

North American Response Centers

# HP 3000 APPLICATION NOTE #19

## STACK OVERFLOWS:

### Causes and Cures For COBOL II Programs



 **HEWLETT  
PACKARD**

December 1, 1986  
Document P/N 5958-5824/2649

**HP Computer Museum**  
**[www.hpmuseum.net](http://www.hpmuseum.net)**

**For research and education purposes only.**

HP 3000 APPLICATION NOTES are published by the North American Response Centers twice a month and distributed with the Software Status Bulletin. These notes address topics, where the volume of calls received at the Centers indicates a need for addition to or consolidation of information available through HP support services. You may obtain previous notes (single copies only, please) by returning the attached Reader Comment Sheet listing their numbers.

<u>Note #</u>	<u>Published</u>	<u>Topic</u>
1	2/21/85	HP 3000 Printer Configuration Guide (superseded by note #4)
2	10/15/85	Terminal Types for HP 3000 HPIB Computers (superseded by note #13)
3	4/01/86	HP 3000 Plotter Configuration Guide
4	4/15/86	HP 3000 Printer Configuration Guide - Revised
5	5/01/86	MPE System Logfile Record Formats
6	5/15/86	HP 3000 Stack Operation
7	6/01/86	Cobol II/3000 Programs: Tracing Illegal Data
8	6/15/86	KSAM Topics: Cobol's Index I/O; File Data Integrity
9	7/01/86	Port Failures, Terminal Hangs, TERMDSM
10	7/15/86	Serial Printers - Configuration, Cabling, Muxes
11	8/01/86	System Configuration or System Table Related Errors
12	8/15/86	Pascal/3000 - Using Dynamic Variables
13	9/01/86	Terminal Types for HP 3000 HPIB Computers - Revised
14	9/15/86	Laser Printers - A Software and Hardware Overview
15	10/01/86	FORTTRAN Language Considerations - A Guide to Common Problems
16	10/15/86	IMAGE: Updating to TurboImage & Improving Data Base Loads
17	11/01/86	Optimizing VPLUS Utilization
18	11/15/86	The Case of the Suspect Track for 792X Disc Drives

