A Printing Solution for a Multimedia Environment

For environments in which users are confronted with a myriad of printers to choose from, HP SharedPrint provides a simple graphical interface that enables users to select a target printer and a set of options without encountering the typical problems associated with this process.

by John Mandler

Most computer users in a networked computer environment take certain aspects of the system for granted—especially shared resources such as disks and printers. Users expect these resources to work flawlessly without any concern or interference on their part. Unfortunately, the ideal is not always possible, especially for printers. Printers often fall far short of users' expectations. Because of this, we placed special emphasis on providing a robust, easy-to-use printing solution for HP MPower.

To understand the HP MPower printing solution it is first necessary to examine the fundamentals of a spooled printing system. Fig. 1 shows the phases and basic operations of a spooled printing system.

The user controls the specification phase by selecting an output device from a list provided by the system. The object (file) to print forms a second part of the the job specification. Finally, any options that control the appearance of the final output should be noted. The options may be used to control when the job is printed, to define the desired feedback, to select some feature of the output device, or to specify the final appearance of the printed output.

The queue and control phase is handled by the native spooling system. It packages the user's job specification, locates the target machine, establishes the communication path, and places the job in a queue. The spooling system eventually dequeues the job, invoking the functions that process the job, and directs the output to the printer. The spooling system also handles query and control requests from the user, providing the feedback required for a robust system.

The processing phase performs the real work in a printing subsystem. This phase is initiated by the spooler, which starts a process that connects a dequeued file to the process input and the printer data port to its output port. The processing phase performs the translation and formatting of the input data stream, applying options as appropriate, and sending out printer control functions to the selected printer when necessary.

Problems

Each of these three steps—specification, queue/control, and processing—has problems that can range from nuisance issues such as the printer is out of paper to a compete breakdown of the printing subsystem (see Fig. 2). Problems in the specification phase include the need for the user submitting a print job to spend time finding the names and capabilities of available printers and spending more time learning how to add special printer control options to a print request. The options give the user more control over the final look of the document.

Queue and control problems include minimal to no feedback on the state of a job or printer. Failures are often not reported to the user, giving the appearance that jobs just disappeared from the queue. Administering the spooling system can be a difficult task. The steps to add a printer are not obvious and modifying a printer's default behavior requires more knowledge about the printer and programming than the user may want to know.

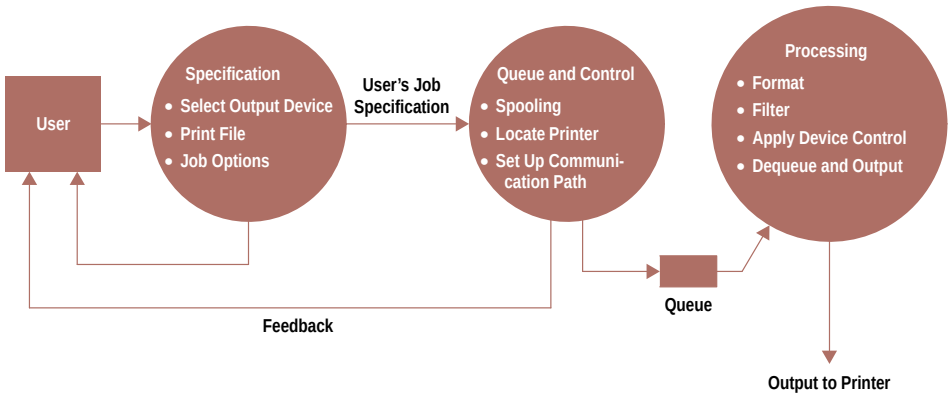

Fig. 1. Basic operation of a spooled printing system.

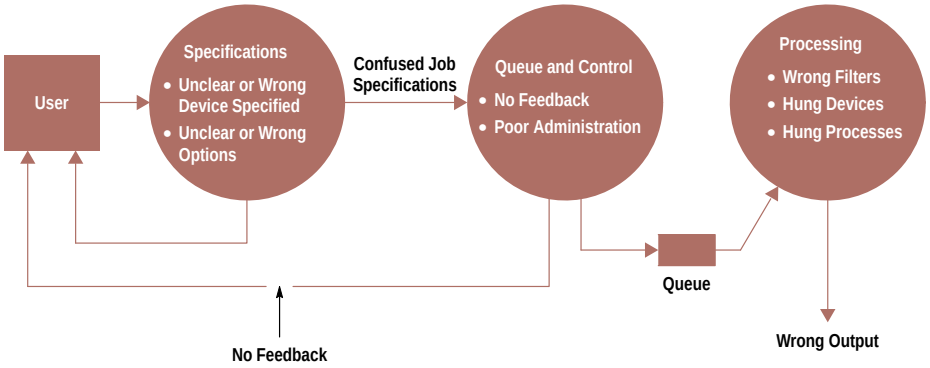

Fig. 2. Problems associated with some spooled printing systems.

The problems encountered during the processing phase can be the most difficult to diagnose and fix. Consider that the print job process is a background process running on a server, which is often remote from the user. If the wrong operations are performed on a print job, the result can be a hung spooling system, no output, or pages of useless output from the printer. Determining the failure requires intimate knowledge of the spooling system and any shell scripts or programs the spool system executes.

