HP 3000 Computer Systems

Special Documentation Package
MPE V/P

Release E/F.00.00

HEWLETT PACKARD

CONVENTIONS USED IN THIS MANUAL

NOTATION	DESCRIPTION

COMMAND

Commands are shown in CAPITAL LETTERS. The names must contain no blanks and be delimited by a non-alphabetic character (usually a blank).

KEYWORDS

Literal keywords, which are entered optionally but exactly as specified, appear in CAPITAL LETTERS.

parameter

Required parameters, for which you must substitute a value, appear in ***bold italics***.

parameter

Optional parameters, for which you may substitute a value, appear in *standard italics*.

[ ]

An element inside brackets is optional. Several elements stacked inside a pair of brackets means the user may select any one or none of these elements.

Example: [A]
          [B]   user may select A or B or neither.

When brackets are nested, parameters in inner brackets can only be specified if parameters in outer brackets or comma place-holders are specified.

Example: *[parm1[,parm2[,parm3]]]*   may be entered as:

    parm1,parm2,parm3   or
    parm1,,parm3   or
    ,,parm3   , etc.

{ }

When several elements are stacked within braces the user *must* select one of these elements.

Example: { A }
          { B }   user must select A or B.

...

An ellipsis indicates that a previous bracketed element may be repeated, or that elements have been omitted.

<u>user input</u>

In examples of interactive dialog, user input is underlined.
Example: NEW NAME? <u>ALPHA1</u>

superscript^c

Control characters are indicated by a superscript^c. Example: Y^c. (Press Y and the CNTL key simultaneously.)

[RETURN]

[RETURN] indicates the carriage return key.

<<COMMENT>>

Programmer's comments in listings appear within << >>.

** Comment **

Editor's comments appear in this form.

CONTENTS

.

Welcome to release E/F.00.00, the first release of MPE V/P. This Special Documentation Package will supplement your MPE IV manuals with descriptions of the new features of MPE V/P; your new MPE V/P manuals will be shipped to you as soon as possible. Please take a moment to complete and mail the form on the last page.

MPE V/P provides the additional operating system performance gained from disc caching, and disc caching enhancements to IMAGE/3000 and OPT/3000. The article "MPE V/P Offers Increased Performance" discusses disc caching and is followed by "Disc Caching Commands" detailing the disc caching commands: :SHOWCACHE, :STARTCACHE, :STOPCACHE, and :CACHECONTROL. Additional articles elaborate on disc caching with IMAGE/3000 and OPT/3000.

On the HP 3000 operating system, the MPE version number has been modified to state the software base V.UU.FF. For your preference this information can be accessed through the logon or :SHOWME command. The article "MPE Reports Its Base V.UU.FF" explains both methods.

You may have expanded your directory since release C/D.01.00. The article "Directory Expanded?" provides procedures to check if your system includes an expanded directory, and how to update the directory for disc caching.

Our documentation plan for the HP 3000 operating system has undergone several changes. Most of the manuals have new part numbers for the MPE V/P operating system. Some manuals have been combined, and other manuals are new. Additional information on documentation can be found in "MPE V Documentation Has a New Look".

New Software Support Service product numbers have been established for MPE V/P. If you currently have a Software Support Service contract (CSS, PICS, SS, or MUS), please contact your HP Sales Representative or System Engineer to ensure that you have the correct support product.

Marissa Kobylenski
Editor

MPE V/P Offers Increased Performance

By Jeff Byrne and Larry Russell

The E/F.00.00 release of MPE V/P adds support for the disc caching MPE product. This innovative product significantly improves the response time and transaction throughput of I/O- and database-intensive applications running on MPE V/P. This release is supported on HP 3000 High Performance Series 39, Series 42, 48, and 68. Disc caching comes standard with the Series 42, 48, and 68. Disc caching field upgrade products are available to upgrade a Series 39 to a High Performance Series 39, a Series 40 or 40SX to a Series 42, a Series 44 to a Series 48, and a Series 64 to a Series 68, with more than 2 Mb of main memory.

