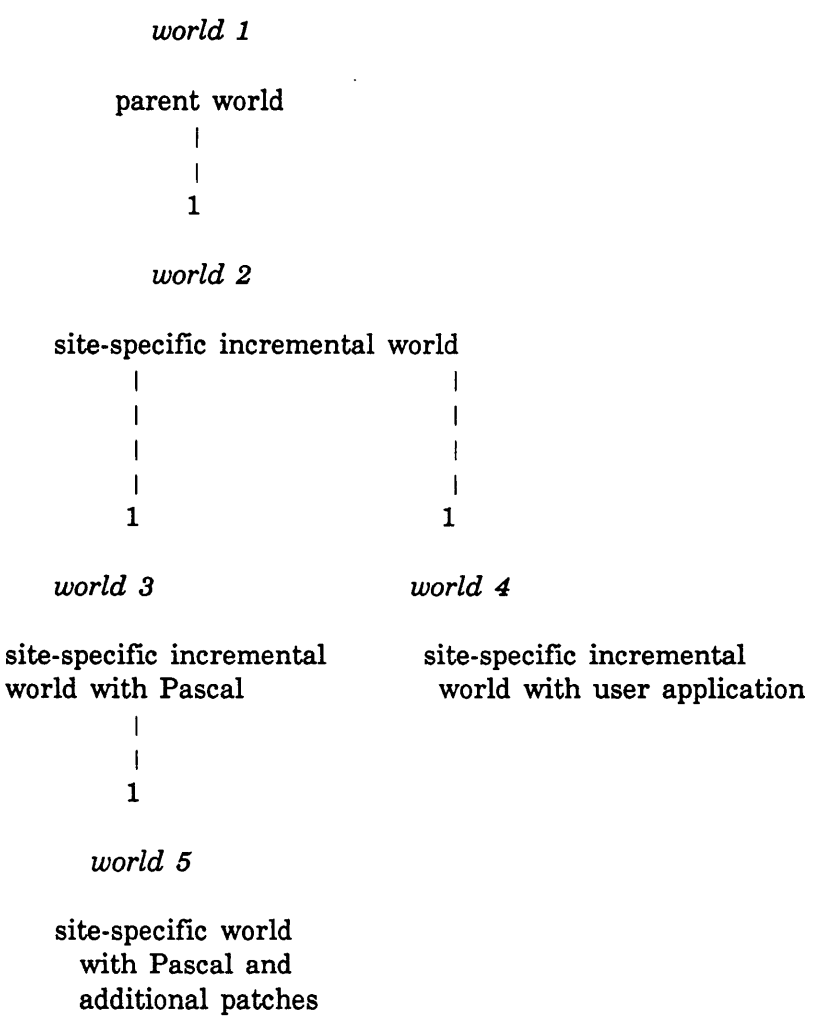

Figure 2. Incremental Worlds

Here is another example that shows how disk space is conserved. Assume you have a distribution world that is 34,000 blocks in size; this is the parent world. You might want to set the site and load site modifications and save the resulting world. Doing a nonincremental disk save results in a world containing the parent and the modifications, which might be 40,000 or more blocks. Doing an incremental disk save would result in a world of perhaps 3,000 to 5,000 blocks that consists of only the new blocks and modified blocks of the parent. You might then boot this incremental world, load a system into it, and incrementally disk save the result, which might again be 3,000 to 5,000 blocks. You now have ready access to three worlds (a distribution world, a world customized for your site, and a world with an additional system) that total perhaps 46,000 blocks instead of 90,000 blocks. You always have the option to do a full disk save on a multigeneration incremental world; for example, you could do a full disk save on the incremental world with the system loaded into it.

Note: It is important to keep all ancestors for all incremental worlds you intend to use. If you delete some ancestor, a descendent will not be usable because it requires blocks that are in the ancestor.

The format of world files (full and incremental) is such that they can freely be moved around within a site and (less practically) between sites. Thus, a site can place the distribution world supplied by Symbolics on all the machines at the site. You can boot this distribution world on one machine, modify it, and incrementally disk-save it. You can then distribute the resulting incremental world to all the other machines. The parent world, as well as any other intermediate incremental worlds, must be available on that machine for this incremental world to boot.

For information about the command that shows the IDS parent worlds for a specified world-load file: See the section "Show IDS Parents Command" in *Genera Handbook*.

For information about the command that shows the IDS children for a specified boot file: See the section "Show IDS Children Command" in *Genera Handbook*.

For information about the command that shows the IDS files for a specified host: See the section "Show IDS Files Command" in *Genera Handbook*.

Using Incremental Disk Save (IDS)

You can enable IDS with the Enable IDS command in your boot file.

The following steps describe the procedure for enabling IDS:

1. Edit the boot file to include the Enable IDS command, as shown in the following example:

```
Clear Machine
Load Microcode microcode-file-name
Load World world-file-name
Set Chaos-Address this-machine's-chaos-address
Enable IDS
Start
```

The Enable IDS command must precede the Start command.

2. Boot the machine using the recently modified boot file.

3. Save the world using either the Save World command or **zl:disk-save**. IDS is enabled by default. To disable IDS, replace the Enable IDS command with the Disable IDS command in your boot file.

   This IDS-enabled world can be copied to other Symbolics computers at your site.

A world that is saved (either completely or incrementally) with IDS enabled retains this characteristic; IDS is enabled in any of the subsequently saved worlds. Therefore, you do not need to use the Enable IDS command again. However, doing so does not harm anything and ensures that IDS is enabled. (Resetting the FEP does disable IDS.)

To do an incremental disk save, proceed normally to modify your world; for example, doing site modifications, loading local site patches, loading any systems that will be used, and so on.

After making the modifications, save the world by using either the Save World command.

Note: Using **gc-immediately** or **si:full-gc** on a world prior to using IDS negates the benefits of using the command. This is because IDS saves only those pages of the world that have been changed; both the Start GC command and **si:full-gc** change the organization of the world. Thus, the size of a world which is garbage-collected and then saved with IDS is not substantially smaller. However, it is recommended that you use the ephemeral-object garbage collector (EGC) on your world. The EGC is enabled by default in Genera 7.0.

If IDS is enabled, you are asked:

```
(Complete or Incremental)
```

Answer Incremental if you want to save an incremental world; answer Complete if you want to save a full image, for example, if you want to condense collected patches. You are then prompted for the pathname for the saved file. A full save gives the same default as before, namely Genera-*major-minor* or System-*nnn-mmm*. An incremental save changes this default in the following ways:

- "Genera-" is prefaced with "Inc-" or "System-" is replaced with "Inc-".

- "-from-" is appended.

- A trimmed version of the loaded world name (the pathname that appears in the first line of the herald) is appended.

- If the result is longer than 32 characters (the limit of filenames in the FEP file system), the filename is truncated and the last 4 characters within the limit are replaced with -etc.

Examples: the pathname for "loaded world" is a possible modification of the *previous* default):

Loaded world	*Action*	*Default*
Genera-7-0	Set Site	Inc-Genera-7-0-from-Genera-7-0
Inc-SITE-7-0-from-Rel-7.-0	Load Patches	Inc-271-99-from-IncSITE-7-0--etc
Inc-SITE2-from-Inc-SITE-7-0-etc	Load Fortran	Inc-271-99-from-IncSITE2-fro-etc
Inc-FORTRAN-from-Inc-SITE2-etc		

Suggestions for Using IDS

For maximum flexibility, you should try using many levels of Incremental Disk Save. For example:

Steps	*World*
0	Distribution World
1 Boot the distribution world Use Set Site Incrementally save	world-1
2 Reboot with world-1 Load site patches Incrementally save	world-2
3 Reboot with world-2 Load a system Incrementally save	world-3

This encompasses a total of four worlds (a distribution world and three incrementally saved worlds) instead of two worlds (a distribution world and one collective save). The most stable modifications should be made first. That way, if the system that was loaded in the third step changes, it is not necessary to redo the Set Site or load the site patches; it is only necessary to reboot the world with site patches (world-2), reload the system, and incrementally save. If the site maintainer wishes to load new local site patches, the second and third steps need to be repeated, but not the first. The users of the result of the third save have the option to use the old site patches (and not get the new site patches) or rebuild their system with the new site patches.

Note that enabling IDS is pervasive across disk-saves. That is, if you turn it on before saving a world load, the resulting world has the facility enabled in it when that world is again booted.

3. Creating More Room on the Local Disk

There are two file systems available on the Symbolics computer: the Lisp Machine File System (LMFS) and the FEP File System (FEP FS). LMFS is a general-purpose, highly flexible file system, suitable for everyday use. Currently, only the Symbolics processor understands how to operate on LMFS files. The FEP FS is a simple, basic file system that both the Symbolics computer and front-end processors understand how to access.

The FEP FS contains two kinds of files.

- *FEP files* are used to store the information that the FEP uses to do things like boot Lisp and manage virtual memory. Examples of FEP files include world load files, netboot cores, microcode load files, paging files and boot files.

- *File system partitions* are very large, specialized FEP files. LMFS stores its structure and data in file system partitions. User files are stored by LMFS in partitions.

Sometimes the Save World or Copy World commands might inform you that you have run out of FEP file system space, and offers you the option of editing your FEP directory. For systems with 167-Mbyte or more of storage, you should delete and expunge old, unneeded world loads, and then resume from the Save World "out of room" error or retry the Copy World operation. You should not delete any world loads from a 140-Mbyte system. See the section "Instructions for Managing Disk Space on the 3640 With a 140 Megabyte Disk", page 36.

It is wise to keep a large (about 40K), noncritical world load or extra paging file on the Symbolics computer's disk, so it will be available for the FEP Disk Restore command to use in case all world loads become nonfunctional.

Sometimes, writing a file out to a LMFS produces an "out of room" error. This means that the present allocation of that particular LMFS is not large enough to accommodate your request for space. It might help to expunge directories with deleted files in them to remove unneeded versions of files, using the Zmacs command Dired (m-X).

If you still do not have enough space after you have deleted and expunged unnecessary files, consider creating an auxiliary file partition. You should only do so, however, on systems that have at least 280 Mbytes of storage. This is because 140-Mbyte systems have no room at all for an auxiliary file partition, and allocating an auxiliary file partition on a 167-Mbyte system can limit your space for large world loads. Even for 280-Mbyte systems, you are trading off world load space for file space when you create auxiliary partitions.

If you are loading worlds from the local disk, be sure to reserve enough FEP file system space for two world loads. World loads are approximately 50,000 blocks

each. The two world loads are your current world and a new world that you are copying. See the section "Booting, Netbooting, and Autobooting", page 2. If you are netbooting, you should reserve an extra paging file approximately the size of a world load, but you won't need the space required for the second world load. See the section "Netbooting", page 8. For details on how to create auxiliary file partitions: See the section "Multiple Partitions", page 181.

Warning: Once you have created an auxiliary file partition, you should never delete it, because you would lose all the data contained in that partition and make the entire Lisp Machine File System unusable.

If you run out of room while writing a LMFS file and then create a new partition to increase the LMFS space, you cannot resume the file operation that failed. Instead, you must abort that operation by pressing c-ABORT, and then retry the operation.

3.1. Allocating Extra Paging Space

Programs that use large amounts of virtual memory might require you to allocate additional paging space, to perform better or to perform at all. Only systems with at least 280 Mbytes of disk storage have enough room to permit additional paging files without adversely affecting the maintenance of worlds on the machine. In order to add an extra paging file to your virtual memory set, you must first create a FEP file using the Create FEP File command. Then, you can activate the paging file from Lisp by using the Add Paging File command. To create a 20-K block paging file on unit 0 type:

```
Create FEP File fep0:>page1.page 20000
```

After creating the extra paging file, any boot files should be modified to use this new paging file. Use the Declare Paging-files command in the boot file to load any paging files you want to use. A typical boot file before inserting the command to load the paging file might look something like this:

```
Clear Machine
Load Microcode >3640-mic.mic.389
Load World >Dist-7-0.load
Set Chaos-Address 401
Start
```

After creating the new paging file, edit your boot file to include the Declare Paging-files command. The new Boot.boot file might look something like this:

```
Clear Machine
Load Microcode >3640-mic.mic.389
Declare Paging-files fep0:page page1
Load World >Dist-7-0.load
Clear Paging
Set Chaos-Address 401
start
```

For information about the Declare Paging-files command: See the section "FEP System Commands: General Usage".

It is safe to delete extra paging files, but only if they are not in active use. You cannot change a paging file that is being used, without booting. To change the paging area you have set up, first boot without adding the paging file to be deleted. Be sure to cold boot by hand, and when you type the Declare Paging-Files command, *do not* specify the extra paging file that you intend to delete. Once you have booted, you can delete the unwanted paging file by editing the FEP directory. Be sure to remove any references to the file from your boot file as well.

3.2. Adding a Paging File From Lisp

If you want to add a paging file from Lisp, use the new command:

```
Command:  Add Paging File
```

Prior to adding the paging file you may have to create the FEP file by using the command:

```
Create FEP File
```

to create the paging file.

Add Paging File Command

Add Paging File *pathname :prepend*

Adds a paging file that has already been created. See the section "Create FEP File Command" in *Genera Handbook*.

pathname — The pathname of the existing FEP file, which becomes the new paging file. The default pathname is the disk unit from which you most recently booted. For example, if you most recently booted from FEP1:>, the default paging file might look like:

```
FEP1:>.page
```

Each paging file must have a unique name.

keywords — :prepend

:prepend — *{yes no}* Yes means to put the paging file at the beginning of the list of swap space to use when new space is needed. This makes the new paging file used almost immediately. No, which is the default, puts the paging file at the end of the list of paging files. Consequently, this new paging file will not be used until the previous swap space is completely used.

4. Instructions for Managing Disk Space on the 3640 With a 140 Megabyte Disk

Since the 140 megabyte disk drive of the 3640 contains a smaller paging file than the 3600 or 3670, you must manage your 3640 FEP file system differently. For a complete description of paging files: See the section "FEP File Types", page 117.

This section describes the different procedures that you follow to manipulate paging space when:

- Loading the world.
- Customizing and saving that world load.
- Saving future world loads.

The disk of your 3640 contains a world load file, a large paging file (called Page.page), and an auxiliary file (called Aux.page) that is the same size as the world load file. You use the auxiliary file in one way for normal operation and in another way when putting a new world on the disk.

WARNING: If your system does not contain an auxiliary file (use the Show FEP Directory command to look for the file named *Aux.page*), call your field representative.

In normal operation, you boot a world load file and use both Page.page and the auxiliary file for paging. In this case, you call the auxiliary file Aux.page.

When you want to create a new world or transfer a new world to the disk, you boot your world load file and use only Page.page for paging. Instead of using the auxiliary file for paging, you rename it and use it to receive a new world you are creating. Once you have successfully created the new world, you rename the old world load file to Aux.page and use it as your auxiliary paging file. For specific instructions for this procedure, see "Installing Genera 7.0 on a 3640 with One 140-Mbyte Disk" in the *Software Installation Guide.*

The auxiliary file is always actively in use, either as:

- A paging file (in normal operation)
- The target file for new world load

4.1. Customizing and Saving the World on a 3640

The shipped configuration assumes the auxiliary file (Aux.page) is to receive your site's customized world load and so contains just one actual paging file (Page.page). **Note:** The size of the distribution world, the Page.page file, and Aux.page file in this example are only examples, since the sizes of these files vary from release to release.

distribution-world.load.1	30,000
Aux.page.1	30,000
Page.page.1	45,000

A customized, normal configuration uses the auxiliary file as a paging file and so contains two paging files:

New-world.load.1	30,000
Aux.page.1	30,000
Page.page.1	45,000

To create your customized world, follow these instructions:

1. Boot the distribution world using the correct microcode. Use only Page.page.1 for paging and reserve the auxiliary file. **Do not boot with the Aux.page file.** You should initially boot by hand rather than use the boot file so that you can set your chaos address and give the correct Add Paging instruction:

   ```
   Clear Machine
   Load Microcode microcode-file-name
   Declare Paging-Files  fep:page
   Load World world-load-file-name
   Set Chaos-Address this-machine's-chaos-address
   Start
   ```

2. Log in by using **si:login-to-sys-host**, for example:

   ```
   (si:login-to-sys-host)
   ```

3. Rename the auxiliary file to whatever name you wish, for example:

   ```
   Rename File FEP:>Aux.page.1 FEP:>New-world.load.1
   ```

4. Customize the booted world and then save it into your new world load file:

   ```
   Save World FEP:>New-world.load.1
   ```

   Since you are asking to save the world into an existing file, you are prompted for an action with which to proceed. The correct answer is Overwrite. Then

you are asked if you want to update the boot file. Answer yes. The Set Chaos line that you manually typed is added to the boot file at this time.

5. Rename the distribution world to be the auxiliary file:

```
Rename File FEP:>distribution-world.LOAD.1 FEP:>Aux.page.1
```

6. At this point, you should edit the boot file, FEP:>Boot.boot, to add the auxiliary file as an additional paging file. Insert this line:

```
Declare Paging-Files  fep0:page aux
```

after the Load Microcode command. For an explanation of the Declare Paging-Files FEP command: See the section "FEP System Commands: General Usage".

Your edited boot file should look like this:

```
Clear Machine
Load Microcode microcode-file-name
Declare Paging-Files  fep0:page aux
Load World world-load-file-name
Set Chaos-Address this-machine's-chaos-address
Start
```

7. Save the edited version.

8. Log out and halt the machine.

9. Boot the new world using the boot file.

4.2. Saving Subsequent Worlds on a 3640

Whenever you wish to create a new world on your 3640 disk, you must follow a procedure similar to that shown above.

1. Boot manually, and do *not* type the Declare Paging-Files fep0:aux command, since you will be saving the latest world into it:

```
Clear Machine
Load Microcode microcode-file-name
Declare Paging-Files  fep0:page
Load World world-load-file-name
Set Chaos-Address this-machine's-chaos-address
Start
```

2. Login by using **si:login-to-sys-host**, for example:

```
(si:login-to-sys-host)
```

3. Rename the auxiliary file to whatever name you wish, for example:

```
Rename File FEP:>Aux.page.1 FEP:>Newer-world.load.1
```

4. Either customize the booted world and save it into your new world load file, or transfer the world from some other machine:

```
Save World FEP:>Newer-world.load.1
```

or:

```
Copy World "other-machine-name"|FEP0:>remote-world-name.load
FEP0:>newer-world.load
```

Since you are asking to save the world into an existing file, you are prompted for an action with which to proceed. The correct answer is Overwrite. Then you are asked if you want to update the boot file. Answer yes.

5. Rename the old world to the auxiliary file to be used for paging:

```
Rename File FEP:>Old-world.LOAD.1 FEP:>Aux.page.1
```

6. Log out and halt the machine.

7. Boot the machine using the boot file.

5. Enabling the Who-Calls Database at Your Site

The **who-calls** database is a cache that maps *names*, which are symbols, to code and variables that use those symbols in some way. A name can be used as a constant, a variable, a function, a macro, an instance variable, a condition, and a few others.

The **who-calls** database is activated when you use the Set Site command during site customization if the database has not already been enabled or disabled. By default, the Set Site command calls the function **(si:enable-who-calls :new)**, which records only the callers in any layered products, special software, or programs loaded into the world. The database is turned on in this way to make it easy to include new software or layered products in the database. **Note: (si:enable-who-calls :new)** does not cause the callers in code in the Distribution world to be recorded.

If you prefer to have all of the Symbolics-supplied software in the database, you can use the function **si:enable-who-calls** with the argument **:all**. Using the argument **:all** has the additional advantage that you can create the database once and then save the database when you save the world. However, creating a full database takes a long time and about 2000 pages of storage. **si:full-gc** remakes the entire database each time it is called.

If you want only explicitly-named files to be in the database, use the function **si:enable-who-calls** with the argument **:explicit**.

If you use the function **si:enable-who-calls** without any arguments, it prompts you for an argument and offers help.

si:enable-who-calls &optional *mode* *Function*

This command takes an argument which is the mode in which the database should be enabled. If you need the full database, use **si:enable-who-calls** during site configuration time. If you do not need a full database you can create an incremental database by choosing a suitable mode.

mode can be one of the following:

:all	Creates a full callers database. This takes many minutes and about 2000 pages of storage. **:all** also queries about the old state.
:all-remake	Creates a full callers database but does not query about the old state. This takes many minutes and about 2000 pages of storage. Use this if you do not want to perform a **si:full-gc**. **si:full-gc** would discard the existing database and remake it, causing extra work.
:new	Creates a callers database that includes only new functions.

:all-no-make	Creates a callers database that includes only new functions until a **si:full-gc** is performed. **si:full-gc** creates the entire database. This takes many minutes and about 2000 pages of storage.
:explicit	Enables items to be added to the callers database explicitly by using **si:add-files-to-who-calls-database** or **si:add-system-to-who-calls-database.**

After you use the function **si:enable-who-calls** with the argument best suited for the type of database you want to create, you can compress the database by using either **(si:compress-who-calls-database)** or **(si:full-gc)**.
(si:compress-who-calls-database) makes the database smaller, thus making the world load smaller if you save the world after using the function. If the database is large, this function may take many minutes to compress the database. Here are examples of the ways in which you can couple these functions:

- If you use the function **(si:enable-who-calls :new)**, you can load any special files (these may be layered products, private systems, the local site system, etc.) and then you use either either **(si:compress-who-calls-database)** or **(si:full-gc)** to compress the database.

- If you want to have the entire body of Symbolics-supplied software in your **who-calls** database, there are two options you can take during the customization of the distribution world. For the first option use this form:

  (si:enable-who-calls ':all-no-make)

  Followed by this form:

  (si:full-gc)

  (si:enable-who-calls ':all-no-make) creates a callers database that includes only new functions. When you do a **si:full-gc** the entire database is created. This takes a long time and about 2000 pages of storage. The world load file, when you save it, will be about 2000 pages bigger than it would be if you used the **:new** mode.

  The second option is to use this form:

  (si:enable-who-calls ':all)

  Followed by this form:

  (si:compress-who-calls-database)

  (si:enable-who-calls ':all) creates a full callers database. This also takes a long time and about 2000 pages of storage. **:all** also queries about the old state. **:all-remake** suppresses the query. **si:compress-who-calls-database** compresses the who-calls database by garbage-collecting the database.

- If you load your own programs into the database and wish to include any specific items in the database, you can use the function:

  (si:enable-who-calls ':explicit)

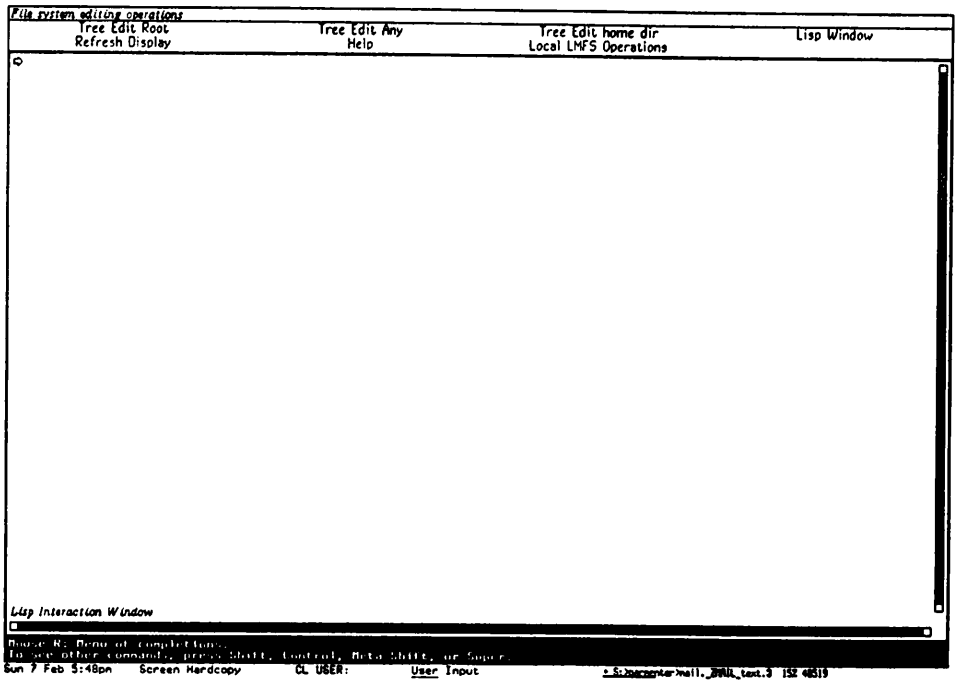

Figure 3. Initial File System Editing Operations Menu

6.1. File System Editing Operations Commands

The following table describes the menu structure of the File System Editing Operations Program. The commands are organized into four levels, in increasing order of potential for causing problems if you don't use them correctly. The table is followed by explanations of the commands at each level. Discussions of how these commands are used in various file maintenance procedures make up the next few chapters of this book.

Level 1	General file system editing operations.
Level 2	Local file system control operations. These are to be used by the person at your site responsible for maintaining the file system and doing backup dumps. This menu is summoned by clicking on [Local LMFS Operations] in the Level 1 menu.

Level 3 — Server and maintenance operations.

These should be used only by someone who is very knowledgeable about the file system. This menu is summoned by clicking on [LMFS Maintenance Operations] in the Level 2 menu.

Level 4 — File system internal data structure editing operations.

You should not attempt to use these commands unless you are an expert in Lisp Machine File System internal data structures. This menu is summoned by clicking on [LMFS Internal Tools] in the Level 3 menu.

Note: *The commands in menus three and four can damage your file system if used incorrectly. If you have any questions about these operations, please call Symbolics Software Support.*

Some commands are described as "typing out" certain information. If the large pane is a *Lisp Interaction Window*, such information is simply displayed on that window; if it is a *File System Editor*, then the information is displayed in a typeout window over the file system editor information. You can flush the typeout window and restore the display of the File System Editor by pressing any character or by clicking on [Refresh Display].

6.1.1. Level 1 Menu

File System Editing Operations:

[Tree Edit Root] — Enter the File System Editor, using the root directory of the local file system as the base directory. See the section "File System Editor" in *Reference Guide to Streams, Files, and I/O.* This puts the large pane into the *File System Editor* state.

[Tree Edit Any] — Enter the File System Editor; it prompts you for the name of the base directory. See the section "File System Editor" in *Reference Guide to Streams, Files, and I/O.* This puts the large pane into the *File System Editor* state.

[Tree Edit home dir] — Enter the File System Editor, using your home directory as the base directory. See the section "File System Editor" in *Reference Guide to Streams, Files, and I/O.* This puts the large pane into the *File System Editor* state.

[Lisp Window] — Put the large pane back into the *Lisp Interaction Pane* state. This is useful for getting out of the File System Editor.

[Refresh Display] — When the large pane is in the *File System Editor* state, and you use one of the commands that "types out" information, the information appears on top of the File System Editor window, and you are told "Type any char to flush:". You can

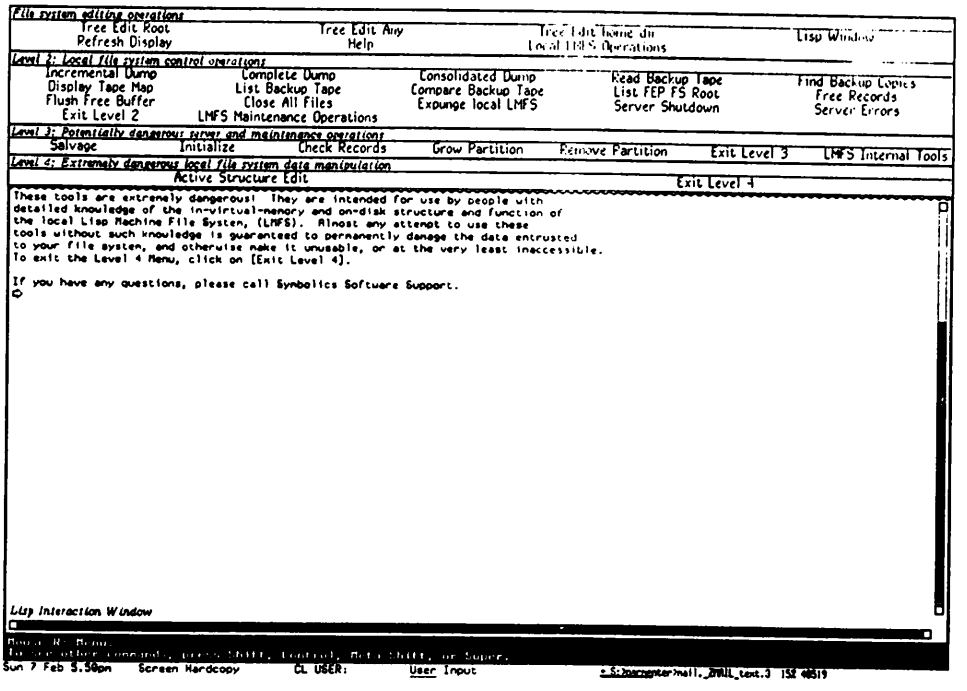

Figure 4. The Four Levels of the File System Editing Operations Menu

use this command to clear the screen and redisplay the File System Editor window without removing your hand from the mouse. You can also use this command to proceed from **MORE** pauses.

[Help] Type out general information about the File System Maintenance program and File System Editor.

[Local LMFS Operations]
Bring up the Level 2 menu Local File System Control Operations.

6.1.2. Level 2 Menu

The commands in this menu are intended to be used by the person at your site who is responsible for maintaining the file system and doing backup dumps.

Local File System Control Operations:

[Incremental Dump] Do an incremental dump of the local root directory. Offers an Accept-Values menu to adjust all parameters. See the section "Dumping, Reloading, and Retrieving", page 140.

[Complete Dump] Do a complete dump of the local root directory. Offers you an Accept-Values menu to adjust all parameters. See the section "Dumping, Reloading, and Retrieving", page 140.

[Consolidated Dump] Do a consolidated dump of the local root directory. Offers you and Accept-Values menu to adjust all parameters. See the section "Dumping, Reloading, and Retrieving", page 140.

[Read Backup Tape] Retrieve single files or reload full tapes. You select the activity you want from a pop-up menu. You can also use [Read Backup Tape] to list or compare tapes, if you change the default operation in the menu by clicking on Compare in the menu.

[Find Backup Copies] Locate files on backup tape. It prompts you for a file specification to locate.

[Display Tape Map] Display a directory listing of what should be on a backup tape. It prompts you for the tape number. You can display a listing of any backup tape directory on any machine connected to your network by giving the machine name and the pathname of the tape directory. Standard pathname defaulting and merging work.

[List Backup Tape] List the contents of a backup tape that you have mounted on a tape drive. It prompts you for the tape specification.

[Compare Backup Tape]
Compare the contents of a backup tape that you have mounted to the contents of the local file system. It prompts you for the tape specification. To compare the tape to another (remote) file system, use [Read Backup Tape].

[List FEP FS Root] List the FEP file system's root directory from the default disk unit. See the section "Show FEP Directory Command", page 4. See the function **zl:print-disk-label**.

[Free Records] Type out information about the number of free records in the local file system. The last line tells you how many records are marked as free, how many are marked as used, and the sum of these numbers, which is the total number of records in the file system.

Clicking middle on [Free Records] prepares a directory-by-directory usage report of record use, indicating how many records are in use by files in each directory. It prompts you for the name of a file in which to place the report.

Clicking right on [Free Records] displays how many records are in use in each partition. This information is necessary for the commands that allow you to change the size of, add, or remove partitions.

See the section "Free Records", page 182.

[Flush Free Buffer] Write the internal pool of free disk records back to the disk. This happens automatically when you log out. After doing this, you can cold boot without losing records. See the section "Free Records", page 182.

[Close All Files] Call **fs:close-all-files**. This has nothing to do with the Lisp Machine File System as such; it closes any open files in use by your machine, whether local or remote. This is occasionally useful for cleaning up after problems occur, but be aware that by using [Close All Files], you can cause new problems for any programs in the machine that are validly using files at the time.

[Expunge Local LMFS]
Expunge all directories on the local file system. It tells you how much space was recovered.

[Server Shutdown] Shut down file servers at a future time, or reschedule or cancel a shutdown.

This command lets you shut down a file server cleanly. You run this command only on a 3600-family computer that is acting as a file server for other users. Clicking left on [Server Shutdown] means that you plan to shut down the file system soon. It asks you for a short message to be sent to people using the file server, which you can use to explain why it is being shut down and when it will return. It also asks you when you want the shutdown to take place; the default is five minutes. All users of the file system are sent periodic messages warning them that the server is going to be shut down. Finally, when the time comes, it closes all Chaosnet servers on the machine, and disables creation of new servers. When servers are shut down, you can cold boot the machine or whatever else you want to do. While the shutdown is "in progress" (the messages are being sent), you can cancel it by clicking middle on [Server Shutdown] or reschedule it by clicking right on [Server Shutdown].

[Server Shutdown] only shuts down the network server; it does not affect the local operation of the file system itself. It shuts down all servers, not just file servers, since anything that requires the file servers to be shut down also requires that all servers be shut down.

[Server Errors] Display all the error messages associated with errors encountered by the file server. When such errors occur, you get a message that begins as follows:

```
[File Server got an error: ...]
```

The message contains descriptive information about the error.

[Exit Level 2] Return from this menu to the top-level menu of General File System Editing Operations.

[LMFS Maintenance Operations]
Bring up the Level 3 menu of Server and Maintenance Operations. **Note:**The operations on the Level 3 menu can damage your file system if misused. They should only be performed by someone who is very knowledgeable about the file system.

6.1.3. Level 3 Menu

The commands in this menu are intended only for users who are thoroughly familiar with the operations of the file system; misuse of these operations can destroy or damage the file system.

Server and Maintenance Operations (Potentially Dangerous):

[Salvage] Run the salvager. See the section "Salvager", page 183.

[Initialize] Create a new file system. Use this tool to add file storage space to your local Lisp Machine File System. This operation asks you several questions and prints out information to make sure you really want to initialize the FEP file that would contain the file system. These verifications are to ensure that you do not accidentally destroy any previous file system residing there. [Initialize] takes about a minute for each four thousand records (a record is four 256-word disk blocks). It queries you about the FEP file size before it initializes the new file system.

Clicking right on [Initialize] presents a menu of initialization. This is how you can add new FEP files to the running file system. For a description of this operation: See the section "Adding a partition to LMFS", page 186.

Note: Added partitions should never be initialized. The partitions are automatically initialized by the system when they are created.

This tool is intended only for manipulating the local Lisp Machine File System, specifically, adding partitions to it for the purpose of storing files in them. If you want to create a FEP file for some other purpose, or if you want to create an additional paging file, you must do so from Lisp, using the Create FEP File command. See the section "Create FEP File Command" in *Genera Handbook*. See the section "Allocating Extra Paging Space", page 34.

Warning: *Attempting to create an additional paging file using [Initialize] will destroy the file system on your machine, permanently and irretrievably.*

[Check Records] Check each record in the local file system for consistency, and notify you about any problems. (This is also available from the salvager.) This option scans the hierarchy, going through the directories, making sure that each directory entry really describes a file that agrees with it, and that each record of each file is validly identified as a part of that file.

[Grow Partition] Increases the number of blocks available in a partition. It offers you a menu to select the partition to grow and prompts for the number of blocks by which to increase the partition.

[Remove Partition] Remove active partitions from the file system and delete them. It does this by walking over the local file system, evacuating files and directories from the partitions to be removed, to other partitions. For medium and large file systems, this operation takes a long time. In order for it to succeed, there must be enough room in other partitions to contain the evacuated files and directories. This tool determines whether or not sufficient room exists for the operation to complete successfully, and queries if it suspects that sufficient room is not available. You can click right on [Free Records] to get a partition-by-partition report. Salvaging might be necessary to properly identify all free records. If you need to do this: See the section "Salvager", page 183.

[Remove Partition] provides the option of deleting the FEP file when all LMFS files have been removed from it. If it detects that a partition is not completely empty, it reports this and allows you to abort the process.

Do not attempt to use this tool to manipulate FEP files for any other purpose, or to manipulate FEP files in use by LMFS in any other way, for example, do not use this tool on these files: lmfs.file, lmfs1.file, or fspt.fspt. Misuse of this tool causes irretrievable destruction of data in your file system.

[Exit Level 3] Return from this menu to the Level 2 menu of Local File System Control Operations.

[LMFS Internal Tools]
Bring up the Level 4 menu of File System Data Manipulation Operations.

Warning:*Editing the internal structures of your file system can result in data being irretrievably lost. You should not use these tools unless you are sure you know what you are doing.*

6.1.4. Level 4 Menu

The commands in this menu are intended only for users who are thoroughly familiar with the internal organization and implementation of the file system; misuse of these operations can destroy or damage the file system.

Local File system data manipulation (Extremely Dangerous):

[Active Structure Edit]
Edit the active file system data structure. This displays "active" internal data structure as a scroll window, and is intended to be used by those debugging local file system problems.

[Exit Level 4] Return from this menu to the Level 3 menu of Server and Maintenance Operations.

7. Maintaining the Namespace Database

After you bring up the new release, you can perform the following steps as part of your site administration activities:

- Register users, hosts, printers, and networks in the site's namespace
- Move the new release to other Lisp Machines at the site
- Install new releases distributed in patch tape format
- Install new releases distributed in world load format
- Install world loads from other sites

You should reflect any changes, such as new users or changes in the site's hardware, in the namespace database. Register new hosts and printers in the namespace database before connecting them to the network or supporting host. Register new users in the namespace database either before they use the new release or as part of the process when they log in for the first time. Whether you are registering new users or new hardware, the process is most easily done by copying and modifying an existing object of the same type using the Namespace Editor (invoked by the Edit Namespace Object command or the **tv:edit-namespace-object** function), and then saving the new object. Once it has been saved, it is part of your site's configuration and all machines running the new release know about the new object the next time they boot, or sooner in some cases. (See the section "Namespace System" in *Networks*.)