#### NOTICE

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

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

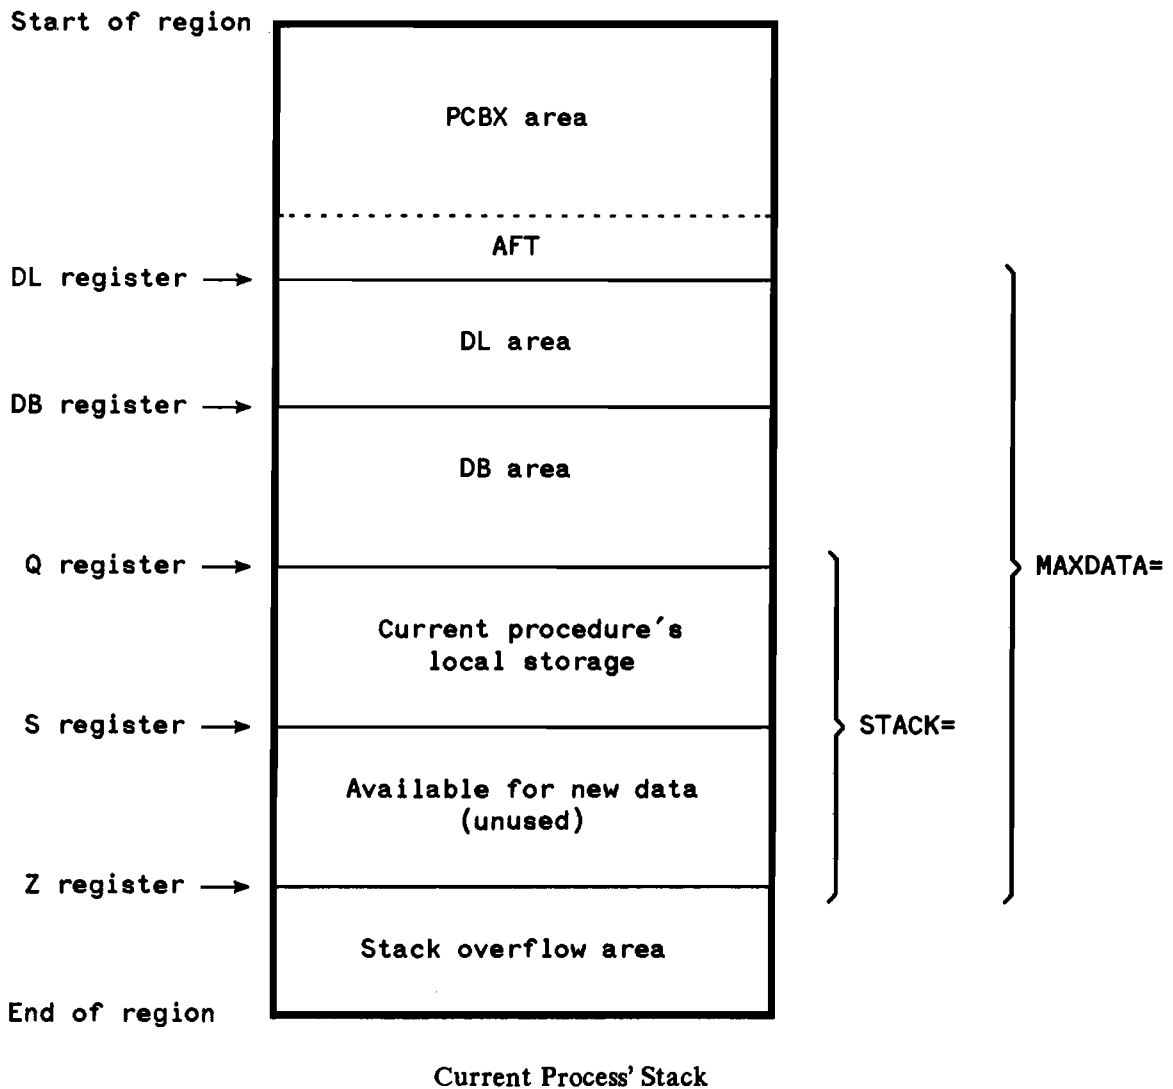
This document contains proprietary information which is protected copyright. All rights are reserved. Permission to copy all or part of this document is granted provided that the copies are not made or distributed for direct commercial advantage; that this copyright notice, and the title of the publication and its date appear; and that notice is given that copying is by permission of Hewlett-Packard Company. To copy otherwise, or to republish, requires prior written consent of Hewlett-Packard Company.

# STACK OVERFLOWS: Causes & Cures for COBOL II Programs

This Application Note discusses Stack Overflows and suggests how you can prevent them. Examples and references are oriented towards the COBOL II/3000 programmer, however, much of the information can be also be applied to the other languages on the HP 3000. For a more detailed discription of how the stack operates please refer to *HP 3000 Application Note #6: HP 3000 Stack Operation*.

## What is a Stack Overflow?

A stack overflow is an error condition that almost every HP 3000 programmer encounters sooner or later. As its name suggests, stack overflows occur when a process' data stack attempts to 'overflow' its boundaries. This happens when the S register, indicating the Top of Stack (TOS), moves beyond the current end of the stack, pointed to by the Z register. MPE will attempt to increase Z, but can only do so when the DL to (new) Z space is less than or equal to MAXDATA (see below). Otherwise, the process is aborted with a STACK OVERFLOW message.



To MPE, a process' stack is simply a portion of a data segment. It, like any other data segment, is limited to a maximum size of 32767 words. Unfortunately, not all of this space is actually available to the user process. MPE itself uses the area below DL, known as the Process Control Block Extension (PCBX), to hold process specific information not contained in other system-wide tables. The Available File Table (AFT), containing entries for the process' open files, also resides in the PCBX. With another 128 words beyond the Z register reserved for the processing of error conditions (such as stack overflows), the amount of space available to the user process is actually only 31232 words. (This number could be reduced further if a large number of files are opened creating a large number of AFT entries).

In general, the key to avoiding stack overflows is to manage the stack as efficiently as possible. Usually, this means using as little of it as you can at any given time.