HP SharedPrint Solutions

The HP Mpower printing subsystem is based on the HP SharedPrint product. It solves many of the printing problems mentioned above by providing a robust, easy-to-use subsystem layered on top of a spooling system. The HP SharedPrint solution addresses the specification phase by providing a graphical user interface tool that includes a list of printers to choose from and a highlighted list of options that can be specified for a particular printer (see Fig. 3).

HP SharedPrint provides tools that enable the system administrator to set up a printer easily and quickly, making customization trivial. Another graphical user interface tool provides real-time status of print jobs and printers, along with control functions to cancel print jobs, add or delete printers, and modify a printer's default behavior (see Fig. 4). Both of these tools are described in detail later in this article.

The HP SharedPrint processing modules eliminate the major problems that cause system administrators and end users so much trouble. An intelligent server provides print job processing based on an algorithm that dynamically builds and executes the elements necessary to handle the job properly. Built-in error detection logic is used to provide a hard-copy message to the user detailing any problems in processing the job.

HP SharedPrint Overview

HP SharedPrint has a client-server architecture. The client components include the two graphical user interface screens

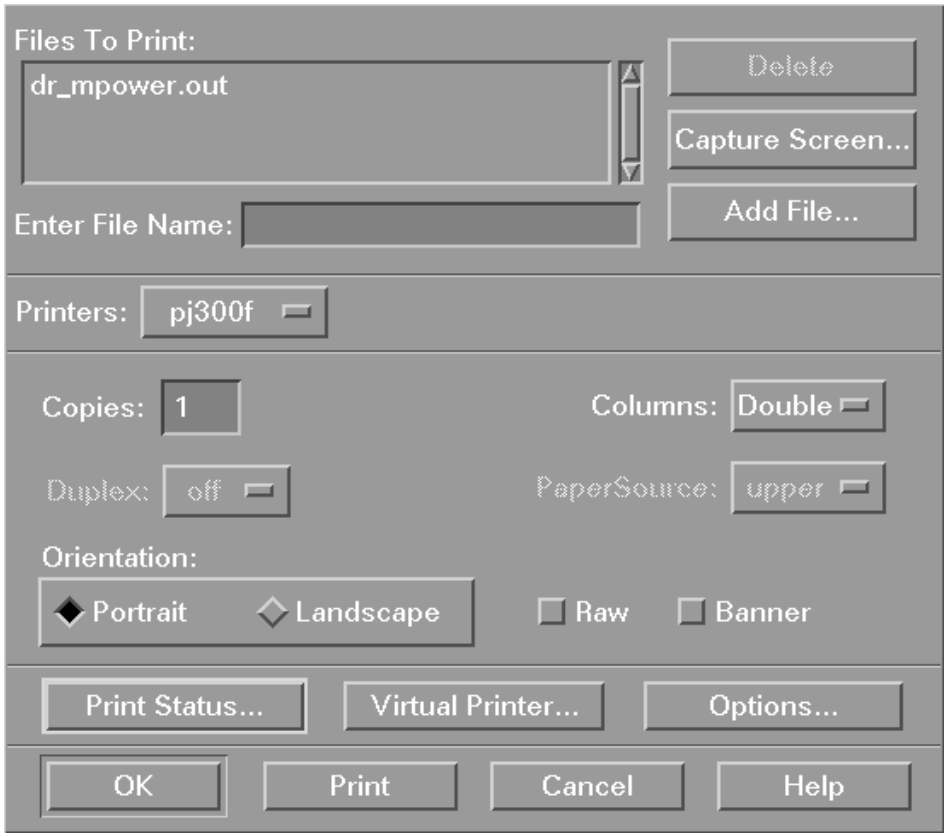

Fig. 3. HP SharedPrint graphical user interface showing a list of available printers.

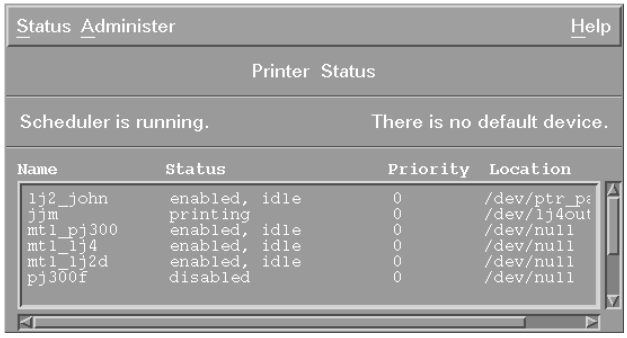

Fig. 4. HP SharedPrint's graphical user interface for system administration.

mentioned above. Fig. 5 shows a high-level view of the HP SharedPrint client/server system. The job specification client communicates with the HP SharedPrint server via the base spooling system. The server provides real-time feedback to the user about what options are valid for a specific job and printer combination.

One of the major features of HP SharedPrint is its ability to print many of the common file types† transparently on any of the supported printers. This functionality eliminates the need for the user to know which file types can be printed on which devices.

The HP SharedPrint server, being highly configurable, can be used to drive smart as well as dumb printers. Also, because of its highly modular design, the server can be layered on any spooling system.