To use the Namespace Editor to create and update objects, click on [Namespace] on the System menu or select a Lisp Listener and type the Edit Namespace Object command or the **tv:edit-namespace-object** function.

See the section "The Namespace Editor", page 52.

7.1. The Namespace Editor

7.1.1. Introduction to the Namespace Editor

The namespace is the means of tying Symbolics sites together. The namespace consists of a number of objects. A namespace object is some entity that is known to all the systems at your site, such as users, hosts, printers, or the site itself.

Each of these objects has a number of attributes associated with it. Some of these attributes are just for general information, but many of them are used by your system and other systems at your site for essential operations. For instance, you must

have a User object in the namespace to be considered a user. Attributes of your user object include your login name and home host, which must be present for you to log in, but also your birthday, which is just general information, although some programs might make use of it.

The Namespace Editor allows you to create, edit, show or delete any namespace object. There are CP commands for each of these actions.

- Show Namespace Object
- Edit Namespace Object
- Create Namespace Object
- Delete Namespace Object

If you go to a Lisp Listener and issue the following command:

```
Show Namespace Object User username
```

naming yourself or some other regular Symbolics user, you can get a good idea of what goes into a namespace object. If you issue

```
Show Namespace Object Host hostname
```

naming one of the server machines at your site, you'll begin to get a picture of the amount of information, useful to individuals as well as to systems on the network, that is stored in the namespace.

See the section "Introduction to the Namespace System" in *Genera User's Guide.*

The Namespace Editor enables you and other users, particularly your site manager, to control the information in the namespace. The Namespace Editor uses dynamic window technology to check for correct input, supply help, and offer completion. In addition, the Namespace Editor checks for errors in Chaos, Internet, and DNA addresses along with checking for nicknames in use by other hosts in the local namespace. Services are checked to see if they are known on the local machine and warnings are printed when there's a problem.

Now you can try the command:

```
Edit Namespace Object User username
```

Move the mouse around to various attributes of the namespace object and try editing them. As long as you don't save the object, you can do as you please, although it's probably a good idea to experiment only with your own object.

7.1.2. Using the Namespace Editor

The Namespace Editor is accessible in several different ways.

- Use the CP command Select Activity and name the Namespace Editor.
- Assign the Namespace Editor to a SELECT key combination by using SELECT =.

- Use the Namespace Editor CP commands.

See the section "Managing the Namespace Database" in *Networks*.

See the section "Software Interface to the Namespace System" in *Networks*.

7.1.2.1. Namespace Editor CP Commands

Show Namespace Object Command

Show Namespace Object *class name keywords*

Shows the information in the namespace database for *name*.

class — The type of object. Possible types are: Host, Site, Namespace, User, Network, and Printer.

name — The name of the object, that is a user-id, the name of a machine, or the name of the site or namespace.

keywords — :Format, :Locally, :Output Destination

:Format — {Normal, Detailed} Whether to show fields that are empty. The default is Normal, to omit empty fields. The mentioned default is Detailed, to show all fields.

:Locally — {Yes, No} Whether to show the information cached in the local machine or to consult the namespace server. The default is No, to consult the server. The mentioned default is Yes.

:Output Destination — {Buffer, File, Printer, Stream, Window} Where to redirect the typeout done by this command. The default is the stream ***standard-output***.

Here is what the namespace object for a user might look like:

```
Command: Show Namespace Object (a namespace object [default Site ACME]) User KJONES

Showing USER KJONES in namespace ACME:
Lispm Name: Kjones
Personal Name: Jones, Kingsley
Nickname: King
Work Address: Building 3-701
Work Phone: 5891
Home Host: ACME
Mail Address: Kjones ACME
Birthday: 19 June
Project: Database
Supervisor: Finklestein
```

Edit Namespace Object Command

Edit Namespace Object *class name keywords*

Modify an object in the namespace database.

To create a new object:See the section "Create Namespace Object Command", page 55.

class {User Printer Network Host Site Namespace nil} The kind of object to create or edit. The default is nil, which runs the Namespace Editor as an activity with no class specified.

name Name of the object to create or edit. The default is any.

keywords :locally

:locally {yes no} Whether to edit only a local copy of the information for the object. The default is no, meaning to edit the object in the central namespace database. The mentioned default is Yes, edit only a local copy.

Create Namespace Object Command

Create Namespace Object *class name keywords*

Adds a new object to the namespace. See the section "Creating a New Namespace Object" in *Genera User's Guide.*

class {File-System, Host, Namespace, Network, Printer, Site, User} The type of object you want to create.

name The name of the new object.

Keywords:

:Copy From {*name*} The name of an object of the same class to provide the initial property list.

:Locally {Yes, No} Whether to edit only a local copy of the object. The default is No, the mentioned default is Yes.

:Property List Initial property list for the object.

Delete Namespace Object Command

Delete Namespace Object *class name keywords*

Removes information about the object *name* from the namespace.

class {File-System, Host, Namespace, Network, Printer, Site, User} The type of object you want to delete.

name The name of the object to be deleted.

keywords :Locally

:Locally {Yes, No} Whether to delete only a local copy of the object. The default is No. The mentioned default is Yes.

Add Service to Hosts Command

Add Service to Hosts *service-triple hosts*

Adds a service attribute in the namespace for one or more hosts.

service-triple A service triple consists of a service, a medium, and a protocol, such as NETBOOT SLAP NETBOOT

hosts The name of one or more hosts to which you want to add the service.

Keywords:

:locally {yes, no}The default is no. A no answer means the actual namespace database on the namespace server is changed. A yes answer means that if and when you save your changes, they are only changed in your Lisp environment. Only you see the changes, and the changes are forgotten when you cold boot unless you save the changed world.

:verbose Print messages for each host modified.

Here is an example adding netboot service to three hosts at once:

```
Command: Add Service To Hosts (A Service Triple service)) netboot
   (medium) slap (protocol) netboot (A sequence of hosts or All)
   HARPAGORNIS, WINTER, TOWHEE
   (keywords) :Locally

Adding service NETBOOT SLAP NETBOOT to hosts (locally).
  Done.
```

Remove Service From Host

Remove Service from Hosts *service-triple hosts*

Removes a service attribute in the namespace for one or more hosts.

service-triple A service triple consists of a service, a medium, and a protocol, such as NETBOOT SLAP NETBOOT

hosts The name of one or more hosts from which you want to remove the service.

Keywords:

:locally {yes, no}The default is no. A no answer means the actual namespace database on the namespace server is changed. A yes answer means that if and when you save your changes, they are only changed in your Lisp environment. Only you see the changes, and the changes are forgotten when you cold boot unless you save the changed world.

:verbose Print messages for each host modified.

Here is an example of removing netboot service from a host:

```
Remove Service From Hosts (A Service Triple (service) netboot
(medium) slap (protocol) netboot (A sequence of hosts or All) HARPAGOR

Removing service NETBOOT SLAP NETBOOT from hosts.
  Done.
```

7.1.2.2. Editing a Namespace Object

First select the Namespace Editor by using the Edit Namespace Object command. If you do not supply the class and object name to Edit Namespace Object, the Namespace Editor window comes up empty and you can click on [Edit Object] or enter the Edit Object command. You are prompted for the name of an object to edit. The current information for the object is retrieved from the namespace database and displayed in the window.

Click Middle on the attribute name for information on the attribute.

The attribute fields are mouse-sensitive. Clicking on an attribute prompts you for information. Mouse clicks have the following meaning:

Left Replace the information in the attribute.

Middle Edit information in the attribute.

Right Menu.

Sh-Middle Delete the information in the attribute.

The window can be scrolled using the SCROLL and m-SCROLL keys, the scroll bar, or the mouse.

Once you have finished editing the information, you have three ways to proceed. You can click on [Quit] without saving the changed information. If you are just practicing using the Namespace Editor, that would be appropriate.

The other two choices are to save the information locally or globally. If you save it globally, the new information is stored in the site's namespace database. If you save it locally, the new information is stored only in your machine's local copy of the namespace; changes save locally affect only your machine.

The initial state of the Namespace Editor is the global mode. When you are in global mode the middle of the screen looks like:

Editing HOST TOWHEE in namespace SCRC

If you have clicked on [Locally], you are in local mode. The middle of the screen looks like:

Editing HOST TOWHEE (locally) in namespace SCRC

You can click on [Locally] to toggle the mode between global and local. When you are ready, click on [Save Object] to save the information. Then click on [Quit] to exit the Namespace Editor.

7.1.3. Namespace Editor Commands and Menu Items

The Namespace Editor commands are available as menu items, or as commands you type to the `Namespace:` prompt. These are the commands:

- Edit Object — Alter an existing object.
- Save Object — Save an object.
- Copy Object — Create a new namespace object similar to an existing object.
- Create Object — Create a new namespace object.
- Delete Object — Delete an object.
- Revert Object — Eliminate the effects of editing the object.
- Locally — Cause changes either to be limited to the current world or to affect the site's namespace database.
- Clear History — Clear the history of objects read into the Namespace Editor.
- Previous Object — Go back to the previously edited object. You can do this with `c-m-L` also. `c-m-L` takes a numeric argument as well.
- Show History — Show all the namespace objects that have been edited. You can do this with `c-0 c-m-L` also.
- Not Modified — Mark the object not modified. You can do this with `m-~` also.

The Edit Object Namespace Editor Command

Edit Object *namespace-object namespace*

Reads in the object from a namespace server, makes it the current object, and allows you to make changes.

If *namespace-object* is described in more than one namespace (is "multi-homed"), Edit Object prompts for a second argument, specifying which namespace you want to edit. In general a namespace object is only described in one namespace, and the *namespace* argument is not required.

Keyword:

:locally — {yes, no}The default is no. A no answer means the actual namespace database on the namespace server is changed. A yes answer means that if and when you save your changes, they are only changed in your Lisp environment. Only you see the changes, and the changes are forgotten when you cold boot unless you save the changed world.

Examples:

```
Edit Object User Washington

Edit Object User Washington :Locally Yes

Edit Object Host Gateway-1 ABC
```

The Save Object Namespace Editor Command

Save Object

Save the current object. If the object is being edited locally, save the changes only to the Lisp world. Otherwise, save the changes to the namespace database.

Keyword:

:force-save — {Yes, No} The default is No, which means do nothing if no changes were made. If the answer is Yes, saves the object even if no changes were made.

The Copy Object Namespace Editor Command

Copy Object *name*

Creates a new namespace object named *name* of the same type as the current object. The initial contents of the new object are copied from the current object. You can now edit any of the contents, and save the new object.

Keyword:

:locally — {Yes, No} The default is No. If the answer is No, the actual namespace database on the namespace server is changed. If the answer is Yes, the changes are made only locally. This means that if and when you save your changes, they are only made in your Lisp environment. Only you see the changes, and the changes are forgotten when you cold boot unless you save the changed world.

See the section "The Create Object Namespace Editor Command", page 60.

The Create Object Namespace Editor Command

Create Object *class name*

Creates a new namespace object of type *class* named *name* and makes it the current object. The object can then be edited and saved.

Keywords:

:copy from	{name of an object of the same class} This will copy the contents of the object named into the newly created object. Duplicate names are removed before the contents are inserted into the new object to avoid conflicts.
:locally	{Yes, No} The default is No. The actual namespace database on the namespace server is changed. A Yes answer means make the changes locally. This means that if and when you save your changes, they are only made in your Lisp environment. Only you see the changes, and the changes are forgotten when you cold boot unless you save the changed world.
:property-list	{a Lisp expression} This specifies an initial property list for the newly created object in the internal form that object properties are specified in. This keyword option will have no effect if the :copy from keyword option is used.

The Delete Object Namespace Editor Command

Delete Object

Deletes the current object from the namespace database. If the current object is being edited locally, the deletion only happens in Lisp virtual memory, not in the real namespace database. This command prompts for confirmation.

The Revert Object Namespace Editor Command

Revert Object

Discards any changes that you made to the current object. If the object is being edited locally, it is reverted to the version in virtual memory on the local machine. If the object is being edited globally, a fresh copy of the object is read from the namespace server.

The Locally Namespace Editor Command

Locally

Changes whether the current object is being edited locally or not. If the current object is being edited locally, change it to be edited globally, and vice versa.

If you are editing an object locally, if and when you save your changes, they are only made in your Lisp environment. Only you see the changes, and the changes are forgotten when you cold boot unless you save the changed world.

If you are editing an object globally, which is the default in most cases, the actual namespace database on the namespace server is changed. This means the information in the change is available to the entire site.

The Clear History Namespace Editor Command

Clear History

Clears all history information from the Namespace Editor. After this command, there is no current object, and no history for the Previous Object command to use.

See the section "The Show History Namespace Editor Command".

The Previous Object Namespace Editor Command

Previous Object

Selects the previous object on the history list and makes it the current object. This command is also bound to c-m-L.

Doing this once turns the current object into the previous object. This means you can use c-m-L to go from one to the other.

Keyword:

:count	{*integer*} Means go back this many objects on the history list. The default is 2. An equivalent numeric argument can be given to c-m-L.

See the section "The Show History Namespace Editor Command".

The Help Namespace Editor Command

Help *topic*

Displays help about *topic* if it is available. Commands, menu items, overviews, and object attributes can all be topics.

The Not Modified Namespace Editor Command

Not Modified

Changes the status of the current object to not modified. This command is also bound to m-~. See the section "The Save Object Namespace Editor Command", page 59.

7.1.4. Attributes in the Namespace Editor

7.1.4.1. User Attributes in the Namespace

Login Name

Specifies the appropriate login name for each of several hosts; a set of pairs. The first element of each pair is a token giving the login name, and the second element is the host object that corresponds to that name. Generally, you should have one **login-name** attribute filled in for every account that you have on a host on the network.

The Symbolics computer uses these login names when it connects to a host to log in a file server or a tape server. **login-name** is not required, but lack of this attribute causes the Symbolics computer to ask for the name to use for each server, which might be inconvenient. Passwords are not stored in the database because it is not secure; the Symbolics computer prompts the user for a password interactively when one is required.

```
Login Name:  GWash VIXEN
Login Name:  GWash PEGASUS
```

Lispm Name

Specifies the user name displayed in the status line; a token (required). Used by the **lispm-finger** service as the user name. The Lisp variable **zl:user-id** is set from this attribute. Typically it is similar to the actual name of the user object, but with upper- and lower-case so it looks better.

```
LispM Name*: GWash
```

Personal Name

Specifies the user's personal name; a token (required).

```
Personal Name*: George Washington
```

Nickname

Specifies a personal nickname; a token. Unlike host nicknames, user nicknames cannot be used to look up the user.

```
Nickname: Sleeper
```

Work Address

Specifies a work (business) address; a token.

```
Work Address: The White House, Washington D.C.
```

Work Phone

Specifies the work (business) phone number; a token.

```
Work Phone: 202-555-1212
```

Home Address

Specifies the user's home address; a token.

```
Home Address: Mount Vernon VA
```

Home Phone

Specifies the user's home phone number; a token.

```
Home Phone: 202-999-1234
```

Home Host

Specifies the user's host machine; a host (required). This is the system from which the user's lispm-init file is read.

```
Home Host*: SHOOFLY
```

Mail Address

Specifies the network mailbox at which the user wants to receive mail; a pair (required). The first element is the mailbox name (a token), and the second element is a host object. Defaults to name@home-host.

```
Mail Address*: GWash@VIXEN
```

Birthday

Specifies the user's birthday; a token.

```
Birthday: February 22
```

The Project User Attribute

Specifies what the user is working on; a token.

```
Project: Father of Country
```

The Supervisor User Attribute

Specifies who the user is working for; a token.

```
Supervisor: The People
```

Affiliation

Specifies the user's group affiliation; a single character. The character is arbitrary and can refer to different sets of users at different sites.

```
Affiliation: Z
```

The Remarks User Attribute

Specifies other relevant information; a token.

```
Remarks: "I cannot tell a lie."
```

The Type User Attribute

Specifies a special user type such as a DAEMON. No value is needed for non-special users.

The User Property User Attribute

Specifies a user-chosen property for this object; a pair whose first element is an indicator (by analogy with property lists) and whose second element is a token denoting whatever the user chooses to associate with that indicator. Several classes of objects have the **user-property** attribute, including users, hosts, printers, sites, namespaces, and networks. This is simply a place-holder where you can store any extra information. For example:

```
User Property: ID-number 123-45-6789
```

7.1.4.2. Host Attributes in the Namespace

Site Host

Specifies the site at which this host is located; a site object (required).

```
Site: SCRC
```

Nickname

Specifies alternate names for the host. You can have several nicknames. You can always add another.

```
Nickname: Junko
Nickname: a name in namespace ACME
```

Short Name

Specifies additional nicknames; a set of names. A short-name is used when a program wants to display a host's name without using up too much space. This is also used in the printed representation of pathnames.

```
Short Name: J
```

Machine Type

Specifies the hardware type of the machine this host is; a global-name. For example:

```
Machine Type: 3653
```

Common values for **machine-type** are:

```
3600
3610
3620
3630
340
3645
3650
3653
3670
3675
vax
pdp10
pdp11
ibmpc
honeywell-dps-8m
alto
pe3230
cadr
```

System Type

Specifies the operating system run on the host; a global-name (required). The Symbolics system uses this information to figure out how to parse pathnames for a given host; be sure to enter this information correctly. For example:

```
System Type*: LISPM
```

Common values are:

Value	*Type and Version of Software*
`lispm`	Symbolics software, any version
`unix42`	UNIX version 4.2BSD and later versions
`unix`	UNIX versions prior to 4.2BSD

vms4.4	VMS version 4.4 and later versions
vms4	VMS version 4.0, 4.1, 4.2, and 4.3
vms	VMS versions prior to version 4
tops-20	TOPS-20 software, any version
alto	ALTO software, any version
its	ITS software, any version
multics	MULTICS software, any version
minits	MINITS software, any version
magicsix	MAGICSIX software, any version
mos	MOS software, any version
ms-dos	MS-DOS software, any version

Address

Specifies the network addresses of this host; a set of pairs. Addresses are always represented as tokens because each kind of network has a different kind of address. The network must be a valid network that has already been created. The address must be be the correct type for that network to be accepted.

For information on the individual network types and their corresponding address conventions: See the section "Network Addressing" in *Networks*. Here is an example of a pair:

```
Address:  CHAOS 401
```

Pretty Name

Specifies a "pretty" version of the name of the host; a token. Unlike the real name, the nicknames, and the short name, this does not count as a name as far as the database system is concerned (you cannot use it to find the host). It appears in the herald.

```
Pretty name: Slate-colored Junco
```

The pretty name normally appears on screen displays and in prompts. That is, it appears where real people need to see the name. The actual name of the object is used by the software.

Finger Location

Describes the physical location of the host itself; a token.

```
Finger-location:  Across the alley from the Alamo
```

Note: This is used by Release 6 systems when they are performing the **lispm-finger** and **show-users** services. Genera 7.0 and later systems use the **console-location** attribute instead, unless it is not filled in.

Location

Specifies a description of the physical location suitable for programs to understand; a pair. The first element is a token that identifies the bu'lding the machine is in. The second element is a token that says what floor of the building the machine is on.

```
Location: Lab 2
```

Note: This is used by Release 6 systems. Genera 7.0 and later systems use the **console-location** attribute instead, unless it is not filled in.

Console Location

Describes the physical location of the host's console; a triple.

You are prompted for the building, the floor number, and a description of the location, but you can, in fact, enter any string values you like, using "quotes" to enter multiword values.

```
Console-location: Empire State 107 Near King Kong
```

Printer

Specifies the preferred printer for this host; a printer object. This printer is used by default when files are hardcopied from this host. If this attribute is not provided, the site's **default-printer** attribute is used.

```
Printer: American
```

Bitmap Printer

Specifies the preferred bitmap printer for this host; a printer object. This printer is used by default when screen images are hardcopied from this host. If this attribute is not provided, site's **default-bitmap-printer** attribute is used.

```
Bitmap Printer: Asahi
```

Print Spooler Options

Specifies options for any print spoolers running on this host; a set of pairs of global-names and tokens. A typical global-name for print-spooler-options is Home-directory; its value denotes the directory where hardcopy requests are stored. The default for Symbolics computers is local:>print-spooler>.

```
Print Spooler Options:  Home-directory local:>print-spooler>
```

Spooled Printer

Specifies printers for which this host provides a spooling service. When you enter a printer object, you are prompted for a Home Directory and pairs giving printer options for that printer.

```
Spooled Printer: PRENSA
  Home Directory: a pathname of a file
  Other Options: zero or more pairs of a global name and a token
```

The home directory is where hardcopy requests are stored. The default for Symbolics computers is local:>print-spooler>.

Service

Specifies services and protocols supported by this host; a list of triples of the form *service medium protocol*. Each triple specifies that the host is capable of providing *service* when you connect to it using *medium* and *protocol*.

```
Service:  FILE CHAOS NFILE
```

For information on services, mediums, and protocols: See the section "Service Attributes in the Namespace Database" in *Networks*.

Server Machine

Specifies whether the object described is a server machine; a token. If the value is Yes, the host is a server machine. If No, which is the default, the host is not a server machine. The default is of No is actually **nil**, or undefined.

```
Server Machine: Yes
```

This attribute only applies to Symbolics computers. Server machines do not automatically enable their services when you boot them. This is to prevent premature creation of servers before the machine has completely initialized. You must enable services yourself, either in your lispm-init file, using **sys:enable-services** or the Enable Services command: See the section "Enable Services Command" in *Genera Handbook*.

File Control Lifetime

Specifies the lifetime of a file control connection; a time interval. See the section "Reading and Printing Time Intervals" in *Programming the User Interface -- Concepts*. When a Symbolics computer connects to this host as a user of the **file** service, it will automatically close its control connection after the specified time interval.

```
File Control Lifetime: 30 minutes
```

Peripheral

Specifies a peripheral device. Click on the peripheral type. You are prompted for the values associated with that peripheral type.

```
Peripheral: None  Graphics-Tablet  Kanji-Tablet Modem Pad  Sdlc-Interface
Serial-Unit: a decimal integer greater than or equal to 0 and less than or equal to 3
  Baud: 300  600  1200  1800  2000  2400  3600  4800  7200  9600  19200  56000
  Address: a string
  Autoanswer: Yes  No
```

Default Secondary Name Server

If Yes, specifies that this host acts as a secondary namespace server for any namespace that has a namespace server at all. This is used in sites that have several namespaces.

A typical configuration is for each primary namespace server to also act as a secondary namespace server for the other namespaces. This is beneficial because that host performs namespace service for any namespace whose server is unreachable.

This configuration can be achieved two ways:

- Each namespace can list all other namespace servers in the **secondary-namespace-server** attribute of its namespace object.
- Each namespace server can set the **default-secondary-namespace-server** attribute to Yes in its own host object.

```
Default Secondary Namespace Server:  Yes
```

Internet Domain Name

The Internet Domain Name associated with the namespace; a token. See the section "Dialnet and Internet Domain Names", page 189.

```
Internet Domain Name: SCRC.Symbolics.COM
```

User Property

Specifies a user-chosen property for this object; a pair whose first element is an indicator (by analogy with property lists) and whose second element is a token denoting whatever the

user chooses to associate with that indicator. Several classes of objects have the **user-property** attribute, including users, hosts, printers, sites, namespaces, and networks. This is simply a place-holder where you can store any extra information. For example:

```
User Property: ID-number 123-45-6789
```

7.1.4.3. Site Attributes in the Namespace

Pretty Name

Specifies a version of the name suitable for people to read; a token.

The pretty name normally appears on screen displays and in prompts. That is, it appears where real people need to see the name. The actual name of the object is used by the software.

```
Pretty Name: The Old Mill Stream
```

Site Directory

Specifies the file name of the directory that holds the Symbolics computer system's site-specific files at this site; a token (required). This is used only to find the files that define the logical hosts, such as sys:. All other site-specific pathnames are managed by logical pathname translations or by the descriptor file attribute of a namespace.

```
Site Directory*: blue:>sys>site>
```

Site System

Specifies the name of a system (in the **defsystem** sense) to be loaded into Symbolics computer worlds at this site; a token. The site's system installers should load this system into worlds to be used at the site; the loading does not happen automatically. A warning is given if a user logs into a machine and the site's site-system is not loaded in the world.

```
Site System: HARV-SPECIFIC
```

Default Printer

Specifies the default printer to use for printing text files at this site; a printer object. This will be used by hosts that do not have their own **printer** attribute.

Default Bitmap Printer

Specifies the default printer to use for printing screen images at this site; a printer object. This attribute is for hosts that do not have their own **bitmap-printer** attribute.

Host For Bug Reports

Specifies the host to which bug reports should be sent (required). The Report Bug CP command, and the commands in the Debugger, Editor, and Zmail for reporting bugs use this attribute.

```
Host for Bug Reports*: blue
```

Timezone

The timezone at this site; a global-name (required).

```
Timezone*: EST
```

Secure Subnets

Specifies an association of networks and secure subnet numbers; a set of pairs. The first element of each pair is a type of network; it must be CHAOS or INTERNET. The second element of each pair is a set of subnet numbers, the interpretation of which depends on the type of the network. For a CHAOS network, the set is represented as octal character strings. For an INTERNET network, the set is represented as decimal character strings. Hosts on these subnets are considered trustworthy.

This attribute controls the subnet security feature of any servers that use the **:trusted-p** or **:reject-unless-trusted** keywords to **net:define-server**. The following Symbolics servers respect the **secure-subnets** attribute: NFILE, QFILE, TCP-FTP, and TFTP.

Dont Reply To Mailing Lists

Specifies a set of names of mailing lists to which Zmail does not reply by default; tokens. This attribute is useful only to those who have not set the PEOPLE NOT TO REPLY TO option in their Zmail init files.

Other Sites In Mail Area

A list of other sites that share the same list of mailboxes as this site.

Host Protocol Desirability

Specifies a tuning factor to be used in the Generic Network System's cost estimates when trying to construct a path to a service; a triple of the form (host protocol desirability), in which host represents a host, protocol names some protocol that host supports, and desirability is a token expressing a floating-point factor for the cost calculations. See the section "Desirability of Network Protocols" in *Networks*.

For example:

```
Host Protocol Desirability: YUKON CHAOS-MAIL 0.75
```

Services and protocols are discussed elsewhere: See the section "Symbolics Generic Network System" in *Networks*.

If you change the value of **host-protocol-desirability**, you must either cold boot or use the following function, to make the change take effect: See the function **neti:recompute-all-namespace-server-access-paths** in *Networks*.

Local Namespace

Specifies the site's local namespace; a namespace object (required). This is the namespace that will be used at the site. Normally, there is exactly one namespace for each site.

```
Local Namespace*: harvard
```

The asterisk (*) in a namespace attribute prompt indicates that you must supply a value for that attribute.

Other Sites Ignored In Zmail Summary

Specifies a set of site objects. Zmail does not display the host names of hosts from the specified sites in its summary window as well as not doing so for this site.

All Mail Addresses Forward

Set to Yes to have mailers on Symbolics machines at the site automatically forward all mail addresses that have the host supplied in the mail address.

Root Domain Server Address

A pair of a network and an address for a root domain server used by the Domain Name System. Here are some examples:

```
Root Domain Server Address: INTERNET 10.0.0.51
Root Domain Server Address: INTERNET 10.1.0.17
Root Domain Server Address: INTERNET 128.213.5.17
```

See the section "How to Install the Internet Domain Names System" in *Networks*.

Standalone

Specifies whether the host at this site is a standalone machine; a token. If the value is the string "yes", then only one host exists at this site and no response to the *who-am-I* network broadcast request at boot time is expected. If the attribute is not present or the value is not "yes", then multiple Symbolics computer hosts exist at this site; when one host is booted, another host answers its who-am-I query.

Validate Lmfs Dump Tapes

Specifies whether the LMFS backup dumper attempts to validate backup tapes; a token. If the value is "yes", then the LMFS backup dumper validates backup tapes. If the value is not "yes" or if the attribute is not provided, no validation is done.

Terminal F Argument

An associate set of specifications for what the various arguments to the FUNCTION F key should do. Each component is a triple consisting of a number (a string of the decimal number) or the string "none", a global name, and a set of hosts. The global names can be:

login	The login file computer.
local-lisp-machines	All Symbolics computers at this site.
all-lisp-machines	All Symbolics computers on the local network.
host	The hosts in the third element of the triple.

For example:

```
Terminal F Argument: NONE LOCAL-LISP-MACHINES
Terminal F Argument: 0 ALL-LISP-MACHINES
Terminal F Argument: 1 HOST VIXEN CUPID COMET
Terminal F Argument: 2 LOGIN
```

User Property

Specifies a user-chosen property for this object; a pair whose first element is an indicator (by analogy with property lists) and whose second element is a token denoting whatever the user chooses to associate with that indicator. Several classes of

objects have the **user-property** attribute, including users, hosts, printers, sites, namespaces, and networks. This is simply a place-holder where you can store any extra information. For example:

```
User Property: ID-number 123-45-6789
```

7.1.4.4. Printer Attributes in the Namespace

Type

Specifies the device type of the printer; a global-name (required). This attribute implies some data formats that are interpreted by the device. For example:

```
Type*: LGP2
```

Common values are:

lgp
lgp2
ascii
press
xgp

Site

The site where the printer is located; a site object. Generally all printers at a site are offered in menus of potential output devices for the destination of a hardcopy request.

```
Site: SCRC
```

Pretty Name

Specifies a name for the printer; a token.

```
Pretty Name: Caspian Sea
```

The pretty name normally appears on screen displays and in prompts. That is, it appears where real people need to see the name. The actual name of the object is used by the software.

Format

Specifies the print formats supported by the device; a set of global-names. These are in addition to those implied by the **type** attribute.

```
Format: LGP
```

Common print formats are:

lgp
lgp2

press
xgp
ascii
tektronix

Interface

Specifies the type of interface by which this printer is attached to its host. Click on the correct choice.

```
Interface: Serial  Elp  Other
```

Interface Options

Specifies parameters of the hardware interface. Click on the correct values.

```
Interface Options:
  Unit: 1
  Baud: 300  600  1200  1800  2000  2400  3600  4800  7200  9600  19200  56000
  Other Options: zero or more pairs of a global name and a token
```

Host

Specifies the host to which the printer is directly connected; a host object (required).

```
Host*: LETHE
```

Protocol

Specifies the protocols to use for direct unspooled printing; a set of global-names. If protocol is not specified, the HARDCOPY service is invoked on the host to which the printer is directly connected.

Default Font

Specifies the name of the font that should normally be used for this printer; a token. If not specified, the default-font is usually determined by the type of printer.

Note: This attribute is used when a Release 6 system requests hardcopy. Genera 7.0 systems use the **body-character-style** attribute instead, unless it is not filled in.

Header Font

Specifies the name of the header font that should normally be used; a token. If not specified, the header-font is usually determined by the type of printer.

Note: This attribute is used when a Release 6 system requests hardcopy. Genera 7.0 systems use the **header-character-style** attribute instead, unless it is not filled in.

Body Character Style

Specifies the name of the character style that should normally be used for this printer; a character style. The first element is the family; the second element is the face; the third element is the size. See the section "Character Styles" in *Symbolics Common Lisp -- Language Concepts*. If not specified, the default character style is usually determined by the type of printer.

```
Body Character Style: SWISS.ROMAN.LARGE
```

DPLT Logo

Specifies the name of the logo printed by DPLT; a global-name.

```
DPLT-LOGO: Symbolics
```

Character Size

Specifies the size of a character in micas; a pair of width and height, in decimal. (A *mica* is 10 microns, or 1/2540 of an inch.)

Page Size

Specifies the size of the page in device units; a pair of width and height, in decimal.

```
Page Size: 135 80
```

Fonts Widths File

Specifies the name of the `fonts.widths` file for this printer; a token. It is best if this is a fully qualified physical pathname instead of a logical pathname, for example:

```
Font Widths File: SCRC|A:>sys>stats>lgp-1>fonts.widths
```

Printer Location

Describes the physical location of the printer; a triple. The first element is a token that identifies the building. The second element is a token that is the floor number. The third element is a textual description.

```
Printer Location: 11CC 3 In joseph's office
```

User Property

Specifies a user-chosen property for this object; a pair whose first element is an indicator (by analogy with property lists) and whose second element is a token denoting whatever the user chooses to associate with that indicator. Several classes of objects have the **user-property** attribute, including users, hosts, printers, sites, namespaces, and networks. This is simply a place-holder where you can store any extra information. For example:

```
User Property: ID-number 123-45-6789
```

7.1.4.5. Namespace Attributes in the Namespace Editor

Primary Name Server

Specifies those hosts that are primary namespace servers for this namespace; a set of host objects. A primary server is an authority regarding its namespace. The namespace data are stored in files controlled by the primary namespace server.

```
Primary Name Server: BLUE
```

Secondary Name Server

Specifies secondary namespace servers for this namespace; a set of host objects. A secondary server is not an authority on a namespace, but can provide a backup in case the primary server is temporarily unavailable. It attempts to keep a copy of the namespace information current by querying the primary server more often than a nonserver machine would.

```
Secondary Name Server: ORANGE
Secondary Name Server: PINK
Secondary Name Server: the name of a host
```

Search Rules

Specifies the search rules, expressed as a set of namespaces (required).

```
Search Rules*: HARVARD YALE
```

See the section "Names and Namespaces" in *Genera User's Guide.*

Descriptor File

Specifies the descriptor file for the namespace; a pathname (required). See the section "Namespace Database Descriptor Files" in *Networks.*

```
Descriptor File*: BLUE:>SYS>SITE>HRV-NAMESPACE.TEXT
```

Internet Domain Name

The Internet Domain Name associated with the namespace; a token. See the section "Dialnet and Internet Domain Names", page 189.

```
Internet Domain Name: SCRC.Symbolics.COM
```

User Property

Specifies a user-chosen property for this object; a pair whose first element is an indicator (by analogy with property lists) and whose second element is a token denoting whatever the user chooses to associate with that indicator. Several classes of objects have the **user-property** attribute, including users, hosts, printers, sites, namespaces, and networks. This is simply a place-holder where you can store any extra information. For example:

```
User Property: ID-number 123-45-6789
```

7.1.4.6. File System Attributes in the Namespace for Statice

File System attributes in the namespace are used by the Statice product.

For more information: See "File System Attributes in the Namespace"

7.2. Registering Users

The easiest registration strategy is to create your entry by copying someone else's entry. To copy another entry, use the Edit Namespace Object command and then click on the namespace editor's [Copy] command. If you are the first person at your site to register, copy the user object for user LISP-MACHINE.

Individual users can run the Edit Namespace Object command for themselves the first time they use a new release. If they have not created an appropriate user object, then logging into the new release fails because the system does not find the user object in the namespace database. Should this be the case, the system offers to create the user object with the Edit Namespace Object command. Click on the namespace editor's [Copy] command and copy the user object for user LISP-MACHINE.

For an overview of changing objects in the namespace database: See the section "Maintaining the Namespace Database", page 52.

7.3. Registering Hosts

To create the new host object in the site namespace, type:

```
Edit Namespace Object Any RETURN
```

Click on [Create] and then specify the name of the new host. Or, use the form:

```
(tv:edit-namespace-object :host New-Host :create t)
```

To determine the service attributes required for your LISPM: See the section "Define Site Command", page 154.

7.4. Registering a Tape Drive in the Namespace

To register a tape drive for a Symbolics machine, use the `Edit Namespace Object` command to add the tape drive to the namespace database. Specify `tape chaos rtape` as the last `Service: Set:` entry in the host object. See the section "**service:** host object attribute", page 68.
For example:

```
Edit Namespace Object :host Janis RETURN
```

pops up namespace menu that displays all the attributes of the host Janis. Add the tape service by clicking on `Set:` of the last `Service: Set:` entry, then type:

```
tape chaos rtape RETURN
```

8. The Front-End Processor

8.1. Introduction to the FEP

Symbolics computers use a front-end processor known as the FEP. The FEP is a small computer inside the processor cabinet, based on a microcomputer chip. It plays several roles in the operation of the system, the most visible being booting (loading and starting the software of) the Lisp system.

This discussion uses FEP EPROM version 127 or higher as an example. Use the Show Version command to determine the FEP version of your machine. See the section "FEP System Commands: General Usage". If you have FEP EPROMS of version lower than 127, please contact Customer Service.

"V127" is an example. On your system, use the version number of the FEP EPROM which your machine is using. This may not be V127. Other FEP EPROMS are G206 and G208.

The FEP system has two parts:

- EPROMS containing the kernel of the FEP system.
- Loadable overlay files (with the extension .FLOD) containing the rest of the FEP software, in the FEP file system, commonly called flods, or flod files.

The FEP system provides a kernel that is independent of all changeable knowledge of Genera and Symbolics machines. This kernel remains constant and does not need to be modified to support new features. Instead, support for new features can be provided by overlay files that are read from disk or cartridge tape. New FEP features are distributed as part of software releases in flod files.

8.2. Using the FEP

The FEP system software implements a small number of commands in EPROM and the rest in the *overlay files*, which have an extension of .FLOD. To be used, a command must be defined in an overlay file that has been scanned. Additionally, the overlay file must be resident in memory. Only one overlay file can be resident in memory at any time.