## **The MAXDATA Parameter**

The most common method of eliminating stack overflows is to use the MAXDATA parameter at RUN and/or at PREP time. It specifies the maximum size to which the stack (the DL to Z area) can grow. Although MPE does NOT reserve this amount of space in memory when the process is loaded, it does reserve space in Virtual Memory (VM) to hold the stack when it is swapped out of memory. In fact, the only way for a stack to expand past its initial limits is by being swapped out to VM and back.

It is usually best to specify the smallest MAXDATA that will work without aborting your program. Although specifying a larger MAXDATA may solve your STACK OVERFLOW problem, if MAXDATA is larger than the stack will ever grow, VM disc space will be wasted; and because MAXDATA specifies the maximum space from DL to Z, a very large MAXDATA will restrict the size to which the PCBX can expand. This could cause a process which opens many files to abort with a "No room left in stack segment for another file entry (FSERR 74)" error indicating no more AFT entries can be created. Therefore, it is best to use the smallest MAXDATA possible. The easiest way to determine this is to experiment and use the smallest value that works.

## **The STACK Parameter**

The STACK parameter is frequently confused with MAXDATA but has a significantly different function. Whereas MAXDATA specifies the maximum size a stack (from DL to Z) can attain, the STACK parameter specifies the amount of stack space you wish the process to start with when it is first loaded into memory. This overrides the initial space allocation computed by the SEGMENTER in the PREP command and, if the space is not needed, it will be wasted.

## **The NOCB Parameter**

If MAXDATA does not solve your overflow problem, the NOCB parameter might. NOCB stands for No Control Block and tells the MPE file system to move its control blocks from the PCBX area of the stack to an extra data segment (XDS). This is done regardless of whether or not space is available in the PCBX. Although using NOCB causes the process to incur slightly more CPU and I/O processing overhead, it sometimes is the only alternative. To run your process NOCB only requires you specify it at RUN and/or PREP time, MPE takes care of the managing the extra data segment holding the file control blocks.

## The VPLUS 'FAST' Forms File Option

VPLUS uses the DL to DB area of the stack for its "COMAREA EXTENSION". This COMAREA EXTENSION holds the largest save field, input and data buffers for the forms file, the largest form definition, function key labels, a save field table, a field status table, a local forms storage directory, a "protect string" (to indicate protected fields), and the forms file directory. The size of the directory, built by the VPLUS forms file compiler, varies depending upon whether the forms file is a regular forms file or a "fast" forms file. A regular forms file contains "source" records, "intermediate" records, and "code" records. A fast forms file contains only "code" records making it smaller than a regular forms file. Since its directory, which resides in the COMAREA EXTENSION, is also smaller, a "fast" forms file uses 800 words less stack space than a regular forms file. The FORMSPEC utility can be used to create your "fast" forms file.

## SORT

When a procedure is called, space is allocated on the stack for its local data items (requiring the stack to expand if there is not enough room). This is also true when a SORT is performed. When the sort is invoked by the SORT verb, all stack space needed for communication with the sort subsystem is allocated. This space is not released until the sort is over. Therefore, if COBOL's INPUT PROCEDURE and OUTPUT PROCEDURE clauses are used, the stack space for any procedures *they* call (IMAGE, VPLUS, user-written, etc) will be allocated *in addition to* that space already being used for/by the sort routines. If your process is nearing the limit of its stack size, an overflow can occur.

As shown in the following example, using the INPUT PROCEDURE and OUTPUT PROCEDURE clauses, any procedure call within INPUT-PROC or OUTPUT-PROC would expand the stack beyond the size already used by SORT.

```
SORT SORT-FILE ASCENDING/DESCENDING KEY IS SORT-KEY
      INPUT PROCEDURE IS INPUT-PROC
      OUTPUT PROCEDURE IS OUTPUT-PROC.
```