Fig. 6 shows a detailed view of the main components that make up the HP SharedPrint client/server system. The HP SharedPrint client contains the graphical user interface programs sprint (HP SharedPrint print) and spadmin (HP SharedPrint administration). Sprint prompts the user for the file to print (① in Fig. 6) and the designated printer and then uses NCS (network computing services) to connect to the HP SharedPrint feedback server ② and ③. The feedback server hands sprint a set of valid options based on the user's inputs ④. The option information is used to validate which options the user can select for the current job.

After completing the job specification, the user selects the OK button in the sprint window to queue the job for printing. At this point the job request moves to the line printer spooler (lp) client ⑤. The lp client packages the print request and connects to the lpsched daemon located on the lp server ⑥. The lpsched daemon places the job request in the printer queue and at the appropriate time starts the HP SharedPrint job processor to process the job ⑦.

The HP SharedPrint server includes the HP SharedPrint job processor, configuration files that specify various mappings, filters, and drivers that perform translation and formatting, and the HP SharedPrint feedback server. Data from the configuration files is used by the feedback server and the job processor. The feedback server uses the data to provide option information to the sprint user interface, and the job processor uses the data to define the processing steps. The job processor invokes the filters and drivers to finish processing a job. The filters translate and format data files going to the selected printer, and the drivers place device-specific control information in the data stream going to the printer.

The Print Client

The sprint print client enhances the basic HP VUE printing model by providing an easy-to-use graphical interface. Sprint does not alter or replace the current HP-UX* spooling system functionality but is layered on top of the spooling system. This leads to a smooth transition from the standard HP-UX printing models to HP SharedPrint. The sprint graphical user interface allows users to:

- Select a file to print via a drag and drop operation from the HP VUE file manager interface
- Select a printer from a displayed list of printers
- Set options based on real-time feedback
- Save groups of options for reuse later on similar print jobs.

† A file type is a special data format required of all printers, such as PCL image files, RTL (raster transfer language), PostScript, or applications (image or text).

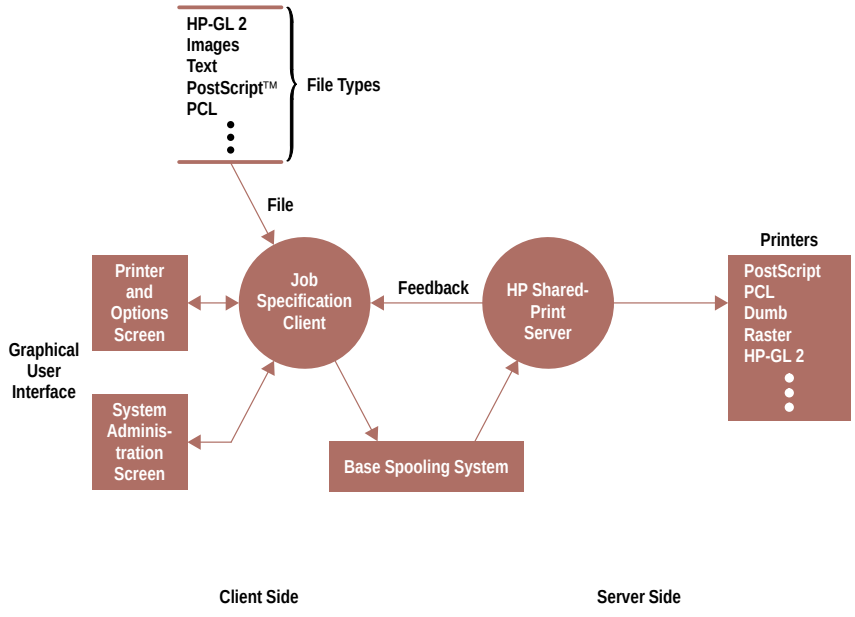

Fig. 5. A high-level view of HP SharedPrint's client/server architecture.

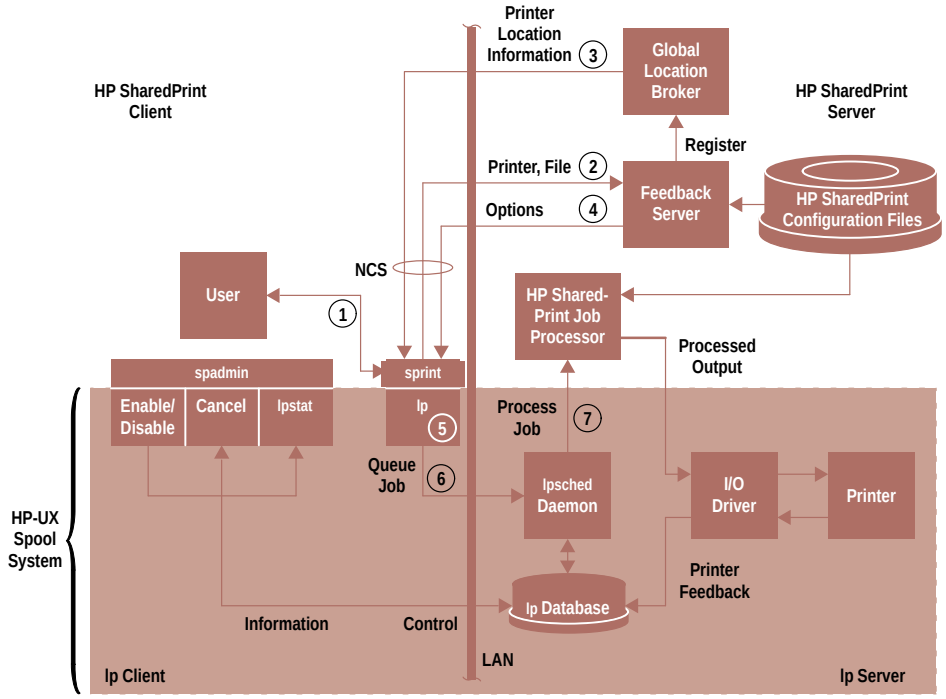

Fig. 6. A detailed block diagram showing the components that make up HP SharedPrint's client/server system.

