To: Distribution

From: R.H. Morrison and A.R. Downing

Date: March 9,1976

Subject: Multics HEALS II, Phase 2

1.0 INTRODUCTION.

The Honeywell Error Analysis and Logging System (HEALS) is a facility for capturing and logging hardware error data, then sorting and analyzing the data, and finally reporting the errors and error rates by type (peripheral device, media, etc.). The purpose is to assist FED to achieve its goal of on call maintenance. The premise is that hardware failures tend to be intermittent before they become solid, and many times intermittent failures are recoverable. Errors, and particularly error rates, are an indication of incipient solid failure and a diagnostic which FED can use to schedule preventive maintenance and improve the dispatching of personnel and parts to a site. Additionally, the amount of machine time required by FED for test and diagnosis will be reduced.

HEALS II is implemented under GCOS, and the Multics HEALS II Product Functional Specification requires that Multics have at least the capabilities of GCOS. This MTB summarizes the implementation under GCOS and proposes an implementation for Multics.

2.0 SUMMARY OF PRODUCT FUNCTIONAL SPECIFICATION (PFS)

The following summarizes reference 1 in Section 6.0

The purpose of HEALS II is to improve system availability by speeding up the recognition and isolation of hardware errors. This helps satisfy the Marketing requirement of high availability of service by reducing unscheduled down time. It is anticipated that overall maintenance costs will be reduced.

HEALS II is intended for use primarily by FED personnel. Detailed error information is captured by the system, and summary reports indicate which hardware and/or media need maintenance. The summaries can indicate an impending failure by flagging significant changes in error rates. Error data is captured while the system is running, whether in BOS, Multics, or Salvager, and whether or not the system is providing normal service to the users.

Whenever an abnormal situation occurs, it should be detected and logged. All program-accessible information pertinent to the situation will be obtained and saved on disk. It shall be possible for a site to disable the entire error logging. (It need not be easy to do this, since the error logging and analysis should be of substantial benefit to Honeywell as well as the customer.) No provisions need be made for a site to be able to easily inhibit portions of the error logging.

When an error is detected and logged, a message is not necessarily sent to the system operator. In general, notify the operator if:

(1) there is something the operator can do;
(2) the operator's near-term future actions should be affected by the error rate;
(3) there is the potential of an upcoming system crash and it is a "notify now or never" situation.

"For your information" messages should go only to the log on disk for later perusal and analysis.

Provision must be made for the System Administrator to control how much disk space is reserved for the system error log. It is recommended that a site set aside sufficient disk space to retain at least two weeks of system error log traffic.

The output from the analysis programs is intended to assist the Field Engineer in localizing problems to specific hardware or media items. The analysis programs shall be one or more Multics procedures that are invoked from command level. The current and planned analysis in GCOS should be used as guides to the detailed design of the Multics analysis programs. Providing a total of the errors in various categories is necessary but not sufficient. Error rates must also be calculated. The detailed error data

shall be summarized amd sorted to provide information to facilitate correlation by the Field Engineer between high error rates and specific hardware functions, hardware equipment, and/or media. Some error categories may have time as the basis for the rate calculations, but most error rates should be expressed as a fraction of the operations attempted. FED and customers can then take action when the error rates exceed predetermined threshold rates.

The general format of the output reports shall be substantially the same for both Multics and GCOS systems. Both Multics and GCOS shall use the same sorting rules, nomenclature, row and column headings, row and column order, and mathematical algorithms. The analysis programs shall be able to display individual error messages in original encoded form, and decode and display them in a formatted and easily understood form. However, the primary audience is FED personnel wno are familiar with hardware and exhaustive English explanations are not required.

Most of the capability described in Section 3 of the Product Functional Specification and summarized here already exists in Series 60, Level 66. Multics must be at least equal to Series 60, Level 66

System throughput shall not be degraded by more than 0.5% by the error logging functions as measured in the standard way in tasks/hour by Systems Support Engineering.

Capturing all of the relevant details when errors occur is more inportant than the amount of programming effort, execution time, memory space, or disk space required.

Providing easily understood summary reports is more important than the amount of programming effort, execution time, memory space, or disk space required.