Descriptions of each FEP command show where the command comes from. Not all commands are available on all machines at all times.

Scanning an overlay file makes all the commands it contains known to the FEP. Scanning inserts all the commands defined in that file into the FEP's list of valid commands. This list is the command tree; it indicates which commands are avail-

able and in which overlay file they reside. Commands remain in the command tree unless the FEP is reset or the machine is powered down. After an overlay file has been scanned, it can be loaded into memory.

When you type a command, the FEP system either dispatches directly to the command (if the command is supported in the kernel, or if the correct overlay is resident) or automatically loads the overlay file first and then dispatches to the command. In the latter case, prompting for the next command argument is delayed a short time while the overlay file is read.

Here is a list of the overlay files and what they contain. The wildcard (*) is replaced by the name of the FEP EPROM on the machine: V127, G206, or G208.

Overlay File:	*Contains:*
*-info.flod	Commands that give information about the machine, for example, the Show Configuration command.
*-loaders.flod	Commands to load the machine, for example, the Load Microcode and Load World commands.
*-lisp.flod	Commands for manipulating Lisp, for example, the Start, Continue, and Show Status, Set LMFS FSPT Unit, and Show LMFS FSPT Unit commands. The last two commands allow you to specify a disk unit for the location of the LMFS's FSPT and to show you this location. This overlay file also contains FEP-based support for the UNIBUS option.
*-debug.flod	The FEP debugger, which is invoked by the Debug command.
*-tests.flod	The Test commands.
*-disk.flod	The Disk Restore and Disk Format commands.

The last two overlay files listed do not belong in the hello.boot file since they are used only during software installation or testing. Use the Scan command on the *-tests and *-disk overlays when necessary, by typing the following to the FEP prompt before issuing the contained commands:

```
Scan v127-disk.flod
Scan v127-tests.flod
```

"V127" is an example. On your system, use the version number of the FEP EPROM which your machine is using. This may not be V127. Other FEP EPROMS are G206 and G208.

Use the Copy Flod Files command to copy flod files to your FEP.

Copy Flod Files Command

Copy Flod Files *keywords*

Copies FEP Eprom overlay files to your machine.

:Disk Unit	{*integer*} Disk onto which flod files will be copied. The default is 0.
:Create Hello File	{Yes, No, Ask} Whether or not to create a Hello.boot file after copying. The default is Ask.
:From Directory	{*pathname*} Directory from which to copy files. The default is `sys:n-fep;` .
:Hosts	{*name*, All} Host(s) to which flod files will be copied. The default is your local FEP.
:Silent	{Yes, No} Display files as they are copied. The default is Yes.
:Version	FEP version. For example, G206 or V127.

8.2.1. FEP System `hello.boot` File

The `hello.boot` file is used to add new commands to the FEP. Each time the machine is reset, the commands that are included in the overlay files must again be made available. Using the FEP command, Hello, to load the `hello.boot` file makes these commands available.

In order to use the FEP system, you must create a hello.boot file in the editor, with a pathname of `FEPn:>hello.boot.`, FEP*n* refers to the disk unit number, in the case where a computer has more than one disk. This file normally contains a sequence of Scan commands, which scan the overlay files containing the standard commands and the Initialize Hardware Tables command.

Here is an example of the sequence of Scan commands to put in the `hello.boot` file:

```
Scan FEP0:>v127-info.flod
Scan FEP0:>V127-loaders.flod
Scan FEP0:>v127-lisp.flod
Scan FEP0:>v127-debug.flod
Initialize Hardware Tables
```

"V127" is an example. On your system, use the version number of the FEP EPROM which your machine is using. This may not be V127. Other FEP EPROMS are G206 and G208.

Make sure you press `RETURN` after `Initialize Hardware Tables`, and then save the file. For an explanation of the Scan commands:See the section "Using the FEP", page 80.

8.2.2. Lisp Utility for the FEP System

The following function writes the distributed overlay files to a cartridge tape in a format acceptable to the FEP's Scan command. Use this function from Lisp:

tape:write-fep-overlay-flods-to-cart *overlay-prefix* &optional (*tape-spec* "**local:cart**") &rest *private-flods* *Function*

overlay-prefix is a string such as "v127", indicating the FEP EPROM version. *tape-spec* defaults to "LOCAL:CART"; if you supply **nil**, *tape-spec* prompts for a tape specification, which must specify a cartridge tape.

Each item of &rest *private-flods* is a string, such as "LOADERS"; you can use these to specify individual overlay files that you want to write to tape. The file sys:n-fep;*overlay-prefix*-private-flod.lisp is among the additional flods written to tape. The pathname of each file is displayed as it is written to cartridge tape.

"V127" is an example. On your system, use the version number of the FEP EPROM which your machine is using. This may not be V127. Other FEP EPROMS are G206 and G208.

For information about what each overlay file contains: See the section "Using the FEP", page 80.

To use **tape:write-fep-overlay-flods-to-cart**, type the following form at a Lisp Listener. In this example, we want to copy FEP EPROM version 127 to tape.

```
(tape:write-fep-overlay-flods-to-cart  "V127")
```

This writes all the flod files for FEP Eprom version 127 from disk to tape.

"V127" is an example. On your system, use the version number of the FEP EPROM which your machine is using. This may not be V127. Other FEP EPROMS are G206 and G208.

We recommend that you use this function to make a backup tape containing the overlay files. Then, this tape will be of use to you if you ever have a disk without enough overlay files on it to boot Lisp. If you are in this situation, and have a tape of the overlay files, load the tape into a tape drive and type the following at the FEP prompt:

```
FEP Command: Mount cart:
FEP Command: Scan Cart:
FEP Command: Scan
.
.
.
```

Repeat the last command, Scan, until you get an "End of File". Now type boot to activate the boot file and boot Lisp, or type each command from a boot file manually if you do not have a boot file. If you type each command from a boot file manually, this is the command sequence:

```
Clear Machine
Load Microcode microcode-file-name
Load World distribution-world-file-name
Set Chaos-Address this-machine's-chaos-address
Start
```

Once you have booted Lisp, copy the overlay files from sys:n-fep; onto the FEP file system. Use the Copy Flod Files command to do this. For example:

```
Copy Flod Files SYS:n-fep;v127-*.flod.newest FEPn:>
```

"V127" is an example. On your system, use the version number of the FEP EPROM which your machine is using. This may not be V127. Other FEP EPROMS are G206 and G208.

See the section "Copy Flod Files Command", page 81.

8.2.3. Hints on using the FEP

The FEP command prompt is displayed when you are at FEP command processor level. It looks like this:

```
FEP Command:
```

The FEP command processor provides defaults and documentation where appropriate. When using it, remember these hints:

- You need type only enough of a FEP command to identify it uniquely, as shown below:

Input	*Completes to*
b RETURN	Boot
l w RETURN	Load World (default is FEP0:>Genera-7-2.load):
st RETURN	Start

- You can press the HELP key for a list of all FEP commands. For example:

  ```
  HELP
  ```

  Prints out:

```
Add ...more...
Attach ...more...
Boot  --  Execute an indirect command file
Clear ...more...
Compute ...more...
Continue  --  Continue running the machine after a halt
Copy ...more...
Debug  --  Enter the Fep Debugger
Declare ...more...
Detach ...more...
Disable ...more...
Dismount  --  Dismount a device
Enable  ...more...
Find ...more...
Fsm ...more...
Hello  --  Execute a "hello" file (default FEP:>Hello.boot) to initialize the FEP
Initialize ...more...
Load ...more...
Netboot -- Netboot a world via loading a netboot core world into 3600 memory
Mount  --  Mount a device
Reset ...more...
Scan  --  Scan a file (module) for commands
Set ...more...
Show ...more...
Shutdown  -- Halt the machine
Start -- Start the machine
```

Some of these commands are used in ordinary booting; others exist primarily to help system maintainers debug unusual problems.

- You can press the HELP key after typing a command name, for a list of all possible completions to that command. For example:

```
set <SPACE> HELP
```

prints:

```
Set Chaos-address  --  Set the chaos address
Set Default-disk-unit  --  Sets the default disk unit
Set Display-string  --  Set the NanoFep's display string
Set Ethernet-address  --  Set the Ethernet address
Set Lisp ...more...
Set LMFS ...more...
Set Monitor-type  --  Sets the monitor type to Philips or Moniterm
Set Prompt  --  Sets the top level command prompt
Set Wired ...more...
Set World-to-netboot  -- Tell the world just loaded to netboot a particular descriptior
```

Note that you must press SPACE after typing a command name and before pressing HELP to receive a list of the command's arguments.

- You can insert parenthetical comments in any white space within or after FEP commands. Such comments make useful documentation for boot files. For example:

```
load world >World1.load (contains geological survey programs)
set chaos-address 401 (Koala)
```

`load world >World1.load` and `set chaos-address 401` are FEP commands, and the parenthetical phrases are user-supplied comments.

Finally, be careful! Some commands are used primarily by system maintainers to debug unusual problems. Among these are Disk Restore and Disk Format. Be careful when giving these latter commands. If you make a mistake, you might destroy the state of the loaded or saved Lisp system.

8.3. FEP System Features

The following sections describe features of the FEP system.

8.3.1. Pathname Completion Is Supported

Commands support two kinds of pathname completion, one of which you activate by pressing the COMPLETE key and the other by pressing the HELP key.

If you press the COMPLETE key the FEP attempts to complete the pathname supplied so far. It replaces your input with as much of a completed pathname as it can without running into a conflict between two similar pathnames in the file system. For example, if you type

```
Load Microcode (default is FEP0:>3640-mic.mic) 3640 COMPLETE
```

the FEP system might respond with

```
Load Microcode (default is FEP0:>3640-mic.mic) FEP0:>3640-
```

if the possibilities are:

```
FEP0:>3640-mic.mic
FEP0:>3640-fpa-mic.mic
```

Similarly, if you type

```
Load World (default is ...) Inc HELP
```

the FEP system would show all the files that begin with Inc, such as the Incremental Worlds, created by Incremental Disk Saves.

8.3.2. Pathname Merging Is Supported

Pathnames given to FEP commands are merged against the default in much the same way as in Lisp. Therefore, you need to specify only those fields that are different from the default. If either the name or the type is given, the default version

is the newest version, not the version of the default. Pathnames are not case sensitive. Examples of pathname merging are:

Default	*Input*	*Merged*
FEP0:>3600-MIC.MIC	3600-FPA-MIC	FEP0:>3600-FPA-MIC.MIC
FEP0:>3600-MIC.MIC	fep1:	FEP1:>3600-MIC.MIC
FEP0:>3600-MIC.MIC	..389	FEP0:>3600-MIC.MIC.389
FEP0:>3600-MIC.MIC.389	3600-FPA-MIC	FEP0:>3600-FPA-MIC.MIC
FEP0:>3600-MIC.MIC	3600-mic..389	FEP0:>3600-MIC.MIC.389

8.3.3. Show Directory Command Understands Simple Wildcards

Examples of the use of simple wildcards are:

Specification	*Lists all*
*.load	world loads
*.boot	boot files
v127*.flod	.FLOD files for FEP version 127
sys.*	files with SYS as a substring of their name
*.mic.417	version 417 microcode files

The Show Directory command does not include either directory wildcards or handling a version of 0 (meaning newest) correctly. Show Directory does not support relative pathnames.

FEP pathnames have a four-character type-field limitation, thus, you cannot use .*lo*d to find all .LOAD and .FLOD files.

8.3.4. The FEP Determines Microcode Default From the Hardware Configuration

The initial default for the Load Microcode command is determined from the hardware configuration. Here are some of the hardware configurations:

Hardware Configuration	*Default*
3600, FPA	FEP:>3600-fpa-mic.mic
3600, FPA, XSQ	FEP:>3600-fpa-xsq-mic.mic
3640, FPA	FEP:>3640-fpa-mic.mic
3640, FPA, XSQ	FEP:>3640-fpa-xsq-mic.mic

This is only the initial default; using the Load Microcode command sets the default to the appropriate pathname if the command completes successfully. You can always reset the default to the initial (computed) default with the Compute Microcode Default command.

8.3.5. Show Directory Shows Detailed Information

The Show Directory command shows detailed information about each FEP file system directory entry in much the same way as the CP Show Directory command does. In addition to showing the pathname, it also shows the following information about each file:

- The number of blocks allocated
- The number of bytes and the byte size (or DIRECTORY if the file is really a directory)
- Flags (such as don't delete, deleted, don't reap)
- Creation time (in Greenwich Mean Time)
- File comment
- File author

8.4. Cold Booting

Cold booting completely resets Lisp. When you are finished using the computer, you can cold boot it to put it into a fresh state for the next user. Avoid cold booting a machine that someone else may be using, though, since the other person might be expecting the machine to remain in its current state.

You can cold boot a world off the local disk or you can netboot a world from a remote netboot server.

Here is the procedure for cold booting:

1. Go to a Lisp Listener by pressing SELECT L.
2. Log out, if you are logged in, by typing the Logout command.
3. Halt the machine by typing the Halt Machine command. (The function **(sys:halt)** can also be used.)

   If you cannot get a Lisp Listener window, or if no Lisp Listener is responding to keyboard input, press h-c-FUNCTION. However, the Halt Machine command is preferred over h-c-FUNCTION for stopping Lisp, because h-c-FUNCTION might interrupt disk I/O operations.

Pressing h-c-FUNCTION does not immediately stop Lisp. Instead, the FEP asks Lisp to stop itself cleanly. If Lisp does stop itself, the FEP prints the message "Lisp stopped itself." If Lisp does not stop itself after about three seconds, the FEP prints, "Waiting for Lisp to stop itself..." If after another three seconds Lisp does not stop itself, the FEP forcibly stops Lisp and prints, "Halting execution of Lisp." The purpose of this behavior is to reduce the chance of halting the computer during a disk write, which might cause ECC errors.

4. When control has returned to the FEP (you see the FEP Command: prompt) you issue the FEP Boot command to cold boot the machine from a boot file:

   FEP Command: Boot *file-name* RETURN

   where *file-name* is a boot file, with an extension of .boot. Its default value is the last file name given the Boot or Show File command. Its initial default value is Boot.boot on the current default disk unit.

   If you are booting a world from the local disk, the boot file must include a Load World command. If you are netbooting a world from a remote netboot server, the boot file must include a Netboot command.

   All FEP commands can also be entered by hand. You can cold boot a machine by typing in all the commands shown for boot files, in the order they are given.

   The following is a typical boot file:

   ```
   Clear Machine
   Load Microcode microcode-file-name
   Declare Paging-files paging-file-names
   Load World distribution-world-file-name
   or
   Netboot world-description
   Set Chaos-Address this-machine's-chaos-address
   Start
   ```

   Alternatively, if the microcode is already loaded and the Chaosnet address is set, you can type the FEP commands manually:

   ```
   Load World distribution-world-file-name
   Start
   ```

   or

   ```
   Netboot world-description
   Start
   ```

   For more information: See the section "Booting, Netbooting, and Autobooting", page 2.

Cold booting from the local disk takes approximately one minute. Netbooting takes a bit longer. It takes another minute or so for Lisp to start. After the system is cold booted, you will probably want to log in.

During this time, the machine might print a message asking you to enter date and time information, if it has no other way to find it. (This behavior is site-dependent.) If so, type it in the following format:

```
09/21/88  15:04
```

Be sure to enter the date and time correctly, as it is important that the file system know exactly when files are created and modified. If the calendar clock has been set, the machine uses the calendar clock reading as the default time for you to type in. If the calendar clock has not been set, the machine offers to set it to the time you specify.

8.5. Resetting the FEP

Resetting the FEP restarts the FEP system, thereby discarding knowledge of the FEP's free storage area. Resetting might be necessary if you unplug the console video cable from either end or turn the console off and on. You also need to reset the FEP if you receive the error message: `No More Memory`.

You can reset the FEP from either the keyboard or the processor front panel. Note that when the FEP is being reset the fault light (located on the front panel of the processor box) is turned on by the hardware. Then, when the FEP finishes initializing itself the FEP turns the fault light off.

- To reset the FEP from the keyboard:

  1. Type the Halt Machine command at a Command Processor command prompt to stop Lisp and give control of the keyboard to the FEP.

     If no Lisp Listener is responsive, press `h-c-FUNCTION` to stop Lisp.

  2. Type the command Reset FEP to the FEP prompt.

  3. Press `Y` to answer the confirmation prompt.

- To reset the FEP from the processor front panel:

  1. Push the red RESET button on the processor front panel.

  2. Press the spring-loaded YES switch to answer the "Reset FEP?" question (This question is asked only if you have a 3600 machine model).

  To reset the FEP from the processor front panel on a 3650:

  ° Turn the key switch momentarily to RESET.

After you reset the FEP, the keyboard is connected to the FEP, not to Lisp. Type the Hello command to the FEP prompt, and then give the Start command and press `RETURN` to warm boot the machine and Lisp, and return control of the keyboard to Lisp.

8.6. FEP Commands

Some FEP commands are involved in normal use of the computer. These include Boot, Show Directory, Show Version, and Start. Other commands are used primarily by system maintainers to debug unusual problems. Among these are Disk Restore and Disk Format. Be careful when giving these latter commands. If you make a mistake, you might destroy the state of the loaded or saved Lisp system.

If your system seems to lack a FEP command, you may need to scan a flod file to make the command available:See the section "Using the FEP", page 80. It is also the case that not all FEP commands are available (or needed) on all models of Symbolics Lisp machines.

8.6.1. Interactive FEP Commands

The FEP commands in this section are commonly typed directly to the FEP prompt. Any FEP command can be typed interactively, but the commands in this section are the most likely to be used in this way.

8.6.1.1. Boot

Boot *file-name*
: Executes the commands specified in *file-name*. *file-name* is the name of a boot file; it defaults to the last file name given the Boot or Show File command. Its initial default is `>Boot.boot`.

This command is available in all FEP EPROM versions.

8.6.1.2. Continue

Continue
: Continues the computer's operation from where it left off. However, if you have stopped the world and loaded new microcode, Continue does not work. Instead, you must warm boot by using the Start command. See the section "Start", page 93.

  This command is loaded when you scan the FEP overlay file `*-lisp.flod`.

Debug

Debug
: Enter the FEP debugger. See the section "Debugging in the FEP", page 136.
  This command is available in all FEP EPROM versions.

  This command is loaded when you scan the FEP overlay file `*-debug.flod`.

8.6.1.3. Hello

Hello
: Takes a pathname (defaulting to `FEPn:>hello.boot`) that normally contains a sequence of Scan commands and the

Initialize Hardware Tables command. The Scan commands scan the flod files containing the standard commands. Run the `hello.boot` file each time the FEP is reset or the machine is powered up by typing Hello to the FEP prompt. When all the commands in the `hello.boot` file are completed, you can boot your machine with the Boot command.

See the section "Using the FEP", page 80.

This command is available in all FEP EPROM versions.

8.6.1.4. Reset FEP

Reset Fep

Restarts the FEP program, discarding the FEP's knowledge of what microcode was loaded, and so on. After the FEP is reset, you must initialize the overlay files by typing Hello to the FEP command prompt. Then, boot the machine by typing Boot to the FEP command prompt. If the FEP is running in RAM, asks whether to switch back to PROM.

This command is available in all FEP EPROM versions.

8.6.1.5. Reset Video

Reset Video

Reloads the console screen's sync program.

This command is available in all FEP EPROM versions.

8.6.1.6. Show Directory

Show Directory *wildcard-directory-spec*

Displays the contents of a directory matching the wildcard in the FEP file system and the associated file comments. *directory-spec* must end in >* and can be preceded by fep: or fep*n*:. For example, Show Directory >* is acceptable, and shows the contents of the directory on the default disk. The default is FEP:>*.* or previous spec. For information about simple wildcards: See the section "Show Directory Command Understands Simple Wildcards", page 87.

Use Show Directory to check whether a file is in the FEP file system directory. For example, to see a directory listing of the contents of the files on FEP1, type the following to the FEP prompt:

```
FEP Command:  Show Directory  FEP1:>*.*
```

This command is available in all FEP EPROM versions.

8.6.1.7. Show File

Show File *file-name* Displays the contents of *file-name*, a file in the FEP file system. *file-name* defaults to the last file name given to the Show File or Boot command. Its initial default is >Boot.boot. For example, to look at the contents of a boot file, type:

```
Show File  FEP0:>boot.boot
```

This command is available in all FEP EPROM versions.

8.6.1.8. Shutdown

Shutdown Halts the FEP. To restart (and reset) it, push the RESET button on the processor front panel. The preferred way to turn off the machine is:

1. Halt Lisp, using the Halt Machine command.
2. Halt the FEP, using the Shutdown command.
3. Power off the processor and console.

On the 3600, the Shutdown command asks "Do you really want to halt the FEP?" Answer `Y` to confirm. It then displays the message "FEP Halted" in the nanofep display.

On all machine models other than the 3600, the Shutdown command asks the question "Do you really want to power down the 3600?" Answer `Y` to confirm. It then lights the fault light on the switch panel.

This command is available in all FEP EPROM versions.

8.6.1.9. Start

Start Starts the loaded Lisp world.

- If the world has just been loaded, this is a cold boot.
- If the world has been netbooted, this is a cold boot from a remote machine.

- If the world had been loaded previously, this is a warm boot.

The Start command checks for an acceptable network address; if none was set, it is read from the computer and you are asked to confirm it.

This command is loaded when you scan the FEP overlay file `*-lisp.flod`.

8.6.2. FEP Commands for Boot Files

The FEP commands in this section can be typed interactively, but they are most commonly included in boot files. There are three kinds of boot files, identified here by their default names:

- `boot.boot` — The common boot file. See the section "Booting a World", page 2.
- `hello.boot` — Boot file used to scan FEP overlay (flod) files. See the section "FEP System `hello.boot` File", page 82.
- `autoboot.boot` — Combines functions of `boot.boot` and `hello.boot` for autobooting.See the section "Autobooting", page 12.

8.6.2.1. Clear Machine

Clear Machine — Clears the internal state of the registers and memories.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

8.6.2.2. Clear Paging-files

Clear Paging-files — Clears the list of remaining pages that Lisp uses, and clears the FEP's list of paging files added by the Add Paging-file command. This command does not clear the list of declared paging files (declared by the Declare Paging-Files command or the Declare More Paging-files command). To clear the list of declared paging files, you must give the Declare Paging-files command without listing any files.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

8.6.2.3. Clear Screen

Clear Screen — Clears the console's screen.

This command is available in all FEP EPROM versions.

8.6.2.4. Declare More Paging-files

Declare More Paging-files *file-names*

Declares *file-names* to be the paging files for all subsequent Load World commands and adds the files to the list of previously declared paging files; it is the same as the Declare Paging-files command except that it does not clear previous declarations. For information about the Declare Paging-files command: See the section "Declare Paging-files", page 95.

This command is of use when you decide to declare additional paging files after you have already used the Declare Paging-files command. It is also useful when you want to declare many paging files, but cannot do so with a single Declare Paging-files command because doing so on a single line can overflow the FEP's line input buffer.

Here is a sample boot file with the Declare Paging-files command and the Declare More Paging-files command:

```
Clear Machine
Load Microcode FEP1:>3640-mic.mic.418
Declare Paging-files FEP0:>Page More-Page Even-More-Page
Declare More Paging-files FEP0:>Yet-More-Page Still-More-Page
Load World FEP1:>Release-7-2.load.1
Set Chaos 401
Start
```

You can also put the Declare More Paging-files command in the hello.boot file, after Initialize Hardware Tables. The Declare Paging-files command and the Declare More Paging-files command do not have to appear in the same boot file, either hello.boot or Boot.boot.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

8.6.2.5. Declare Paging-files

Declare Paging-files *file-names...*

Declares *file-names...* to be the paging files for all subsequent Load World commands until a new Declare Paging-files command overrides it. *file-names* is a list of files separated by **spaces, not commas**. The default pathname (directory and file extension) for the first file is always FEP0:>.page. The default for subsequent files is the previous pathname without the filename. For example, if you specify FEP1:>abc.xyz, the default for the next file is FEP1:>.xyz.

The Declare Paging-files command is the same as the Declare More Paging-files command, except that the latter command does not clear previous declarations. For information about the Declare More Paging-files command: See the section "Declare More Paging-files", page 95.

For example, this command:

```
Declare Paging-files fep0:>page aux fep1:>page2 page3
```

declares the files: FEP0:>page.page, FEP0:>aux.page, FEP1:>page2.page, and FEP1:>page3.page.

The command does not declare duplicates. For example:

```
Declare Paging-files fep0:>page aux page
```

does not declare fep0:>page.page a second time.

The command checks to see if each paging file actually exists. If a file does not exist, a warning is issued. The file is still added to the list of declared files in case the file is created some time in the future.

To undeclare all paging files, use Declare Paging-files without specifying any files.

If there are no declared paging files, the Load World command simply loads the paging file called FEP:>Page.page. If there are declared paging files, it adds them in the order in which they were declared. If there is a problem with a file, it warns you and goes on to the next file.

In .boot files, the Declare Paging-files command should be before the Load World or Netboot command. Here is a sample boot file with the Declare Paging-files command:

```
Clear Machine
Load Microcode FEP1:>3640-mic.mic.418
Declare Paging-files FEP0:>Page More-Page Even-More-Page
Load World FEP1:>release-7-2.load.1
Set Chaos 401
Start
```

You can also put the Declare Paging-files command in the hello.boot file, after Initialize Hardware Tables. The Declare Paging-files command and the Declare More Paging-files command do not have to appear in the same boot file, either hello.boot or Boot.boot.

This command is loaded when you scan the FEP overlay file *-loaders.flod.

8.6.2.6. Enable IDS

Enable IDS

Informs Lisp that the Incremental Disk Save facility should be enabled. This command must be issued after the Load World command and before the Start command. This command is pervasive; you do not need to execute it for a world that has been saved with IDS enabled.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

8.6.2.7. Initialize Hardware Tables

Initialize Hardware Tables

Initializes hardware tables inside the FEP. This command reads the ID proms of the boards in the machine. Before this command is executed for the first time, the FEP has no knowledge of the type or location of the memory boards in the machine. This command is not strictly necessary, because any command that requires that the hardware tables be initialized automatically executes it. However, as the initialization function might have to print diagnostic messages on the console, it is useful to run it from the Hello.boot file; that way,any problems are identified at a predictable time, rather than spontaneously in the middle of some other activity.

This command is loaded when you scan the FEP overlay files `*-loaders.flod` or `*-loaders.flod`.

8.6.2.8. Load Microcode

Load Microcode *file-name*

Loads microcode memory and other high-speed memories from the specified file. The default value of *file-name* is the last file name given to the Load Microcode command. Its initial default is `>Microcode1.mic`, which is determined by the hardware configuration: See the section "Compute Microcode Default FEP Command", page 102. Give the Clear Machine command before the Load Microcode command, if the computer was just powered on.

```
FEP Command:  Load Microcode FEP1:>3640-mic.mic.418
```

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

8.6.2.9. Load World

Load World *file-name*

Restores enough of the saved world in the computer so that

you can start up the machine. It prints both the desired microcode version for this world and the currently loaded microcode version. The default value of *file-name* is the last file name given to the Load World command. Its initial default is >Released-World.load.

```
FEP Command:  Load World FEP1:>genera-7-0.load
```

The Load World command can initiate netbooting in certain circumstances. See the section "Netbooting IDS Worlds", page 12.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

8.6.2.10. Mount

Mount *pathname*

Mounts the device (and unit) specified in the device field of the pathname.

To mount disk unit 2, use the following:

```
Mount FEP2:
```

Whenever a device is supplied in any new FEP system command (for example, CART: or FEP3:), the device is mounted automatically. The Lisp system expects the FEP to inform it about all usable disks at the time the FEP starts Lisp running; thus, this command is necessary only when a disk that Lisp needs to know about has not been mentioned in any other interaction with the FEP. See the section "Dismount", page 105.

This command is available in all FEP EPROM versions.

8.6.2.11. Netboot

Netboot *world-description*

Netboots the world described by *world-description*. Polls the netboot servers on the local subnet for worlds that match *world-description* and netboots the most recent of those worlds.

Here is an example:

```
Netboot chip-poker
```

In effect, the Netboot command replaces Load World if you are netbooting.

See the section "Netbooting", page 8.

See the section "World Description for Netbooting".

The Load World command can initiate netbooting in certain circumstances. See the section "Netbooting IDS Worlds", page 12.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

Scan

Scan *pathname*
: Adds (or updates) the commands in the specified file to the command tree. This command is normally included in `hello.boot` files. See the section "FEP System `hello.boot` File", page 82.
This command is available in all FEP EPROM versions.

8.6.2.12. Set Chaos-address

Set Chaos-address *octal-value*
: Sets the Chaosnet address. The default value of *octal-value* is the previous Chaosnet address. It is set to zero when the FEP is started.

The FEP checks for an acceptable Chaosnet address before starting Lisp. If none is specified as argument to this command, it warns you, asks whether the current setting is acceptable, and allows you to change it if necessary. Here is what you type to the FEP prompt:

```
FEP Command:  Set Chaos-Address 401
```

This command is loaded when you scan the FEP overlay file `*-lisp.flod`.

8.6.2.13. Set Default-disk-unit

Set Default-disk-unit *unit*
: Sets the default disk unit. *unit* becomes the default for most subsequent disk references. However, within a command file executed by a Boot command, the default disk unit is the one on which the command file is located. Here is what you type if you want the default disk unit to be FEP1.

```
FEP Command:  Set Default-disk-unit  1
```

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

8.6.2.14. Set Display-string

Set Display-string *string*

Displays the string in the nanofep display of the Symbolics 3600. The length of the string is limited to 12 characters, the number of characters in the nanofep display. If more characters are used, the string is truncated. This command can be used in a `.boot` file.

This command is loaded when you scan the FEP overlay file `*-lisp.flod`.

Set Ethernet Address

Set Ethernet Address

Sets the Ethernet address. Use this command if you are using DNA.

This command is loaded when you scan the FEP overlay file `*-lisp.flod`.

Set LMFS FSPT Unit

Set LMFS FSPT Unit

The Lisp Machine File System (LMFS) allows the use of multiple partitions residing on one or more disk drives. The selection of partitions to be used by LMFS is determined by a database called the *file system partition table* (FSPT). It is contained in a FEP file named `>fspt.fspt` on a boot drive.

If any machine at your site has more than one disk, it may be difficult to find the disk location of the FSPT. In order to make finding the location of a FSPT easy, insert the Set LMFS FSPT Unit command in your `Hello.boot` file. This command causes LMFS to look for the file named `>fspt.fspt` on the disk unit specified. For example, if you put your FSPT on disk unit 2, put the following in your `Hello.boot` file:

```
Set LMFS FSPT Unit 2
```

For more information about LMFS and the FSPT: See the section "Multiple Partitions", page 181.

This command is loaded when you scan the FEP overlay file `*-lisp.flod`.

8.6.2.15. Set Prompt

Set Prompt *string* Sets the FEP prompt to *string*.

For example:

```
FEP Command: Set Prompt "In the FEP Again? "
In the FEP Again?
```

This command is available in all FEP EPROM versions.

8.6.3. FEP Commands for System Maintainers

The FEP commands described in this section are generally of interest only to system maintainers. Most users do not need these commands.

8.6.3.1. Add World File

Add World File *pathname*

Explicitly adds the world specified in the pathname to the internal world database. For example, type the following to add an IDS world named Inc-Release-7-2-from-Rel-7-2.load, which resides on FEP0:

```
FEP Command:  Add World File FEP0:>Inc-Release-7-2-from-Rel-7-2.load
```

The FEP's internal world database keeps track of the relationships between IDS worlds on the local disk and also helps determine which world is the best choice to use as a netboot core. You cannot add a netboot world description to the internal world database.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

Clear Color Background Screen

Clear Color Background Screen

Clears the regular screen in a color console system. Since the FEP writes to the overlay screen, it may be easier to read the overlay screen after you clear the regular screen.

This command appears is available only on color console systems with FEP EPROM Versions V127 or G206.

Clear Command Tree

Clear Command Tree

Removes all commands defined by the overlay files,

This command is available in all FEP EPROM versions.

8.6.3.2. Clear World Files

Clear World Files Clears the internal world database. You can use this, for example, if the FEP has been running continuously for a long time and the database is cluttered with worlds that no longer exist.

The FEP's internal world database keeps track of the relationships between IDS worlds on the local disk and also helps determine which world is the best choice to use as a netboot core.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

Color

Color

If your system has both color and monochrome consoles, this command switches you to the color console.

This command is available only on color console systems with V127 or G206 FEP EPROMs.

You must either warm boot or cold boot after issuing this command.

8.6.3.3. Compute Microcode Default FEP Command

Compute Microcode Default

Computes the default microcode name from the hardware configuration and resets the Load Microcode command default to the computed name. Some examples:

Hardware Configuration	*Default*
`3600, FPA`	`FEP:>3600-fpa-mic.mic`
`3600, FPA, XSQ`	`FEP:>3600-fpa-xsq-mic.mic`
`3640, FPA`	`FEP:>3640-fpa-mic.mic`
`3640, FPA, XSQ`	`FEP:>3640-fpa-xsq-mic.mic`

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

For a complete list of microcode types: See the section "Genera 7.2 Microcode Types".

8.6.3.4. Detach Graphics Tablet

Detach Graphics Tablet

Allows a graphics tablet to be disconnected from a serial port. This command supports software distributed by the Symbolics Graphics Division.

This command is supported on the following models only: 3600, 3640, 3645, 3670, 3675.

To attach a graphics tablet see: See the section "Attach Graphics Tablet FEP Command".

This command is loaded when you scan the FEP overlay file `*-lisp.flod`.

8.6.3.5. Disable IDS

Disable IDS

Informs Lisp that the Incremental Disk Save facility should not be enabled. This must be issued after the Load World command and before the Start command. Disable IDS is pervasive. It is included for completeness; Symbolics does not at this time know of any situations that would warrant its use.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

8.6.3.6. Disable Load-to-paging Migration

Disable Load-to-Paging Migration

Disables the automatic migration of world load file pages to the paging file. When Load-to-Paging Migration is disabled, pages are written out to the paging file only if they are modified. Unmodified pages remain in the world load file.

This command should be executed after the Load World command and before the Start command.

The advantage of having Load-to-Paging Migration enabled is that it provides better paging performance. The advantage of having it disabled is that the roughly 80 percent of the world load that is never modified does not need to have paging space allocated for it, so the effective available paging space is increased by roughly 80 percent of the size of the world load. Load-to-Paging Migration is disabled by default.

The state of Load-to-Paging Migration is irrelevant to netbooting.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

8.6.3.7. Disk Restore

Disk Restore

Restores the file system or files from cartridge tapes to disk. **Note:** Before using Disk Restore, ensure that memory is clear and the appropriate microcode is in place, by giving the commands Clear Machine and Load Microcode Tape:. Note the trailing colon (:) after Tape.

This command is destructive in its operation and should not be used lightly. It is used primarily by system maintainers to debug unusual problems. Be careful when giving this command. If you make a mistake, you might destroy the state of the loaded or saved Lisp system.

Disk Restore displays two questions:

1. `Have you used Set Disk-type for all units that do not have valid label blocks?`

   The disk type must be known before the first reference to it, in case the label block is not yet written.

   A tape can contain information in either *image restore* (*block restore*) format or *file restore* format. In both cases, up to 1152 characters of information are displayed, describing the contents of the section of the tape. Usually the information is supplied by the person producing the tape.

   - *image restore*: Data recorded in this format can be either an initial file system or raw disk blocks from a source disk. The tape-generating program writes out block numbers normalized to block 0 (and writes the number of the original starting block into the header information) so that they can be written to the new disk at a different location if desired. File systems *must* be restored to block 0; the header information reminds you of this.

   - *file restore*: Data recorded in this format must be a file. The header includes the name of the source file, its length, and a comment supplied by the writer of the tape. You are asked for a destination pathname for the data; the default disk unit is assumed unless another is specified. Both a file system and the specified destination file must already exist on the unit, and the destination file must be large enough to hold the tape data. If the data pages of the destination file are not contiguous (because of bad blocks, say, or lack of contiguous space), the restored data is fragmented also.

     When the file restoration is completed, a special restoration block is read, containing the length of the file, the author, the creation date, and a comment.

     Large files (for example non incremental worlds) cross tape boundaries. But since all block numbers are relative to the beginning of the file, the second reel of tape is logically continuous with the first, and file restoration proceeds as for single-reel files.

This command is loaded when you scan the FEP overlay file `*-disk.flod`.

2. `Do you want to restore it?`

   You are prompted with Y (Yes), N (No), S (Skip microcodes), and F (Find microcode). If you answer "no", it skips the current tape restore section and searches for the next one. If you answer "skip microcodes" it stops when it reaches the world load, if you answer "find microcode", it asks what microcode to find. The default is the current microcode.

8.6.3.8. Dismount

Dismount *pathname*
Forcibly dismounts the device (and unit) that is the device field of the pathname.

See the section "Mount", page 98. This command is available in all FEP EPROM versions.

8.6.3.9. Enable IDS

Enable IDS Informs Lisp that the Incremental Disk Save facility should be enabled. This command must be issued after the Load World command and before the Start command. This command is pervasive; you do not need to execute it for a world that has been saved with IDS enabled.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

8.6.3.10. Enable Load-to-paging Migration

Enable Load-to-Paging Migration
Enables the migration of referenced pages from the world load file to the paging file. This copies each page that has been read from the world load file to the paging file. All future reads of that page then come from the paging file rather than from the world load file.