The top level interface for sprint is shown in Fig. 3. The dimmed buttons represent unavailable items. For example, the menu item PaperSource is dimmed to indicate that the selected printer does not support multiple paper trays.

The button marked Options will take the user to the screen shown in Fig. 7. This screen provides another set of options. These options are sent to sprint by the feedback server.

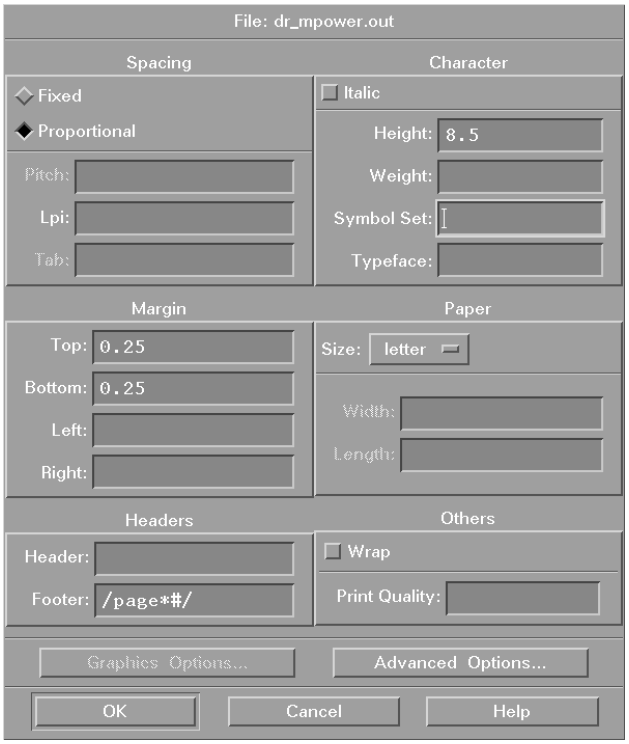

Fig. 7. The Options screen that shows the options available for a certain printer.

The Administration Client

The HP SharedPrint administration (spadmin) graphical user interface is shown in Fig. 4. Selecting the printer icon from the HP VUE front panel activates the spadmin client. Print request or printer status information is displayed and updated every 30 seconds. The Administer menu enables more complex tasks such as editing printer configuration files, adding and deleting printers, and making other changes to the print spooler state.

Core Server

The HP SharedPrint core server components are shown in Fig. 8. The core server performs the processing for a print job and provides the options feedback to the user. Two groups of configuration files are used for processing a job. One group contains information that defines job processing and feedback rules. The other group contains printer configuration information that defines the default settings and

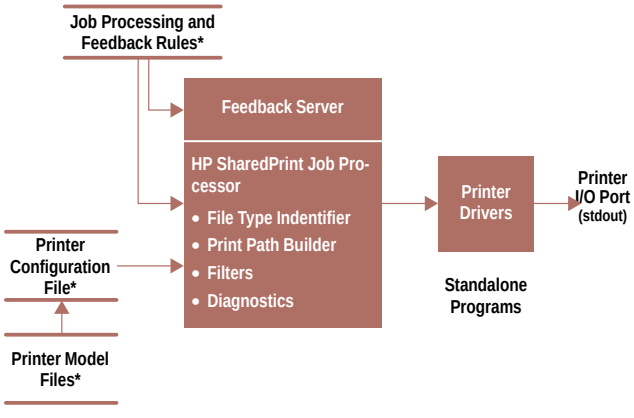

* HP SharedPrint Configuration Files

Fig. 8. HP SharedPrint's core server components.

behavior for a particular printer. There is one printer configuration file for each printer connected to the system. Filters and drivers are standalone programs invoked by the core server to process a print job.

The core server consists of two separate processes that share a common set of data and some common functions. These processes are the job processor and the feedback server. The job processor is invoked by the print daemon lpsched with the arguments:

job_id user_name title copies options files

The first four arguments represent the job number, the user who queued the job, the title of the job, and the number of copies to print. The options argument contains zero or more tokens, each representing an option specified by the user queueing the job. For example –charheight 8 –orientation landscape are typical tokens in the options argument. These options are set in the sprint client or on the lp command line by prepending the letter –o to the option string. The files argument is a list of one or more files to be printed.

The job processor serves all the supported printers by dynamically configuring itself for each print job. The information needed for this operation comes from the printer configuration file. Each printer requires a unique printer configuration file, which is created when a printer is added to the system. This configuration file can be modified at any time via the HP SharedPrint client spadmin or by using any text editor.

The job processor uses some control programs to read the configuration files to determine which set of filters to execute to print a job properly. Details of these steps are described later.