DISC CACHING ADVANTAGES

Disc caching enhances system performance by using main memory as a high speed buffer for disc resident information. Disc caching places those most frequently-accessed portions of files and directories in available regions of main memory where they can be accessed repeatedly by an executing program. Thus, instead of going immediately to the disc to get requested information, disc caching first checks to see if the file information is already in main memory (in the disc cache). When the information is located in the disc cache, I/O performance is improved in two principal ways. First, several time-consuming disc accesses are eliminated. Second, information may be accessed in main memory 10 to 100 times faster than it may be accessed from disc; disc caching greatly reduces the time required for an I/O transaction to complete. This improvement means better response and higher throughput for I/O-intensive applications.

While most users of I/O-intensive applications will experience significant performance gains with disc caching, it is most effective in applications in which multiple users share the same information. Programs which share data bases, such as IMAGE/3000, benefit greatly. Disc caching not only reduces the contention for these data bases, but it also reduces the lock service times of the data base control blocks because I/O requests can be serviced more quickly.

MEMORY CONSIDERATIONS

The MPE Memory Manager controls the operation of disc caching, and handles file information in the cache dynamically and transparently, in the same manner as a data segment would be treated. As a result, special areas of main memory need not be dedicated to the disc cache. However, users who wish to gain significant performance benefits from disc caching must install an adequate amount of main memory on their systems. Your local SE or Performance Specialist SE can help you decide how much additional memory will be required, if any.

SYSTEM MANAGER OR SYSTEM SUPERVISOR INTERACTION

Although the operation of disc caching is fully transparent to the user, the System Manager or System Supervisor can enable, disable, monitor, and control caching through designated commands. Since disc caching can be enabled on a device-by-device basis, only those disc drives which will benefit from caching need to have caching enabled. Full guidelines on the use of all the disc caching commands are provided in the article "Disc Caching Commands" in this issue.

Disc Caching Commands

By Al Kondoff

Many MPE users will have a significant improvement in system performance as a result of disc caching. Currently disc caching is available only with MPE V/P on High Performance Series 39, 42, 48, and 68. This new feature speeds up disc file reads by bringing larger pieces of the disc files into memory when needed by programs, and by retaining these pieces in memory so that they are available for repeated references. These repeated references avoid the overhead of additional input requests to the disc. Writes to disc files will also have improved performance, since cached writes need only be moved into cached memory, rather than waiting for the physical disc updates to complete. This allows programs to continue running while the physical disc updates occur.

Disc caching is most useful in systems with excess main memory and processor capacity. Available unused memory is consumed by cached disc domains. Disc references to these domains are resolved with CPU MOVE instructions rather than physical disc accesses.

The operation of disc caching is transparent to the programs running in a cached system. Process accounting on a caching system differs from process accounting on an uncached system, because processes will be charged for CPU time associated with disc caching. In general, CPU utilization increases and disc utilization decreases in a cached system. The system and user processes in particular, spend much less time waiting for disc file input/output to complete.

To control disc caching, four new commands have been added to the operator command set. These commands enable or disable disc caching per device, and monitor and control disc caching. In addition to these new commands, the FSETMODE intrinsic has been extended to allow refined control over disc caching management on a file basis, as opposed to a system-wide basis. The new commands and new intrinsic are described below:

STARTCACHE

```
:STARTCACHE ldev
```

This command is used to enable caching on a single disc and requires System Supervisor (OP) or System Manager (SM) capabilities. Each disc will require a separate invocation of the command. The parameter *ldev* is the logical device number of the device to be cached. This device must be a disc, and the system must have the caching option installed. When the system is started, caching is disabled by default. Caching can be started by :STARTCACHE commands in the System Operator's startup UDC.

STOPCACHE

```
:STOPCACHE ldev
```

This command is used to disable caching on a single disc and requires System Supervisor (OP) or System Manager (SM) capabilities. Again each disc will require a separate invocation of the command. The *ldev* is the logical device number of the device to have caching disabled.

CACHECONTROL

```
                {SEQUENTIAL=sequentn}
                {RANDOM=randomn      }
:CACHECONTROL                         [,...]
                            {YES}
                {BLOCKONWRITE={NO }  }
```

PARAMETERS

sequentn

Assigns the number of sectors read from disc of a sequential file, if the information is not found in cache. The disc read will stop at extent boundaries and will fetch at least the amount requested. The number of sectors must be greater than 1, but less than 96 (inclusive). Default is 96 sectors.

randomn

Specifies the number of sectors read from disc of a random access file of a cache read, if the information is not found in cache. The disc read will stop at extent boundaries, and will fetch at least the amount requested. The number of the sectors must be greater than 1, but less than 96 (inclusive). Default is 16 sectors.