This command should be executed after the Load World command and before the Start command.

The advantage of having Load-to-Paging Migration enabled is that it provides better paging performance. The advantage of having it disabled is that the roughly 80 percent of the world load that is never modified does not need to have paging space allocated for it, so the effective available paging space is increased by roughly 80 percent of the size of the world load. Load-to-Paging Migration is disabled by default.

The state of Load-to-Paging Migration is irrelevant to netbooting.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

8.6.3.11. Find World Files

Find World Files *pathname*

Examines all files that are incremental worlds in the specified FEP directory. Files that appear to be valid incremental worlds cause the internal database of IDS files to be updated. As the IDS files are found, the names of the files and generations of the disk save are displayed.

The FEP's internal world database keeps track of the relationships between IDS worlds on the local disk and also helps determine which world is the best choice to use as a netboot core.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

Load Color Sync-program

Load Color Sync-program

Load a color sync program from disk. This command is used only with color consoles.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

8.6.3.12. Load Sync-program

Load Sync-program *file-name*

Loads the specified file (of type `.sync`) into the sync program memory of the I/O board and clears the screen. This is used for machines with monitors that require different sync programs than the one preprogrammed into the FEP. The default value of *file-name* is the last file name given to the Load Sync-program command. Its initial default is `>Sync.sync`.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

8.6.3.13. Reset Device

Reset Device *pathname*

Performs a device-dependent reset of the device (and unit) specified in the device field of the pathname.

```
FEP Command: Reset Device FEP1:>
```

This command is available in all FEP EPROM versions.

8.6.3.14. Reset Most

Reset Most — Resets the processor clock, the Lbus, the sequencer, the video, and the disks. If you think that the internal state of the computer is inconsistent, try Reset Most before power-cycling.

This command is loaded when you scan the FEP overlay file `*-lisp.flod`.

8.6.3.15. Reset Sequencer

Reset Sequencer — Resets the sequencer data paths.

This command is supported on the following models only: 3600, 3640, 3645, 3670, 3675.

Retension Cartridge-tape

Retension Cartridge-tape
Retensions the tape by rewinding the cartridge tape.

For the 3610, 3620, 3630, 3650, and 3653, this command is part of the FEP PROM.

For the 3600, 3640, 3645, 3670, and 3675 it is loaded from the FEP overlay file, `*-disk.flod`.

Set Color Monitor-type

Set Color Monitor-type *monitor-type*
Identifies the color monitor to the FEP.

This command is available only on color console systems with FEP EPROMs V127 or G206.

Set Console

Set Console {color | monochrome}
Sets the console to either color or monochrome.

This command is available only on color console systems with FEP EPROM G208.

See the section "Monochrome FEP Command".See the section "Color", page 102.
You must either warm boot or cold boot after issuing this command. You must also reload the correct microcode.

Set Monitor Type

Set Monitor Type {color | monochrome}
{Philips | Moniterm | Amtron | Tektronix | Sony}

Identifies the console type.

This command is available only on systems with FEP EPROM G208.

This command replaces Set Monitor-type and Set Color Monitor-type for these systems.

8.6.3.16. Set Monitor-type

Set Monitor-type *monitor-type*

Specifies the monitor type. The Set Monitor-type command ensures that the sync program used is for the monitor type requested. *monitor-type* can be either **Moniterm** or **Philips**; the types can be abbreviated to their first letter, **m** for **Moniterm** and **p** for **Philips.**

Set Monitor-type is used if the monitor is changed at a site and the ID prom is not changed accordingly. This command can be used in a boot file.

The following examples show two valid uses of the same command.

```
Set Monitor-type Moniterm

set mon m
```

This command is available only on systems with V127 or G206 FEP EPROMs.

8.6.3.17. Set World-to-netboot

Set World-to-netboot *world-description*

Replacement for the Netboot command in certain rare circumstances. Usually, the Netboot command selects the "best" world to use as a netboot core. If the "best" core does not, for some reason, seem to work, you can use Load World to load another world to use as a netboot core and then use this command to netboot using that world rather than selecting the "best" core. In other words, this command uses the currently loaded world to netboot, rather than searching for a "better" one. See the section "World Description for Netbooting".

This command is loaded when you scan the FEP overlay file *-loaders.flod.

Show Command Modules

Show Command Modules

Displays the currently active overlay files, whether or not they are currently loaded, the commands in the overlay files, and the command descriptions. For example, if the overlay file FEP0:>v127-loaders.flod is not currently loaded in your environment, the display looks like this:

```
Pathname FEP0:>v127-loaders.flod, not loaded
```

"V127" is an example. On your system, use the version number of the FEP EPROM which your machine is using. This may not be V127. Other FEP EPROMS are G206 and G208.

This command is loaded when you scan the FEP overlay file *-info.flod.

Show Command Tree

Show Command Tree

Displays the current command tree. Show Command Tree shows whether a command came from an overlay file or a module. It also shows the address within FEP memory at which the command resides (or would reside if the overlay file were loaded).

This command is loaded when you scan the FEP overlay file *-info.flod.

8.6.3.18. Show Configuration

Show Configuration Displays the hardware configuration, scans the backplane, and describes the boards on the bus.

This command is loaded when you scan the FEP overlay file *-info.flod.

8.6.3.19. Show Disk Label

Show Disk Label *unit*

Displays the label of *unit*, the specified disk unit. This is done independently of the unit's being mounted, so you can tell

what the label contains. The default for *unit* is the current default disk unit. For example, to see the disk label of FEP1, type:

```
FEP Command:  Show Disk Label 1
```

This command is available in all FEP EPROM versions.

Show Ethernet Address

Show Ethernet Address

Displays the Ethernet address, if set. Use this command if you are using DNA. See the section "Set Ethernet Address", page 100.

This command is loaded when you scan the FEP overlay file `*-info.flod`.

Show LMFS FSPT Unit

Show LMFS FSPT Unit

Show the current location of the *file system partition table* (FSPT). It is contained in a FEP file called named `>fspt.fspt`.

For more information on the FSPT: See the section "Set LMFS FSPT Unit", page 100.

Show Serial

Show Serial

Show status of a serial unit.

This command is loaded when you scan the FEP overlay file `*-lisp.flod`.

This command is available on the following hardware models: 3600, 3640, 3645, 3670, 3675.

8.6.3.20. Show Status

Show Status

Displays the internal status of some machine registers.

For information on interpreting the output of this command: See the section "Finding Out Why Your Machine Crashed", page 127.

See the section "FEP Show Status Command Output", page 128.

The CP command equivalent is called Show Crash Data. See the section "Show Crash Data Command", page 132.

This command is loaded when you scan the FEP overlay file `*-lisp.flod`.

8.6.3.21. Show Version

Show Version — Displays the version number of loaded FEP software.

This command is available in all FEP EPROM versions.

8.6.3.22. Show World Files

Show World Files — Shows the internal world database in an understandable format. It starts by showing each file that does not have a parent. This usually means full disk saved worlds, but can also mean incremental worlds whose parents have been deleted, or netboot cores. It then displays the incrementally saved worlds in depth-first fashion. Each description line shows the generation number, the file, the timestamps (IDs) of the file, and the timestamps of the parent (parent IDs). The display also shows which files can be used as netboot cores. The FEP's internal world database keeps track of the relationships between IDS worlds on the local disk and also helps determine which world is the best choice to use as a netboot core.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

8.6.4. FEP Commands for System Internals Maintenance

The FEP commands in this section require a serious knowledge of Lisp machine internals for successful use. Many of these commands are destructive in their execution and should be used with caution.

8.6.4.1. Copy File FEP Command

Copy File *from-pathname to-pathname*
Copies a file from disk to tape on that same machine.

This command is available in all FEP EPROM versions.

8.6.4.2. Load FEP

Load Fep *file-name* — Loads and starts loadable FEP programs. The names of the FEP programs are usually of the form V127-*name*, where V127 is the number of the FEP version on which the program runs and *name* is the name of the program.

"V127" is an example. On your system, use the version number of the FEP EPROM which your machine is using. This may not be V127. Other FEP EPROMS are G206 and G208.

This command is available in all FEP EPROM versions.

Test A-memory

Test A-memory — Tests all locations in A memory. This takes a couple of minutes.

This command is loaded when you scan the FEP overlay file `*-tests.flod`.

Test All

Test All — Runs all FEP tests. See the section "Test Main Memory", page 112.See the section "Test Simple Main Memory", page 112.See the section "Test A-memory", page 112.See the section "Test Location", page 112.See the section "Test Disks", page 112.

This command is loaded when you scan the FEP overlay file `*-tests.flod`.

Test Location

Test Location — Tests a single location in main memory.

This command is loaded when you scan the FEP overlay file `*-tests.flod`.

Test Disks

Test Disks — Tests the disks on the system.

This command is loaded when you scan the FEP overlay file `*-tests.flod`.

Test Simple Main Memory

Test Simple Main Memory
Runs a version of Test Main Memory with less functionality that takes less time to run.

See the section "Test Main Memory", page 112.
This command is loaded when you scan the FEP overlay file `*-tests.flod`.

Test Main Memory

Test Main Memory — Tests all locations in main memory. This takes a long time.

This command is loaded when you scan the FEP overlay file `*-tests.flod`.

Enable Trap Handling

Enable Trap Handling

Enable trap handling in case it got turned off. This command is mainly for use by system developers.

This command is loaded when you scan the FEP overlay file `*-lisp.flod`.

Set Lisp Release

Set Lisp Release — Set the Lisp release version intended.

This command is loaded when you scan the FEP overlay file `*-lisp.flod`.

8.6.4.3. Set Wired Addresses

Set Wired Addresses *%wired-virtual-address-high*

Sets values for wired addresses. This solves the following problem. If there are local or Symbolics-distributed patches to the wired system, and if these patches cause an internal limit to be exceeded, an error is signalled stating that the variable **sys:%wired-virtual-address-high** needs to be increased and gives a suggested new value. This command makes it easy to set the necessary variables.

Note: This command must be executed after the Load World command and before the Start command.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

Load Complete World

Load Complete World

Load all of a world into memory. This command is mainly for the use of system developers.

This command is loaded when you scan the FEP overlay file `*-loaders.flod`.

8.6.4.4. Show Disk Types

Show Disk Types — Lists the names of possible disk types along with information about each disk's cylinder, head, sector, and format gap sizes. Before you can issue this FEP command, you must first use the FEP command Scan V*xxx*-disk.

This command is loaded when you scan the FEP overlay file `*-disk.flod`.

8.6.4.5. Set Disk Type

Set Disk Type *unit type pack-id*

Tells the FEP that disk *unit* is of type *type* and has pack id *pack-id*. Disk Restore might need this information if the disk has no label block or if the label block contains incorrect information. Give Set Disk-type after an implicit or explicit Mount command.

This command is loaded when you scan the FEP overlay file `*-disk.flod`.

8.6.4.6. Disk Format

Disk Format

Formats the disk. This command overwrites all data on the disk. **This command is destructive in its operation and should not be used lightly.** If you make a mistake, you might destroy the state of the loaded or saved Lisp system. This command is primarily for the use of system maintainers in debugging unusual problems.

When you give the command, the FEP asks several questions; it expects answers in the following form:

Questions	*Valid answers*
Of type	M2284, T306, M2284, M2294 M2312, M2351A, XT1140, XT1150 XT2190, D2257, P807
On unit	Disk unit number
With pack id	0
From cylinder	Cylinder number; inclusive lower bound
Through cylinder	Cylinder number; inclusive upper bound

The answers to these questions, such as pack id, can be obtained by using the FEP command Show Disk Label.

This command is loaded when you scan the FEP overlay file `*-disk.flod`.

8.6.4.7. Add Disk-type

Add Disk-type

Lets you declare an arbitrary disk type to the FEP. Use this command if you need to format and restore a disk which is not yet supported by Symbolics. You can declare up to four disk types before you have to give the Clear Disk-types command. Add Disk-Type is needed only to format and restore disks. It is not needed for normal operation of any validly formatted disk with a FEP file system.

Add Disk-type has the following arguments, for which it prompts with the argument names in parentheses:

name	The textual name by which this disk type is known
cylinders	The number of cylinders supported by the drive
heads	The number of heads on the drive
sectors	The number of sectors
gap1	The length of "gap1"
gap2	The length of "gap2"
gap3	The length of "gap3"
fast	0 for slower disks, 1 for faster disks

These numbers require careful computation and involve some restrictions of the computer hardware. The calculation should be done by Symbolics Customer Service.

This command is loaded when you scan the FEP overlay file `*-disk.flod`.

8.6.4.8. Clear Disk-types

Clear Disk-types Clears all disk types declared with the Add Disk-type command.

This command is loaded when you scan the FEP overlay file `*-disk.flod`.

Clear Disk-counters

Clear Disk-counters Clears the registers that keep track of what has been happening on the disk. There is no Show Disk-counters command.

This command is loaded when you scan the FEP overlay file `*-disk.flod`.

8.7. FEP File System

The Symbolics computer disk has a file system called the *FEP file system*. The entire disk is divided up into *FEP files* (that is, files of the FEP file system). FEP files have names syntactically similar to those of files in the Symbolics computer's own local file system. However, the FEP file system and the Lisp Machine File System (LMFS) are completely distinct.

The *FEP file system* manages the disk space available on a disk pack, grouping sets of data into named structures called *FEP files*. All the available space on a disk pack is described by the FEP file system. A single FEP file system cannot extend beyond a single disk pack; each disk pack has its own separate FEP file system.

The FEP file system supports all of the generic file system operations. It also supports multiple file versions, soft deletion and expunging, and hierarchical directories.

Although "FEP" is an acronym for *front-end processor*, the FEP file system is managed by the main Lisp processor. It is called the FEP file system because the FEP can read files stored in the FEP file system. For example, the FEP uses the FEP file system for booting the machine and running diagnostics.

Disk streams access FEP files. A disk stream is an I/O stream that performs input and output operations on the disk. (For information about streams: See the section "Types of Streams" in *Reference Guide to Streams, Files, and I/O*. See the section "Stream Operations" in *Reference Guide to Streams, Files, and I/O*. When disk streams are opened with a **:direction** keyword of **:input** or **:output**, the disk stream reads or writes bytes, respectively, buffering the data internally as required. When the **:direction** is **:block**, the disk stream can both read and write the specified disk blocks. Block mode disk streams address blocks with a block number relative to the beginning of the file, starting at file block number zero. This *file block number* is internally translated into the corresponding disk address. The checkwords of all disk blocks contained in the FEP file system are reserved for use by the FEP file system, so block mode transfers should not use the checkwords stored in the disk array. See the section "3600-Family Disk System Definitions and Constants" in *Internals, Processes, and Storage Management*.

The FEP file system is also used by the system for allocating system overhead files, such as the paging file. See the section "FEP File Types", page 117. This section lists some of these files and what they are used for.

The need to allow the FEP to access FEP files, and also to allow the system to use them imposes some constraints on the design of the FEP file system. The internal data structures of the file system must be simple enough to permit the FEP to read them, and a small amount of concurrent access by both the FEP and Lisp must be tolerated. A FEP file's data blocks should have a high degree of locality on the disk to minimize access times. And the FEP file system must be very reliable, since the FEP needs to use the file system for running diagnostics and for booting the machine.

Note: Because of these constraints, the FEP file system is not intended to be a replacement for LMFS. (See the section "Lisp Machine File System" in *Reference Guide to Streams, Files, and I/O*.) Allocating new blocks for FEP files is slow, so creating many files, especially many small files, might impair the performance of the FEP file system, and ultimately the virtual memory system, if paging files or world load files become highly fragmented.

8.7.1. Naming of FEP Files

The FEP filename format is similar to the LMFS filename format. See the section "Lisp Machine File System" in *Reference Guide to Streams, Files, and I/O*. There are differences, however. Here are the format details of a FEP filename:

host The name of the FEP file system host. The format for a FEP host is *host*|FEP*disk-unit*, where the *host* field specifies which machine's FEP file system you are referring to, and *disk-unit* specifies the disk unit number on the machine. The *host* field defaults to the local machine if you omit it and the terminating vertical bar (|). If you omit both the *host* and *disk-unit* fields, the FEP host defaults to the disk unit the world was booted from on the local machine. For example:

`Merrimack|FEP0`	The FEP file system on Merrimack's unit 0.
`FEP2`	The FEP file system on the local machine's unit 2.
`FEP`	The FEP file system the booted world load file resides on.

directory The name of the directory. The FEP file system supports hierarchical directories in the same format as in LMFS. Each directory name is limited to a maximum of 32 characters; there is no limit on the total length of a hierarchical directory specification.

name The name of the FEP file, which cannot exceed 32 characters.

type The type of the FEP file, which cannot exceed 4 characters.

version The version number of the FEP file, which must be a positive integer or the word "newest".

FEP files can be renamed. For example, if you save a world containing MACSYMA, you might want to rename the world file to >macsyma.load or >macsyma1.load. Be sure to update your boot file if you intend this to be the default world.

8.7.2. FEP File Types

By convention, the following file types are used by the FEP file system for files used by that system.

boot The file contains FEP commands that can be read by the FEP's Boot command. *boot* files are text files, and can be manipulated by the editor. See the section "Configuration files", page 119.

load The file contains a world load image, or *band*, that is used to boot the system.

mic The file contains a microcode image, plus the contents of other internal high-speed memories that are initialized when the computer is booted. For example, >3640-mic.mic.389 contains version 389 of the microcode for the 3640 and 3670.

fspt — The file contains a LMFS partition table. It tells LMFS which FEP files to use for file space. For example, >fspt.fspt.newest is the default partition table used by LMFS.

file — The file contains a LMFS partition which holds the machine's local file system. The entire Symbolics computer local file system normally resides inside one big file of the FEP file system. For example, >lmfs.file.newest is the default LMFS file partition.

page — The file contains disk space that can be used by the virtual memory system. To increase the effective size of virtual memory, you can add additional paging files. See the section "Allocating Extra Paging Space", page 34. For example, >page.page.newest is the default file used by the virtual memory system as storage for swapping pages in and out of main memory.

flod — The file contains a FEP Load file. FEP Load files contain binary code the FEP can load and execute.

fep — The file contains binary information used by the FEP file system. These files should not be written to by user programs. Some examples of these files are:

>root-directory.dir — This is the root directory for the FEP file system.

>free-pages.fep — Describes which blocks on the disk are allocated to existing files.

>bad-blocks.fep — Owns all the blocks that contain a media defect and should not be used.

>sequence-number.fep — Contains the highest sequence number in use. The FEP file system uses sequence numbers internally to uniquely identify files. This is to assist in rebuilding the file system in case of a catastrophic disk failure.

>disk-label.fep — Contains the disk pack's physical disk label. The label is used to identify the pack and describe its characteristics.

dir — The file contains a FEP directory. For example, fep0:>root-directory.dir.newest contains the top-level root directory. The directory file for fep0:>dang>examples> would reside in fep0:>dang>examples.dir.1.

8.7.3. Configuration files

Configuration files contain FEP commands tailored for a particular Symbolics computer configuration. The commands are executed if you specify the file as argument to a Boot command when cold booting the machine. See the section "FEP Commands", page 91.

The configuration file >Boot.boot usually contains FEP commands to:

- Clear the internal state of the machine
- Load the microcode
- Load a world
- Set the Chaosnet address
- Start the machine

To change the selection of microcode and world loads that are booted by default, simply use Zmacs to edit the file FEP*n*:>Boot.boot, where FEP*n* is the disk unit. Be careful to avoid typographical errors; otherwise, you might have to type in the commands manually in order to boot the machine. Also, be sure that the last command in the file is followed by RETURN.

8.7.4. FEP file comment properties

Comment properties supply additional information about the contents of FEP files. They are listed inside square brackets, where the reference or expunge date appears for other file systems. You can list the contents of the FEP file system by using the Show FEP Directory command. The Zmacs command Dired (m-X) of fep*n*:>*, or the form (dired "fep:*n*>*") (where *n* is the disk unit) invokes the directory editor on the FEP file system. An example of the Show FEP Directory command output is shown in figure 5.

8.7.5. Accessing FEP Files

FEP files are accessed by open disk streams. A disk stream is opened by the **open** function. (See the section "Accessing Files" in *Reference Guide to Streams, Files, and I/O*. That section contains more details on accessing files.) If a FEP file system residing on a remote host is referred to, a *remote stream* is returned with limited operations, as specified by the remote file protocol.

In addition to the normal **open** options, the following keywords are recognized:

:if-locked This keyword specifies the action to be taken if the specified file is locked. This keyword is not supported by the remote file protocol.

- **:error** Signal an error. This is the default.
- **:share** Open the specified file even if it is already locked, incrementing the file's lock count. This mode permits multiple processes to

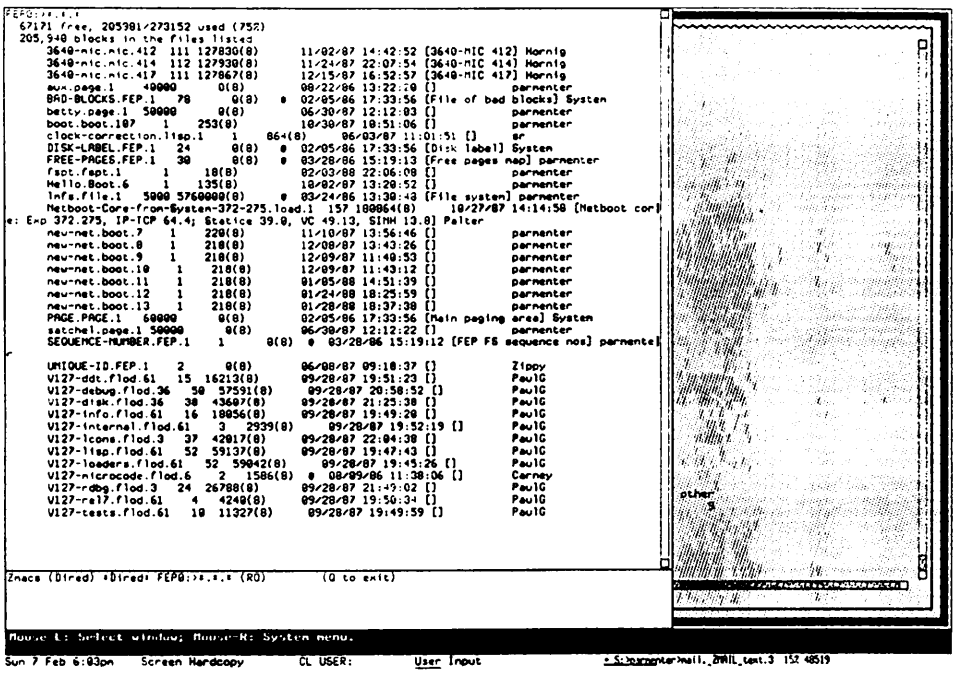

Figure 5. FEP File Comment Properties

write to the same file simultaneously. (See the section "FEP File Locks", page 124. That section contains more information on file locks.)

:number-of-disk-blocks

The value of this keyword is the number of disk blocks to buffer internally if the **:direction** keyword is **:input** or **:output**. This keyword is ignored for other values of **:direction** or for files on remote hosts. The default **:number-of-disk-blocks** is two.

8.7.6. Operating on Disk Streams

All disk streams to a local FEP file system handle the following messages:

:grow &optional *n-blocks* &key **:map-area :zero-p** *Message*

This message allocates *n-blocks* of free disk blocks and appends them to the FEP file. The value of *n-blocks* defaults to one. If **:zero-p** is true the new blocks are filled with zeros; otherwise, they are not modified. The return value of **:grow** is the file's data map (the format of the data map is described in **:create-data-map**'s description below). The value of **:map-area** is the area to allocate the data map in, which defaults to **default-cons-area**.

:allocate *n-blocks* &key **:map-area :zero-p** *Message*

This message ensures that the FEP file is at least *n-blocks* long, allocating additional free blocks as required. Returns the file's data map (the format of the data map is described in **:create-data-map**'s description below). **:map-area** specifies the area to create the data map in, and defaults to **default-cons-area**. The newly allocated blocks are filled with zeros if **:zero-p** is true. **:zero-p** defaults to **nil**.

:file-access-path *Message*

This message returns the disk stream's file access path.

For example, you can find out what unit number a FEP file resides on as follows:

```
(send (send stream :file-access-path) :unit)
```

:map-block-no *block-number grow-p* *Message*

This message translates the relative file *block-number* into a disk address, and returns two values: the first value is the disk address, and the second is the total number of disk blocks, starting with *block-number*, that are in consecutive disk addresses. *grow-p* specifies whether the file should be extended if *block-number* addresses a block that does not exist. When *grow-p* is true, free disk blocks are allocated and appended to the FEP file to extend it to include *block-number*. Otherwise, if *grow-p* is false, **nil** is returned if *block-number* addresses a block that does not exist.

:create-data-map &optional *area* *Message*

This message returns a copy of the FEP file's data map allocated in area *area*, which defaults to **default-cons-area**. A FEP file data map is a one-dimensional **art-q** array. Each entry in the file data map describes a number of contiguous disk blocks, and requires two array elements. The first element is the number of disk blocks described by the entry. The second element is the disk address for the first block described by the entry. The array's fill-pointer contains the number of active elements in the data map times two.

:write-data-map *new-data-map disk-event* *Message*

This message replaces the file's data map with *new-data-map*. *disk-event* is the disk event to associate with the disk writes when the disk copy of the file's data map is updated. This message overwrites the file's contents and should be used with caution.

8.7.7. Input and Output Disk Streams

Input and output disk streams are buffered streams. In addition to the standard buffered stream messages, local input and output disk streams also support the messages described elsewhere: See the section "Operating on Disk Streams", page 120.

Input disk streams read bytes of data starting at the current byte position in the FEP file, updating the byte position as the data is read. Output disk streams write bytes of data in the same way.

The bytes of data are stored in buffers internal to the stream. The **:number-of-disk-blocks open** keyword controls how many disk blocks the internal buffers can hold. When the current pointer moves beyond a disk block boundary, the buffered disk block is written to the file for an output stream, or the next unbuffered block is read in from the file for an input stream. Output streams also write out all the buffered disk blocks when the stream is sent a **:close** message without an **:abort** option.

8.7.8. Block Disk Streams

Block disk streams can both read and write disk blocks at specified file block numbers. A file block number is the relative block offset into the file. The first block in the file is at file block number zero, the second is at file block number one, and so on.

Block disk streams do not buffer any blocks internally and are not supported by the remote file protocol.

See the section "Operating on Disk Streams", page 120. In addition to the messages described in that section, block disk streams support the following messages:

:block-length *Message*

The **:block-length** message returns the length of the FEP file in disk blocks.

:block-in *block-number n-blocks disk-arrays* &key **:hang-p :disk-event** *Message*

The **:block-in** message causes the disk to start reading data from the disk into the disk arrays in *disk-arrays*, starting with the file block number *block-number*, and continuing for *n-blocks*. *disk-arrays* can be a disk array or a list of disk arrays. The value of *n-blocks* is the number of disk blocks to read. When *n-blocks* is greater than one, each disk array is completely filled before using the next disk array in *disk-arrays*. The checkwords stored in the disk arrays are reserved for use by the FEP file system. See the section "3600-Family Disk System Definitions and Constants" in *Internals, Processes, and Storage Management*. Unused disk arrays or portions of disk arrays remain unmodified.

When the value of **:hang-p** is true, which it is by default, **:block-in** waits for all the reads to complete before returning. If the value of **:hang-p** is false, **:block-in** returns immediately upon enqueuing the disk reads without

waiting for completion. In this case, all *disk-arrays* and the *disk-event* must be wired before sending the **:block-in** message, and must remain wired until the disk reads complete.

If the **:disk-event** keyword is supplied, its value is the disk event to associate with the disk reads. Otherwise the **:block-in** message allocates a disk event for its duration. A **:disk-event** must be supplied when **:hang-p** is false.

:block-out *block-number n-blocks disk-arrays* &key **:hang-p** **:disk-event** *Message*

The **:block-out** message causes the disk to start writing the data in the disk arrays in *disk-arrays* onto the disk, starting with the file block number *block-number*, and continuing for *n-blocks*. The arguments to the **:block-out** message are identical to those of the **:block-in** message.

8.7.9. FEP File Properties

In addition to having a name and containing data, FEP files also have properties. These properties store information about the file itself, such as when it was last written and whether it can be deleted or not. File properties are read by the **fs:file-properties** function, and modified by the **fs:change-file-properties** function. The **fs:directory-list** function also returns the file properties of several files at once. (See the section "Accessing Directories" in *Reference Guide to Streams, Files, and I/O.*)

The following file properties can be both read and modified:

:creation-date The universal time the file was last written to. Universal times are integers. (See the section "Dates and Times" in *Programming the User Interface -- Concepts.*)

:author The user-id of the last writer. The user-id must be a string.

:length-in-bytes The length of the file, expressed as an integer.

:deleted When **t** the file is marked as being deleted. A deleted file can then be marked as being undeleted by changing this property to **nil**. The disk space used by a deleted file is not actually reclaimed until the file is expunged.

:dont-delete When **t**, attempting to delete or overwrite the file signals an error. **nil** indicates the file can be deleted or written to.

:comment A comment to be displayed in brackets in the directory listing. The comment must be a string.

The following file properties are returned by the **:properties** message, but cannot be modified by **:change-properties**:

:byte-size The number of bits in a byte. The value of this property is always 8.

:length-in-blocks The block length of the file expressed as an integer.

:directory If **t**, the file is a directory; otherwise **nil**.

8.7.10. FEP File Locks

A FEP file is *locked* for the interval from when it is opened for reading or writing until it is closed. If the **:direction** keyword is **:input**, the file is *read-locked*; if the **:direction** keyword is **:output** or **:block**, the file is *write-locked*.

When the **:if-locked** keyword is **:error**, which is its default, a file that is read-locked can still be opened for reading but signals an error if opened for writing; a file that is write-locked cannot be opened for reading or writing. This permits multiple readers to access a file concurrently, while prohibiting writing to the file being read.

When the **:if-locked** keyword is **:share** in an open call for write, it succeeds in opening the file even if it is already read- or write-locked.

An expunge operation on a file that is either read- or write-locked does not expunge the file. If expunging a directory fails to expunge a file, the file must be closed and the directory expunged again.

8.7.11. Installing microcode

Use the Copy Microcode command to retrieve any new microcode from the file system of the sys host.

Copy Microcode Command

Copy Microcode *{version or pathname} destination keywords*

Installs a version of microcode.

version or pathname
Microcode version number or pathname to copy. *version* is a microcode version number (in decimal). *pathname* rarely needs to be supplied. It defaults to a file on FEP*n*:> (where *n* is unit number of the boot disk) whose name is based on the microcode name and version. (The file resides in the logical directory sys:l-ucode;.) The *version* actually stands for the file *appropriate-hardware*-MIC.MIC.*version* on FEP*n*:>. (See the Section "Genera 7.2 Microcode Types" in *Software Installation Guide*)

destination FEP file specification. The pathname on your FEP*n*:> directory. The default is created from the microcode version.

keywords :update boot file

:update boot file {*FEP-file-spec*, None, Query}. The pathname of the boot file you want it to update. The default is the current default boot file name.

Initially, the Symbolics personnel who install your system establish these microcode files for you.

8.7.12. Using a Spare World Load for Paging

You can reuse FEP file space for paging files. You may have a spare world load file, which you can transform into a paging file. For example, once you have successfully installed a new software release, you can rename the old world load to be a paging file. **Note:** Do not use the world load you are currently running for a paging file, as this action overwrites the previous contents of the specified file.

If you are considering this step, you may want to take a look at netbooting. See the section "Netbooting", page 8. If your old world load is Release-7-1.load, is resident on FEP0:>, and is, say, 35,000 blocks in size, and you want to create a new paging file called FEP0:>page2.page (also with a block size of 35,000), follow these steps:

1. Rename the file FEP0:>release-7-1.load to FEP0:>page2.page using the Rename File command. For example, type:

   ```
   Rename File FEP0:>release-7-1.load  FEP0:>page2.page
   ```

   Now the world load has been renamed to a paging file.

2. Use the Add Paging File command to initialize the paging file from the Lisp environment.

3. Edit your FEP*n*:>Boot.boot file to declare the new paging file. Use the Declare Paging-files command in your boot file to do this. For information about the Declare Paging-file command: See the section "FEP System Commands: General Usage".

You can also create new FEP files and use them for extra paging space:See the section "Allocating Extra Paging Space", page 34.

8.7.13. Adding a spare world load as LMFS file space

Partitions can be added to LMFS by following these steps:

1. Create the partition you wish to add to LMFS prior to entering the File system editing operations program. In addition, when you add a new partition or a partition on another disk, the disk should be free of errors and properly initialized and formatted.

2. Press SELECT F to select the File system editing operations program.

3. Click on [Local LMFS Operations] to invoke the second level of the File System Maintenance Program.

4. Click on [LMFS Maintenance Operations] to invoke the third level menu, which is a menu of file-system maintenance operations.

5. Click right on [Initialize] to invoke a menu of initialization options, which offers [New File System] and [Auxiliary Partition] as choices. Clicking on [New File System] is similar to clicking left on [Initialize]; it initializes a partition to be the basis of a file system.

6. Click on [Auxiliary Partition] to add another partition.

7. Enter the pathname of the FEP file to be used as the new partition. The default presented, which is correct for [New File System], is never correct for adding a partition.

8. Click on [Do It]. The system then performs much verification and error checking, roughly as much as when initializing a new partition. It must not be interrupted while performing these actions.

9. When finished, the File system editing operations program adds the partition and edits the FSPT automatically.

8.8. Disk Handling

You can include a disk specification of the form FEP*n*: (where *n* refers to disk unit *n*) as the first field of file and directory references to the FEP. A specification of fep: (with no unit number) refers to the disk unit from which the current Lisp world was booted, that is, the unit containing the world load file. If fep*n*: is omitted entirely, the default disk unit, set by Set Default-disk-unit is assumed.

For information on FEP commands for working with disks: See the section "FEP Commands for System Maintainers", page 101.See the section "FEP Commands for System Internals Maintenance", page 111.

8.8.1. Multiple Disk Units

Each Symbolics computer can access more than one local disk. The following conditions apply:

- The FEP can access any disk at all. Currently, the hardware allows a maximum of eight disks.

- You can boot a Lisp world from any disk by using the FEP command Load World. Also, you can add paging files from any disk by using the FEP command Declare Paging-files in your boot file. To load paging files from Lisp, you can use the Add Paging Files command.

- The form FEP: refers to the disk from which you booted the current world. If this is disk 0, then FEP: is equivalent to the form FEP0:. However, it is also possible to specify another disk explicitly, using such forms as FEP1: or FEP2:.

- World loads are not specific to a type of disk. This means that if one Symbolics computer has a T306 disk and another has an M2284 disk, world loads can be transferred back and forth between the disks.

8.8.1.1. Disk Types

The FEP currently supports the following types of disks:

Century Data T306 (300 megabytes unformatted capacity — removable)

CDC EMD368 (368 megabytes unformatted capacity)

CDC EMD515 (515 megabytes unformatted capacity)

Fujitsu M2284 (168 megabytes unformatted capacity)

Fujitsu M2294 (335 megabytes unformatted capacity)

Fujitsu M2351A (474 megabytes unformatted capacity)

Maxtor XT1140 (140 megabytes unformatted capacity)

Maxtor XT2190 (190 megabytes unformatted capacity)

NEC D2257 (167 megabytes unformatted capacity)

Priam P807 (335 megabytes unformatted capacity)

8.9. Finding Out Why Your Machine Crashed

When your machine crashes, using the FEP Show Status command can give you useful information for diagnosing the cause of the crash. For an outline of the information that Show Status prints: See the section "FEP Show Status Command Output", page 128.

The Show Status output section "3600 program counters" includes the macro PC, the CPC, and the 16 OPCs. The macro PC is the address of the current instruction of compiled Lisp code. The CPC is the address of the current microinstruction. The OPCs are the addresses of the 16 most recently executed microinstructions; OPC+0 is the most recent, OPC+17 the earliest. An arrow points at either the CPC or the first OPC, depending on the error condition that stopped the machine. This is the microinstruction that was executing when the event occurred that was the proximate cause of the machine's stopping itself.

8.9.1. FEP Show Status Command Output

The register contents and program counters displayed by the Show Status command give some information on machine states causing the FEP message "Lisp stopped itself". They are generally not useful for interpreting wired-ferror halts. Show Status merely prints the contents of certain hardware registers, decoding the bits symbolically. The FEP does not interpret these contents, so some output might not be meaningful. The following cautions apply:

- You must interpret some bits depending on the value of other bits.
- Some registers listed below are printed only if they contain "useful" information.

The most important registers are *Sequencer status* and *MC error status*.