The other process in the core server, the feedback server, is run on every system where an HP SharedPrint printer is installed. This server provides the sprint interface with feedback on what options are valid for a specific file and printer combination.

The feedback server is automatically started when the first HP SharedPrint printer is added to the system. Other HP SharedPrint printers on the same system use the same HP SharedPrint server.

Configuration Files. The configuration files shown in Fig. 8 are text files that provide all the information necessary to process any of the supported file types on any of the supported printers. Users can modify any of these files using an editor, or in the case of the printer configuration file, the HP SharedPrint system administration utility spadmin. The configuration files include:

- Printer model files. These files contain information for specific printer models. For any particular printer model, this file includes the options supported, default values for any filters, and the file types the printer model can handle.
- Printer configuration file. This file contains information about a specific printer. The file is created from one of the printer model files when a printer is added to the system. Users can modify this file to reflect changes made to the printer, such as swapping paper trays or inserting font cartridges. Changes take effect at the next print job.
- Object name extension file. This file contains a list of entries that map file extension values to file types. For example, files with the extension .c (C source files) would be mapped to an ASCII file type.
- Pipeline file. This configuration file specifies a filter pipeline process for each combination of input and output file type. Each of the filter specifications can include options and white space.
- Options map file. This file lists supported options for each filter or driver.
- Options file. This file lists known options and aliases.
- Types file. This file lists known file types and aliases.

These last five configuration files make up the files labeled as job processing and feedback rules in Fig. 8.

Filters. Filters are programs that transform a data file from one format to another. Filters perform either translation or formatting. Translators change the encoding of the information in the data file without changing its appearance. For example, in converting ASCII text to PostScript format the text remains the same, but the file is expanded to contain PostScript commands. Formatters modify how the information in the data file will be rendered on the page. For example, a formatter might rotate or scale an image, add commands for multicolumn text printing, or add footers and headers to a document. These filters often have options that the user can set to control the final appearance.

Filters are invoked by the job processor. Each filter reads from standard input and writes to standard output with error messages being sent to standard error. The filters supported by HP SharedPrint include:

- txpcl. This filter formats and translates text documents into PCL (printer control language) format. The formatting options allow users to change point size, print portrait or landscape, perform double column printing, add headers and footers, and change the typeface and symbol-set mapping.
- txps. This filter formats and translates text to PostScript. It also supports most of the options included in the text-to-PCL filter txpcl.
- cgmhpgl2. This filter reads binary CGM (computer graphics metafile) format and produces an HP-GL 2 byte stream. It is used to obtain 2D graphics output from StarBase or HP-PHIGS CGM files.
- psrip. This filter allows users to print a PostScript file on a non-PostScript printer. The filter reads a level-1 PostScript file and creates a raster image for each page. The raster image is fed to a device-specific formatter which adds the appropriate control or command codes for the target device.
- pclrip. This filter performs the same function as the psrip filter on PCL files. The PCL level can be level 1 through 4 and can include the level-5 PCL scalable typefaces. PCL 3+, used by the HP PaintJet, PaintJet XL and DeskJet 500C printers, is not supported by this translator. The HP-GL 2 extension to level-5 PCL is also not supported by this translator.
- ilFilter. This is an image library filter that converts all HP MPower supported image formats, a level-1 PostScript file, or a PCL 4 file into a Postscript, RTL, or PCL file. The filter is composed of three stages: a producer that converts the input object into an image library format, one or more filters that

manipulate the raster image, and a consumer that converts the image library output to a PostScript, PCL, or RTL file.

Drivers. The function of a driver is to add job control information to the data stream produced by the filters. It does not interpret or change the data stream, but adds printer-specific commands to the byte stream. Think of the driver as handling printer-specific functions, whereas filters handle language-specific functions.

The job control information depends on the target printer. Each job control feature in the printer can be selected or set by the system administrator when the printer is configured or specified by users in their personal options files.

In most cases the driver will send a header with any job control information then cat the input stream from the filters to the output stream. In some cases, the driver will have to examine the first few bytes of a file to see if there are any reset control characters that might interfere with the driver header stream. For instance, PCL files may contain an ESC E that resets the printer state. The driver will have to strip these bytes if a header must be sent.

Processing Steps

For each print job, the HP SharedPrint job processor executes the algorithm outlined in Fig. 9. Each of the steps in Fig. 9 represents a module or inline code involved in processing a job.

Parse Command Line. This step is the interface to the base spooling system. Each spooling system passes the job data to the processing modules in a well-defined form. For the HP-UX spooling system, the job processing module is called via the lpsched daemon with the six command line arguments described earlier. The command line information, which includes the user name, printer name, job options, and files to print is collected and stored in the job processor for use as appropriate.

The main purpose of this module is to enclose the specifics of a spooling system in one location so that the process of porting HP SharedPrint to a new spooling system would involve changing only this module.

Get Printer Type. A key step in properly processing a job is to know the characteristics of the target device. One such printer characteristic is the type of languages or file types the printer understands.

The printer file types are listed in the printer configuration file, which is created when a printer is added to the spooling system. This step involves parsing the configuration file for the target printer. If no file-type entry exists in the configuration file, or the configuration file cannot be located, an error page is sent to the printer, the printer is disabled, and a mail message is sent to the user who queued the job.