BLOCKONWRITE

Specifies whether or not to block the process until posting of cache buffers to disc completes. The BLOCKONWRITE must be equal to YES or NO. If BLOCKONWRITE=YES, enabling the Serial Write Queue or disabling BLOCKONWRITE for a file is overridden. Default is NO.

The CACHECONTROL command is used to tune the performance of caching on a running system. This command requires System Manager or System Supervisor capabilities. The :CACHECONTROL requires that at least one keyword be provided. Any two or all three keywords can be used in a single command, but each keyword can only be used once. The order in which the keywords are specified is not significant. If a keyword is specified more than once, the system will give a warning and the last value specified will be used.

SHOWCACHE

```
:SHOWCACHE
```

The :SHOWCACHE command provides a brief description of caching performance. No capabilities are required. For example enter:

:SHOWCACHE

DISC LDEV	CACHE REQUESTS	READ HIT%	WRITE HIT%	READ%	PROCESS STOPS	K-BYTES	% OF MEMORY	CACHE DOMAINS
1	105542	89	63	86	9897	69	1	51
2	191487	91	54	55	10794	714	17	225
3	99868	85	87	21	14895	577	14	63
4	39416	90	45	55	2371	117	2	25
Total	436313	90	67	55	37957	1477	37	364

Data overhead is 94K bytes, 49% of user I/Os eliminated.
Sequential fetch quantum is 96 sectors.
Random fetch quantum is 16 sectors.
Blockonwrite=no

The following descriptions explain the above summary, since disc caching was enabled on the device.

COLUMN HEADING	MEANING
DISC LDEV	Logical device number(s) of the disc(s) which is being cached. Read across the row to see the other entries on a per disc basis. The bottom row is a total across all discs.
CACHE REQUESTS	Tally of the number of logical I/O requests against caching. This is a double word counter, and will reset to zero automatically when caching is started or when the counter overflows.
READ HIT%	Cumulative percentage of logical disc I/O in the system for which the request was a read request, and the needed data was in a main memory cache domain when the read request occurred.
WRITE HIT%	Cumulative percentage of logical disc I/O in the system for which the request was a write request, and the needed data was in a main memory cache domain when the write request occurred.
READ%	Percentage of logical disc I/Os in the system which were read requests.
PROCESS STOPS	Number of times a process was stopped in order for a cache I/O request to complete. Process stops occur due to read misses, or due to write hits in which the cached region is currently being posted to disc. This is a single word counter which will reset to zero, if it overflows.
K-BYTES	Count of the number of thousands of bytes (1024-byte groups) which are currently being used for main memory cache domains.
% OF MEMORY	Percentage of the total memory in the machine which is currently allocated for use by cache domains.
CACHE DOMAINS	Number of cache domains in memory.

5

Data overhead is the current dynamic data structure overhead used to support disc caching. Data overhead does not include the portions of disc, as reflected in the K-BYTES column, in main memory.

FSETMODE INTRINSIC

The FSETMODE intrinsic has been expanded to include the BLOCKONWRITE and "Serial Write Queue" parameters. The following syntax illustrates the FSETMODE intrinsic as it specifically relates to disc caching. For additional information on the FSETMODE intrinsic, please refer the to the MPE V Intrinsics Reference Manual (32033-90007).

SYNTAX

```
                  IV        LV
     FSETMODE (filenum,modeflags);
```

PARAMETERS

filenum
integer by value
A word identifier supplying the file number of the file to which the call applies.

modeflags
logical by value
A 16-bit value that denotes the access mode options in effect, as described below.

Bit (14:1) - BLOCKONWRITE for this file. This bit the following settings:

=1 Enabled (YES).

=0 Disabled (NO). Default.

Bit (15:1) - Serial Write Queue. This bit has the following settings:

=1 Enabled for writes (YES).

=0 Disabled for writes (NO). Default.