FEP buffer status

Bit	*Meaning*
Spy DMA Enb	Spy bus being used by FEP to access disk or net (means spy bus being used for normal functions)
Write to dev / Read from dev	Spy DMA direction
Drive busy	Spy DMA mode (who controls busy line)
Int Enb	Spy DMA enable to interrupt FEP
Count up / Count down	Spy DMA address increment direction
Busy	Spy DMA busy (inside FEP)
Spy DMA busy	Spy DMA busy line (actual line on backplane)
DMA setup	[meaning unknown]

FEP Lbus control

Bit	*Meaning*
ECC Diag	Normal memory error correction logic disabled; instead, FEP can read or write the 8 extra bits of main memory
Doorbell Int Enb	Doorbell (Lisp-to-FEP signal) interrupt enabled
Use Uncorrected Data	FEP unaware of corrected Lbus data if single-bit-error
Ignore Double ECC Error	FEP does not get bus error if uncorrectable Lbus error (either double-bit error or nonexistent memory)
Task 3 Req	FEP trying to wake up microtask 3
Doorbell	Doorbell ringing (Lisp-to-FEP signal)
Lbus Buffer Busy	[self-explanatory]
Lbus Buffer Some Parity Error	[self-explanatory]

FEP Board ID control

Bit	*Meaning*
Continuity	Read-back of random signal that checks board presence
Lbus ID Req	Lbus reading board IDs, not doing normal functions
Half Speed	Main processor clock running at half speed

FEP Proc control

Bit	*Meaning*
Lbus Power Reset	Reset all Lbus devices due to power turn-on or turn-off
Lbus Power Reset (on bus)	Same as above, but actually read back from the bus
Lbus Reset	Reset all Lbus devices
Lbus Reset (on bus)	Reads back the Lbus Reset
Clear Errors	Bit that clears FEP error registers (not an error)
FEP Int Enable	FEP interrupt enable (not an error)
Kept Alive	FEP died and was reset by nanofep
FEP Ram Par Err	Parity error in dynamic ram on FEP board

Sequencer error status

(Status of the SQ board and the main error status bits that can halt the machine)

Bit	*Meaning*
Microcode-halted	A "halt" microinstruction was executed, for one of the following reasons: • A call to the %HALT function (due to a wired-ferror or a call to HALT) • A fatal error, such as an error while entering the error handler or an error in wired code (page fault, disk handlers) • Executing an undefined macroinstruction (running too old a version of microcode or executing bad macrocode) • Failure of a microcode consistency check (stack frame too large, stack overwritten)

Self-explanatory hardware errors:

Bit	*Meaning*
Spare-error-bit	[never happens, unless manually wired to some signal]
GC-Map-parity-error	GC MAP ram on DP board
Type-map-parity-error	TYPE MAP ram on DP board

Page-Tag-parity-error	PAGE TAG ram on FEP board
A-memory-parity-error	AMEM ram on DP board
B-memory-parity-error	BMEM ram on DP board
MC-error (map, ifu, or main mem)	Error on MC or IFU board; see MC / IFU error status
AU-error	Error on AU (FPA) board (if the machine has one)
Task-state-memory-parity-error	TSKM ram on SQ board (doesn't always halt machine)
Control-memory-parity-error	CMEM ram on SQ board (also for microcode breakpoints if an L-Console program is cabled up for debugging)

Hardware "errors" that are not always errors:

Bit	*Meaning*
CTOS-low-parity-error	CSTK ram on SQ board (low half of output register)
CTOS-high-parity-error	CSTK ram on SQ board (high half of output register)

(Note: If CTOS-came-from-IFU is true, above two bits have no meaning.)

Sequencer miscellaneous status

(Status bits that are not errors)

Bit	*Meaning*
CTOS-came-from-IFU	CTOS register holds macroinstruction dispatch address from IFU (or TMC) rather than contents of CSTK ram
TSK-STOP (sequencer stopped)	Machine is stopped for some reason
Errhalt-Sync	Some error bit is on (stops machine)
MC Wait	Microinstruction waiting for memory control to allow the instruction to continue
Task Switch	Switching to a different microtask

MC / IFU Error status

Bit	*Meaning*
Double bit error	An uncorrectable error in main memory, or a reference to a nonexistent Lbus address. Further information under the heading, ECC syndrome.
Map A parity error, Map B parity error	Parity errors in the map caches on the MC (TMC, IFU) board.
Hit in both map A and map B	Both map caches claiming to map the same address. Could be the map hardware, or some hardware or microcode problem causing map to be written with bad data.

IFU op parity error	Parity error in internal IFU operation.
IFU arg parity error	Parity error in internal IFU argument.

ECC syndrome

(An octal number followed by an address with x's in it)

This register contains the most recent main-memory read-error correction status. The error can be caused by a read by the processor, a read by the FEP, or a read by a DMA I/O device. The events that set this register include nonexistent memory reference, single-bit error correction, and double-bit error detection. Nonexistent memory and double-bit error halt the processor (even if it was the FEP or an I/O device that got the error). Currently, the FEP disables itself from getting a bus error if it references nonexistent Lbus memory or gets a double-bit error in Lbus memory.

One other event that can set this register is a bug in the FEP's code for examining the machine's status. In this case, the first two digits of the address are usually 77.

The address is the physical address of the location referenced. Only bits 23-18 and 1-0 are valid (the rest are x'ed out). These are sufficient bits to determine which Lbus slot (bits 23-19) and which of the 8 banks within a memory board are being referenced. To convert the address to an Lbus slot number, consider the one or two digits at the left of the x's to be an octal value, and divide it by 2. This is a logical slot number, as printed (in decimal) by the Show Configuration command. It is not related to the numbers printed on the machine chassis. Slot 0 is at the left, as seen from the front of the machine.

The syndrome codes are as follows:

	0	1	2	3	4	5	6	7
000	okay	36	37	2-bit	38	2-bit	2-bit	3
010	39	2-bit	2-bit	6	2-bit	8	9	2-bit
020	40	2-bit	2-bit	13	2-bit	15	16	2-bit
030	2-bit	18	19	2-bit	20	2-bit	2-bit	unused
040	41	2-bit	2-bit	23	2-bit	25	26	2-bit
050	2-bit	28	29	2-bit	30	2-bit	2-bit	31
060	2-bit	33	34	2-bit	35	2-bit	2-bit	unused
070	NXM	2-bit	2-bit	unused	2-bit	unused	unused	2-bit
100	42	2-bit	2-bit	0	2-bit	1	2	2-bit
110	2-bit	4	5	2-bit	7	2-bit	2-bit	10
120	2-bit	11	12	2-bit	14	2-bit	2-bit	unused
130	17	2-bit	2-bit	unused	2-bit	unused	unused	2-bit
140	2-bit	21	22	2-bit	24	2-bit	2-bit	unused
150	27	2-bit	2-bit	unused	2-bit	unused	unused	2-bit
160	32	2-bit	2-bit	unused	2-bit	unused	unused	2-bit
170	2-bit	unused	unused	2-bit	unused	2-bit	2-bit	unused

3600 program counters

Label	*Meaning*
Macro PC	The address of the current instruction of compiled Lisp code. This is prefaced with either (Odd) or (Even) since there are two instructions per word.
Current micro PC (CPC)	The address of the current microinstruction.
Old PCs (OPC)	The addresses of the 16 most recently executed microinstructions. OPC+0 was executed most recently, OPC+17 least recently.

Show Crash Data Command

Show Crash Data *keywords*

Obtains from the FEP the most recent output of the FEP's Show Status command and makes it available in your Lisp world. This information includes some hardware state, the compiled code program counter (macro PC), virtual memory address (VMA), stack pointer and frame pointer (SP and FP), and the 16 most recent microcode program counters (OPCs).

The information displayed by the Show Crash Data command from Lisp is similar to the information displayed when you issue the Show Status command from the FEP. The difference is that the Show Crash Data command's display is a replay of the machine's status when it crashed sometime in the past, and the FEP's Show

Status command reveals the machine's status at the time when the command is issued.

The Show Status command displays the internal status of some machine registers. For more information about the Show Status command: See the section "FEP Show Status Command Output", page 128.

If you want to put all of the information from the Show Crash Data command in the current editor buffer or in a mail message, use the Editor command Insert Crash Data (m-X). This makes it easy to retain information about a Lisp crash that put you into the FEP. (This data is preserved even if you cold boot the machine.)

The information retrieved from the Show Crash Data command is available until the next time Lisp halts, until the FEP is reset, or the machine is powered down. Therefore, even if you are unable to warm boot your machine, you can cold boot and then make a report with the data of the last crash.

Here is a possible sequence of events:

- The machine crashes and puts you into the FEP. The FEP issues the Show Status command automatically.
- Use the FEP command Start to warm boot the machine. If this does not work, use the Boot command to cold boot the machine.
- Use the command Show Crash Data to retrieve from the FEP the most recent output of the FEP's Show Status Command.
- Use the Insert Crash Data (m-X) command in the editor to save the information in a buffer.

keywords :Interpreted, :Output Destination

:Interpreted Interprets the data in the current world.

:Output Destination
Directs the output of this command to a specified location. The location can be the pathname of a buffer or file, a printer, stream, or window.

8.9.2. Decoding micro PCs

Use the functions **dbg:decode-micro-pc** and **dbg:decode-micro-pcs** to decode the microcode PCs printed by the FEP command Show Status.

dbg:decode-micro-pc *pc* &optional (*name* **sys:%microcode-version**) (*version* **(sys:microcode-version-number sys:%microcode-version)**) *Function*

dbg:decode-micro-pc is useful for investigating why a machine crashed. It decodes an octal microinstruction address printed by the FEP command Show Status. To use this function you should first write down the Show

Status output. You can then either warm boot the machine using the Start command or call **dbg:decode-micro-pc** on another machine. To decode more than one octal microinstruction address: See the function **dbg:decode-micro-pcs**, page 134.

dbg:decode-micro-pcs *pcs* &optional *(name* **sys:%microcode-version)** *(version* **(sys:microcode-version-number sys:%microcode-version))** *verbose (load-symbol-table* **:ask)** *Function*

dbg:decode-micro-pcs is useful for investigating why a machine crashed. It decodes the octal microinstruction addresses printed by the FEP command Show Status. To use this function you should first write down the Show Status output. You can then either warm boot the machine using the Start command or call **dbg:decode-micro-pcs** on another machine. To decode only one octal microinstruction address: See the function **dbg:decode-micro-pc**, page 133.

pc is an address in the microcode, taken from the CPC or OPC information printed by the Show Status command. Show Status prints these numbers in octal; if your default radix is decimal, precede *pc* by **#o**. Normally the number in the Show Status output with the arrow (→) pointing to it is the relevant number, but sometimes, decoding all of the numbers gives you additional clues.

name and *version* are optional; they specify the version of the microcode that was running at the time of the crash. You can omit these arguments if you call **dbg:decode-micro-pc** or **dbg:decode-micro-pcs** while using the machine that crashed and while running the same microcode version as at the time of the crash. You can also omit these arguments if you call this function from another machine that has a software *and* hardware configuration that is *identical* to that of the machine that crashed. To find the microcode version name and number that a machine is running, use the command Print Herald with the keyword :detailed or take the name and version number of the microcode file in the machine's boot file (normally fep0:>Boot.boot). Microcode version numbers are decimal; include a period at the end of the number if your default radix is octal.

8.9.2.1. dbg:decode-micro-pc(s) Examples

Example of **dbg:decode-micro-pc** function call:

```
(dbg:decode-micro-pc #o44552 "3600-mic" 389.)
```

Example of **dbg:decode-micro-pcs** function call:

```
(dbg:decode-micro-pcs (#o44552 #o12345) "3600-mic" 389.)
```

dbg:decode-micro-pc and **dbg:decode-micro-pcs** print information that depends on the microinstruction:

Microinstruction	*Information printed*
Halt instruction	The reason it halts the machine. An example is "error in the error handler". These reasons are

constant strings in the microcode source program and do not represent any dynamic analysis of the state of the machine.

Signaller of a Lisp error
: The internal form of the error message. This is not the same form of error message you would ever see otherwise; normally Lisp software translates these messages into conditions and signals them, and the conditions define more readable error messages. This is useful mainly in decoding OPCs earlier than the one with the arrow, when the machine halted because of "error in the error handler".

Handler for a macroinstruction in compiled Lisp code
: The name of that macroinstruction. A halt here might be caused by running a world together with an incompatible microcode, such as a microcode from an earlier release, that does not implement an instruction used by that world.

If all else fails, the function offers to load the microcode symbol table (from the sys:l-ucode; directory) and then prints the symbolic name of the microinstruction. Loading the microcode symbol table takes a few minutes. Microinstruction symbolic names can sometimes be clues to help in figuring out what the machine was doing at the time it crashed.

Two types of symbolic names exist: those with and without parentheses.

If the name includes parentheses, it is a list of the name of a microcode routine and the path through that routine to reach the microinstruction in question. *Beware of a pitfall!* These names are not unique; the same microinstruction can be reached by multiple paths from different microcode routines. For example, a microinstruction named (FTN-AR-1 3) might also be part of the microcode for the CAR instruction; you cannot assume too much from the name if it contains parentheses. It is only a clue.

If a symbolic name is just a symbol and has no parentheses, it is unique and names the first microinstruction of a microcode routine.

Beware of assuming too much. If the reason Lisp stopped itself is not "microcode halted", the information that **dbg:decode-micro-pc** and **dbg:decode-micro-pcs** print is not likely to be helpful, though it might be useful to people who understand the hardware.

For more information about crash data: See the section "Show Crash Data Command", page 132.

8.9.3. Decoding macro PCs

To decode the macrocode PC printed by the FEP command Show Status, warm boot or go to another machine running identical software and call the function **sys:%find-structure-header** on the number printed by the FEP. This is an octal number; use **#o** if necessary. It should return a compiled-function object, which is the function that was executing at the time. To find the exact place in the function that was executing, note the difference between the number printed by the FEP and the address in the printed representation of the compiled-function object. You can use **sys:%pointer-difference** to compute this difference. Multiply this by 2, and add 1 if the FEP said the PC was odd (not even). The result is the instruction number of the current instruction; disassemble the compiled function to see it.

Example:

```
FEP Command:  Show Status
...
3600 program counters:
  Macro PC/ (Odd)1244531
  ...
FEP Command: Start
...
(%find-structure-header #o1244531)
#<DTP-COMPILED-FUNCTION EQUAL 1244530>
(%pointer-difference #o1244531 *)
1
(1+ (* * 2))
3
(disassemble ***)
  0 ENTRY: 2 REQUIRED, 0 OPTIONAL
  1 PUSH-LOCAL FP|0              ;A
  2 PUSH-LOCAL FP|1              ;B
  3 BUILTIN EQL STACK
  ...
```

Instruction 3 (EQL) is the one that halted.

8.10. Debugging in the FEP

The release tapes include some files provided as an extra debugging aid. These files can be used to enter a debugging mode in the FEP. This mode is especially useful for problems that cause control to return to the FEP, making it impossible to use the debugging methods normally used in Lisp.

These files have names of the form: v*n*-debug.flod, where *n* is the FEP version number. The files should now reside on your sys host in the directory with the logical pathname SYS:N-FEP;V127-DEBUG FLOD, where 127 is the version of FEP software.

"V127" is an example. On your system, use the version number of the FEP EPROM which your machine is using. This may not be V127. Other FEP EPROMS are G206 and G208.

To use these files, you should copy the appropriate file to the FEP file system *before* you need to use it. To copy the file, first find out which version of FEP software is installed in your machine. You can do this by either using the Show Herald command in Lisp, or by typing the Show Version command at the FEP level. To copy this file to your FEP file system, use the Copy File command.

For example, if you are using FEP version 127 software, you would use the following command to copy the .flod file to the FEP file system:

```
Copy File sys:n-fep;v127-debug.flod fep0:>v127-debug.flod
```

The flod file cannot be used on any other FEP version; trying to use one on a different FEP version has no effect.

After you have copied the file to the FEP file system, you can enter the debugging mode by loading the file with the Load FEP command, as shown in the following example:

```
FEP Command: Load Fep >v127-debug.flod
```

This puts you into a debugging mode very similar to the Debugger in Lisp, whereby you can move up and down the stack to examine the state of the machine and determine the source of the problem. The HELP key lists the available commands.

One particularly useful command, when the machine has crashed during paging, is c-m-S. This command allows you to switch between the auxiliary stack (where paging code runs) and the normal stack (where user code runs). If the machine crashed while executing on the auxiliary stack, user stack frames will not be found until c-m-S is executed.

If you need to stop the execution of Lisp and give control to the FEP, use the Halt Machine command.

Halt Machine Command

Halt Machine

Stops execution of Lisp and gives control to the FEP. You can now enter FEP commands, for example, to warm or cold boot the machine.

If you need to know the machine model, use the function **(si:machine-model)**.

si:machine-model *Function*

This function returns a keyword symbol designating the model number of the current 3600-family computer.

Possible return values are as follows:

:unknown	The model number cannot be determined (usually indicating lack of some ID prom)

:/3600 or :|3600|	(The keyword whose print-name is "3600".) The machine is a Symbolics 3600.
:|3670|	The machine is a Symbolics 3670.
:|3675|	The machine is a Symbolics 3675.
:|3640|	The machine is a Symbolics 3640.
:|3645|	The machine is a Symbolics 3645.

If you want to look at your machine's hardware configuration, use the Show Machine Configuration command.

Show Machine Configuration Command

Show Machine Configuration *host keywords*

Shows the board-level hardware information about any 3600-family machine on the same network as your machine.

host	The name of a 3600-family machine. The default is your machine.
keywords	:Output Destination

:Output Destination
{Buffer, File, Printer, Stream, Window} Where to redirect the typeout done by this command. The default is the stream ***standard-output***.

This information is useful for service personnel. You might be asked for the machine serial number (in line 3) if you call Symbolics Software Support. The display from Show Machine Configuration looks like this:

si:machine-model *Function*

This function returns a keyword symbol designating the model number of the current 3600-family computer.

Possible return values are as follows:

:unknown	The model number cannot be determined (usually indicating lack of some ID prom)
:/3600 or :|3600|	(The keyword whose print-name is "3600".) The machine is a Symbolics 3600.
:|3670|	The machine is a Symbolics 3670.
:|3675|	The machine is a Symbolics 3675.
:|3640|	The machine is a Symbolics 3640.
:|3645|	The machine is a Symbolics 3645.

```
Chassis (PN 170219, Serial 381) in Chassis or NanoFEP:
    Manufactured on 4/2/85 as rev 1, functions as rev 1, ECO level 0
    Machine Serial Number: 4415
    FEP Version Number: 127
Datapath (PN 170032, Serial 3163):
    Manufactured on 3/7/85 as rev 3, functions as rev 3, ECO level 0
Sequencer (PN 170042, Serial 2572):
    Manufactured on 4/21/85 as rev 4, functions as rev 4, ECO level 0
Memory Control (PN 170052, Serial 2741) in Memory Control or IFU:
    Manufactured on 4/14/85 as rev 5, functions as rev 5, ECO level 0
Front End (PN 170062, Serial 3392) in FEP:
    Manufactured on 4/4/85 as rev 5, functions as rev 5, ECO level 0
512K Memory (PN 170002, Serial 2638) in LBus slot 00:
    Octal Base address: 0
    Manufactured on 3/1/85 as rev 2, functions as rev 2, ECO level 0
512K Memory (PN 170002, Serial 515) in LBus slot 01:
    Octal Base address: 2000000
    Manufactured on 9/20/83 as rev 2, functions as rev 2, ECO level 0
512K Memory (PN 170002, Serial 3357) in LBus slot 02:
    Octal Base address: 4000000
    Manufactured on 6/20/85 as rev 2, functions as rev 2, ECO level 0
IO (PN 170157, Serial 1062) in LBus slot 03:
    Octal Base address: 6000000
    Manufactured on 5/7/85 as rev 6, functions as rev 6, ECO level 0
512K Memory (PN 170002, Serial 358) in LBus slot 04:
    Octal Base address: 10000000
    Manufactured on 7/16/83 as rev 2, functions as rev 2, ECO level 0
FEP Paddle Card (PN 170069, Serial 912) in FEP -- PADDLE side:
    Manufactured on 3/21/85 as rev 1, functions as rev 1, ECO level 0
IO Paddle Card (PN 170245, Serial 480) in LBus slot 03 -- PADDLE side:
    Manufactured on 3/25/85 as rev 1, functions as rev 1, ECO level 0
    Ethernet Address: 08-00-05-03-00-0F
```

Figure 6. Display of a Machine's Configuration

9. Dumping, Reloading, and Retrieving

A file system can be damaged or destroyed in any number of ways. Users can delete files by accident. To guard against such a disaster, it is wise to *dump* the file system periodically, that is, write out the contents of the files, their properties, and the directory information onto magnetic tapes. If the file system is destroyed, it can then be *reloaded* from the tapes. Individual files can also be *retrieved* from tapes, in case a single file is destroyed, or just accidentally deleted (and expunged). Dump tapes can also be used to save a copy of all the files on a system for archival storage.

In a *complete dump*, all of the files, directories, and links in the file system are written out to tapes. This, obviously, saves all the information needed to reload the file system. However, a complete dump can take a long time and use a lot of tape, especially if the file system is large. In order to make it practical and convenient to dump the file system at short intervals, a second kind of dump can be done, called an *incremental dump*.

In an incremental dump, only those files and links that have been created or modified since the last dump (of either kind) are dumped; things that have stayed the same are not dumped. (All directories are always dumped in an incremental dump.) Now, if the file system is destroyed, you reload it by first reloading from the most recent complete dump and then reloading each of the incremental dump tapes made since that complete dump, in the same order in which they were created. Therefore, you do not need to retain incremental dump tapes that were made *before* the most recent complete dump was done; you can reuse those tapes for future dumps.

Since all tapes containing incremental dumps done since the last successful complete dump must be reloaded in order to restore the file system, doing a complete dump regularly makes recovery time faster. Doing complete dumps also lets you reuse incremental dump tapes, as described above. The more incremental dump tapes you must load at recovery time, the longer it takes to recover, and thus the more chance there is that something will go wrong. Thus, it is advantageous to perform complete dumps periodically.

A *consolidated dump* is like an incremental dump, in that it only dumps files that have been created or changed recently. However, a consolidated dump backs up only those files that have been created or changed since a specified *consolidation date*. A consolidated dump is the appropriate kind to take if some event destroys recent incremental dump tapes, or they are found to be unreadable. If a complete dump extends through several days, it is wise to take an incremental dump between tape stopping points as appropriate.

9.1. Performing Dumps

To perform a dump, follow these steps:

1. Mount a magnetic tape on an available and usable tape drive which need not be connected to a machine on which you are performing the dump.

2. Press SELECT F to select the File System Editing Operations Program and click on [Local LMFS Operations].

   This invokes the second level of the File System Editing Operations Program, which is called *Local File System Control Operations*.

3. Choose either [Incremental Dump], [Complete Dump] or [Consolidated Dump] by clicking on the menu.

   These commands respond with an Accept Values menu that lets you set the parameters of the dump; the only difference among the commands is the initial value of some of the parameters in this window.

4. Change the values in this window as needed.

Figure 7 shows the second level of the File System Editing Operations Program.

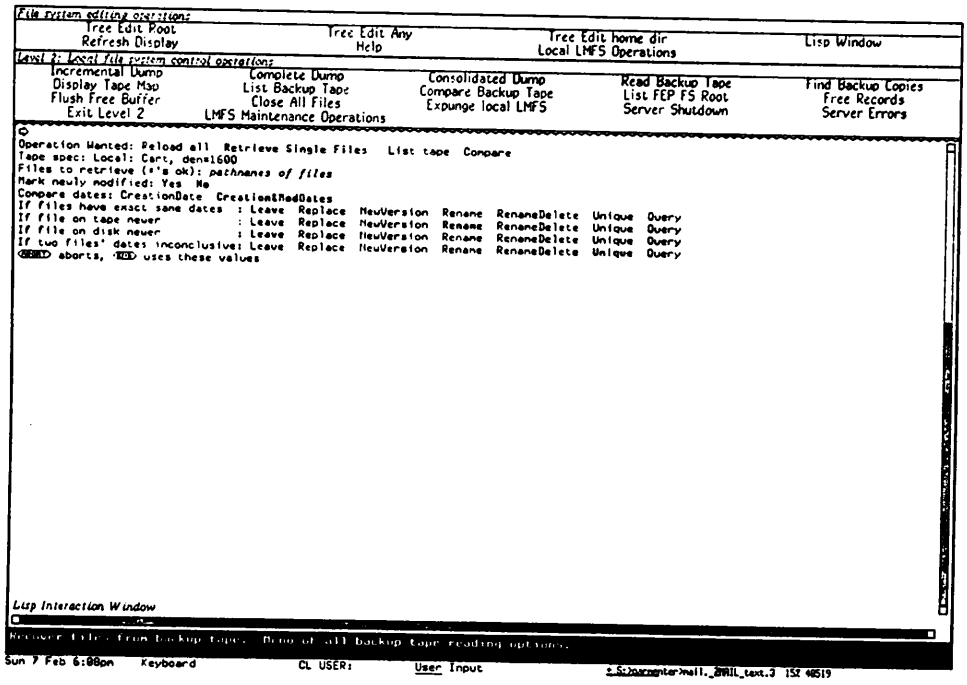

Figure 7. Performing a Dump

Here is an explanation of the parameters offered for modification in this window:

Dump Type

There are three possible types of dump: *incremental*, *consolidated*, and *complete*. The three File System Maintenance commands initialize this field according to their own requirements; the [Dump] command initializes it to *complete*.

Pathnames

The pathname, or pathnames, specifying what is to be dumped. If there is more than one pathname, they are separated by commas. This value controls what files and directories are inspected for dumping. The type of the dump (complete, incremental, or consolidated) and the status of the individual files controls what subset of these files are actually dumped. For information about the different types of dumps: See the section "Dumping, Reloading, and Retrieving", page 140. Names of single files or links can be used to dump single files or links.

Wildcard specifications can also be used: this is the normal way to dump many files from one directory, or from a subtree. Subtrees are dumped via recursive (**:wild-inferiors**, "**") directory wildcards. The pathname you type is merged with "local:>**>*.*.*".

To dump the whole file system, which is the normal default, the appropriate pathname is:

```
>**>*.*.*
```

To dump all the files in directory >foo>bar, and all of its inferiors, the appropriate pathname is:

```
>foo>bar>**>*.*.*
```

To dump all the latest Lisp files in directory >abel>baker, but not any of its inferiors, the appropriate pathname is:

```
>abel>baker>*.lisp.newest
```

See the section "Naming of Files" in *Reference Guide to Streams, Files, and I/O*. See the section "LMFS Pathnames" in *Reference Guide to Streams, Files, and I/O*.

Tape Reel ID

Every reel of tape produced by the dumper must have a Tape Reel ID, which is a string of up to eight characters. You must explicitly supply a value for this option. The reel ID is used to identify this reel of tape to the backup system; it appears in the dump maps and in any messages about the tape. The **:complete-dump-tape** or **:incremental-dump-tape** property of any file dumped is set to this value, as well. The Tape Reel ID should be written with a pen onto the label of

the tape so that the tape can be identified by sight. *Note that you must supply a Tape Reel ID.*

Tape Drive Spec — The host name of the machine on which the tape drive to be used appears. This is initialized to a reasonable default, which is usually Local: if there is a tape drive present on the local host.

Dump deleted files — This can be either *Yes* or *No*, and says whether files marked as deleted but not yet expunged should be dumped on the backup tape. The default value is *No*; deleted files normally are not dumped.

Tape when done — This controls what the dumper does with the tape when it has finished passing over all the specified pathnames. These are the available options:

Offline	Rewinds the tape and puts it offline. It declares the dump finished. This is the default.
Rewind	Rewinds the tape without putting it offline. It declares the dump finished. It facilitates listing or verifying the tape contents.
Leave	Leaves the tape positioned at the end without rewinding it or putting it offline. It declares the dump finished. It facilitates more dumping later, by leaving the tape in the correct position for using "Append to tape" in a later dump invocation.
Query	At the end of the dump, all of these options are presented, and you can choose whether to rewind and set the tape offline, rewind it, leave it at end, or dump some more files. If you choose to dump more files, the dumper menu is offered again, and the new files are appended to this tape. The dump is not declared finished until you click on "Abort".

Person operating — The identification of the person doing the dump. This is entered into the backup map, and sets this person as the file author of that map. Normally, this is the same as the login ID of the user performing the dump, and that is its default value. However, if the user who is performing the dump is not logged in to the machine from which the dump is

invoked, this field should be filled with that user's name. It is important for documentation purposes and site record-keeping.

Consolidate from

This field is only used during a consolidated dump. It is a date and time in the past, entered in any acceptable Lisp Machine format. The consolidated dump dumps all, and only, files that have been created or modified since this date.

Set date dumped

When the dumper finishes writing a tape, it marks all the files it has dumped as having been dumped on that tape at this time, and creates a *tape directory*, as described below. These measures allow the file to be retrieved later, and indicate that the file no longer needs to be dumped in incremental dumps. This is the default action, which corresponds to a value of *Yes*. **Note:** The LMFS dumper should not be used to move software between sites as it is far too general, and system-independent; use the carry system and the distribution tape system instead. However, if you do use the dumper to make tapes that are not part of your site's backup, such as for moving software between machines, you do *not* want to indicate that the files were dumped, or to make a tape directory. Select *No* in this case.

Restart pathname

The purpose of this feature is to allow restarting of complete dumps that are interrupted by any sort of failure. When the dumper finishes a tape, it prints out the pathname of the last file dumped. Although this is recorded in the dump map, *the pathname and the name of the tape should be recorded on paper by the person doing the dump*, especially if a complete dump is being done.

To restart a dump, fill out the menu as usual, but type in the pathname of the last file known to have been dumped as the value of [Restart Pathname]. The dumper scans the sub-hierarchy indicated, but does not dump files already dumped, or even progress down, seeking files to be dumped, into directories that the restart pathname indicates have already been processed. The skipping of files and directories already dumped is based on sorting order, not whether the file has actually been dumped. Thus, if files A, C, E, G, and I exist, and files A through E get dumped one day, and the dump is interrupted and restarted from E the next day, a D created in the interim is not dumped.

Comment

A string, of arbitrary contents, written on each reel of the dump and in the dump map. This might say why the dump was performed, or any other special information about this dump.

When you are done filling in values, press END; if you decide not to do a dump after all, press ABORT. If there is something wrong about the set of parameters you have specified, the program displays a message and presents you with the Accept Values window again. Otherwise, it displays a message saying that the dump has started successfully, and proceeds. While the dump is in progress, the name of the file that is being dumped is shown in the far right-hand field of the status line (this is the field that normally shows you the names of files that are being read or written).

The dumper creates a file called the *dump map*. The dump map is a character file, giving a complete description of what has been dumped, directory by directory and file by file, including the time of dumping, the tape on which the file was dumped, the tape reel ID of the previous tape on which the file appears (if any), and so on. The dump map is created in the >dump-maps directory. Its name is constructed from the type of dump and the date and time at which it was started; the file type is **map** and the version is **1**. A typical dump map might have the pathname:

```
>dump-maps>complete-3/15/86-9:02.map.1
```

The dumper puts all information about the dump, the operator, the time of day, the options, and so forth, in the map. It also puts error recovery information there, and descriptions of tape-changings, as well as the number of files dumped on each tape. The dumper performs a **:finish** operation on the map file at the end of each tape, so that if the system crashes during a multi-tape dump, information about previous tapes is guaranteed to be intact and accessible.

The dumper also creates a file called the *tape directory* for each separate reel of each dump. This is a binary file saying what is on the tape, with more or less the same information as the dump map. You use this file when you try to locate dumped copies of a file. See the section "Finding Backup Copies of Files", page 151. The tape directory is also created in the >dump-maps directory. Its name is the tape reel ID of the tape, its file type is **directory**, and its version is **1**. A typical directory map might have the pathname:

```
>dump-maps>INC00001.directory.1
```

The dumper dumps files successively to tape, and at the end of each tape, rewinds and unloads the tape, asking for a new tape if there are more files to be dumped. It is only after it has done this that it sets backup dates for the files and makes the dump directory.

If the dumper gets an irrecoverable error while writing a tape, it attempts to write end-of-file marks on the tape and inform you of what has happened. It gives you the option of either considering the files on that tape to have been validly dumped, in which case the dump continues on the new tape, or discarding the tape, in which case it redumps all the files that it had dumped on the bad tape onto the new tapes. The problem and its chosen recovery are described in the dump map.

By default, the dumper tries to read each tape before writing on it. This is to avoid accidentally overwriting valuable tapes. For tapes to be appended to, this is necessary. For other tapes, it is desirable. It often takes a long time to attempt to read blank tape, to prove that a new tape is really new. The dumper explains and queries if it is not confident that the tape being written on is the right one.

Some sites may want to waive this checking. This is necessary when tape hardware is in use that cannot time out while reading blank tape, and therefore reads the whole tape when a new tape is checked, with no way to stop it. The site option **validate-lmfs-dump-tapes**, an attribute of the Site object for a site, which is normally elected, enables the suppression of this checking.

9.2. Reloading and Retrieving

Reloading is the process of moving all the files on a backup tape into a local file system. *Retrieving* is the process of moving selected files.

Reloading and retrieving can load files onto any LMFS.

Two other functions are related to reloading and retrieving: listing the contents of backup tapes with the List tape option, and verifying the contents of the dump with the Compare option. The *reloader* program implements all four of these functions.

To invoke the reloader:

1. Press SELECT F to get to Level 1 of the File System Editing Operations program.

2. Click on [Local LMFS Operations] to invoke Level 2.

3. Click on [Read Backup Tape].

Figure 8 shows the Accept-Values menu that appears when you click on [Read Backup Tape].

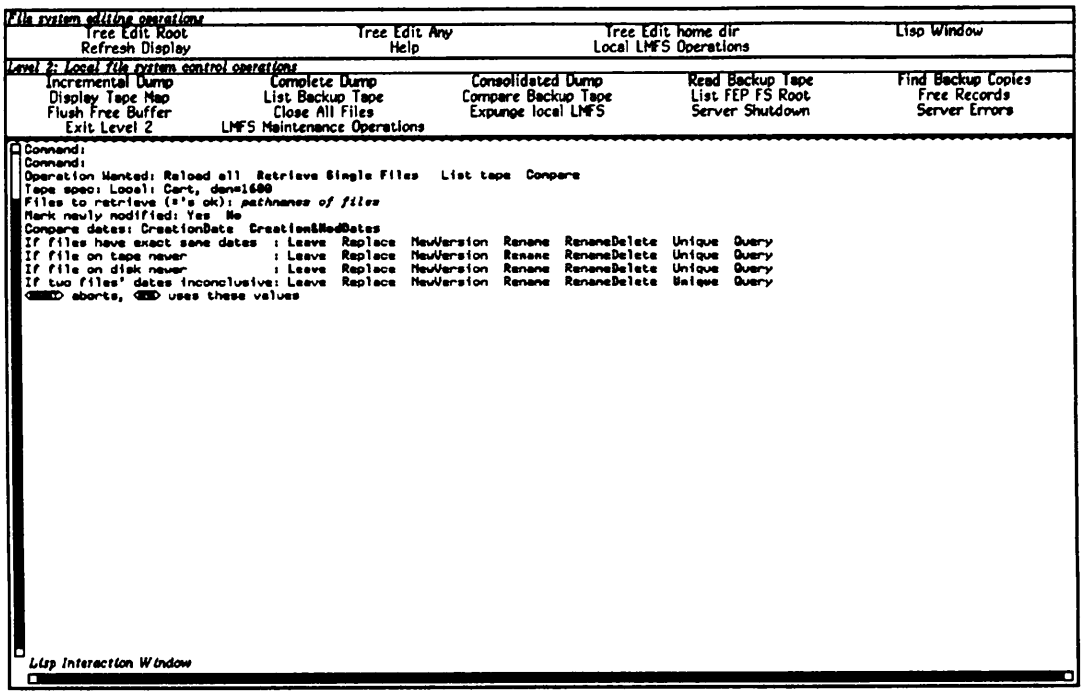

Figure 8. The Read Backup Tape Menu

The reloader Accept-Values menu contains the following items:

Operation Wanted This selects which function the reloader is to perform.

These are your operation choices:

Reload all
: The reloader is to read the tape, and reload all files on it. It will not reload files that are already present on the local LMFS.

Retrieve Single Files
: The reloader will expect a list of pathnames, to tell it which files to reload. It will not reload files that are already present on the local LMFS. You specify these pathnames by clicking on the menu item [Files to retrieve]. You can specify wildcard pathnames. For example, a wildcarded pathname for a single file looks like this:

```
E:>trees>.*
```

List tape
: The reloader will read the tape and display a description of its contents on the screen. It lists the dumps appearing on the tape, and which files are on the tape. This does *not* verify that the files were dumped correctly; use "Compare" for that.

Compare
: The reloader will read each file on the tape, look for the same file in the file system, and compare the two, bit for bit, reporting any discrepancies. Use this option to verify that a dump tape just created contains good data.

After selecting the operation to be performed, you supply the following information:

Tape Spec:
: The tape spec. Normally, for a 3600 with an attached cartridge tape the default is **Local: cart**, which is an acceptable expression of the local machine. Otherwise, a reasonable default will be chosen, based on available tape servers at your site. The tape drive is a character string identifying the tape drive to be used on the tape host. You only need to supply this value if the machine selected has more than one tape drive, and needs this information to select one.

Files to retrieve (*'s ok):
: One or more pathnames, which are usually wildcard pathnames. This is used when you click on [Retrieve single files]. Any file matching one of these wildcard pathnames will be reloaded, unless it is already there. For more information, see the following sections:

See the section "Naming of Files" in *Reference Guide to Streams, Files, and I/O*.See the section "Wildcard pathname mapping" in *Reference Guide to Streams, Files, and I/O*.See the section "LMFS Pathnames" in *Reference Guide to Streams, Files, and I/O*.

Mark newly modified: Yes No

Specify yes to mark the newly reloaded files as *not yet dumped* so that they will get dumped by the next incremental or consolidated dump. Choose no if you don't want the files to be dumped in the next dump. The default is No.