3.0 GCOS HEALS II IMPLEMENTATION

3.1 GCOS HEALS II Structure

The structure of GCOS HEALS II is a concatenation of a number of semi-independent programs operating on several data files. The programs are GESEP, CORTLR, TAPSUM, DISKRP, and HEALS. HEALS is a standard GCOS system which contains the subprograms HEAL, MPCD, and ECFR. The individual programs were independently developed over a period of years; HEALS II does not do much more than run them in sequence and provide some interfacing and transitioning for the first four programs. Most of the sequencing and option selection is done by the Job Control Deck for HEALS II.

Three data files are used to log system error information. First is the Statistical Collection File (SCF) which is a catch-all file for system operations data. I/O errors are logged in this file as type 3 records. Activity and job accounting statistics, including some I/O statistics, are logged as type 1 records. The type 3 record is the principal data source for peripheral error reports, with some incidental information coming from the type 1 report.

The second and third data files are the Error Collection File (ECF) and the Error Summary File (ESF) which contain the logging data for mainframe errors, MPC controller statistics, and device statistics. Logging to these files is done by HEALS; Heal is the logging subprogram which logs mainframe errors and gathers and logs MOS memory ELAC data and MPC statistics. (ECFR and MPCD generate and write HEALS output reports.) HEALS also performs several supervisory functions not related to data capture, logging, analysis, and reporting.

3.2 GCOS HEALS II Reports

There are fifteen reports produced by the GCOS HEALS II programs. (This section is abstracted from Reference 2.)

1. The I/O Error report summarizes all of the type 3 I/O error records found on the system accounting file and details the data found on those records. It many times will be used as a final reference when more specific data is needed after first analyzing

other HEALS II reports.

2. The Activity Summary Report summarizes the statistics accumulated on the SCF for all activities executed on the system for the current accounting period. This report can be used by the customer as an indication of total system throughput for this accounting period.

3. The Fault Summary Report details the faults for all programs for this accounting period. The report is more meaningful to the customer than to FED because it will give him an indication of the amount of throughput that is non-productive on his system. Most faults that occur are the results of programs that are being debugged.

4. The Job Abort Summary Report details the Aborts for all jobs executed on the system for the accounting period. This is more meaningful to the customer because he can obtain an indication of the number of jobs which were unsuccessful.

5. The Core Utilization Summary Report illustrates for the customer the size of activities that are executed on this system. Using this report he can obtain statistics which relate to the typical activity size, time, and memory storage used.

6. The Reel Error Statistics Report lists the reel numbers of the first 512 tapes reporting errors for this accounting period. They are sorted by descending order of the total number of Data Alerts logged against those reel numbers. When normal tape device maintenance is being performed regularly, such as cleaning and unit repair, this report will indicate which tape reels may need maintenance by the tape librarian. Further considerations may be necessary if the previous report shows that excessive errors are occurring on a particular tape unit.

7. The Tape Errors by Handler and Command Report tallies all tape errors by handler device number and tape subsystem command. It will allow the field engineer to determine which tape device may need additional diagnosis and direct him to the subsequent reports. It will also allow him to quickly determine whether the tape subsystem may be experiencing excessive read or write failures. Data Alerts totals displayed on this report for each Tape Unit reflect alerts encountered ONLY when a reel serial number was present.

8. The Tape Errors by Reel-Number/Unit Report will illustrate that a tape reel is failing on multiple devices. This will assist field engineering in determination of media versus device problems. When write errors are occurring on one reel number for several devices, the indication is that the tape reel needs attention by the tape librarian.

9. The Tape Errors by Unit/Reel-Number Report will illustrate errors that may be occurring when different tape reels are mounted on the same tape device. If a tape device were experiencing excessive write errors on several tape reels, then the device could be defective.

10. The Tape Unit Variance Report can be used by Field Engineering to quickly determine which device is experiencing the most Data Alerts with respect to connects for the entire tape subsystem. It also will eliminate the Data Alerts caused by the worst tape reel that had been mounted on the unit and recompute the error ratios/percents using the error free connects and Data Alerts remaining. The entries in this report are sorted by the percent FAIL column in descending sequence. When there is a large number of connects and the percent FAIL column is also a large number, the probability of a bad tape unit is increased. Only Data Alerts will be used to construct this report.

11. The Disk Error Statistics Report summarizes the type 3 I/O error SCF records for system mass storage errors that have occurred during the current accounting period. The continuous binary seek address is converted to its device specific decimal equivalent so that the Field Engineer might relate the failure to a specific physical characteristic of the device. All read, write, or seek errors will be reported.