The FSETMODE intrinsic has been extended to allow access to the MPE global Serial Write Queue and BLOCKONWRITE parameter on a file-by-file basis. By enabling the Serial Write Queue on a file, all write requests are guaranteed to be performed in the order they were issued. The Serial Write Queue does this to preserve integrity while allowing the performance benefits of uninterrupted process execution. (BLOCKONWRITE=NO). Since this option may have an effect on total system performance, please consult your SE or SEO (Service Engineering Organization) before enabling it. The Serial Write Queue is globally disabled if the :CACHECONTROL BLOCKONWRITE=YES was specified, since this causes all processes on the system to wait for their physical disc I/O to complete.

BLOCKONWRITE management can also be controlled on a file-by-file basis with the FSETMODE intrinsic. If BLOCKONWRITE is locally enabled, the user's application will not be notified of a disc

6

write completion until the physical write operation has completed. Therefore, the BLOCKONWRITE option guarantees "commitment" of the transaction to disc prior to completion notification to the user's application. The FSETMODE setting of an individual file will be overridden if the :CACHECONTROL BLOCKONWRITE=YES is specified.

A combination of these two options in FSETMODE provides complete integrity and transaction commitment notification with disc caching, with minimal performance impact. Many disc writes can be performed with BLOCKONWRITE=NO, serial write queue enabled, and guarantee integrity. On the final disc write of a transaction, both localized BLOCKONWRITE=YES and Serial Write Queue enabled can be specified to provide both integrity and total transaction commitment notification.

ADDITIONAL CONSIDERATIONS

With disc caching enabled, the system can have many writes pending against the discs, as opposed to several writes in non-cached MPE. In the event of a system failure, the potential for losing writes to disc is increased. Again, it is possible to globally enable :CACHECONTROL BLOCKONWRITE=YES to limit the write exposure, so it is equivalent to non-cached MPE, but at the expense of some caching performance.

MPE-managed subsystems, such as IMAGE and KSAM, have been enhanced to utilize disc caching synchronization and transaction commitment notification. If other user-managed disc files are sensitive to the order in which writes are performed, two options are available. Application programs can be modified to call FSETMODE for these sensitive files, or the global :CACHECONTROL BLOCKONWRITE=YES can be specified after caching is enabled. The :CACHECONTROL option will negate a portion of the benefit provided by disc caching, but will provide total compatibility with non-cached MPE.

7

IMAGE/3000 Changes for Disc Cache

By Doug Griffin

IMAGE/3000 version B.04.00 has been expanded to support larger Process Identification Numbers (PINs), larger Extra Data Segment numbers, and disc caching.

To facilitate future MPE support of the larger PINs, IMAGE/3000 has expanded the size of the PIN in its internal tables. Because of future larger PINs, each lock descriptor in the lock area of the Data Base Control Block will be 2 to 4 words longer. This increase in the size of lock descriptors may cause more occurrences of IMAGE/3000 status 62 :"Lock area is greater than 4096 words" or "Data Base Control Block size is greater than 31000 words."

Changes have also been made to support larger Extra Data Segment numbers. IMAGE/3000 now uses a System Data Base Control Block (SDBCB) to hold the Extra Data Segment numbers of all the data bases on the system. The first word of the Base-ID (returned by DBOPEN) will no longer contain the Extra Data Segment number of the Data Base Control Block (DBCB). Instead, the Base-ID will contain the access path number in the first 8 bits and an index into the SDBCB where IMAGE/3000 will locate the Extra Data Segment number of the specific data base. Programs which use the Base-ID for purposes other than IMAGE/3000 intrinsics may need modifications.

When IMAGE/3000 logging is enabled for a data base, DBEND modes 1 and 2 will always wait for the log record to be posted to disc before allowing a program to continue. This will ensure that logical transactions which have completed will always be recovered during roll forward recovery, plus ensure posting of completed transactions to disc before MPE disc caching allows the program to continue. To ensure all writes of a transaction are posted to disc before the DBEND log record is written, a program only waits for the DBEND log record to post to disc. IMAGE/3000 serializes all writes by using disc caching's "System Serial Write Queue".

MPE Reports Its Base V.UU.FF

By Jon Cohen

By using the :SYSDUMP command, any user with System Manager (SM) or System Supervisor (OP) capabilities can generate a modified version of the MPE system tape. Usually, this system tape is simply a backup of the system. Occasionally, users build their own versions of the system software with :SYSDUMP. For example, a "patched" system tape is often generated with :SYSDUMP.