The reloader will move all the files from the backup tape into the local file system unless files of the same name, type, and version exist. If this is the case, the reloader compares the creation and or creation and modification dates of these files according to guidelines you specify here. These are the guidelines you can choose:

Compare Dates: CreationDate Creation&ModDates

Compares the creation dates or creation dates and modification dates of the files you are trying to reload, depending on what you choose.

If files have exact same dates:

Your choices are to: Leave, Replace, NewVersion, Rename, RenameDelete, Unique, Query. The default is Leave. The meaning of these choices is explained below.

If file on tape newer:

Your choices are to: Leave, Replace, NewVersion, Rename, Unique, Query. The default is Rename. The meaning of these choices is explained below.

If file on disk newer:

Your choices are to: Leave, Replace, NewVersion, Rename, Unique, Query. The default is Leave. The meaning of these choices is explained below.

If you choose the creation and modification dates for the Compare Dates option in the first menu item, then this option is presented:

If two files' dates inconclusive:

Your choices are to: Leave, Replace, NewVersion, Rename, Unique, Query. The default is Unique. The meaning of these choices is explained below.

Here is the meaning of each choice for the questions above:

Leave	Leaves the file in the file system and does not reload it.
Replace	Loads the file from tape, completely replacing the file in the file system.
NewVersion	Loads the file from tape as the newest version.
Rename	Renames the file uniquely, and loads the file from tape.
RenameDelete	Renames the file uniquely, deletes it, and loads the file from tape.

Query	Stops the reload and asks you what to do when the tape is ready.
Unique	Loads the file from tape under a unique name in the proper directory.

When you have made your choices from the reloader Accept-Values window, press END to begin the reloader. If you do not wish to proceed with the reload, tape list, or compare, press ABORT.

The reloader prints information about each file it reloads or about every file on the tape, if you have selected [List tape]. When it gets to the end of the tape, it stops. If you want to continue reloading, you must mount another tape and restart the reloader. Tapes can be reloaded in any order, but choose your reload options carefully. Generally you should load tapes in the opposite order of creation, and not replace any newer files. This prevents writing over files, such as mailboxes.

The reloader does not delete or expunge files. If you are reconstructing a file system from backup tapes that include many incremental backups, you must occasionally intervene to delete and expunge unwanted old files that are reloaded, in order to ensure adequate file space.

9.2.1. Comparing Backup Tapes

A few notes about the comparer: Occasionally, hardware problems can cause bad backup tapes to be written, without any error having been detected by the dumper. You should always verify a backup by using the "Compare" option of the reloader to *verify* the backup tape. The comparer verifies each bit of every file on the tape, and, for files that still exist, reports any discrepancy.

If you find that an incremental backup tape is deficient, and you decide that the files must be redumped, you must perform a consolidated dump, with a consolidation date equal to the time the bad incremental dump started. You must delete the backup dump tape directory ("*tape-name*.directory" in >dump-maps) by hand, if the the backup-copy finding mechanism is to be aware that the tape has been abandoned.

The reloader produces an error log file in the local directory >reloader-logs.

The dumper assumes that no undetected problems occurred. Thus, it does not run the comparer automatically. The dumper sets backup times when it finishes each tape.

9.3. Finding Backup Copies of Files

In order to retrieve individual files, or groups of files, rather than reloading an entire tape, you must know on what tape the files were dumped. The LMFS backup mechanism provides a way to search the binary tape directories, which are produced by the dumper, to find backup copies of files. This is the Lisp function **lmfs:find-backup-copies**, which you can invoke by clicking on [Local LMFS Oper-

ations] in Level 1 of the Files System Editing Operations program. Then click on [Find Backup Copies]. This prompts you for files pathnames.

lmfs:find-backup-copies searches all the binary tape maps at the site for the tape locations of all backup copies of a file. It prompts for names, which it parses with respect to the local host unless you name an explicit host in the pathname. All of the files requested must be from the same host. This applies only to Lisp Machine hosts. It looks at the binary tape maps on the specified host in the directory >dump-maps.

Normally, you specify wildcard pathnames for this function to match. A non-wild pathname is considered to be a valid degenerate case of a wildcard pathname. The default with which all of the names are merged is "local:>**>*.*.*". This function uses the same pathname interpretation conventions as the reloader.

Here is a sample interaction with **lmfs:find-backup-copies**:

```
Enter file pathnames for which to search, separated by commas.
Wildcard names are allowed.  [default LOCAL-LISPM:>**>*.*.*:]
Paths: f:>sys>io>*.lisp, f:>bsg>*.init, f:>lisp>*.q*
Searching for:
  F:>sys>io>*.lisp.*
  F:>bsg>*.init.*
  F:>lisp*>**>*.q*.*
-----
F:>LISPM>COLDLD.QBIN.39, created 2/17/88 00:26:58, on tape fscns001
  (Backup dump of 2/24/88 13:05:44)
F:>LISPM>COLDUT.QFASL.57, created 2/18/88 19:06:36, on tape fscns001
  (Backup dump of 2/24/88 13:05:44)
F:>LISPM>UTILS>NEW>EVAL.QBIN.13, created 2/19/88 04:45:17, on tape fscns001
  (Backup dump of 2/24/88 13:05:44)

....and so on
----- Scan complete.
```

10. Setting and Defining Sites

The location of your machines is called a site. When you boot new software (distributed as a world load), your site is not set. The herald gives the machine name as *DIS-LOCAL-HOST*. In order to begin operating your computer at a site, you must configure the new software at a site by using the Set Site or the Define Site command.

If you are configuring software at a site that already exists, you must use the Set Site comand. This gives your machine access to the network capabilities of an already existing site.

If you are configuring software at a site that does not yet exist, you must use the Define Site command. This creates a new site for your machine. Later, you can expand the site to include other users, hosts, printers, and networks. For more information: See the section "Customizing and Saving Worlds", page 14.

10.1. Set Site Command

Set Site *site-name*

Starts a dialogue to set the current site to be *site-name*. It should be the first thing you type to your machine after booting a new distribution world. The command makes the identity of all objects included in the site's namespace known to your machine.

In order for the Set Site command to work:

1. Your machine must be declared as a host in the site's namespace. See the section "Registering Hosts", page 79.

2. The namespace database for each site is stored on a machine called the namespace server. If the site's namespace server is not the local host, you must know the name and chaosnet address of the namespace server. For more information: See the section "Choosing Chaosnet Addresses".

10.2. Set Site Dialogue

Here is an explanation of each line of the Set Site dialogue for a site. Following each line of the Set Site dialogue is the explanatory text.

```
Command: Set Site (site name [default Get from network]) downunder
```

First you must issue the Set Site command. If you accept the default (called "Get from network"), your machine automatically queries other network machines to learn their site name, and uses that site name.

If you wish to change your site, or if your site information is unavailable, you must enter a site name manually. In this example *downunder* is the site name.

When you enter the Set Site command, the system prompts:

```
What host is a namespace server for DOWNUNDER
(default: Local):
```

This question prompts for the name of the namespace server for your site.

1. If the site's namespace server is the local host (the machine on which you are typing), you must press RETURN to accept the default. **Do not type the name of the local host.** The machine then asks for the location of the *descriptor file*. The *descriptor file* holds namespace data files. The default value is the same one offered when defining a site:

   ```
   Where is the descriptor file for DOWNUNDER
   (default local:>sys>site>downunder-namespace.text):RETURN
   ```

   The machine is now on the network at the site *downunder*.

2. If the site's namespace server is not the local host, then you must provide the name of the namespace server for the site. The machine then asks for the server's chaosnet address, and asks you to confirm the server's name. In this example, *sydney* is the name of the server.

   ```
   What host is a namespace server for DOWNUNDER
   (default: Local): sydney
   Chaosnet address for sydney: xxxxx
   Host responds as SYDNEY, ok? (Y or N) Yes.
   ```

   The site's namespace determines all other relevant information, and the machine is now on the network at the site *downunder*.

Your machine now has access to all objects located at the site, and you are ready to login and use the new world load. You may want to save the new, site-specific version of your world. For more information:See the section "Customizing and Saving Worlds", page 14.

10.3. Define Site Command

Define Site *site-name*

Starts a dialogue to create a new site called *site-name*.

The default values for the site's namespace server, SYS host, computer storing the namespace data files, and host for bug reports, are all the local host. If you provide a non-local host for any of these, you must know the non-local Chaosnet address(es) and operating system(s). For more information:See the section "Choosing Chaosnet Addresses". Here are the default values for the Define Site command:

System Type*: LISPM
Service: CHAOS-STATUS CHAOS-SIMPLE CHAOS-STATUS
Service: SHOW-USERS CHAOS NAME
Service: TIME CHAOS-SIMPLE TIME-SIMPLE
Service: UPTIME CHAOS-SIMPLE UPTIME-SIMPLE
Service: LOGIN CHAOS TELNET
Service: LOGIN CHAOS SUPDUP
Service: LOGIN CHAOS 3600-LOGIN
Service: SEND CHAOS CONVERSE
Service: SEND CHAOS SEND
Service: NAMESPACE CHAOS NAMESPACE
Service: NAMESPACE-TIMESTAMP CHAOS-SIMPLE NAMESPACE-TIMESTAMP
Service: LISPM-FINGER CHAOS-SIMPLE LISPM-FINGER
Service: FILE CHAOS NFILE
Service: FILE CHAOS QFILE
Service: CONFIGURATION CHAOS CONFIGURATION
Address: CHAOS *nnnnn*

In the Address attribute line, *nnnnn* represents a valid 5-digit octal Chaos address.

10.4. Define Site Dialogue

Here is an explanation of each line of the Define Site dialogue for a site. Following each line of the Define Site dialogue is the explanatory text.

```
Command: Define Site (site name) downunder
```

First you issue the command and give the site name. This might be the name of your company, or, if you are more whimsical-minded, it might be related to the machine names you have chosen. In this example, *downunder* is the name of an example site.

```
Define a new site named DOWNUNDER (as opposed to looking for an
existing definition of DOWNUNDER on disk)? (Yes or No) Y
```

verifies that you want to create a new site rather than setting your site to an existing one.

```
What host is to be a namespace server for DOWNUNDER
(default LOCAL): RETURN
```

This question prompts for the name of the machine where the namespace database is stored. If the namespace server is not the local host, then you can provide the name of the namespace server for the site. **If the namespace server is the local host, do not type the host's name, instead use the default by pressing** RETURN.

```
What host is to be the SYS host for DOWNUNDER
(default LOCAL): RETURN
```

This question prompts for the name of the machine where Symbolics-supplied

source and documentation files are to be stored by default. If the SYS host is not the local host, then you can provide the name of the SYS host for the site. For more information: See the section "Logical Pathnames and the SYS Host".

```
What Symbolics computer will store the namespace data
files for DOWNUNDER (default LOCAL): RETURN
```

This question prompts for the name of the machine that stores the namespace data files. Although this machine must be a Symbolics computer, it does not have to be the same machine that is the namespace server, but we strongly suggest that it be the same machine.

```
What host is to be used for bug reports for DOWNUNDER
(default LOCAL): RETURN
```

This question prompts for the name of the machine that receives bug reports. This will most likely be the machine that you use for your mailer. For information about the mailer: See the section "Installing and Configuring the Mailer", page 203.

```
What is the real name of the local host: koala
```

This question prompts for your machine's name at the site. Machine names can be different at different sites. In this example, *koala* is the name of the local host. For more information: See the section "Why do you name machines and printers?" in *Genera User's Guide*. The default values for the namespace server, the SYS host, the computer storing the namespace data files, and the host for bug reports, are all the local host. If you provide a non-local host for any of these, the computer prompts for the non-local chaosnet address(es) and operating system(s)

```
What directory on KOALA will hold the namespace data files
(default >sys>site>): KOALA:>sys>site>
```

This question prompts for the location of your namespace data files. These files are accessed each time the site is set. We strongly suggest that you use the default value to avoid possible confusion.

```
What directory on KOALA corresponds to SYS: SITE;
(default >sys>site>): KOALA:>sys>site>
```

This question prompts for the directory that holds your systems translations files. These are necessary for logical pathname translations. We strongly suggest that you use the default value to avoid possible confusion.

```
What account should be used for the system to login to KOALA
(default: LISPM): wombat
```

This question prompts for a name the system uses automatically when it needs to access files. In this example we are using the name wombat. Do not supply your own name or that of another real user. When the system needs to log in automatically (for example, to read the namespace data files for the first time) the namespace server logs in with this user name.

```
What is the local timezone [default EST]: EST
```

This question prompts for the timezone. The default is Eastern Standard Time (EST) or Eastern Daylight Time (EDT) depending on whether Daylight Savings Time is in effect. For more timezone information: See the section "Specifying a

Time Zone for Your Site".

```
Is DOWNUNDER a standalone site (there are no servers to respond
to a Who-am-I request)? (Y or N) Yes.
```

This question prompts for whether your site is a standalone site consisting of one Symbolics machine, or if it has many Symbolics machines.

Several notification messages follow and confirm the definition of your site.

You are now ready to login and use the newly configured software. The site of your local host is automatically set at the newly defined site. You may want to save the new, site-specific version of your world. For more information: See the section "Customizing and Saving Worlds", page 14.

11. Using Tape Facilities on Symbolics Computers

Symbolics supports different tape formats for different purposes. Each format is specific to one Lisp Machine tool. Many people wonder which tool is appropriate for which application.

The *LMFS dumper and reloader* are used for backing up files from a local Lisp Machine file system and reloading those files on the same local Lisp Machine file system, in the same place, at a later date. The intended use of LMFS backup is to reload files onto the same machine from which they were dumped.

The *Distribution dumper and loader* are intended to distribute transportable systems and libraries, defined by system declarations on logical hosts, from one site to another. This tool is best used when transporting many files within a system, rather than transporting just a few unrelated files. The distribution system specializes in finding appropriate source and object versions of systems, appropriate patch files, and so forth. Its use of logical pathnames allows it to create a tape which can be easily loaded into the filesystem of a foreign host.

TAPEX is a format for transferring character (source) files between hosts. It is the only one of these formats which can be read or written by other than a Lisp Machine. TAPEX programs exist for TOPS-20 and Multics, as well as for the Lisp Machine (**tape:tapex**). It cannot deal with any type of file except character files, which are written in standard ASCII. Each individual dump or load requires an interaction.

Carry tape is the most general tool, and is used to dump individual files and sets of files (specified by wildcard filespec) and load them at any site. Recent improvements to the wildcard facility make this a very powerful and easy-to-use tool.

FEP Tapes:

- *IFS tape* is the Initial File System tape. A different tape is shipped with each disk, and is used in case of dire emergency to restore the structure of the FEP file system. This tape should not be used without consulting with Software Support. The FEP reads the IFS tape, and discards all of the data on the disk. The FEP then creates a new root directory, and **empty** files to hold world loads and microcode files. After you initialize the disk with the IFS tape, you read a FEP tape containing the world loads and microcode files.

- *FEP-tape* is a format which allows the user to read and write world loads and microcode files both from and to cartridge tape. Tapes written with this program can be read with the FEP commands Load Microcode CART: and Disk Restore.

11.1. Tape Facilities and Their Uses

Here are the possible tape facilities you can use with a Symbolics computer. Below are the purposes, restrictions, and capabilities of each tape.

Note: The 3600 uses a Cypher cartridge tape drive, which is capable of writing cartridge tapes of only 20 megabytes. The other machine models use an Archive cartridge tape drive, which is capable of writing cartridge tapes of more than 20 megabytes. If you write a cartridge tape containing more than 20 megabytes (on any machine model excluding the 3600) you will not be able to read it on a 3600. Since only two tape utilities (see chart below) currently limit the data written to tape to below 20 megabytes, you should be aware of this potential incompatibility.

If you use a machine which does not have a local tape drive, make sure that the machine with the tape drive has the **rtape** option in its namespace object. For further information see: See the section "Registering a Tape Drive in the Namespace", page 79.

LMFS

Purpose: Backing-up the file system on Symbolics computers.
Restrictions: Use only for LMFS files; may only restore to Symbolics computers.
Tape drive must be local? No
Can span multiple tapes? Yes
Is tape verifiable? Yes
Can write more than 20 megabytes per cartridge tape? Yes. The default is to write as much as the cartridge tape drive in use allows.
Because of this, if you write a LMFS cartridge tape on an Archive
cartridge tape drive, and it is longer than 20 megabytes, it cannot be read
on a Cypher cartridge tape drive.
Tapes readable by Lisp or the FEP? Lisp

Distribution

Purpose: Distributing software which is layered (not world loads).
Restrictions: Requires definition of logical pathnames.
Tape drive must be local? No
Can span multiple tapes? After a fashion. If you create multi-tape distributions, each resulting tape is an individually restorable single tape. You must enforce correct ordering when restoring the tapes.
Is tape verifiable? No
Can write more than 20 megabytes per cartridge tape? No, unless you bind the default value of the special parameter, **(tape:+cart-4track-max-length+)**, to a higher value.

Note that the default limits all Distribution cartridge tapes to less than 20 megabytes, so that all tape drives can read tapes written by Distribution.
Tapes readable by Lisp or the FEP? Lisp

Carry

Purpose: Moving files between hosts.
Restrictions: More difficult to use when many pathnames are needed.
Tape drive must be local? No
Can span multiple tapes? No
Is tape verifiable? Yes
Can write more than 20 megabytes per cartridge tape? Yes. The default is to write as much as the cartridge tape drive in use allows.
Tapes readable by Lisp or the FEP? Lisp

FEP-Tape

Purpose: Writing and reading world load and microcode files.
Restrictions: World load or microcode files only.
Tape drive must be local? No, if reading or writing a tape from Lisp; Yes, if reading a tape from the FEP.
Can span multiple tapes? Yes
Is tape verifiable? Yes
Can write more than 20 megabytes per cartridge tape? Optionally.
The default is to write cartridge tapes of less than 20 megabytes.
Tapes readable by Lisp or the FEP? Lisp and FEP.

TAPEX

Purpose: Making tapes in format for compatibility with TOPS-20 and Multics systems.
Restrictions: Use only for Source files.
Tape drive must be local? No
Can span multiple tapes? No
Is tape verifiable? Yes
Can write more than 20 megabytes per cartridge tape? Yes

11.2. Distribution Subsystem

The distribution subsystem is used at Symbolics to write and restore distribution tapes. Customers use the subsystem mostly for loading the distribution tapes they receive from Symbolics; that is, they restore the files on these tapes to their own file systems. However, users can also use this facility to write their own tapes.

The distribution tape facility runs on a Lisp Machine and requires a tape drive. The tape drive can exist on that machine, or it can be part of another machine that is accessible through the network and has a remote tape server program.

To restore data from tape to disk, use the Restore Distribution command. To write systems to tape for distribution, use the Distribute Systems command. Note that the distribution subsystem is limited to distributing the .newest version of each subsystem.

Distribute Systems Command

Distribute Systems *systems keywords*

Writes systems to tape for distribution. Expects the input to be one or more systems. The default is Null. After you confirm the command, Distribute Systems lists the systems to write to tape, and asks if you want to perform the operation. Your choices are Y, N, Q, or S. Type `Y` for Yes, `N` for No, `Q` for Quit, or `S` for Selective. If you choose Selective, each file is listed, and you are asked if you want to distribute that particular file. You can select as many or as few files as you want. After you enter this information, you are prompted for a tape spec.

systems is a list consisting of items separated by commas, each item being either (1) a system name or (2) a system name followed by a space and a version number.

keywords :Default Version, :Distribute Patch Sources, :File Types, :Include Components, :Include Patches, :Included Files Checkpoint, :Menu, :Output Destination, :Query, :Source Category, :Use Cached Checkpoint, :Use Disk

:Default Version {Released, Latest, Newest, *version-designator*} Version of the system to distribute if not individually specified in Systems. The default is Released.

:Distribute Patch Sources
{Yes, No} Whether to include patch sources for system patches. The default is No. The mentioned default is Yes.

:File Types {Sources, Binaries, Both, Patches-Only, Default} What file types to distribute. The default leaves it to the specifications in individual DEFSYSTEMs.

:Include Components
{Yes, No} Whether to include any component systems of the systems being distributed. The default is Yes.

:Include Patches {Yes, No, Selective} Whether to include the patch files for the systems being distributed. The default is Yes.

:Included Files Checkpoint
{Patch, Release, None} Limit distributed files to those after this patch number or release name, or None (do not limit). The default is None.

:Menu {Yes, No} Whether to use a menu interface to specify details of the distribution. Choosing Yes presents a Distribute Systems frame to select which files are distributed. For detailed

information about this frame: See the section "Distribute Systems Frame", page 163.The default is No. The mentioned default is Yes.

:Output Destination

{Buffer, File, Printer, Stream, Window} Where to redirect the typeout done by this command. The default is the stream ***standard-output***.

:Query

{Everything, Yes, Confirm-Only, No} Whether to ask about distributing each file. The default is Confirm-Only. The mentioned default is Everything. Everything queries you about each file being distributed, and again for each system being distributed. Yes queries you about each file being distributed. Confirm-Only queries you about each system being distributed. No does not query.

For queries about individual files, the possible responses are:

Y Yes, distribute this file.
N No, do not distribute this file.
I Include the remaining files in this system.
B Bypass (do not include) the remaining files in this system.
D Directory. Show the directory containing this file.
E Edit this file.
S Source compare this file.

For queries about systems, the possible responses are:

Y Yes, distribute all files listed in this system.
N No, do not distribute any files listed.
Q Quit. Do not distribute any files listed in this system.
S Selective. Query about each file in this system individually.

:Source Category {Basic, Optional, Restricted, Optional-only, Restricted-only} Indicates which source category or categories to write to tape for distribution. The default is Basic.

:Use Cached Checkpoint

{Yes, No} Use the last checkpoint gathered for this system. Using the cached checkpoint information, if there is any, saves time. But it is safe to use only if you are sure no more patches have been made since the cached information was computed. The default is No. The mentioned default is Yes.

:Use Disk

{Yes, No} Writes the input to disk (test mode), rather than to tape. Test mode writes a special file which is an image of what would be written to tape. Use this when you are preparing a distribution and want to see what files would be written to tape. The default is No. The mentioned default is Yes.

11.2.0.1. Distribute Systems Frame

When you use the Distribute Systems command, and specify the keyword :Menu with the value Yes, you invoke the Distribute Systems frame. This frame enables you to define and edit the specifications for one or more systems to be written on a distribution tape. You can add specifications for new systems one at a time, make changes to the details of the existing systems, or delete them entirely. You can also access the Distribution Systems frame by using the command Select Activity Distribution System.

Figure 9 shows the Distribute Systems frame.

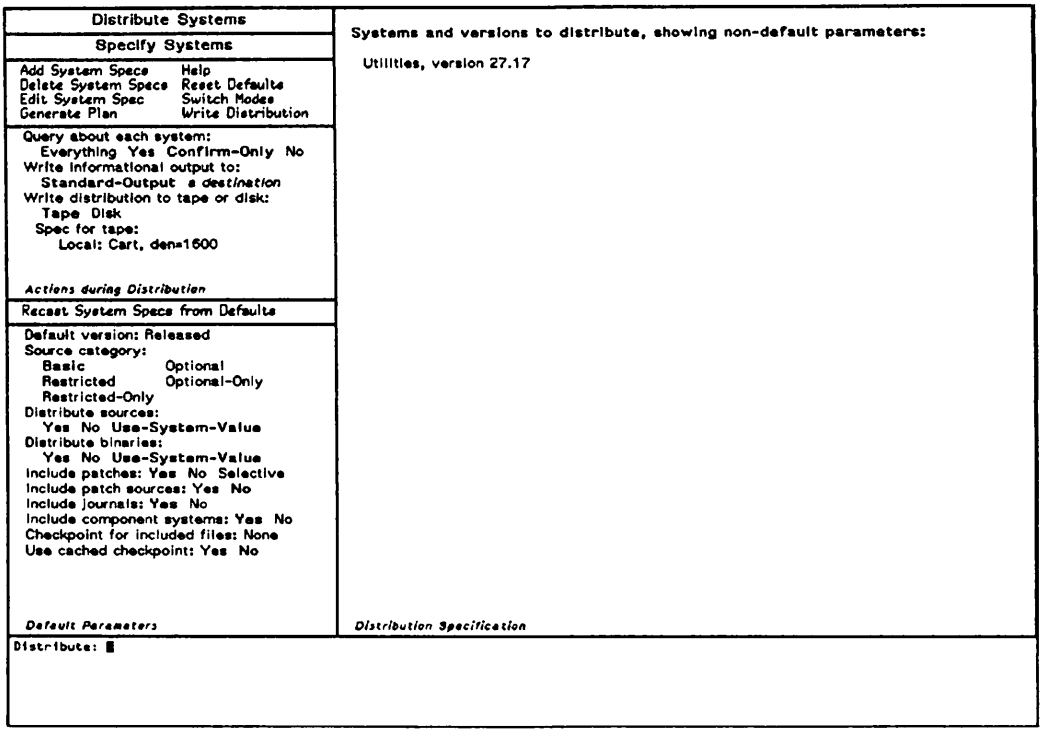

Figure 9. Distribute Systems Menu

Overview:

The entire Distribution System frame consists of six different sections: one distribution display specification pane, one pane indicating your status in the process of distributing systems, two command menus, and two panes for setting parameters.

While using the Distribute Systems frame, you are in different phases of the process, depending on which activity are performing. The top-left pane always displays the name of the activity you are performing. Initially, the top-left pane displays the words Specify Systems, since this is the first activity you perform.

The first phase occurs when you make your system specifications. These specifications appear in the large Distribution Specification pane if you enter the Distribute Systems frame with the Distribute Systems command, and specify the keyword :Menu, with the value Yes. When you are in this first phase, the top-left pane dis-

plays the words Specify Systems, to indicate your activity in the process of distributing systems.

After you finish specifying systems, you click on the Generate Plan command in the command menu. The top-left pane then quickly displays the words Generating Plan.... The distribution plan then appears in the Distribution Specification pane, and the words Editing Distribution appear in the top-left pane. This signifies that you are now in the phase of editing the distribution plan.

Then, when you are satisfied with the distribution plan, you click on Write Distribution in the command menu, and the words Writing Distribution appear in the top-left pane, signifying that that you are writing the distribution. While writing the distribution, the process prints the progress log in a typeout window, and also in another location if you specify it to do so.

Here is a description of the different panes:

Distribution Specification Pane -- the right-side pane

This pane displays the system specs, as you have defined them up to now.

Specify Systems Pane -- the top-left pane

The top-left pane always displays the name of the phase you are in. The name changes as you move from one activity to another.

The Command Pane -- the second-left pane

The second-left pane offers a command menu for setting parameters of the systems you want to distribute. You can choose one of the commands by clicking on the command name.

Here is an explanation of these commands.

Command	*Description*
Add System Specs	Adds specifications for new systems.
Delete System Specs	Deletes a specification in the list of systems. You can specify individual systems or **all**.
Edit System Spec	Enables you to edit the detailed parameters of an individual specification already in the list of parameters.
Generate Plan	Generates the plan so you can examine it before writing the final distribution. Computes the exact list of files to write, according to your specifications. When it finishes, it switches to the Edit Distribution Plan phase, and displays the list of files.
Help	Gives information about Specify Systems commands.
Reset Defaults	Enables you to reset the actions and default parameters to their initial values.

Switch Modes — Enables you to switch back and forth between the Specify Systems and the Edit Distribution Plan phases. (This command does not change the state of the frame in any other respect).

Write Distribution — Creates the distribution. Clicking on this writes the distribution to the file or tape device specified in the Actions during Distribution pane.

To add specifications for new systems, you can use Add System Specs in the Specify Systems menu, or you can click left on the title line of the system spec display. The parameter values initially offered in the pop-up editor are from the current set of Default Parameters.

To edit the detailed parameters of an individual specification already in the list, you can use the Edit System Spec in the Specify Systems menu, or you can click left on the displayed line for that system specification.

To delete specifications for a system you can use Delete System Specs, or you can click middle on the displayed line for that system specification.

Actions during Distribution Pane -- the third-left pane

The third-left pane is a parameters pane for changing the settings that control the way Distribute Systems writes the distribution, such as where to direct the information lines that tell which files are being written. You can change one of these parameters by clicking on the value.

This menu enables you to set four parameters:

Query about each system:

Presently this parameter is not needed by the Distribute Systems frame. It is included for use in future enhancements to this frame.

Write Informational output to:

Standard-Output — {Buffer, File, Printer, Stream, Window} Where to direct the typeout done by this command. The default is to write the information on a typeout window on the screen.

Write distribution to tape or disk:

Tape or Disk. — If you choose tape, it prompts you for the tape spec:

```
Spec for tape: Local: Cart, den=1600
```

If you choose disk, asks for the pathname for a dummy tape file on disk.

Recast System Specs from Defaults Pane -- the fourth-left pane

The fourth-left pane is a command menu for setting the parameters of the systems to distribute. You would click on the command in this pane after you have changed some parameters in the bottom-left pane, Default Parameters.

This command is useful if you have changed some system specs in the Distribution Specification frame, and you wish to revert all the specs to the defaults displayed in the Default Parameters pane. When you click on Recast System Specs from Defaults, it changes all the specs for each system to the ones specified in the Default Parameters pane.

Default Parameters Pane -- the bottom-left pane

The bottom-left pane is a parameters pane, containing the default values for the details of how Distribute Systems write each system. When you create a new system spec, Distribute Systems uses these defaults as the initial settings of the parameters for that system spec. Once you create a system spec, it contains a full set of parameter values. These parameter values are displayed in the Default Parameters pane. Note that the Distribution Specification pane only displays the settings that do not correspond to the default parameters. If you change a default parameter, the display of the parameters of the existing system specs changes to reflect that one of their values is no longer the default value. You can change one of these parameters by clicking on the value.

This menu enables you to set ten parameters:

Default Version:

The system version, one of the following:

Released	The version designated as released in the journal file.This is the default.
Latest	The most recent version recorded in the journal file.
Newest	Newest means to ignore the versions in the journal file and just find the newest files.
A positive integer	The version of the system you want to distribute.

Source Category:

Basic	Distribute only sources marked Basic.
Restricted	Distribute only sources marked Restricted, Optional, or Basic.
Restricted-Only	Distribute only sources marked Restricted.
Optional	Distribute only sources marked Optional, or Basic.
Optional-Only	Distribute only sources marked Optional.

Distribute Sources:

Yes	All sources in the specified source category distributed.
No	No sources distributed.
Use-System-Value	Use the parameter set in the DEFSYSTEM for each system.

Distribute Binaries:

Yes	All binary files included.
No	No binary files included.
Use-System-Value	Uses the parameter set in the DEFSYSTEM for each system.

Include patches:

Yes	All patches included.
No	No patches included.
Selective	Prompts you for which patches to include.

Include Patch Sources:

Yes	All patch sources included.
No	No patch sources included.

Include journals:

Yes	All journals included.
No	No journals included.

Include component systems:

Yes	All component systems included.
No	No component systems included.

Checkpoint for included files: Use cached checkpoint:

Yes	Use the last checkpoint gathered for this system. Using the cached checkpoint information, if there is any, saves time. But it is safe to use only if you are sure no more patches have been made since the cached information was computed.
No	Do not use the last checkpoint gathered for this system.

After Setting All of the Parameters

After you set all of the parameters to the desired values, use the Generate Plan command in the command menu, to compute the exact list of files to write, according to your specifications. When it finishes, it switches to the Edit Distribution Plan phase, and displays the list of files. You can edit the distribution plan it you like. When you are done editing the distribution plan, or are satisfied with it, click on the Write Distribution command to write the distribution.

How to Scroll the Screens

In the screens with scroll bars, you can scroll the screen display with the SCROLL key, or by typing m-SCROLL. In addition, you can use m-< to move to the beginning of the display, and m-> to move to the end of the display.

Restore Distribution Command

Restore Distribution *keywords*

Restores data from tape to disk. This command reads the directory listing of the files on tape, and then restores the files according to the pathname and property information on the tape. For each file it restores, Restore Distribution checks to see if the file already exists in the file system. If it does not exist, it restores the file. If the file does exist, the default behavior is that Restore Distribution states that the file already exists in your file system and skips the file.

keywords :Menu, :Output Destination, :Skip Restoration, :Use Disk

:Menu {Yes, No} Whether to use a menu interface to specify details of the restoration. Choosing Yes presents a Restore Distribution frame to select which files are restored. The default is No. The mentioned default is Yes.

:Output Destination {Buffer, File, Printer, Stream, Window} Where to redirect the typeout done by this command. The default is the stream ***standard-output***.

:Skip Restoration {Yes, No} Skips restoration of files that already exist in the file system. If the file exists in the system and :Skip Restoration is No, Restore Distribution states that the file it was supposed to restore already exists in your file system, and asks where to restore it instead. You supply a pathname. A pathname of Nowhere means skip this file. The default is Yes. The mentioned default is Yes.

:Use Disk {Yes, No} Reads the input from disk (test mode), rather than from tape. The default is No. The mentioned default is Yes.

11.2.0.2. Restore Distribution Frame

When you use the Restore Distribution command, and specify the keyword :Menu with the value Yes, you invoke the Restore Distribution frame. This frame enables you to examine the list of systems and files on a distribution tape, select the ones you want, and then causes them to be restored to your file system. You can also enter the Restore Distribution frame using the command Select Activity Restore Distribution.

The entire Restore Distribution frame consists of four different sections: the Restore Distribution Command menu, the Files to Restore frame, the Actions during Restore Distribution pane, and the Systems to Restore pane.

Figure 10shows the Restore Distribution frame.

Restore Distribution Pane -- the top-left pane

The top left pane offers offers three commands. To invoke the command, click on the command name.

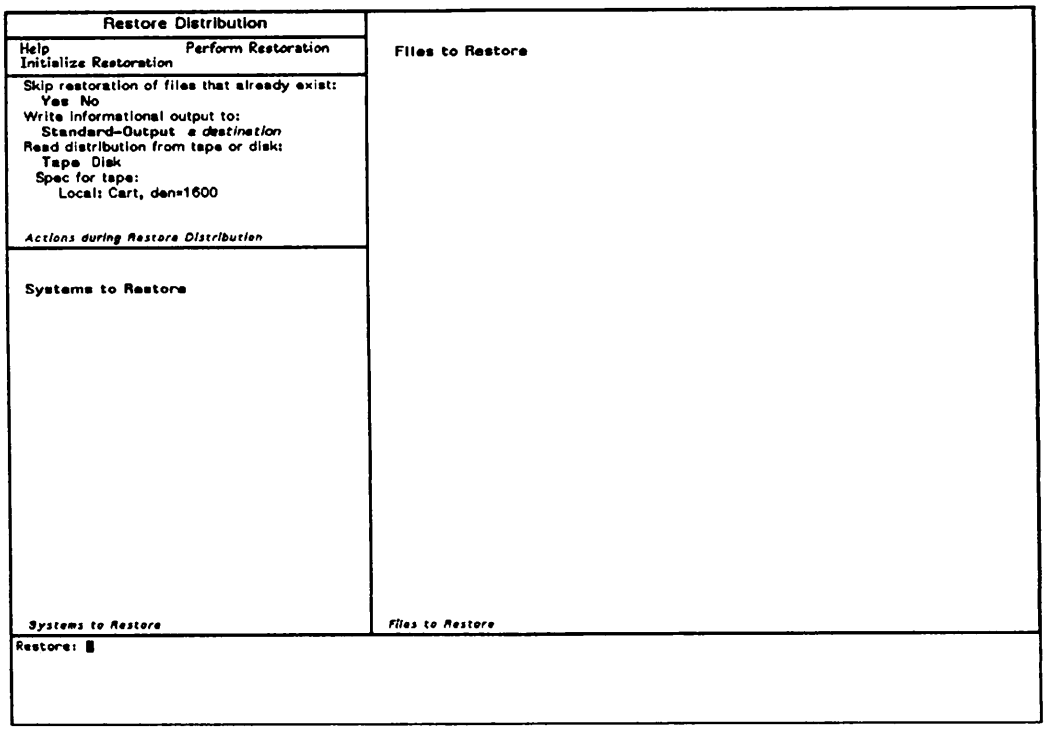

Figure 10. Restore Distribution Frame

Help — Displays information about using the Restore Distribution frame.

Initialize Restoration — Reads the directory of all the systems and files on the tape and displays them for your inspection.

Perform Restoration — Restores all files selected in the Files to Restore pane.

Files to Restore Pane -- the right-side pane

The right-side pane displays a list of files, with system header lines.

Actions during Restore Distribution -- the middle-left pane

The middle-left pane enables you to set the parameters for the restoration, such as where to print information about the restoration.

You can change one of these parameters by clicking on the value.

This menu enables you to set four parameters:

Skip restoration of files that are already exist:

Yes — Restores all files, even if that file already exists on disk.

No — Does not restore a file if that file already exists on disk.

Write Informational output to:

Standard-Output — {Buffer, File, Printer, Stream, Window} Where to direct the typeout done by this command. The default is to write the information on a typeout window on the screen.

Read distribution from tape or disk:

Tape or Disk. — If you choose tape, it prompts you for the tape spec:

```
Spec for tape: Local: Cart, den=1600
```

If you choose disk, asks for the pathname for a dummy tape file on disk.

Systems to Restore pane -- the bottom-left pane

The bottom-left pane lists the systems on the tape.