Build Options List. Each job includes a list of options and associated values. The options list is built from default values defined in the printer configuration file, values passed by the user, or values added by the filters. User-specified options override options defined in the printer configuration file or the filters.

User-specified options can be passed from sprint via the lp command line. Since the HP SharedPrint option names are

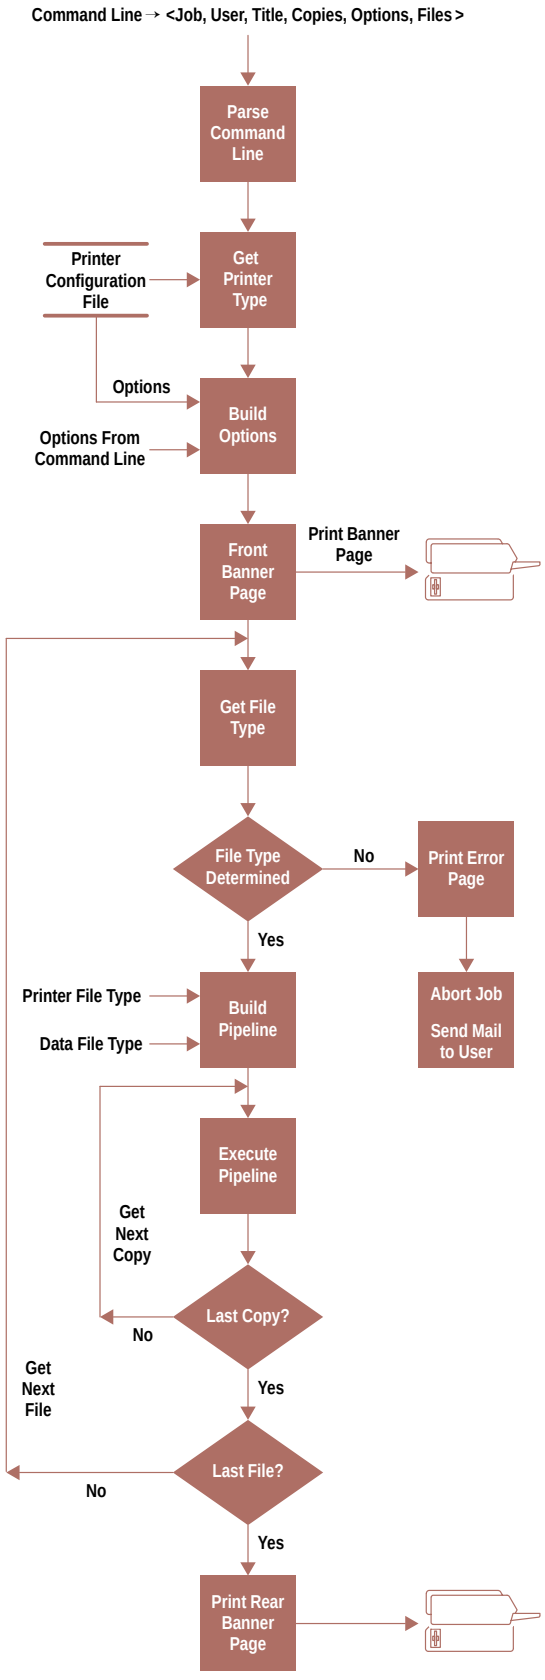

Fig. 9. Job processing algorithm.

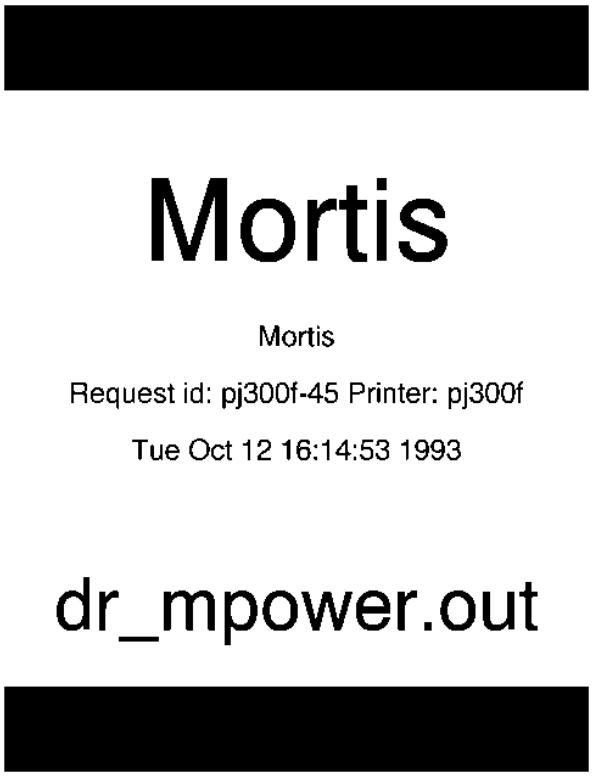

Fig. 10. PostScript banner page.