One of the modifiable aspects of MPE is the Version, Update, and Fix number associated with the system software. It is reported in the form *V.UU.FF* when the user logs on, and with the :SHOWME command.

However, changing the *V.UU.FF* can cause a problem for HP support personnel. If a user reports a problem on a non-standard *V.UU.FF*, and the problem is sent to the factory for resolution, we have no idea which version of the software to examine. With this new software release, MPE now reports both the user-defined *V.UU.FF* number and the official base *V.UU.FF* of the HP version of the operating system software. (Note that the official *V.UU.FF* is not modifiable by a :SYSDUMP.) The base *V.UU.FF* will agree with the version number on the Printing History page of the applicable MPE Manual.

Starting with this software release, the logon banner looks like:

```
HP3000 / MPE V.  X.YY.ZZ (BASE E.00.00) TUE, NOV 1,...
```

where *X* is the user-defined version, *YY* is the user-defined update number, and *ZZ* is the user-defined fix number. In the sample output of the :SHOWME command below, E.00.00 is the base *V.UU.FF*:

```
USER: #S1269 JON.OSE,COHEN      (IN PROGRAM)
MPE VERSION: HP32033X.YY.ZZ.      (BASE E.00.00)
CURRENT: THU, OCT 27, 1983,  4:54 PM
LOGON:   THU, OCT 27, 1983,  8:53 PM
CPU SECONDS: 9          CONNECT MINUTES: 482
$STDIN LDEV: 68          $STDLIST LDEV: 68
```

Directory Expanded?

By Don Darnell

To ensure that installation of MPE V/P will be successful, customers should verify the directory size is less than 6,000 sectors. Follow this procedure:

1. The current size of the directory can be verified by performing a "SYSDUMP $NULL". When the question "ANY CHANGES?", is displayed, enter:

   <u>YES</u>

2. Press (RETURN) for each additional question until "DIRECTORY SPACE ALLOCATION?" is displayed. Enter:

   <u>YES</u>

   SYSDUMP will output:

   DIRECTORY USED = $xxxx$, MIN = $yyyyy$, MAX = $zzzz$.?

3. If the MAX=$zzzz$ value is more than 6000 sectors, do not update to MPE V/P. If you have any questions please call your SE or SEO (Service Engineering Organization). If the value is 6,000 sectors or less, the customer can proceed to update to MPE V/P.

MPE V/P Documentation
Has A New Look

By Brian B Egan

In keeping with the spirit of innovation present in the release of MPE V/P, the MPE Documentation Group is proud to announce the release of a new manual set for MPE V. Figure 1 gives an overview of the manuals in the set.

The Commands, Intrinsics, and Utilities manuals, and the Communicator, have all been completely revised for MPE V/P, and have been given new part numbers. (The MPE IV versions will continue to be available.)

In response to reader surveys and comments, the System Manager/System Supervisor manual and the Console Operator's Guide have been combined into an all-new manual, the MPE V System Operation and Resource Management Reference Manual (32033-90005). The new manual has been extensively rewritten to clarify old concepts and to add new material.

A new manual is currently in progress the Guide for the New User (32033-90009). This book, written in a friendly, easy to use format, is intended to introduce the HP 3000 to those who have had no previous experience. We plan on releasing this book in the next year.

We are also currently producing a new "Pocket Guide": the MPE Quick Reference Guide, intended for publication next year. This book will cover both MPE IV and MPE V, and will be produced in a handy, easily-updated format.

Finally, two MPE IV manuals have been obsoleted for MPE V. The Index to MPE Documentation will be replaced by a comprehensive second index in every MPE V manual. The System Error Message and Recovery Manual will be replaced by enhancements to the system message catalog and HELP facility, presenting the user with recovery information on the screen at the same time that an error message occurs.

The Segmenter, Debug/Stack Dump, and File System manuals will continue in their present forms as joint MPE IV/MPE V/P documents. They will be revised as soon as possible to include the latest information.

The MPE Documentation Group is committed to providing you with the most accurate, timely and usable documentation possible. We welcome your comments and suggestions; please submit them via the postpaid Reader Comment Sheets found at the back of each manual, or by a Service Request (SR).