Initializing a Restore Distribution

To begin the Restore Distribution process, give the specification for the tape device in the *Actions during Restore Distribution* pane. Then, click on the command Initialize Restoration to cause the frame to read the directory on the distribution tape.

Controlling Selection of Systems and Files

Initially, Restore Distribution selects all systems and files for restoration. You can control the selection of a file by clicking left on it to toggle whether it is selected. When a file is deselected, it displays in a smaller font.

You can control the selection of a system by clicking left on it in the Systems to Restore pane. When a system is deselected, it is displayed in a smaller font, and its files are completely removed from the file display.

You can deselect all the files of a system (without deselecting the system itself) by clicking middle on the header line for that system in the file display. You can select all the files of a system by clicking left on the header line. This kind of operation is useful, for example, for selecting just a few of a large number of files in a system by first deselecting them all, and then clicking on the few desired files.

You can deselect all of the systems on the tape by clicking middle on the title line in the Systems to Restore pane. Select individual systems by clicking left on the individual systems names. This operation is useful for selecting just a few of a large number of systems.

After Selecting The Systems and Files You Want

When you have selected the correct systems and files (all the files you want to restore are displayed in large letters) click on the command Perform Restoration. The frame checks the distribution tape to make sure it is the same one from which it was initialized. Then it reads the selected files and restores them to the file system.

While the restoration executes, the informational lines that the operation produces print out on the typeout window over the large pane. If you give some other output destination for the informational output, Restore Distribution writes the output to that place, as well.

How to Scroll the Screens

In the screens with scroll bars, you can scroll the screen display with the SCROLL key, or by typing m-SCROLL. In addition, you can use m-< to move to the beginning of the display, and m-> to move to the end of the display.

Show Distribution Directory Command

Show Distribution Directory *keywords*

Describes the contents of a distribution tape, or a dummy distribution file. For each file in the distribution, this command displays the file's logical pathname, and the physical pathname to which it would be restored.

keywords :Output Destination, :Use Disk

:Output Destination {Buffer, File, Printer, Stream, Window} Where to redirect the typeout done by this command. The default is the stream ***standard-output***.

:Use Disk {Yes, No} Whether to find the directory on disk (test mode) rather than from tape. The default is No. The mentioned default is Yes.

```
Distribution listing, 12/01/87 18:08:31

Systems on this tape:
  LMFS, TAPE, UTILITIES, SERVER-UTILITIES, HARDCOPY

Files on this tape:                                    Translated pathnames:
 In system LMFS:
    1  SYS:LMFS;LMFS.LISP.105                          Q:>rel-7>sys>lmfs>lmfs.lisp.105
    2  SYS:LMFS;PATCH;LMFS.SYSTEM-DIR.83               Q:>rel-7>sys>lmfs>patch>lmfs.system-dir.83
    3  SYS:LMFS;PATCH;LMFS-94.COMPONENT-DIR.1          Q:>rel-7>sys>lmfs>patch>lmfs-94>lmfs-94.component-dir.1
    4  SYS:LMFS;PATCH;LMFS-94.PATCH-DIR.1              Q:>rel-7>sys>lmfs>patch>lmfs-94>lmfs-94.patch-dir.1
 In system TAPE:
    5  SYS:LMTAPE;TAPPKG.LISP.47                       Q:>rel-7>sys>lmtape>tappkg.lisp.47
    6  SYS:LMTAPE;TAPE.SYSTEM-DIR.127                  Q:>rel-7>sys>lmtape>tape.system-dir.127
    7  SYS:LMTAPE;TAPE-74.COMPONENT-DIR.1              Q:>rel-7>sys>lmtape>tape-74>tape-74.component-dir.1
    8  SYS:LMTAPE;TAPE-74.PATCH-DIR.3                  Q:>rel-7>sys>lmtape>tape-74>tape-74.patch-dir.3
    9  SYS:LMTAPE;TAPE-74-1.BIN.1                      Q:>rel-7>sys>lmtape>tape-74>tape-74-1.bin.1
 In system UTILITIES:
   10  SYS:SYS;UTILITY-SYSDCL.LISP.51                  Q:>rel-7>sys>sys>utility-sysdcl.lisp.51
   11  SYS:PATCH;UTILITIES.SYSTEM-DIR.55               Q:>rel-7>sys>patch>utilities.system-dir.55
   12  SYS:PATCH;UTILITIES-20.COMPONENT-DIR.1          Q:>rel-7>sys>patch>utilities-20>utilities-20.component-d
                    .
                    .
                    .
```

Figure 11. Part of the display from Show Distribution Directory

dis:copy-distribution-tape &key *:use-disk* *Function*

Copies an entire distribution tape. An Accept-Variable-Values menu appears, and you can specify a tape to copy, and a destination. The distribution to be copied can be either a tape or the pathname of a dummy source tape on a disk. The destination must be a tape device.

:use-disk specifies to copy from a dummy distribution file rather an actual tape. The file specified must have been created by using Distribute Systems :use disk Yes.

This function copies only tapes in the Genera 7 Distribution format. If you try to copy a Release 6 Distribution-format tape, it prints a special message with suggestions. If you attempt to copy a tape that is in a format other than Genera 7 or Release 6 Distribution format, this function prints an explanatory comment about the tape.

11.3. Carry-Tape System

The carry-tape system provides a means of dumping selected files or sets of files to magnetic tape (cartridge or industry-compatible) and loading them at a later time, possibly at a different site. Using the Carry-Tape System, you can dump files from any host or set of hosts, and reload them to any place on any host.

The carry-tape system provides a standard, system-independent interchange medium for exchanging single programs and files between sites. It is meant to fill in a gap between the LMFS backup dumpers and the distribution tape system. It does not require you to prepare files or declarative forms in advance.

The carry-tape system has three components:

- The carry-tape dumper
- The carry-tape loader
- The carry-tape lister

11.3.1. The Carry-Tape Dumper

tape:carry-dump *file-or-files* &key *tape-host density reel* (*report* **t**) *Function*

Dumps a file or set of files to a carry-tape. You can dump any type of file. Character files are dumped and reloaded using the Genera character set as an interchange medium. Binary files are dumped and reloaded with the proper byte size as long as either of the following is true:

- The file is of one of the system's known canonical types.
- The operating system on which the file resides knows and can supply the byte size.

file-or-files — a pathname, filespec, or list of pathnames and/or filespecs. Wildcard pathnames or filespecs may be used. Recursive ("accordion") wildcards may be used to dump subtrees on those hosts that support them. An example of a pathname which has recursive wildcards is:

```
E:>trees>**>*.*.*
```

tape-host — a host object or the name of a host object to use for tape access. **:local** specifies the local tape drive. If you do not specify a host, the dumper uses the standard tape host prompt and defaulting mechanism.

density — a fixnum, specifying tape density, which may be used when the applicable default is not appropriate.

reel — can be a string, specifying tape reel name for tape servers that need this information (none of the currently supported ones do).

report — tells the carry-tape dumper to report its progress as it dumps files. A value of **nil** tells it not to. A value of **t** tells it to report. The default is to report to ***standard-output***. Any value besides **nil** or **t** is expected to be a stream to which the reports will be written.

Currently, carry dumps must fit on one tape.

The carry-tape dumper starts by finding out all available information about the files to be dumped, verifying their existence. It then asks for confirmation, and proceeds to dump all the files specified, without intervention.

Here is an example of using the Carry-Tape Dumper:

```
(tape:carry-dump "swanee:>minerals>*.d*")
To be dumped:
swanee:>minerals>*.d*: 7 files
Is this right? (Y or N) Yes.
Type name of tape host (default (CR) = POINTER): scrc
Tape mounted on drive mta0:.
Dumping swanee:>minerals>abel.data.3 (5-bit bytes)
Dumping swanee:>minerals>abel.patch-directory.7
.....
End of dump.
```

11.3.2. The Carry-Tape Loader

The carry-tape loader loads files from a carry-tape. The loader makes no attempt to copy any file properties, including author and creation date. It copies only file contents, and provides reasonable defaults for the target file name.

tape:carry-load &key *host density reel* (*report* **t**) *Function*

host a host object or the name of a host object to use for tape access. **:local** specifies the local tape drive. If you do not specify a host, the loader uses the standard tape host prompt and defaulting mechanism.

density a fixnum, specifying tape density, which can be used when the applicable default is not appropriate.

reel can be a string, specifying tape reel name for tape servers that need this information (none of the currently supported ones do).

report tells the carry-tape loader to report its progress as it loads files. A value of **nil** tells it not to. A value of **t** tells it to report. The default is to report to ***standard-output***. Any value besides **nil** or **t** is expected to be a stream, to which the reports will be written.

These arguments are rarely needed.

The carry-tape loader begins its operation by reporting the pathnames given to the dumper, and asks if you wish to load all of the files dumped. If only one filespec or pathname was given, it is assumed that you want to load it all, and no question is asked:

```
(tape:carry-load)
Type tape host or spec (default (CR) = APS0): beta
Tape mounted on drive mta0:.
Carry dump made by DCF.
Dump taken at 6/13/86 09:05:22.
Dumped on machine EAGLE.
Dumped: e:>trees>apple.orchard
```

The set of files dumped as a result of each pathname given to the dumper is called a *group*. If many groups were dumped, the loader lists the pathname of each group at the start of its operation, and asks for instructions about which groups are to be loaded (selectively) and which groups are to be skipped:

```
The following groups of files were dumped:
e:>trees>apple.orchard
e:>animals>whales>tails.tales
e:>baseball>runs>foul.*
-----
Load all these files? (ABORT to get out) (Y, Q, or M)
```

The possible responses are:

`Y` Yes Ignore distinctions of group, and proceed as described below.

`Q` Query Query about each group. The options are:

	N No	Don't load the group.
	Y Yes	Load the group.
	P Proceed	Load this and all succeeding groups.
	Proceed as below for those groups that are selected for loading.	
M Menu	Same as Q, but present a multiple-choice menu instead of querying for each group.	

If you do not want to load anything, you can press ABORT at any time to stop the loader.

The carry-tape loader can either query for the target location of each file to be loaded, or proceed in semi-automatic mode, in which the host and directory from which each file was dumped are used as a key to target loading of subsequent files from that host and directory. The name, type, and version of each file to be loaded are developed automatically from the name, type, and version of the file that was dumped, by means of the same mechanism used by ordinary file copying.

The normal action of the carry-tape loader is to query for each file, with a query of the following form:

```
Load SWANEE:>minerals>rock5.data.6 into BULLWINKLE:/usr2/jones/rock5.data?
        (Y, N, O, or A)?
```

The following responses apply:

Y, SPACE	Load the file into the place specified. The host and directory shown remain the default target directory for all files from this host and directory at the site writing the tape.
N	Do not load this file at all. The host and directory shown remain the default target directory for all files from this host and directory at the site writing the tape, in spite of this.
O	Prompt for another place in which to put this file. The host and directory into which this file is then loaded become the default for all subsequent files from the same host and directory at the site writing the tape. You are queried again at the time such files are encountered.
A	Load the file into the place specified. All further files from the same host and directory at the site writing the tape are then automatically loaded into the same host and directory as this file without querying you.

11.3.3. The Carry-Tape Lister

The carry-tape lister describes what is on a carry-tape. Once started, it does not interact in any way.

tape:carry-list &key *tape-host density reel* (*report* **t**) *verbose* *Function*

host	a host object or the name of a host object to use for tape access. **:local** specifies the local tape drive. If you do not specify a host, the lister uses the standard tape host prompt and defaulting mechanism.
density	a fixnum, specifying tape density, which may be used when the applicable default is not appropriate.
reel	can be a string, specifying tape reel name for tape servers that need this information (none of the currently supported ones do).
report	tells the carry-tape lister to report its progress as it lists files. A value of **nil** tells it not to. The default is to report to ***standard-output***. Any value besides **nil** or **t** is expected to be a stream, to which the reports will be written.

These arguments are rarely needed.

11.4. FEP-Tape System

The FEP-Tape system provides a means of writing and reading the cartridge tapes used to distribute world loads and microcode files. World loads and microcode files are distributed using the FEP-tape format because the FEP commands "Load Microcode CART:" and "Disk Restore" can only read this format. Consequently, tapes written with the FEP-tape system can be read by a machine that does not have a working lisp world, although ordinarily they are read by Lisp. The only requirement is that the machine have an initialized FEP file system and a working set of FEP overlay files. For more information about the FEP system and overlay files: See the section "Introduction to the FEP", page 80.

11.4.1. The Contents of a Tape Made with the FEP-Tape System

Tapes made with the FEP-Tape system are created in *sets*. A FEP-Tape tape set is a series of one or more cartridge tapes. Ordinarily, a FEP-Tape set contains a set of microcode files and a world load file. There can be additional world-load files or other files on a tape set. However, there is no particular advantage in writing unrelated world load files onto the same tape.

The first tape in the set may begin with one or more *initial microcode files*. These are files written in cartridge tape stream mode which can be read by the FEP command Load Microcode CART:.

You only have to put initial microcode files on a tape if it has to be readable by a machine that does not have any microcode loaded. As of Release 7.0, there are no known reasons to write initial microcode files. At some time in the future they might be necessary.

The initial microcode files are usually followed by the microcode and world load files and can continue onto additional tapes. These files are read by the FEP command Disk Restore or by the FEP-Tape application itself. When the FEP reads these files, it restores each one into a previously existing FEP file. **The FEP cannot create new files.** Thus, in order to read a FEP-Tape from the FEP, you must either pre-create a suitable set of files while running Lisp, or write over some other existing files.

11.4.2. How Much Will Fit on a FEP-Tape Tape?

There are two kinds of cartridge tape drives on Symbolics computers, Cypher (4-track) and Archive (9-track). Some 3600s (only the 3600 model, not the series of 3600 machines) have Cypher four-track drives. These drives can read or write the first 20MBytes of a cartridge tape. By default, the FEP-Tape application writes tapes of no more than 19MBytes. This allows the tapes to be read by Cypher tape drives. The FEP-Tape application assumes that the tapes it reads were written for this type of drive.

Other types of Symbolics computers of the 3600 series, such as the 3640 and the 3670, have Archive nine-track cartridge tape drives. Currently, the FEP-Tape application writes or reads 39MBytes using an Archive tape drive, if, and only if, you give the optional argument to the Write Tape commands. The optional argument is *:Full Length Tapes Yes*. This 39MBytes is less than the full capacity of the cartridge tape.

You should only write 39MByte tapes when you are absolutely sure all of the machines that will read the tape have 9-track cartridge tape drives.

11.4.3. Invoking the FEP-Tape System

To invoke the FEP-Tape system, type this command to the Command Processor prompt:

```
Select Activity  FEP-Tape
```

This invokes a frame that consists of a command menu, a file display pane, and a listener pane with the prompt:

```
FEP-Tape Command:
```

The command menu frame lists all possible command options. You type commands in the listener pane at the FEP-Tape Command: prompt. When preparing to write a tape, the FEP-Tape application lists files to be written to the tape in the file display pane, as shown in Figure 12.

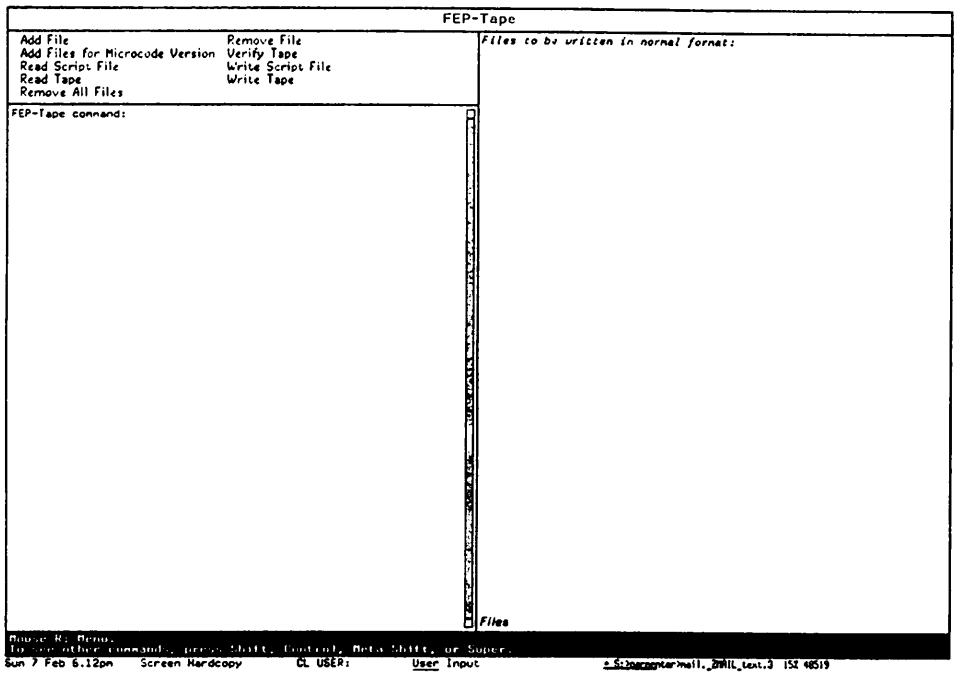

Figure 12. FEP-Tape Menu

The following command menu options are available:

Command	*Definition*
Add File	Adds a single file to the list of files to be written to tape. The command takes arguments that specify: • The pathname. • Whether the file should be written as an initial microcode file, an other file, or both. • A comment to be associated with the file. Only files "other" than the initial microcode files can have comments. Add Files for Microcode Version Adds all the microcode files for a particular version of standard microcode to the tape. The command takes arguments to specify the microcode version, whether to write row-major (Genera 7) and column-major (Release 6) array microcode, and whether to set up the initial microcode files.
Read Script File	Reads a script file containing a file set that you have prepared in advance. Use this command to read in a list of files that you have selected to write to tape.
Read Tape	Reads a FEP-Tape tape. This command prompts for the name of each file, and you specify where to put the file or if you want to skip reading it all together.
Remove All Files	Removes all of the files listed for writing to tape.
Remove File	Removes one file from the set of files listed for writing to tape.
Verify Tape	Reads a newly-written tape set and compares the content of files on disk to the files on the tape.
Write Script File	After specifying your tape-set, use this command to save the list of files in a script.
Write Tape	Writes the files to tape. This command takes arguments to specify the host with the tape drive to use, (default is local) and whether to write 4-track or 9-track tapes.

11.4.4. Writing a FEP-Tape Tape Set

1. **Select the set of files to be written on the tape.**

   There are two ways to do this. One way is to read in a script file containing a file set that you have prepared beforehand. You can use the command Read Script File to read the list of files.

Another way to select the files to be written to tape is to use one or more FEP-Tape commands to specify the files to be written on the tape. You should use the Add File command or the Add Files for Microcode Version command.

The Add File command adds a single file to the list of files to be written to tape. The command takes arguments that specify the pathname, whether the file should be written as an initial microcode file, an other file, or both, and a comment to be associated with the file. Only files other than initial microcode files can have comments.

The Add Files for Microcode Version command adds all the microcode files for a particular version of the standard microcode (non-Prolog) to the tape. The command takes arguments to specify the microcode version, whether to write row-major (Genera 7.0) or column-major (Release 6.0) array microcode, and whether to set up the initial microcode files.

If the script file containing a file set has some files you do not want on your FEP-Tape tape, use the Remove File or the Remove All Files command to remove some files. The Remove All Files command removes every file listed for writing to tape, and the Remove File command removes just one file form the set of files listed for writing to tape.

2. **Save the list of files in a script file.**

   You can do this with the Write Script File command.

3. **Write the files to tape.**

   Use the Write Tape command to do this. This command takes arguments to specify the host with the tape drive to use , (default is local) and whether to write 4-track or 9-track tapes. The argument to specify the host is *:host*, and the argument to specify the length is *:Full Length Tapes*.

4. **After writing the files to tape you can verify with the Verify Tape command.**

   This command reads a newly written tape set and compares the content of files on disk to the files on the tape.

11.4.5. Reading a FEP-Tape Tape

To read a FEP-Tape tape, use the Read Tape command. This scans the tape. For each file, it prompts with the name of the file, and you specify either where to put the file in the file system or that you want to skip reading it all together. This command takes the arguments *:Full Length Tapes* and *:host*.

The FEP command Disk Restore also reads FEP-Tapes.

12. Multiple Partitions

The Lisp Machine File System (LMFS) allows the use of multiple partitions residing on one or more disk drives. It utilizes one or more files of the FEP file system as the vessels in which it stores its files and directories. These FEP files are called *partitions*. Normally, there is one large partition, usually called LMFS.FILE.1. All the files created by LMFS actually reside inside this FEP file, but the existence of these files is known only to LMFS, whose purpose it is to manage them; they are not known to the FEP file system. Since FEP files are limited to one particular disk drive, if a LMFS file is to utilize the space available on multiple drives, partitions must be created on each drive on which it is desired that LMFS store files. Then, LMFS must be instructed to use these partitions.

The selection of partitions to be used by LMFS is determined by a database called the *file system partition table* (FSPT). It is contained in a FEP file named >fspt.fspt on a boot drive. The FSPT is optional. If it is not present, LMFS uses lmfs.file on the FEP boot drive. The FSPT is a simple character database containing the actual pathnames (in the FEP file system) of the partitions to be used for file system access.

If any machine at your site has more than one disk, it may be difficult to find the disk location of the FSPT. In order to make finding the location of a FSPT easy, insert the Set LMFS FSPT Unit command in your Hello.boot file. This command causes LMFS to look for the file named >fspt.fspt on the unit (disk unit) specified. For example, if you put your FSPT on disk unit 2, put the following in your Hello.boot file:

```
Set LMFS FSPT Unit 2
```

Each partition in the file system knows how many partitions make up the file system. Only the FSPT, which is used only at LMFS startup time, indicates the locations of these partitions. That is, the file system databases in the actual partitions do not contain drive and partition numbers or FEP pathnames. Thus, when LMFS is down, partitions can be moved around using Copy File (m-X); as long as the FSPT is edited to indicate their new locations, LMFS comes up (when required) using the moved partitions. **Note:** Since the Copy File (m-X) command copies files according to byte size, you may need to edit the byte count of the partition for the copy file command to work. To do this, multiply the number of blocks by 1152, since partitions were previously created with a byte size of 0. For example:

```
1152 * number of blocks
```

The FSPT is edited only to move partitions around or to add a partition. When you add partitions to the file system, the file system automatically rewrites the FSPT database to include the locations of new partitions.

12.1. Free Records

The basic unit of allocation in the Lisp Machine File System is the *record.* A record is 1152 32-bit words, or four disk blocks. Each file system object is made from an integral number of records. At any time, each record is *in use* (representing an existing file system object) or *free* (not representing anything and free to be used in new objects). When the file system needs to find a new free block to create or grow an object, it does not search through the records looking for a free one, because that would require many disk operations and be very slow. Instead, the file system uses a redundant data structure called the *free record map*, kept in several blocks in a known location in the file system partition. The map has one bit for each record in the file system; this bit marks whether the record is free or in use. The file system can find a free record quickly by examining this map.

If the file system crashes, or something else goes wrong, the contents of the free record map can become inconsistent with the contents of the file system itself. For each record, two different errors are conceivable.

- The record might actually be in use, representing part of an object, but marked as "free" in the map. The system is designed so that this cannot happen, but hardware problems might cause it to happen anyway.

- The record might actually be free, but marked as "in use" in the map.

The first error is much worse than the second; the file system might use the record for a new object even though it is currently representing some existing object, which could destroy the existing object. If the second error occurs, the record simply is not allocated even though it could be. Such a record is said to be *lost.*

The file system is written so that a crash can only cause the second kind of error. While the file system is operating, it maintains a *free buffer* in its data structures in virtual memory. The free buffer is a pool of records that are not actually in use, but are marked as being in use in the free record map on the disk. When it needs to allocate a record, it draws on one of these; when it frees up a record, it adds the record to this buffer. When the buffer gets too big, some records are removed from the buffer and marked as "free" in the map on the disk; when the buffer runs low, more records are marked as "in use" in the map on the disk, and are added to the buffer. So, if the machine is cold booted, or the file system crashes, the records that are in the buffer are lost, but no errors of the first kind are caused. The size of the buffer is maintained at about 30 records, so each crash loses 30 records. To recover, log out of the machine or use the [Flush Free Buffer] command to flush the entire free buffer and mark the records as "free" in the map on the disk. To use the [Flush Free Buffer] command, press SELECT F to enter the File System Editing Program. Click right on [Local LMFS Operations] to invoke the second level of the progam, where you can click on [Flush Free Buffer]. After the buffer has been flushed, you can cold boot the machine without losing any blocks.

Lost records can be found again by the salvager. See the section "Salvager", page 183.

You can check the number of free records in the file system by using the File System Editing Operations program. First, press SELECT F to select the program. Then, click right on [Local LMFS Operations], to invoke the second level of operations. In the second level, if you click left on [Free Records], the program displays a line for each block of the file map, telling you which records are covered by that block, the number of such records, and how many are marked as free. It also tells you how many free records (marked as "in use" in the map) are in the free buffer, and finally displays a grand total of the number of free, used, and total records in the file system.

To find out how many records are actually in use, click middle on [Free Records] to prepare a printable report of record use throughout the file system. This has to pass over every object in the file system, and so it takes some time, especially on large file systems. The discrepancy between the answer of this function and the answer you get when you click left on [Free Records], tells you how many lost records there are; if there are a lot, you might want to run the salvager.

Clicking right on [Free Records] displays how many records are in use in each partition. This information is necessary for commands such as [Grow Partition] that allow you to change the size of partitions, add partitions, or remove partitions.

12.2. Salvager

The *salvager* is a program that reads every LMFS record of the file system and finds and fixes certain inconsistencies and errors. It can fix two classes of problems.

- It can see which records are in use and which are free, and update the free record map to reflect the current state of the file system. This is how you recover lost records.

- It can find objects that are stored in a file system partition but are not referenced by any directory. Such objects are called *orphans*; they exist only if some problem has occurred, such as a file system crash during the creation of a file, or an unanticipated failure of some sort. The salvager finds such objects and puts them back into the directory hierarchy (*repatriates* them).

12.2.1. Using the Salvager

To run the salvager, press SELECT F to select the File System Editing Operations program. Click on [Local LMFS Operations] to invoke the second level of the program. Next, click on [LMFS Maintenance Operations] to invoke the third level of the program. Now click on [Salvage] to obtain a menu of options. If you have a local file system of multiple partitions (occupying multiple FEP files), you are pre-

sented with a menu of partitions to process. This menu, which is an Accept-Values menu, also includes questions about salvager operations. In addition to listing the partitions to be salvaged, the menu offers you the options as shown in figure 13.

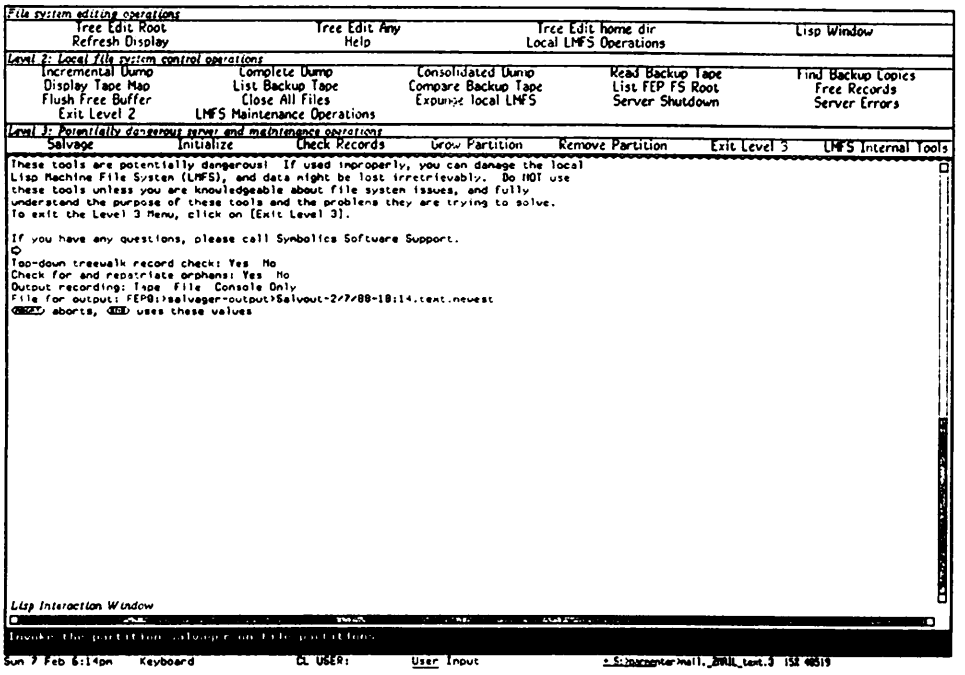

Figure 13. Salvager Options

Here are the options:

`Top-down treewalk record check: yes no`

`Check for and repatriate orphans: yes no`

`Output recording: Tape File Console only`

`File for output:`

The first items on the menu constitute a list of partitions you can select for processing by the salvager. You can choose some or all of the partitions for processing.

The second menu option, `Top-down treewalk record check`, offers to scan all of the directories and files in the local LMFS and report any damaged records (hardware or software), disappeared files, or any other problems. This search starts at the

root and goes through all of the file system, directory by directory, and is performed after all other salvaging activity.

Note: If you deselect any partition for repatriation, then the next menu item, which offers to check for and repatriate orphans, disappears. This happens because it is impossible to construct an accurate model of the hierarchy if each partition is not scanned.

The third menu option, `Check for and repatriate orphans`, offers to find orphaned objects and put them back into the directory hierarchy. During this scan, the salvager also replaces bad directory records with good ones.

The fourth menu option, `Output recordings`, offers to log the salvager output either to tape, in a file, or only to the console.

- If you choose the `Tape` option for output recording, every message goes onto the tape as soon as it is produced because of a special format that is used. Using an industry-compatible tape ensures that all messages appear on tape. If you use a cartridge tape, this is not fully guaranteed. The following forms can be used to view the tape produced in this way:

  lmfs:print-salvager-output-tape &optional *tape-spec* *Function*
  (*stream* **zl:standard-output**)

  Prints the contents of the tape created by the salvager. If you do not supply any arguments you are prompted for a tape spec, and output prints to the console.

  lmfs:copy-salvager-output-tape-to-file &optional *Function*
  tape-spec pathname

  Prints the contents of the tape created by the salvager to a file. You are prompted for any arguments not given.

- If you choose the `File` option for output recording, you must supply a file name. The default file is a FEP file, on boot unit 0. Every time a message is written to file a **:Finish** is done to the file, so that even if the system crashes, the file is intact with all the messages up to the point of failure.

  If you decide to put the output recording in a FEP file, make sure there is enough room, probably about 100 blocks. If you have your output recording sent to another host, choose a host that you are sure will stay on the network during the logging process.

  Warning: If you only have one Symbolics computer or one file server, you cannot use the `File` option because you may not put the output recording in a local LMFS file.

  There are currently no tools for automatically processing a file containing a log of salvager output.

- If you choose the `Console only` option for output recording, note that this is not usually the device of choice. You should choose this option when there is no other means of logging available.

- If any problems occur while the log is running, such as a file closing or a disrupted network connection, a menu appears. This menu asks what to do about continuing the salvager's log. If you enter the debugger while the log is being recorded you are offered restart options for discontinuing or re-selecting log options.

12.2.2. What the Salvager Does

The salvager always reconstructs the free record maps. Running the salvager takes about two minutes per thousand records of file partition.

When the salvager is repatriating an orphan and it cannot find the directory in which the orphan is supposed to reside, it creates a new directory as an inferior of the directory >repatriations, with a name like lost-1 or lost-2. After a repatriating salvager run, you should examine these directories. When the salvager repatriates an object, it types out a message saying that it did so. One of these messages might cause a **MORE** pause. If you plan to leave your console unattended while the salvager is working, you might want to disable **MORE** pauses before you start it.

Note: The salvager always considers storage occupied by orphans to be "in use" for purposes of the free record map, even if it is not repatriating the orphans. Thus, if many orphans existed, they could use up a great deal of disk space. But normally, orphans do not occur at all. When the salvager repatriates it also "fixes" disk errors and misplaced records or directories by replacing them with fresh, empty ones. By nature of the repatriation process, no files are lost in this way.

12.3. Adding a partition to LMFS

You can add partitions to LMFS by using the File System Editing Operations program. First, press `SELECT F` to select the menu for that program. Click on [Local LMFS Operations] to invoke the second level of the menu. Then, click on [LMFS Maintenance Operations] to invoke the third level of the menu. In the third level of the menu, clicking right on [Initialize] yields a menu of initialization options, which offers [New File System] and [Auxiliary Partition] as choices. Choosing [New File System] is similar to clicking *left* on [Initialize]; it initializes a partition to be the basis of a file system. Clicking left on [Initialize] prompts for an initial LMFS partition location, offering FEP0:>lmfs.file as a default location.

When you add a new partition or a partition on another disk, the disk should be free of errors and properly initialized and formatted, and the partition should exist.

To add another partition, choose [Auxiliary Partition]. Enter the pathname of the FEP file to be used as the new partition. (The default pathname presented, which

is correct for [New File System], is never correct for adding [Auxiliary Partition].) Then choose [Do It]. The system then performs much verification and error checking, roughly as much as when initializing a new partition. It must not be interrupted while performing these actions. When finished, it adds the partition and edits the FSPT automatically.

13. Herald Functions and Variables

The herald is the multiline greeting message that is displayed when you cold boot or warm boot, or after you call the Show Herald command, the Save World command or the FEP Disk Restore command. When you cold or warm boot, your machine displays a full-screen herald, which gives some basic information about your machine and how to operate it. When you display the herald with the Show Herald command, you see an abbreviated herald.

Show Herald Command

Show Herald *keywords*

Displays the herald message. The herald is part of the screen display on a cold booted machine. It shows you the name of the FEP file or partition for the current world load, any comment added to the herald, a measure of the physical memory and swapping space available, the versions of the systems that are running, the site name, and the machine's own host name.

keywords :Detailed, :Output Destination

:Detailed {Yes, No} Whether or not to print the version information in full detail. The default is No. The mentioned default is Yes.

:Output Destination
{Buffer, File, Printer, Stream, Window} Where to redirect the typeout done by this command. The default is the stream ***standard-output***.

The following functions and variables are not actually used in printing the herald, but provide the same kind of information as does the Show Herald command.

sct:get-system-version &optional (*system* "**System**") *Function*

Returns three values. The first two are the major and minor version numbers of the version of *system* currently loaded into the machine. The third is the status of the system, as a keyword symbol: **:experimental**, **:released**, **:obsolete**, or **:broken**. *system* defaults to **System**. This returns **nil** if that system is not present at all.

sct:print-system-status-warning &optional (*system* "**system**") *Function*

If *system*'s status is **:experimental**, prints out a warning reminding the user to load patches. If *system*'s status is **:broken**, prints out a warning cautioning the user that the system may not work. Otherwise, it does nothing.

14. Installing Symbolics Dialnet

14.1. Introduction to Dialnet

Symbolics Dialnet is the component of the generic network system that supports the international dial network. The function of Dialnet is to provide a reliable transport medium over possibly unreliable common carrier facilities. The primary uses of this transport medium are mail transfer and remote login. Mail transfer is handled by the Symbolics mail reading and sending program (Zmail) and by the Symbolics store-and-forward mailer. Remote login is handled by the Terminal program.

14.2. Dialnet and Internet Domain Names

Internet Domain Names are part of a tree-structured naming scheme used by the Internet community for distributed administration of a very large namespace. Dialnet also uses the Internet Domain Names scheme of naming and addressing.

This section discusses how Dialnet uses the Internet Domain Names capability. For a complete discussion of Internet Domain Names: See the section "Internet Domain Names" in *Networks*. For installation instructions: See the section "How to Install the Internet Domain Names System" in *Networks*.

Symbolics networks are represented under the second-level domain name Symbolics.COM, and the Symbolics dial namespace is a third-level domain named DialNet.Symbolics.COM. The namespaces of individual Symbolics customer sites are usually represented as fourth-level domains, subdomains of the Symbolics Dial Network. For example, the dial namespace host named CSNY-Young would be in the CSNY.DialNet.Symbolics.COM domain.

The dialnet registries have a mechanism for associating Internet Domain Names with individual hosts, using the **:domain** keyword. This specifies that a given host serves as a mail gateway for a given Internet Domain. See the section "Contents of a Dialnet Registry", page 196.

14.2.1. Internet Domain Names Namespace Attribute

During installation the customer specifies the Internet Domain Name to be associated with the namespace in which local hosts are registered, by editing the **internet-domain-name** attribute of the namespace object that represents the local namespace itself. All hosts that are named within that namespace then inherit the Internet Domain Name that is entered in the namespace object.

For example, the SCRC namespace object might have this attribute:

```
Internet Domain Name:  SCRC.Symbolics.COM
```

SCRC|JUNCO is a host named Junco in the SCRC namespace. Junco inherits the Internet Domain Name of its namespace, so its Internet Domain name is:

```
Junco.SCRC.Symbolics.COM
```

In some cases a host in that namespace is not in the same Internet domain. An individual host can override the Internet domain of its namespace by entering a value in the **Internet-Domain-Name** attribute of its host object. In this example the host SCRC|GRACKLE has the Internet Domain Name Grackle.MIT.EDU.

```
Internet Domain Name:  Grackle.MIT.EDU
```

The **Internet Domain Name** attribute of the host object is used solely to override the attribute of the namespace object.

14.3. Installing Dialnet

Here is an overview of the procedures for installing Dialnet.