An increasing number of users are choosing to dedicate a specific disc pack or group of disc packs to certain customer runs. The SNUMB is therefore displayed here because it could relate to a specific media problem. The pack label is not currently being reported on the type 3 I/O records. The report entries are sequenced by unit address. Devices are printed first by IOM, then by device. All units on IOM-0 will be printed first and in device number sequence regardless of the channel number.

12. The Error Collection File Report formats and prints History Register dumps. The first page of this report is always the History Register Legend. This legend defines the abbreviated mnemonics that are used in History Register dumps reported on subsequent pages.

13. The Error Summary File Report summarizes the MOS and Core Storage error or error correction information, and the Processor data saved on the ECF. The ESF is initialized after it has been deleted as described in the operation section of Reference 2. This report will therefore summarize all errors that have occurred subsequent to the file initialization.

14. The System Abort Summary File Report maintains a history of system aborts. This report displays several of the parameters that can be captured at the time of the system abort, but it doesn't attempt to supply the information necessary to resolve the cause of the abort.

15. The MPC Statistics Report displays the statistical counters for Tape and Disc MPC subsystems. The counters are updated by the application firmware for every event being logged. The HEAL logging progran will periodically save these counters on the Error Summary File.

The display represents valuable statistics, including accurate counts of device usage and certain abnormal conditions. Statistics of particular interest include counts of marginal conditions and errors successfully recovered by the firmware. These statistics are lost whenever an MPC is rebooted or powered off, and accuracy will at times be questionable since some of the counters may theoretically roll over more often than the HEAL logging program sample period. The statistics which come from the ESF will be an accumulation of those maintained in the MPC and will be zeroed after execution of the MPCD program. This is therefore the most accurate tally of statistics available for the accounting period. Those statistics that are reported directly from the MPC are only valid from the last MPC boot, power on, or counter roll-over, and therefore do not represent the best sample for the current accounting period.

Each channel and device address is displayed on this report. When there is more that one logical channel or physical channel address for a device, the statistics will be reported for each, and therefore will be duplicated.

3.3 Logging to the Statistical Collection File

The GCOS SCF has become a general purpose file to which a number of events are logged by various GCOS modules and programs. However, only the type 3 and type 1 records are used by HEALS II.

The type 3 records are written (in effect) by the Interrupt Handler routine of the I/O System (IOS). It checks the status

return words on each interrupt. If the status was other than "channel ready", it performs several tests to decide whether or not to write a type 3 record to the SCF. The decision and status are sent to the appropriate channel module which can further analyze the status and over-ride the decision. If the channel module concurs, the type 3 record is written. (Note, however, that this will become somewhat more complex when extended status is appended to the type 3 record.)

The type 1 activity and job accounting records are prepared and written to the SCF by the GCOS termination modules.

3.4 Extensions to GCOS HEALS II

HEALS II was extended for SR2/H to include the new peripheral devices supported by this release. There is also planned a FW 552 supplementary release which will include logging extended status in addition to regular status for I/O errors on devices which have extended status. This release also allows remote accessing of the error logs and reports via a TSS/IDS approach.

In the long term, there are tentative plans to completely re-do HEALS II as a unified facility. The time frame for this is tentatively mid 1978 in SR5.0.

4.0 MULTICS HEALS II DESIGN CONSIDERATIONS

The design of Multics HEALS II is required by the PFS to have at least the functional capability of GCOS HEALS II. However, the implication of the PFS is that the functional capability be with respect to an error analysis and logging system for FED purposes. There are two areas where a direct re-implementation does not seem useful. These are:

(1) The functions outside the scope of an error analysis and logging system.

(2) A number of the output reports from GCOS HEALS II are based on GCOS job and activity numbers. The Multics process is the nearest thing to a GCOS job, but it is not sufficiently close to convince one that it would be of value to substitute in the reports.

For the Multics implementation it is proposed that HEALS II should be limited to the basic functions of hardware error data capture, logging, analysis, and reporting. In particular, such functions of GCOS HEALS II as managing instruction retry, managing cache memory, etc. will not be a part of Multics HEALS II. Furthermore, functions of GCOS HEALS II which are so specialized to GCOS that no reasonable equivalents exist in Multics (e.g., reporting by job SNUMB and activity number, system abort summary, etc.) will likewise not be implemented.