INTRODUCTORY LEVEL:

GENERAL INFORMATION Manual 5953-7553	GUIDE FOR THE NEW USER 32033-90009

STANDARD USER LEVEL:

MPE V COMMANDS Reference Manual 32033-90006	MPE V INTRINSICS Reference Manual 32033-90007	MPE V UTILITIES Reference Manual 32033-90008
SEGMENTER Reference Manual 30000-90011	DEBUG/STACK DUMP Reference Manual 30000-90012	FILE SYSTEM Reference Manual 30000-90236

ADMINISTRATIVE LEVEL:

MPE V SYSTEM OPERATION & RESOURCE MANAGEMENT Reference Manual 32033-90005

SUMMARY LEVEL:

MPE QUICK REFERENCE GUIDE 30000-90049

Figure 1. MPE V/P Manual Set

12

OPT/3000 Disc Caching Enhancements

By Bruno Morosin

The new version A.00.12 of OPT/3000 is offered on MPE V/P. Several enhancements have been made to OPT/3000 as a result of disc caching, the new feature available with MPE V/P. As a result, many displays have been modified to illustrate the use of disc caching on various system resources.

In this article, the various modifications to OPT/3000 are summarized. It is recommended that this section be removed from the Special Documentation Package and appended to the current OPT/3000 Reference Manual (32238-90001) until the manual is updated.

TERMINOLOGY

Disc caching:
: The process of placing the most frequently accessed portions of files and directories into available regions of main memory.

Cached disc domain:
: A type of segment, analogous to a code or data segment, which may be referenced by MPE disc caching. The maximum attainable size of a cached disc domain is 32K words.

Mapped disc domain:
: A cached disc domain which is currently being referenced, or is in transit between main memory and disc. The domain also has a mapped entry in the Cache Directory table.

Unmapped disc domain:
: A cached disc domain which is neither currently being accessed, nor in transit between the disc and main memory. This domain may be swapped out of memory when space is needed.

MEMORY CONTEXT

Several new displays have been added to the Memory Context to reflect the changes due to disc caching.

MEMORY REPORT DISPLAY. A new field has been added to the ALL MEMORY percentage bar. The " X:Cache Domain" represents the amount of linked memory occupied by cached disc domains.

MEMORY CONTENTS DISPLAY and **BANK CONTENTS DISPLAY**. Both displays have added the character "X" to their legends to represent the size of each individual cached disc domain. In A.00.12 and previous versions of OPT/3000, upper case letters represented user segments, and lower case letters represented system segments. However, in the case of disc caching, an upper case " X " will denote a mapped disc domain, and a lower case " x " will denote an unmapped disc domain.

13

GLOBAL CONTEXT

The Global Resource Usage Report contains changes to the MEMORY USAGE percentage bar and the DISC I/O ACTIVITY bar. The MEMORY USAGE percentage bar is now subdivided into five rather than four different states. The added state is " X:Cached Disc Domains". This field defines the percentage of linked memory occupied by cached disc domains. The DISC I/O ACTIVITY bar has also been modified to reflect disc caching. The " M:Memory Manager " field has been replaced by two new fields: " X:Cached I/O", and " S:Segment Manager ". The " X" state represents the ratio of disc I/Os completed on the behalf of the disc cache management. The " S" state represents the disc I/O activity completed on behalf of the segment management.

CPU MANAGER MEMORY CONTEXT

The CPU MANAGER MEMORY CONTEXT has been changed to reflect disc caching.

CPU/MEMORY MANAGEMENT REPORT and CPU USAGE DISPLAY. The CPU STATE percentage bar contains nine rather than eight CPU states. The " M:Memory Allocation" state has been replaced by two new states: " X:Mem Alloc (Cached Management)" and " M:Mem Alloc (Seg)". The CPU time spent on memory allocation for cache domains is the " X" state, and the CPU time spent on the allocation of segments, is represented by the " M" state.

MEMORY MANAGEMENT ACTIVITY DISPLAY. A new rate bar has been added for disc caching events. The REL CACHE DOMAINS bar represents the number of cache domains released from main memory. The MM I/O READ/WRITE bars have been redefined as SEGMENT READ and SEGMENT WRITE bars representing the number of segment management disc reads/writes performed.