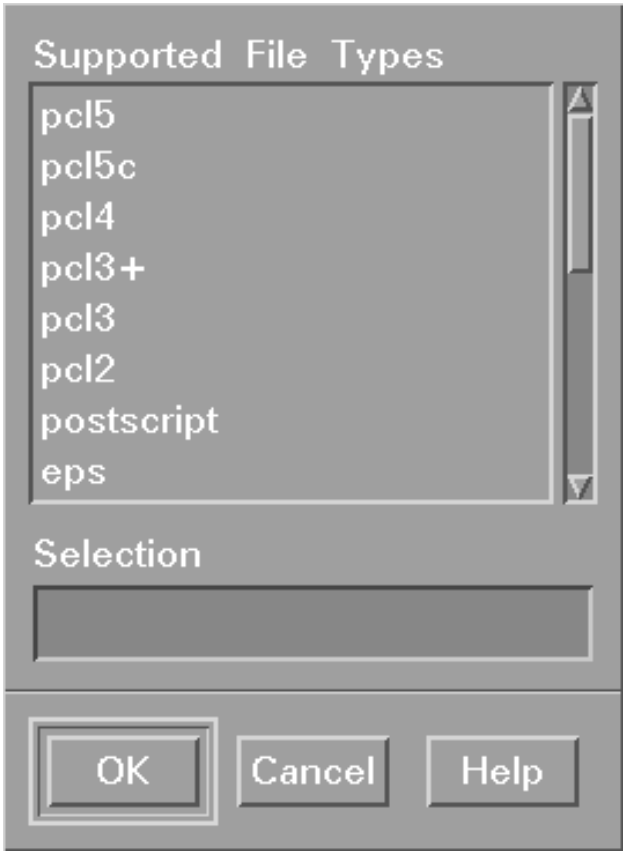

Fig. 11. The sprint dialog box asking about file type.

full-length strings, a set of aliases, which are defined in the options configuration file, can be used on the lp command line. The build-options-list operation will expand the aliased strings when the options list is built.

Front Banner Page. This step involves building and sending a banner page to the printer. This function creates a banner page file, then recursively calls the job processor to send it to the printer. In this way, any type of banner page can be sent to any printer. HP SharedPrint includes a banner page program for PostScript, PCL 5, and ASCII text. Fig. 10 is an example of the PostScript banner page.

Get File Type. This step determines the spooled file's file type. The job processor must know the file type of a job to determine if any processing is necessary to format the spooled file for the target printer. For instance, a PostScript file must be converted to PCL to print on an HP LaserJet printer. The file-typing process uses the following three methods (in the order given) to determine the file type for a job:

- Command line switch
- File type reader
- File extension value.

When a file type is found from any of these methods, the file-typing process terminates. If the file type cannot be determined from any of the above methods, an error page is printed, and the job is aborted.

- Command line switch. With this method if the user includes the file-type option flag (–file_type) on the command line, the file-type value will be set to the value following the –file_type flag. In this case HP SharedPrint assumes the user has some knowledge of the file's content that HP SharedPrint cannot determine. The sprint program will also prompt the user for the file type if the job processor cannot determine the file type. In this case sprint will pop up the dialog box shown in Fig. 11, asking the user to select one of the values.
- File type reader. This method compares the spooled file's contents with a predefined set of patterns for specific file types. These patterns are coded as C programs or Korn shell scripts. Each one of these programs or scripts expects a file as input. If the file type matches the defined pattern, the program or script writes the type string to standard output and returns an exit code of 0 to the calling routine. If there is no match, the program or script returns an exit code of 1.

  A special file-type program, based on the image library, is used for identifying image files. The image library includes a set of functions that allow programmers to open a file and determine if it is a file type that is supported by the image library. If the image library can process the file, the image file-type program writes the string "bitmap" to standard output and exits with a status code of 0.

  The module that determines a job's file type executes each of the file-reading programs and scripts located in a directory called filetypes until a match is found. If no match is found after all the programs have been invoked, the file extension mechanism is used to determine the file type.
- File extension value. File extensions are used by HP VUE to map file types to actions. For example, the extension .au is associated with the audio server, a process that plays or records audio files. HP SharedPrint uses many of the same extension values, mapping them to a file type in the object name extension file. The extension value is extracted from

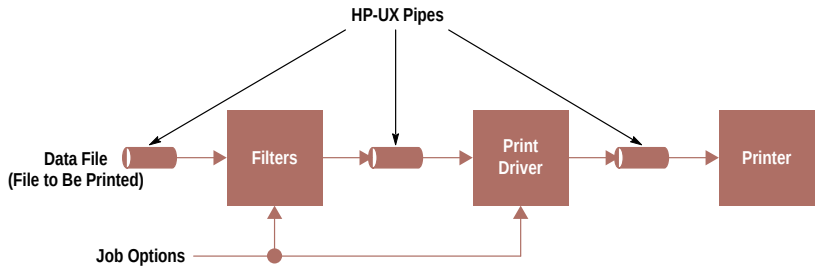

Fig. 12. Pipeline stages created during the pipeline building step in the job processing phase.

the leaf name of the original file that is passed by sprint in the lp spooler title option.

Build Pipeline. The processing pipeline performs the actual work of getting a file printed on the target device. The pipeline consists of one or more filters and a driver, with the output of each stage piped into the input of the next stage as shown in Fig. 12.

The pipeline is derived from the printer file type and the data file type. The pipeline builder scans the pipeline configuration file looking for file type matches. For each match, one or more filters is specified to perform the translation. If no match is found for the specified file types, an error page is printed and the job is aborted.