The control of Multics HEALS II will have little resemblance to that in GCOS, again because of the characteristics of the operating system. Most of the control can be obtained as arguments to the procedures implementing HEALS II. The obvious exception is the control of whether or not the error data is to be captured and logged.

In general, the PFS requirement of capability equal to GCOS HEALS II can be satisfied if the same hardware error data is captured. Working from this data base, it should then be a reasonably straightforward task to analyze the data and produce the reports of GCOS HEALS II, excepting those that are reported in terms of GCOS job flow.

Multics HEALS II should not be limited by the GCOS HEALS II as it now exists; instead the latter should be considered a minimum requirement for the present. There are two reasons for this: first, there are errors on hardware unique to Multics (e.g., the associative memory) that should be logged, and there are events unique to Multics that could be detected and logged; and second,

GCOS HEALS II is still evolving and it will acquire additional features which, from the FED viewpoint, will be applicable to Multics as well.

It is expected that the use of HEALS II would be similar to its use in GCOS. The reports would routinely be generated daily and be interpreted by FED personnel. FED would determine the need for scheduling maintenance activities on a particular unit based on the HEALS II report diagnosis. For closer monitoring of units, one or more of the reports would be generated more frequently on demand, and perhaps be limited to the unit or units of interest.

5.0 PROPOSED MULTICS HEALS II IMPLEMENTATION

5.1 Data Capture

Four types of error and operations data should be captured. These are:

1. I/O error records (including bulk store)
2. Processor error records
3. Device and MPC operations statistics
4. MOS memory EDAC statistics

For the I/O and Processor error records, data must be captured at the time of the event and logged. The statistics, however, are captured and buffered by the hardware from which they must be captured and logged.

The data to be captured for I/O error records is listed in Figure 5-1. An I/O error record potentially should be generated whenever the status return word is not "channel ready". If it is determined that the status is not an error (i.e., is the expected status under the circumstances) the error record should be suppressed (conceivably, it could be written, and a second record written to cancel it).

The data to be captured for processor error records is listed in Figure 5-2. A processor error record should be generated each time the history registers are locked by a fault (Op Not Complete, Lockup, Parity, Command, Store, Illegal Procedure, and Shutdown). (1)

MPC controller and device statistics are captured and buffered by the MPC in counters in the MPC read/write memory. The controller counters can be accessed by the Read Controller Main Memory command, and reset by the Write Controller Main Memory command. (2) The device counters can be accessed with a series of Read Control Register commands, one command addressed to each device.

---

(1) See Reference 7, Section 3.6.2.

(2) See Reference 10, Section 2.2.6.1 and Reference 11, Section 2.5.4.

The counters can be reset with Write Control Register commands. (3) The error correction data for the DSS190 and DSS191 devices can be accessed with a series of Read EDAC Register commands, one command addressed to each device. (4) The data to be collected and logged is shown in Figure 5-3.

The MOS Memory EDAC syndrome data is captured and buffered in the System Controller. The EDAC data can be accessed by a series of Read General Register commands (RSCR processor instructions), one command addressed to each memory unit in each system controller. The data can be reset with Write System Controller General Register commands. (5) The data to be collected and logged is shown in Figure 5-4.

---

(3) See Reference 12, Sections 3.14 and 3.15, and Reference 13, Section 7.5.

(4) See Reference 12, Section 3.16.

(5) See Reference 9, Section 3.4.13, and Reference 14, Sections A2.6 and A2.9.

Primary Extract and Sort Fields:

- Record type (and length)
- Calendar clock time of interrupt
- Device type code
- Tape reel serial number (if tape)
- Disk pack serial number (if disk and if available)

Secondary Extract and Sort Fields:

- Software version ID.
- Installation ID.
- Process ID.
- Calendar clock time of connect

Error Data:

- IO Status
  - Sync bit
  - Power bit
  - Major status
  - Substatus
  - Lost interrupt flag
  - Initiation/Termination interrupt flag
  - IOM Error
  - Record residue
- IO Command (second if DS/DR)
  - Device command
  - Device number
  - IOM number
  - IOM command
  - IOM channel number
  - Record count
- Extended status (when available)
- Number of errors on device (excluding this I/O) from last bootload
- Number of connects on device (including this I/O) from last bootload
- Error ratio (number of errors/64 connects)
- Seek address (if disk)
- Density (if tape)

Figure 5-1. Data Items in I/O Error Record

Primary Extract and Sort Fields:

Record type (and length)
Calendar clock time
CPU number

Secondary Extract and Sort Fields:

Software version ID.
Installation ID.
Process ID.