The following code in which the USING/GIVING clauses are used will accomplish the same result. In this case, however, INPUT-PROC writes each record to TEMP-FILE instead of RELEASE-ing them to SORT-FILE. Likewise, OUTPUT-PROC reads the sorted records from TEMP-FILE instead of RETURN-ing them from SORT-FILE. Using this technique, procedure calls in INPUT-PROC and OUTPUT-PROC do not consume stack space in addition to that used by SORT. Note, however, that in saving stack space, you are increasing (maybe substantially) execution time of the process.

```
PERFORM INPUT-PROC.
      SORT SORT-FILE ASCENDING/DESCENDING KEY IS SORT-KEY
      USING TEMP-FILE
      GIVING TEMP-FILE.
PERFORM OUTPUT-PROC.
```

## Subprograms and \$CONTROL DYNAMIC

As mentioned above, each call to a procedure or subprogram causes space to be allocated on the stack, which may cause the stack to be expanded. Obviously, a program which has many nested calls to subprograms can use a lot of stack space and may cause a stack overflow, unless each subprogram releases the space it uses after it finishes executing.

This can be achieved by compiling subprograms with `$CONTROL DYNAMIC`. This causes stack space to be allocated on an as-needed basis in the dynamic area (Q to Z area) for the procedure each time the call is made. When execution finishes, the stack space used is returned and is now available for use by other subprogram's.

Compiling subprograms with `$CONTROL SUBPROGRAM` (the default) causes data items to be allocated in the DB area (global area) of the stack; the space is not released when execution finishes. Thus, subsequent subprogram calls to other procedures cannot use any of this space and additional stack space must be obtained.

The only reason to use `SUBPROGRAM` rather than `DYNAMIC` is if you must retain subprogram data storage from one call to the next. `SUBPROGRAM` ensures that its variables will remain intact for the next call to another subprogram by not releasing stack space after the subprogram is executed.

Another `$CONTROL` option for subprograms is `ANSISUB`. This allows for space to be obtained and released like the `DYNAMIC` option. However, between calls to the subprogram variables are retained in a disc file (i.e. when the subprogram is exited the variables are written to disc). When the subprogram is called again, these variables are restored from the file. The additional I/O involved with this option makes it the most inefficient option discussed and it should only be used when you must have stack space for subsequent calls and must retain variables between calls.

## Program Design Considerations

Common sense tells us that if we use huge arrays or other large storage areas we increase the probability of a stack overflow. In such cases, it is very important to design your program's storage as efficiently as possible. For instance, when a large table is needed for data validation, consider the use of a KSAM file to contain this information. Extra Data Segments (XDS) can also be used to hold portions of a larger table. With the `DMOVIN` and `DMOVOUT` intrinsics you can dynamically swap these portions in and out of memory, allowing a smaller stack resident table to be used.

If the preceding suggestions do not eliminate the stack overflow, you may need to reconsider the design of your application. One of the ways to reduce the needed stack space is to divide the program into two or more subprograms, with each subprogram having access to only the data that it absolutely needs to accomplish its task.







# READER COMMENT SHEET

North American Response Centers  
HP 3000 Application Note #19: Stack Overflows: Causes & Cures for COBOL II Programs  
RC Questions & Answers (December 1, 1986)

We welcome your evaluation of this Application Note and attached RC Questions & Answers Sheet. Your comments and suggestions help us to improve our publications. Please explain your answers under Comments, below, and use additional pages if necessary.

- Are these documents technically accurate?
- Are the concepts and wording easy to understand?
- Are the formats of the documents convenient in size, arrangement and readability?
- Comments and/or suggestions for future Application Notes:

**AppNote**

Yes  No

Yes  No

Yes  No

**RC Q&A**

Yes  No

Yes  No

Yes  No



This form requires no postage stamp if mailed in the U.S. For locations outside the U.S., your local HP representative will ensure that your comments are forwarded.

**FROM:** \_\_\_\_\_

Name \_\_\_\_\_ Date \_\_\_\_\_

Company \_\_\_\_\_

Address \_\_\_\_\_

\_\_\_\_\_

\_\_\_\_\_

\_\_\_\_\_

FOLD

FOLD



NO POSTAGE  
NECESSARY  
IF MAILED  
IN THE  
UNITED STATES



# BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 1710, SANTA CLARA, CA

POSTAGE WILL BE PAID BY ADDRESSEE

Application Note / RC Q & A Comments  
Hewlett-Packard Western Response Center  
3300 Scott Boulevard  
Santa Clara, CA 95054



FOLD

FOLD

 **HEWLETT  
PACKARD** **RESPONSE CENTER QUESTIONS & ANSWERS**

HP 3000 Questions Commonly Received by the North American Response Centers

**Q.** When I try to direct the work file from SORT/3000 (SORTSCR) to a private volume disc, SORT displays DEVICE UNAVAILABLE (FSERR 55). Is it possible to direct SORT's work file to a private volume?