The pipeline is terminated by adding a driver module to the filter pipeline. The driver is determined from the printer configuration file. The pipeline string at this point will have the format: filter1 | filter2 | driver. Note that each filter might be a filter name followed by some options that control how the filter is executed.

The pipeline is completed by adding the appropriate options to each filter and the driver. The filter options file maps options that apply to each filter or driver so that the correct arguments are passed to the filters. The final pipeline string might look like:

```
txpcl –charheight 8 –orientation landscape | lj3.sh –copies 1
```

Execute Pipeline. The pipeline created in the build-pipeline stage is now executed. The output from the last stage of the pipeline is sent to standard output, which will be the I/O connection to the printer.

If any errors occur in the pipeline execution phase, they are collected and placed on an error page, which is printed at the end of the job. The error page function is called when some unrecoverable event occurs. A special log file contains all the output from the print script. Each filter or driver writes all error or warning messages to standard error. The job processor redirects this text to the error log.

At the completion of the job, the job processor scans the error log. If any information or messages are found in the error log, the job processor builds a text page containing the error messages and some debug information and then prints the text page on the printer.

Diagnostic information is included in the error page to aid in the analysis of the problem.

Final Steps. The last-copy step executes the pipeline again if another copy is needed, and the last-file step cycles the script back through processing the next file if there are multiple files to print. When recycling is required, the initial state of the options, printer configuration file, and output type remain the same and do not need to be recalculated. Finally, the rear banner step, which is identical to the front banner page, may be used to print just a short trailer, or in the case of printers that print face up, the banner page should be printed here.

Feedback Server

The feedback server helps users specify job parameters that optimize the hard-copy output for a given job. Without this feedback, users can be easily overwhelmed by the large number of available options. By enabling only those options that apply to the current job, the feedback server guides the user through the job specification process.

The feedback process shown in Fig. 13 begins with the user selecting a file to print on a specific printer. The sprint program uses the network computing services (NCS) to connect to the server and ask for a list of valid options. The feedback server executes the algorithm shown in Fig. 13 and returns a list of valid options to the client. The sprint program then enables the user to select only the options returned by the feedback server, disallowing (by dimming selections on the menu) other options to the user.

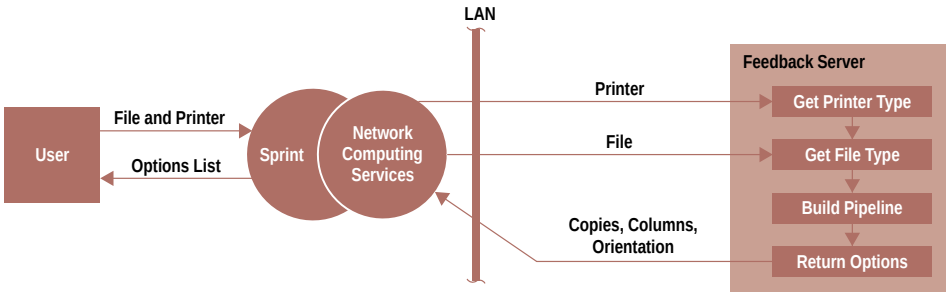

Fig. 13. The feedback server action and data flow.

HP SharedPrint Options

The options that apply to a given print job are a function of not only the target printer, but also any filters that may process the job. In the simplest situation, a file whose type matches the printer type can be sent without any filtering. In this case the set of valid options will be a function of the physical features of the printer, such as duplex or paper tray selection.

Even in the case in which the file type matches the printer, filters can be inserted to perform some preprocessing or formatting. For example, a PostScript document can be passed through a filter that places more than one logical page on a physical page.

HP SharedPrint precomputes the filters that will be used to process a job. This information is then used to extract the supported options from one of the files containing the job processing and feedback rules. It is this set of options that the user sees in the graphical user interface.

Along with each option is the value attached to that option. The filter processing the job must be fed both an option and a value. The source of the value comes from one of three locations: the user, the printer configuration file, or the filter, in descending order. For example, if the user specifies an option and value, this pair is passed to the filter. Otherwise if a value exists in the configuration file, it is used. Finally, if an option and value are not passed to a filter, the filter uses a built-in default value for the option. In this way, the behavior of a filter is controlled first by the user, then the system administrator (via the configuration file), and finally the programmer.

One problem with options is that there is no unified naming scheme. Two filters may provide the same functionality and use different strings to denote an option. HP SharedPrint addresses this by providing the options file to alias filter option names to a set of names defined by HP SharedPrint.

Conclusion

HP SharedPrint provides HP MPower with a robust printing solution. Users can drag and drop any HP MPower object to any printer and use the HP SharedPrint graphical user interface to control the printing process. The graphical user interface decreases user frustration, providing an easy-to-use printing interface. The intelligent server lowers the system administration burden by eliminating catastrophic failures and providing the flexibility to customize a printer's behavior.

PostScript is a trademark of Adobe Systems Incorporated which may be registered in certain jurisdictions.

HP-UX is based on and is compatible with UNIX System Laboratories' UNIX* operating system. It also complies with X/Open's* XPG3, POSIX 1003.1 and SVID2 interface specifications.

UNIX is a registered trademark of UNIX System Laboratories Inc. in the U.S.A. and other countries.

X/Open is a trademark of X/Open Company Limited in the UK and other countries.