- Install the Domain Name Server software.
- Choose a machine to be your Dialnet host.
- Install the hardware that makes Dialnet possible.
- Update the local namespace so that the host knows about the Dialnet hardware.
- Create the private Dialnet registry.

14.3.1. Installing the Domain Name Server for Dialnet

This installation step has two parts: loading the domain name server software, and updating the namespace database.

Loading the Domain Name Server Software

The instructions for loading the domain server software are included elsewhere in the documentation: See the section "How to Install the Internet Domain Names System" in *Networks*.

Updating the Namespace

Dialnet hosts need to have an Internet Domain name associated with them. See the section "Internet Domain Names Namespace Attribute", page 189.

Most Dialnet sites belong in subdomains of the DialNet.Symbolics.COM domain. Use the namespace editor to edit the local namespace object. Enter a value for the **Internet Domain Name** attribute. For example, a site named Acme would have this entry in its namespace object:

```
Internet Domain Name:  Acme.Dialnet.Symbolics.COM
```

Some sites are already using the Internet domain names and have such names assigned to their hosts. Those sites should continue to use the Internet domain names that they are currently using, and enter the domain name of the namespace to the **internet domain name** attribute of the namespace object.

Any host that needs to override the default **internet domain name** stored in the namespace object can enter a different value for the **internet domain name** attribute of their host object. The **Internet Domain Name** attribute of a host object is used solely to override the attribute of the namespace object.

14.3.2. Choosing a Machine to Be Your Dialnet Host

The machine you choose to be your Dialnet host should be one that is on the dial network most of the time. This is important since this machine receives mail for later distribution to other Dialnet hosts and to other local hosts. If the host is only occasionally connected to the dial network, it cannot do a reliable or speedy job of delivering mail to Dialnet hosts. The Dialnet host must be a Symbolics computer, since the Dialnet software currently only runs on this machine.

We recommend that your Dialnet host be the same machine as your namespace/file/mail server, since that host is supposed to be up most of the time as well. Also, you should choose a machine that has a local file system (LMFS), since sending and receiving mail over Dialnet requires use of the Mailer on your Dialnet host and the Mailer requires a local file system.

14.3.3. Installing the Dialnet Hardware

The first step in connecting your Symbolics computer to the dial network is to configure a Symbolics machine at your site with a modem. Modems may be ordered from Symbolics.

The mechanics of the physical connection of your host to the dial network depend on the model of the processor.

- If you have a Symbolics 3600 processor, follow these instructions:

  Many Symbolics 3600 processors contain a Vadic 3450 modem mounted inside the processor cabinet. Bring both a modular jack and a male EIA connector out to the I/O bulkhead. The modular jack (labeled **MODEM TELCO**) accepts a standard modular plug from the data circuit provided by the telephone company. Connect the male EIA connector (labeled EIA 4) (via a short cable, which you can make yourself or order from Symbolics, see below) to any one of the three female EIA connectors that provide access to the serial lines (labeled EIA 1, EIA 2, and EIA 3). EIA 1 corresponds to serial unit 1, EIA 2 to serial unit 2, and EIA 3 to serial unit 3.

The cable between EIA 4 and EIA 1, EIA 2, or EIA 3 should convey the following signals on the pins given below:

° Pin 2 (TXD ; Transmitted Data)

° Pin 3 (RXD ; Received Data)

° Pin 4 (RTS ; Request To Send)

° Pin 5 (CTS ; Clear To Send)

° Pin 6 (DSR ; Data Set Ready)

° Pin 7 (SG ; Signal Ground)

° Pin 8 (CXR ; Carrier Detect)

° Pin 20 (DTR ; Data Terminal Ready)

Terminate this cable with one male connector and one female connector. If you prefer not to build such a cable yourself, you can order it from Symbolics.

If your 3600 has been upgraded to support audio and phase-encoded video [via UPGR-SY70], then the gender of the serial ports is male. Construct the cable as described above except that both connectors on the cable should be female. Again, if you prefer not to build such a cable yourself, you can order it from Symbolics.

Earlier versions of the FEP proms installed in 3600 processors do not support all the features of the 3600 serial lines. Dialnet requires FEP version 127 or higher. The easiest way to determine what version of FEP proms is installed is to use the Show Herald command with the keyword :detailed.

```
Show Herald :Detailed Yes
```

One line of the resulting display shows the Fep version.

- If your Symbolics computer is any machine model, aside from a 3600, or does not contain a Vadic 3450 modem, follow these instructions:

Symbolics computers, other than the 3600 model, have no internal modem, but instead expect an external Vadic 3451 or CDS-224 modem to be connected to one of the three serial ports.

Each modem has two electrical connections (excluding the power cord). One of the electrical connectors is similar to the connector the telephone company puts on telephones. Plug this into the telephone jack. The other electrical connector

has a 25-pin connector which is standard for RS-232 serial lines. Plug this into one of the serial line ports on the back of your Symbolics computer. For example, if you have a Vadic 3451 modem, which has a pre-installed phone line, you simply connect this to the telephone jack. If you have a CDS-224 modem, you first connect the modular jack to the **line in** on the modem and then to the telephone jack.

Bring the serial lines out to the I/O bulkhead and terminate them in male EIA connectors. On the 3640 and 3645 processors, these connectors are labeled SERIAL 1, SERIAL 2, and SERIAL 3. On the 3670 and 3675 processors, these connectors are labeled EIA 1, EIA 2, and EIA 3. On the 3610, 3620, and 3650 processors, these connectors are labeled SERIAL 1 and SERIAL 2.

The cable connecting the modem to the serial port conveys all the signals described above for 3600 cabling; only the gender of the serial port connector differs. If you prefer not to build this cable yourself, you can order it from Symbolics.

14.3.4. Updating the Local Namespace to Know About the Hardware

You must record the physical connection of your machine to the dial network in the namespace database so that the generic network system can decide how best to establish connections over the dial network. You can represent this connection by adding a **peripheral** attribute to the *local view* of the Dialnet host object.

To edit the namespace database use the Edit Namespace Object command, choosing to edit the host object.

The term *local view* is used here in contrast to the dial namespace view of the object. If you are editing a host that has *already* been identified as being in both the local and the dial namespace view namespaces, you are asked to choose (via a small pop-up menu) between the local and the dial namespace view of the object before you actually begin to edit the object. Choose the local view; there are no servers for the dial namespace, so you would have no way to save your changes.

When editing the host object, use the **peripheral** attribute to represent the modem that connects this host to the dial network. There are five relevant indicators for the peripheral attribute: **unit**, **model**, **baud**, **phone-number**, and **autoanswer**.

unit — Corresponds to the serial port number on the I/O bulkhead of the machine. The value should be a number between 0 and 3, inclusive. The serial port for the console is numbered 0.

model — Corresponds to the type of modem attached to this serial port. The value should be one of the following: va212, va3450, va3451, or cds-224.

baud Should be 1200.

phone-number Corresponds to the telephone number of the telephone trunk to which this modem is attached. The phone-number associates this modem with a particular dial network address. (The address, in turn, associates this peripheral with a host in the dial namespace. This latter mapping is done via the Dialnet registry.)

autoanswer Corresponds to the ability of this modem to receive incoming calls from other sites. If you wish to receive calls from other sites, then this indicator should have a value of **yes**. If it has a value other than **yes**, incoming calls are not answered and communication with other sites can only be initiated by the local site.

In general, if two sites wish to communicate over the dial network, at least one of the sites needs to enable **autoanswer**.

An example peripheral attribute might be:

```
peripheral modem unit 2 model va3450 baud 1200 phone-number 16175777348 autoanswer yes
```

14.3.5. Creating Dialnet Registries

The private Dialnet registry contains information describing hosts with which you intend to communicate, plus your own local Dialnet host. You declare the phone number of your local host here, as well as in the host's namespace object. This file lives in sys:site;private-dialnet-registry.lisp.

Dialnet *registries* are files that represent the shape of and possible connections on the dial network. These files really contain namespace information in a form suitable for periodic distribution by separate administrative groups. The following registries exist:

The public dialnet registry

This registry is maintained and distributed by Symbolics. It contains information on publicly accessible Symbolics hosts and domains, on Telenet PADs (the GTE Telenet equivalent of the Arpanet TIP), and on dialing conventions for the international dial network. For example, information in this registry might be the addresses of Symbolics software support groups, all the GTE Telenet PAD access numbers, and enough information about the international phone system to help you find the cheapest way to place calls to the Symbolics hosts described in the registry. This registry might also contain the domain of and address for the users' group, so that new customers could easily contact the group.

This file is sys:site;public-dialnet-registry.lisp.

The users' group dialnet registry

This registry is maintained and distributed by the Symbolics users' group, and contains whatever host and domain information the group's members see fit to enter and distribute.

For example, the users' group might distribute the addresses and domains of all members that wanted to share information via the dial network.

This file is sys:site;users-group-dialnet-registry.lisp.

The private dialnet registry

You maintain this registry at your site. It contains local dialing conventions and any private host and domain information for the site.

A private dialnet registry might contain the names, addresses, and domains of the following:

- Other sites in the organization
- Other organizations that did not want to be published by the users' group
- Common carriers (for example, Tymnet)
- Subscriber data services (for example, Dow Jones Information Services)
- Gateways to other domains (for example, .ARPA)

This file is sys:site;private-dialnet-registry.lisp.

Information moves into the users' group registry when someone at the local site contacts the group, and the group distributes a registry. Information moves into the public registry when someone at the local site contacts someone at Symbolics, and Symbolics next distributes a registry.

Because of delays in the distribution of registries, the private registry may be useful as a repository for advance copies of public or users' group information. When two registries contain differing information about the same object, any conflict is resolved as follows:

1. The public registry is assumed to be least current.
2. The users' group registry is assumed to be more current than the public registry.
3. The private registry is assumed to be more current than either the public or users' group registry.

14.3.5.1. Example of a Private Dialnet Registry

Here is a sample private Dialnet registry:

```
;;;The public dialnet registry contains the common case where the
;;;  phone is just a standard outside line.
;;;ACME is on a Centrex-system phone for which you have to dial 9
;;;  to get an outside line.
(:SUBNET "1901482yyyy>1zzzwwwwwww" :DIAL "yyyy" :cost "0")
    ;;Local extensions
(:SUBNET "1901482yyyy>1zzzwwwwwww" :DIAL "yyyy" :COST "0")
(:SUBNET "1901482yyyy>1800zzzzzzz" :DIAL "91800zzzzzzz" :COST "2")
(:SUBNET "1901482yyyy>1xxxzzzzzzz" :DIAL "9zzzzzzz" :COST "1")

;;; Acme's Dialnet host
(:HOST "ACME-GENIUS" :ADDRESS "19014823002" :LOCAL-NAME "ACME|GENIUS")

;;; Acme's Dialnet domain
(:DOMAIN "Acme.DialNet.Symbolics.COM" :host "ACME-GENIUS")

;;; Acme needs to communicate with Coyote Enterprises.
(:HOST "COYOTE-LISPM-1" :ADDRESS "19035492047")

(:DOMAIN "Coyote.Dialnet.Symbolics.COM" :host "COYOTE-LISPM-1")
```

14.3.5.2. Contents of a Dialnet Registry

A Dialnet registry contains lists of alternating keywords and values. The first keyword in a form (the car of the list) indicates the type of information being conveyed and is one of the following:

:subnet Specifies the cost of and dialing information for a subnet of the dial network. The keywords and values for this form are described elsewhere. See the section "Dial Network Addressing" in *Networks*.

```
(:subnet "1xxxyyyyyyy>1800zzzzzzz" :dial "1800zzzzzzz" :cost "1")
```

:host Specifies a host on the dial network and its address. The name specified by the **:host** keyword is the name of the host in the dial namespace. To avoid duplicate names in this namespace, we use the site name as a prefix. As an example, consider a site named Trilogy, one of whose hosts with a dial network address is named Cerberus. The unique dial namespace name of this host is then trilogy-cerberus.

```
(:host "trilogy-cerberus" :address "14151515151")
```

The address for the host is a string representing the address of the host on the international dial network. In our examples, drawn typically from hosts in the United States, the country code is 1, the area code is a three-digit number whose second

digit is always 0 or 1, and the rest of the address is the familiar local seven-digit number. The address is the concatenation of these three fields.

If a host is present in a local namespace other than the dial namespace, its local name should be noted with the :local-name keyword. For example, the following specifies a host whose name (in the dial namespace) is USMC-Gomer, whose address on the dial network (that is, the network named "dial" in the dial namespace) is 14155551212, and whose name in the local namespace (here, the namespace named USMC) is Gomer.

```
(:host "USMC-Gomer" :address "14155551212"
 :local-name "USMC|Gomer")
```

Each host in the dialnet registry is assumed to support the following mail-related services. When the dialnet registry is loaded, host objects are automatically created or updated to contain the appropriate service attributes for these services.

- STORE-AND-FORWARD-MAIL (over the DIAL medium, using the SMTP protocol)
- MAIL-TO-USER (over the DIAL medium, using the SMTP protocol)
- MAIL-PROBE (over the DIAL medium, using the MAIL-PROBE protocol)
- EXPAND-MAIL-RECIPIENT-SERVICE (over the DIAL medium, using the SMPT protocol)

:domain

Specifies an Internet mail domain and the name of the associated gateway host. Internet mail domains are used by Zmail and the store-and-forward mailer in conjunction with the domain server to direct mail around the dial network, in the absence of other namespace information.

Hosts on the Internet must use their valid domain name as the domain entry in the dialnet registry. For example, for a site with an existing domain name of

```
wilson.smith.nh
```

the valid domain name in the dialnet registry is:

```
(:domain "wilson.smith.nh" :host "wilson|hostname")
```

Note: Dialnet hosts on the Internet should not have dialnet registry entries for other dialnet hosts that are also on the Internet. If you do so, the domain system cannot be guaranteed to obtain the correct information about these hosts.

For a complete discussion of Internet Domain Names and installation instructions: see the section "Internet Domain Names Capability" in *Genera 7.0 Release Notes.*

If your local host is not on the Internet, use the following format in the user's group registry

```
(:host "CSNY-Young" :address "12121234567")
(:domain "CSNY.DialNet.Symbolics.COM"
 :host "CSNY-Young")
```

If the format above is in the user's group registry, mail to the following addresses would be routed via DialNet to Young.CSNY.DialNet.Symbolics.COM for further distribution:

```
Neil@Young.CSNY.DialNet.Symbolics.COM
Graham@Nash.CSNY.DialNet.Symbolics.COM
```

:telenet-pad Specifies the name and address of a GTE Telenet PAD. These gateways to the Telenet network can be an economical way of routing traffic across the dial network and are also useful in their own right as access to higher level GTE Telenet services such as TeleMail. The Symbolics Dialnet implementation uses Telenet automatically when it determines that by doing so it can make a cheaper connection.

```
(:telenet-pad "boston-telenet-pad"
 :address "16172920662")
```

14.3.5.3. Loading a Dialnet Registry

All the information in the three Dialnet registries (public, users' group, and private, if present) can be loaded into the local host with the **dial:load-dialnet-registry** function. You might use this function to load the names and addresses of all Telenet PADs, for example.

```
(dial:load-dialnet-registry)
```

Note that some system programs also load the Dialnet registries. The Store-and-Forward mailer does this when the mailer is started, and the Internet Domain Name Server program does this when that server is started.

14.3.6. Testing Your Dialnet Installation

After you have set up your private Dialnet registry, you should make sure Dialnet is working. To test it, you must load your Dialnet registry. Do this by using the function **(dial:load-dialnet-registry)**. Then, try to connect to a host whose dial-up number is known to you, over the dial network, using the terminal program:

Select the Terminal window by pressing SELECT T. Then, try to dial a host whose phone number you know by typing:

```
Connect to host: dial|dial|13125678901
```

You are told that the given host does not know how to support LOGIN, and then asked if you meant any of the possible protocols the Symbolics Computer knows how to use over a dial-up line. Answer *yes* when it asks about a protocol named TTY-LOGIN, which uses the DIAL-RAW medium. It should then dial up and allow you to log in as if you were sitting at a normal dial-up console. You can customize your console somewhat by pressing NETWORK X and click on the appropriate terminal emulation mode.

14.4. Using the Terminal Program with the Dial Network

Once you have set up the hardware and namespace information that describes how your host is connected to the dial network, you can use that host to dial up other hosts. This is an excellent test of the hardware and software configuration, even if you don't usually dial up other hosts. And, of course, it can be very useful in its own right, providing access to hosts accessible only via dial-up lines.

Press SELECT T to get to the Terminal program, then type a host name at the Connect to Host: prompt. Host names are of the form dial|dial|16175777348. To break that down a bit, that's the host at address 16175777348 on the network named dial|dial, which in turn is the network named dial in the dial namespace. If you need to make the same call frequently, you can add hosts to your own local namespace (not the dial namespace) with addresses on the dial|dial network. In addition to an address attribute, you will probably want to give such a host a service attribute of:

```
Service:  LOGIN DIAL-RAW TTY-LOGIN
```

If this host is in some other domain (this is likely if it is at some other site), you may want to give that host its own Internet Domain Name which corresponds to the Internet Domain Name established by that host's administrator. Use the namespace editor to add a value for the **internet domain name** attribute of the host object for that host.

Telenet PADs have names in the dial namespace such as dial|boston-telenet-pad. Telenet PADs are listed in the public dialnet registry, so to dial up a PAD you first have to load the dialnet registry. See the section "Loading a Dialnet Registry", page 198.

14.5. A Sample Dialnet Installation

This section details a sample installation at site NYC. The goal of this installation is to bring up this site on the dial network, both to exchange electronic mail with other sites and to use the Terminal program to access other hosts over dialup lines. There are five machines at site NYC, named Bronx, Queens, Manhattan, Brooklyn, and Staten-Island. Bronx (a 3600) was the first machine installed at the site, so it bears the title of "site support system" and has an inboard Vadic modem. In this example it also has an attached Kanji tablet for Japanese language

work, and supports an LGP1 printer. When Manhattan (a 3670 with a large disk) later arrived, it was made the SYS host and the namespace server, and Bronx was given over to being a file server for user files, in addition to its continued use as a Japanese workstation and spooler for the LGP1 printer.

1. **Install the Domain Name Server.** This installation step has two parts: updating the namespace database to include the **internet domain names** attribute of the namespace object for the local namespace, and loading the domain server software. See the section "Installing the Domain Name Server for Dialnet", page 190.

2. **Find the modem.** In this example, Bronx has an internally mounted Vadic VA3450 modem by virtue of being a 3600 machine model.

3. **Connect the modem to a serial port.** The serial port, of course, should be otherwise free. In this example, serial port 1 is already in use (for example, it could be supporting a tablet or a printer), so you will be attaching the modem to serial unit 2. A short cable should be used to connect EIA 2 and EIA 3, where EIA 2 is the connector for serial port 2 and EIA 3 is the connector for the inboard modem. For details on how to make this cable (if such a cable wasn't shipped to you by Symbolics) or on how cabling should be done for all machine models except the 3600: See the section "Installing the Dialnet Hardware", page 191.

4. **Connect the modem to the phone company's data circuit.** This requires a telephone line with modular connectors on both ends, of sufficient length to reach from the telephone jack to the **MODEM TELCO** jack on the 3600 I/O bulkhead. (Models other than the 3600 use the same type of cable but the connection is from the telephone jack to the **LINE** jack in the CDS modem). In this example, the data circuit is on a private branch exchange (PBX) that supports direct inward dialing, allows other extensions to be called by dialing their 4-digit trunk numbers, but requires that a 9 be dialed first when making calls outside the PBX extensions. The sample phone number for Bronx is area code 212, phone number 765-4321.

5. **Register the modem in the namespace database.** Use the namespace editor to add a **peripheral** attribute to the local view of the host to which the modem is attached. In this example, you are attaching the modem to serial unit 2, the modem is a Vadic VA3450, the phone number is 1-212-765-4321 (1 (US) country code, 212 area code, 765 exchange, 4321 trunk). Furthermore, this example supposes you want other sites to dial your site and assume some of the communications costs themselves, so you want to enable the autoanswer feature. Thus the peripheral attribute you add and save looks like:

```
peripheral modem unit 3 model va3450 phone-number 12127654321 autoanswer yes
```

The **peripheral** and **modem** indicators must be first and second. Past the **modem** indicator, the order of indicator/value pairs is irrelevant.

For further information on registering the modem in the namespace database:
See the section "Creating Dialnet Registries", page 194.

6. **Create a private dialnet registry.** You need to enter into the private dialnet registry information concerning:

   - Local dialing conventions
   - Hosts connected to the dial network and their addresses
   - Locally supported Internet mail domains

   You do this by creating a private dialnet registry. [You should edit the file sys:site;private-dialnet-registry.lisp and enter the following information. To create the mode line, use this sequence of commands:]

   ```
   Lisp Mode (m-X)

   Set Package (m-X)  enter Dial to the prompt

   Set Lisp Syntax (m-X)

   Update Attribute List (m-X)
   ```

   After you use these commands you are asked the following question. Answer Y.

   ```
   Set attribute for the file and attribute list, too?
   ```

   Or, you can manually type the attribute line as it appears in this example and then use command Reparse Attribute List (m-X).

```
;;; -*- Mode: LISP; Package: DIAL; Base: 10 -*-

;;; one PBX number to another: just dial the extension.
;;; local external phone number: dial 9, then the number.
;;; different area code: dial 9, then the number.
;;; WATS number: dial 9, then the number.
;;;      Note that WATS is cheaper than long-distance.
(:subnet "1212765ssss>1212765dddd" :dial "dddd" :cost "0")
(:subnet "1212765ssss>1212xxxdddd" :dial "9xxxdddd" :cost "1")
(:subnet "1212765ssss>1aaaxxxdddd" :dial "9aaaxxxdddd" :cost "2")
(:subnet "1212765ssss>1800xxxdddd" :dial "91800xxxdddd" :cost "1")

;;; modem is on DIAL|NYC-Bronx, phone (212) 765-4321,
;;;  and the local name of the host is Bronx.
(:host "NYC-Bronx" :address "12127654321" :local-name "NYC|Bronx")

;;; DIAL|NYC-Bronx will run a store-and-forward mailer,
;;; servicing the domain named NYC.DialNet.Symbolics.COM.
(:domain "NYC.DialNet.Symbolics.COM" :host "NYC-Bronx")
```

For related information:

See the section "Symbolics Dialnet" in *Networks*.
See the section "Creating Dialnet Registries", page 194.
See the section "Dialnet and Internet Domain Names", page 189.
See the section "Overview of the Mailer" in *Communicating With Other Users*.

7. **Load the dialnet registry.** Type the following form:

```
(dial:load-dialnet-registry)
```

8. **Test the dial network connection with the Terminal program.** This step is optional, of course, since you might not have a dial-accessible host handy. This example assumes a timesharing host named NYC-Tammany, with a single dialup line at (212) 666-1040. You select the terminal program with SELECT T and enter the address of NYC-Tammany, in this case dial|dial|12126661040. The number is dialed, the connection is made, the screen clears and any subsequent communication takes place with NYC-Tammany. After the session with the foreign host is complete, the connection can be broken by pressing NETWORK L. For more information on using the Terminal program to access hosts via dialup lines: See the section "Using the Terminal Program with the Dial Network", page 199.

15. Installing and Configuring the Mailer

When you want to offer store-and-forward mail service on a Symbolics host (or on several hosts), you must update the host object(s) corresponding to the host(s) that will offer mail service.

For an overview on changing objects in the namespace database: See the section "Maintaining the Namespace Database", page 52.

1. Use the Edit Namespace Object command, with a namespace class of Host, and answer the prompt with the name of the host that is to provide mail service. In this example, for a host named Red, what you type is underlined:

   ```
   Edit Namespace Object <space> (a namespace object or Any
   [default Any]) Host <space> RED <RETURN>
   ```

2. If this host is going to send and receive mail over telephone lines, and if you have loaded a Dialnet registry that included this host, the namespace editor asks whether you want to edit the DIAL or the local namespace view of this host. Choose the local view. For more information on Dialnet: See the section "Symbolics Dialnet" in *Networks*.

3. Add the following service triples to the host object definition:

   ```
   MAIL-TO-USER CHAOS CHAOS-MAIL
   STORE-AND-FORWARD-MAIL CHAOS CHAOS-MAIL
   ```

   These tell the host to deliver mail to local users and forward mail to non-local users

   If the host has an address on the DIAL network, also add the corresponding Dialnet triples; otherwise, skip this step:

   ```
   MAIL-TO-USER DIAL SMTP
   STORE-AND-FORWARD-MAIL DIAL SMTP
   ```

4. Click on [Save] to save your entries in the host's Namespace Object.

5. Click on [Edit] to edit a different namespace object. When the machine prompts for the class of object to edit, click on [Site]. Type the name of your site when the machine prompts for it.

6. Add the following property pair to the site object definition:

   ```
   All-MAIL-ADDRESSES-FORWARD YES
   ```

   This tells the mailers at your site to process any mail addressed to any of the hosts in the site.

7. After you have added the services in the previous steps, click on [Save] to save the changes.

8. Wait for confirmation that the changes have been saved, and then click on [Quit] to exit from the namespace editor.

The mail services have now been registered in your namespace database, for the host you specified.

The next procedure installs the Mailer software onto the machine that is going to be the mail server, which is the same host as specified above.

This installation is actually a further customization of the new release world for your site; that is, the particular world that is run on the host providing mail service is different from the worlds running on the other machines at the site. The difference is that the host which provides the mail service should also be a namespace server and file server. You should backup the files on this machine, since its file system is being used.

1. You should free up enough disk space in the FEP file system of the server machine to accommodate the following material:

   - There should be enough space for a world about 10% larger than the world already there. (The world with the Mailer loaded is saved back onto the disk at the end of this procedure.) You may also want to load the mailer into a world and save it with Incremental Disk Save (IDS). For information about IDS worlds: See the section "Using Incremental Disk Save (IDS)", page 29. If mail service is to be provided by a Symbolics 3640: See the section "Instructions for Managing Disk Space on the 3640 With a 140 Megabyte Disk", page 36.

   - The Mailer sources and binaries occupy about 200 LMFS blocks.

   - The Mailer databases each require a minimum of 20 LMFS blocks. The size of a database increases approximately linearly with the number of messages stored by the mail server (the host); an "average" piece of mail occupies about 2 LMFS blocks.

   You can use the Dired (m-X) Zmacs command, the Delete File command, or SELECT F (the File System Editor) to delete unused or excess files in the same way that you freed up enough space for the customized new release world.

2. Using the Create Directory command, create the following directories on the mailer host:

```
>Mail>
>Mail>Static>
>Mail>Archive>
```

3. Create the two local customizations files, Mailboxes.Text and Options.Text, which are both in the directory >Mail>Static> on the mailer host.

   a. The Mailboxes.text file contains information which allows the Mailer to know where to deliver mail. This file defines all the local recipients of mail, including individuals and mailing lists.

      Using the editor, edit the file >Mail>Static>Mailboxes.text. An example follows:

```
; -*- Mode: Lisp; Package: Mailer; Lowercase: Yes -*-

;; This file belongs on White only.

;; The mail addresses "Postmaster" and "File-Server"
;; are resolved with respect to individual hosts.  Edgar
;; gets all mail addressed to these "non-people" users,
;; will be the person who takes care of dead letters, etc.

(define-local Postmaster Edgar)
(define-local File-Server Edgar)
(define-local Mail-Server Edgar)
(define-local Lisp-Machine Edgar)

;;; People who get mail delivered on White
(deliver-local Robert Edgar Henry Malcolm Peterson Smiley Nelson)

;;; The basic mailing list for bugs.  Each person listed
;;; gets a copy of each bug report, and a copy goes to the files.
(define Bug-LispM
    (list Edgar Davies Henry
          "white:>Mail>Archive>Bug-Reports.text"))

;;; Miscellaneous mailing lists

;;; Good to have a mailing list that allows you to reach everyone
(define seventh-crisis
    (list Robert Edgar Henry Malcolm Peterson Davies Richard Smiley Nelson))

;;; Simple mailing list, commented to show particular preferences
(define sports-followers
       (list Richard       ;; golf
             Smiley        ;; gymnastics
             Nelson)       ;; tennis

;;; Another simple list.
;;; Put yourself on this list if you want to hear about tape bargains
(define tapers
       (list Richard Henry))
```

Write out the file, and stay in the editor for the next step.

Here are some notes about the mailboxes.text file above:

- It is necessary to appoint someone to be the *Postmaster* at your site. This is the person who receives any mail messages that cannot be delivered and cannot be returned. Also, if anyone outside your site needs to communicate with you about the operation of your mailer, it

is conventional for them to send mail to Postmaster. In the example above, the

```
(define Postmaster Edgar)
```

tells the Mail Server to place mail addressed to Postmaster in Edgar's inbox.

- The *deliver-local* list is the list of people who receive mail on this mail server. For example,

```
(deliver-local Edgar)
```

means that mail addressed to Edgar at your site is put into the file Local:>Edgar>Mail.text, which is where Zmail expects to find it.

- Use *define* to make synonyms. For example, *(define Edgar Hoover)* delivers mail addressed to Edgar to whatever mailboxes "Hoover" implies. The second "argument" to define may also be a *list* form, which is how you create a mailing list.

- You probably do not want to use the *define-local* list. Names on this list do not go into the the forward files distributed to the rest of the site.

b. The Options.Lisp file contains various customizations for the operation of the Mailer. This file contains Lisp forms which load when the mailer is started or when you use the function **(mailer:update-options)**. If you make a change to this file while the Mailer is running, type **(mailer:update-options)** to a Lisp listener on the Mailer host, to update the mailer.

Using the editor, edit the file >Mail>Static>Options.lisp. Set it up to look like this:

```
;;; -*- Mode: Lisp; Package: Mailer; Base: 10; Syntax: Common-Lisp -*-

(setq mailer:deferred-delivery-times '("7:45am" "5:15pm" "11:15pm"))
(setq mailer:deferred-receipt-times  '("7:45am" "5:15pm" "11:15pm"))
(setq mailer:deferred-receipt-hosts
      '("Riverside.SCRC.Dialnet.Symbolics.COM"
        "LispM-1.Coyote.Dialnet.Symbolics.COM"))
```

After you finish editing the file, save it.

The variables **mailer:deferred-delivery-times** and **mailer:deferred-receipt-times** tell the mailer how often to try to deliver and pick up mail from the Dialnet hosts. The value of these variables may be **nil**, **t**, a specific time, or a list of times. If you specify **nil** for the value of one of these variables, it means *do not do this*; if you specify **t**, it means *do this as frequently as possible.* If you specify a time interval, such as **"three hours"**, this is the frequency with which to pick up and

deliver mail, in this case, every three hours. If you specify a list of times to pick up and deliver mail, for example, *("7:45am" "5:15pm" "11:15pm")*, this specifies exact delivery and pick up times.

The variable **mailer:deferred-receipt-hosts** is a list of Mailer hosts from which the local Mailer attempts to retrieve mail. This variable is only useful on Dialnet mailers. You should probably always include "Riverside.SCRC.Dialnet.Symbolics.COM" on this list, since Riverside does not dial out for any mail delivery.

When you test the mailer, you may want to set **mailer:deferred-delivery-times** to **t** until you receive your first mail message.

For a description of the options: See the section "Files and Directories Used by the Mailer" in *Communicating With Other Users.*

4. To install the Mailer system, you should create a special world that includes the Mailer. As is true when creating any other world, it is recommended that you do as little as possible to the environment prior to world creation. In this situation, you should do all of the steps listed above, including editing the namespace and creating the Mailer files, and then cold boot the machine. Once you are in the process of building a new world, do not switch windows or do anything that causes an unnecessary allocation of storage.

   First, cold boot the host. After this, disable services, log in to the sys host, and then use the Load System command to load the Mailer system. After using the Load System command, save the world with the Save World command.

   The procedure for installing the mailer system is shown in the following example:

```
Command: Halt Machine RETURN
Fep Command: boot RETURN
     .
     .
     .
Command: Disable Services RETURN

Command: (si:login-to-sys-host)

Command: Load System Mailer

Files to be loaded:

.
.
.

Command: Save World (Complete or Incremental) name-of-file-to-save-it-in
```

When you save the world, you may want to use the Save World Incremental command, which allows you to have an incremental world built on the site-customized distribution world containing the Mailer. For information about this command: See the section "Using Incremental Disk Save (IDS)", page 29. You have now configured the newly saved world to be a mail server for the site.

5. Boot the new world using the FEP Boot command. When services are enabled to the mail server, the mailer starts automatically. Shortly thereafter, you can press SELECT O to bring up the Mailer Log window. Then test the mailer by sending messages to and from various machines at the site.

15.1. Testing and Registering the Mailer

The best way to test the Mailer is to attempt to send mail to various people. Send mail to Postmaster at your local site, to make sure that works. Then, you should send mail to Symbolics so we know about your site and are able to communicate with you. Send mail to "HOSS@Riverside.SCRC.Dialnet.Symbolics.COM" that contains the following information:

- The name of your site, for example, "Acme"
- The name of your mailer machine (Fearless), for example, Fearless.Acme.Dialnet.Symbolics.COM
- How you can be contacted such as your name, telephone number, and hours during which we can call you, in case we receive your mail and cannot respond.

If you are on Dialnet, then include these steps:

- The Dial network address of your mailer machine, for example 16175371234
- Whether you wish to have your telephone number published in the Users' Group or Public Dialnet registries (See the section "Creating Dialnet Registries", page 194.) Or, you can specify if you only want Symbolics to send you mail.

To follow what the Mailer is doing, press SELECT 0 on the Mailer host and watch the log output.

15.2. Configuring Large Sites for Multiple Mail Servers

The Store-and-Forward Mailer supports *forwarding tables* to help coordinate mail delivery at sites with several mail servers.

One particular mail server is configured to be in charge of forwarding-table maintenance. The forwarding tables themselves are written by the this host to the file systems of all the other mail servers at the site. This asymmetry is, in a sense, a further customization of the particular mail server that writes the forwarding tables. The customization is usually done by placing a **setq** of the variable **mailer:forwarding-table-hosts** in the options.text file. For example:

```
(SETQ MAILER:FORWARDING-TABLE-HOSTS
   '("MANFRED" "NATASHA" "BORIS"))
```

Here, the hosts Manfred, Natasha, and Boris receive new forwarding tables from the host to which this init file belongs. The forwarding table for a given host is written in the file >Mail>Dynamic>Forwarding.text on that host's local LMFS.

If you want to run the identical init file on all the server machines at a site, the following example may be instructive. Here, a SYS host (Fearless) runs the Mailer and is responsible for writing out the forwarding tables. The File-Server init file, which all file servers use, includes the following lines:

```
(DEFMACRO FILE-SERVER-ONLY-ON (HOSTS &BODY BODY)
  `(WHEN (OR ,@(LOOP FOR HOST IN HOSTS
     COLLECT `(SEND NET:*LOCAL-HOST* :PATHNAME-HOST-NAMEP ,(STRING
       HOST)))) ,@BODY))

(FILE-SERVER-ONLY-ON (FEARLESS)
  (SETQ MAILER:FORWARDING-TABLE-HOSTS
     '("MANFRED" "NATASHA" "BORIS")))
```

The file Mailboxes.text on Fearless contains the names of all the mailing lists for this network. In addition to the usual forms defining mailing lists, the file contains forms like the following:

```
;; What follows is a global table of mail addresses for the network.
;; There is one entry for each host, listing all of the mail addresses to be
;; forwarded to that host.  Each entry is broken into two sublists,
;; one for mailing lists and one for individuals.  This is the only
;; place in which this table should be edited.
;; The forwarding tables for all other hosts are generated from this one.

(DELIVER-TO NATASHA

  ;;; The mailing list file on Natasha is >mail>static>mailboxes.text.

        ;; Individuals
        Andy Bob Charles David Edgar

        ;; Lists
        ASAS Audio Audiophiles
        BBoard Bikers Bleeding-Hearts Bridge
        ...)
```

Similar **deliver-to** forms are supplied for Boris and all other hosts with store-and-forward-mailers.

When the Mailer on Fearless is booted, (that is when **mailer:start-mailer** is called, either automatically when you enable services, or minimally, by the user) or when the Mailer notices that the local mailboxes.text file has changed and has been stable for at least 10 minutes, it reads in its mailboxes.text file (Fearless will never have a forwarding.text file) and then writes out forwarding.text files on all the other Mailer hosts. Those hosts eventually read in the new forwarding.text files and their own mailboxes.text files, merge the two sets of definitions, and carry on.

The forwarding.text file that Fearless generates for Boris includes forms for hosts with Store-and-Forward Mailers, such as Natasha :

```
; Mailbox forwarding table for BORIS.
; Written 6/10/86 15:33:54 by Mail-Server running on FEARLESS.
; From F:>Mail>Static>Mailboxes.text created on 6/10/86 15:25:24.
; This table is automatically generated by a program.  Do not edit it.

(DELIVER-TO NATASHA

        Andy Bob Charles David Edgar
        ASAS Audio Audiophiles
        BBoard Bikers Bleeding-Hearts Bridge
```

If Boris gets incoming mail for these individuals or lists, the mail is forwarded to Natasha. There is no entry for Boris in this list, since those entries come from the mailboxes.text file on Boris.

Index

B

B B

C

C C

D

E

E

E

F F F

G

G

G

H

H

H

K

K K

L

L L

M

N N N

O O O

P

P

P

Q

R

S

S

S

T

U U U

V

W

X

X X

Y

Y Y

Z

Z Z