I/O CONTEXT AND ACTIVITY

The I/O Activity Report has new rate bars indicating disc I/O completion rates. The " M: Memory Manager " field has been replaced by two new fields: " X:Cached I/O", representing the number of disc I/Os completed on behalf of cached disc domains, and " S:Segment Manager I/O", representing the disc I/Os completed on behalf of segment management.

OPT/3000 LOG FILE

New information has been added to the OPT/3000 log file summary records. Programs referencing earlier versions of the log file will not be affected by the added events. All summary records of the log file, except 0, have been modified in the following manner:

WORD (OCTAL)   EVENT

26-27	Total time (in milliseconds) spent on the user's stack processing caching code since the start of summary report generation.
30-31	The total time (in milliseconds) spent on the Interrupt Control Stack (ICS) processing caching code since summary report generation started.
32-33	The total CPU time spent on disc caching memory allocation.
34-35	Number of disc cache read hits.

14

36-37	The number of disc cache write hits.
76-77	The number of disc cache domain releases.

ADDITIONS AS OF RELEASE A.00.10 OF OPT/3000

134-135	All discs with queue length = 0.
136-137	All discs with queue length = 1.
140-141	All discs with queue length = 2.
142-143	All discs with queue length = 3.
144-145	All discs with queue length = 4.
146-147	All discs with queue length = 5.
150-151	All discs with queue length > 5.

ADDITIONS AS OF RELEASE A.00.12 OF OPT/3000

152-153	The number of disc cache reads for all discs.
154-155	The number of disc cache writes for all discs.

EXISTING FIELDS

156-157	The number of memory manager disc reads (Disc 1).
160-161	The number of memory manager disc writes (Disc 1).
162-163	The number of control operations (Disc 1).
164-165	The number of user disc reads (Disc 1).
166-167	The number of user disc writes (Disc 1).

ADDITIONS AS OF RELEASE A.00.10 OF OPT/3000

170-171	Number of times queue length = 0 (Disc 1).
172-173	Number of times queue length = 1 (Disc 1).
174-175	Number of times queue length = 2 (Disc 1).
176-177	Number of times queue length = 3 (Disc 1).
200-201	Number of times queue length = 4 (Disc 1).
202-203	Number of times queue length = 5 (Disc 1).

| 204-205 | Number of times queue length > 5 (Disc 1). |

ADDITIONS AS OF RELEASE A.00.12 OF OPT/3000

| 206-207 | The number of disc cache reads (Disc 1). |
| 210-211 | The number of disc cache writes (Disc 1). |
- •
- •
- •

The following sequence is repeated for each additional disc. The numbers will begin at 212.

EXISTING FIELDS

The number of memory manager disc reads (Disc n).

The number of memory manager disc writes (Disc n).

The number of control operations (Disc n).

The number of user disc reads (Disc n).

The number of user disc writes (Disc n).

ADDITIONS AS OF RELEASE A.00.10 OF OPT/3000

Number or times queue length = 0 (Disc n).

Number of times queue length = 1 (Disc n)

Number of times queue length = 2 (Disc n).

Number of times queue length = 3 (Disc n).

Number of times queue length = 4 (Disc n).

Number of times queue length = 5 (Disc n).

Number of times queue length > 5 (Disc n).

ADDITIONS AS OF RELEASE A.00.12 OF OPT/3000

The number of disc cache reads (Disc n).

The number of disc cache writes (Disc n).

IMPORTANT

Your new MPE V/P manuals are in production now, and they will be mailed to you as soon as they are complete. In order to ensure that you receive them as quickly as possible, please complete the following form, remove this whole page from the Special Documentation Package, and mail it.

This form requires no postage stamp if mailed in the U.S. For locations outside the U.S., please give it to your local Hewlett-Packard representative, who will forward the form to us.

Name: _____

Company Name: _____

Address: _____

_____

City/State: _____

Country: _____

BUSINESS REPLY MAIL

FIRST CLASS   PERMIT NO. 1070   CUPERTINO, CALIFORNIA

POSTAGE WILL BE PAID BY ADDRESSEE

Publications Manager
Hewlett-Packard Company
Computer Systems Division
19447 Pruneridge Avenue
Cupertino, California 95014

HEWLETT
PACKARD