**A.** Yes. The secret is that the work file's group must be homed to that volume set. You can either log on to such a group, in which case the work file will be created on the volume set automatically, or you can issue a file equation which specifies a group residing on the volume set you'd like to use. For example,

`:FILE SORTSCR=SORTSCR.pvgroup`

In this case, you must have SAVE access to the group specified.

**Q.** What is the caching controller on the 7933/35XP Disc Drives? Is this the same as the disc caching on my system?

**A.** No. The 7933/35 drives offer a hardware controller cache which is separate from MPE's software caching. Controller cache provides a random access memory (RAM) cache board which resides in the disc drive's controller. When a physical request is made to the disc drive, more data is read than asked for. The additional data is stored in physical memory (the cache board). Future accesses may have the data requested already in the cache area, thus eliminating the need for a physical read.

MPE disc caching uses main memory as a high speed buffer for disc. For example, disc file reads are sped up 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 from disc.

Refer to the *MPE V/E Communicator* (Part No. 5955-1770, Volume 2, Issue 1) for a more complete description of MPE's disc caching.

**Q.** How can I control and check the cache on my 7933/35XP Disc Drives?

**A.** This is done with CS80UTIL version A.02.02 or above.

Several new commands have been added to support the 7933/35XP Disc Controller Cache:

- 1) - CACHEON - Turn cache on
- 2) - CACHEOFF - Turn cache off
- 3) - RESET STATS - Reset cache statistics
- 4) - CACHE STATS - Display cache statistics

You can invoke these commands by using the INFO parm at run time. For example:

`:RUN CS80UTIL;INFO="CACHEON 2"` to turn disc controller cache on for LDEV 2.

more...

:RUN CSBOUTIL;INFO="CACHE STATS 1" to display the cache statistics table for LDEV 1.

With this new function you could set up UDC's (User Defined Commands) to enable or disable caching as follows:

```
CACHON ldev
RUN CSBOUTIL.CS80.TELESUP;info="CACHEON !ldev"
***
CACHEOFF ldev
RUN CSBOUTIL.CS80.TELESUP;INFO="CACHEOFF !ldev"
***
```

Note, when the disc is powered on, disc controller caching is enabled.

There is no indication from the front panel to tell if disc controller caching is ON or OFF. The CACHE STATS command will give you this information.

If the front panel display shows NO CACHE 1 (where 1 is the disc unit number) then the internal disc diagnostics detected a problem in the cache hardware and disabled disc controller caching. This indicates that service is required on the disc.

**Q. I have an Image data base and have enabled User Logging. When the data base is DBOPENed the condition word contains -110 and the next word in the status array contains a 7. In Appendix A, "Error Messages", of the Image manual (Part No. 32215-90003) it states that I should contact Hewlett Packard for information. What does this error indicate?**

**A. This error occurs when a user accesses an Image data base that has transaction logging enabled and that user does not have LG capability. It is actually the MPE OPENLOG intrinsic that fails in this case, as described in the MPE V Intrinsic Manual (Part No. 32033-90007).**

LG capability can be added using the ALTACCT and the ALTUSER commands, as described in the Systems Operations and Resource Management Reference Manual (Part No. 32033-90005).

**Q. How do you change the size of the CSTBK Table and what is it used for?**

**A. The CSTBK (CST Block) table contains a pointer to the block of Code Segment Table Extension (CSTX) entries for each program currently running or allocated. Each CSTBK entry is one word long.**

The size of the table is determined by the MAXIMUM NUMBER OF CONCURRENT RUNNING PROGRAMS allowed on the system as set in the SEGMENT LIMIT CHANGES portion of the SYSDUMP dialogue. The table size is set to the maximum number of concurrent running programs you specify, plus 8.