Error Data:

Faulted Instruction
Operand Pair in error
Fault code
Reason code
Retry count
Processor registers
Fault register
Mode register
Configuration switches
Instruction counter and indicator register
Control Unit History registers
Operations Unit History registers
Decimal Unit History registers
Appending Unit History registers
Pointers and lengths

Figure 5-2. Data Items in CPU Error Record

Primary Extract and Sort Fields:

Record type (and length)
Calendar clock time
IOM number
Channel number
MPC number

Secondary Extract and Sort Fields:

Software version ID.
Installation ID.

Data:

MPC controller counters for each MPC
Device counters for each device
EDAC data for each disk (when available)
(Reference MPC and controller EPS-1's for lists of counters)

Figure 5-3. MPC Statistics Record

Primary Extract and Sort Fields:

Record type (and length)
Calendar clock time
SCU number
Store unit number

Secondary Extract and Sort Fields:

Software version ID.
Installation ID.

Data:

MOS Memory EDAC data for each store unit in each system controller (See References 9 and 14)

Figure 5-4. MOS Memory EDAC Record

5.2 Data Logging

Because I/O and processor error records require data to be captured at the time of the error, it is reasonable also to log them immediately. It is proposed that the syserr mechanism be used to do this.

The MPC statistics data is cumulative and buffered in the MPC, and there is no need to log them immediately. The MOS memory EDAC syndrome data is not cumulative and therefore should be collected frequently.

Since some of the statistical data can only be gathered by a privileged process, it is proposed that the initializer be responsible for gathering data on a regular basis and entering it in the syserr_log. The frequency of this data copying will depend on how often it is necessary to copy data without having any data lost due to such events as counter roll over. There will be a set of control commands with which the system operator can alter the sampling rate and other parameters. The first step in report generation will be a request to the initializer to update the ring 4 copy of the syserr_log.

5.3 Data Reduction

The HEALS II error records logged by syserr will be on the ring 4 copy of syserr_log along with all other syserr logged records. Periodically (for example, every hour) and prior to the generation of HEALS II reports, the syserr_log copy should be updated and scanned for new records, and Heals II error records extracted and merged in the heals_log. The main reason for this is to facilitate the generation of the output reports which will involve repeated sorting of the error records over variable time spans. The error records may also be re-formatted to be more convenient for this purpose. In addition, the heals_log will satisfy the PFS requirement to save the error data for some period of time independently of the time the syserr_log is saved. During the extraction and merging of error records, the records can be processed to develop error threshold and trend data for timely output as console messages to the operator.

5.4 Reports

The HEALS II reports listed in Section 3.2 fall into three classes: (1) those that are hardware error oriented and are most

useful to FED, and (2) those that are system performance oriented and are most useful to the customer, and (3) all others. Reports in the second class are organized by GCOS job and activity number and thus present some difficulties.

The first class contains the following reports:

I/O Error
Reel Error Statistics
Tape Errors by Handler and Command
Tape Errors by Reel Number/Units
Tape Errors by Unit/Reel Number
Tape Unit Variance
Disk Error Statistics
Error Collection File
Error Summary File
MPC Statistics

These Reports will be produced by Multics HEALS II in the same format as the GCOS reports.

The second class contains the following reports:

Activity Summary
Fault Summary
Job Abort Summary
Core Utilization Summary

unless it becomes clear that reports equivalent to these but in terms of Multics interactive or absentee processes (or something) have some real use and do not overlap reports from other Multics performance metering facilities, it is proposed that these reports be eliminated from Multics HEALS II.

The System Abort Summary Report is in the third class. It is proposed that this report not be included in Multics HEALS II.

When a report is needed, the requesting person will use a command (for example, "heals_report") to initiate his request. The argument list to this command will contain the name of the report required and information concerning the output of the report. It is proposed that the final report be left in some proper part of the hierarchy for perusal by the requestor. Options to this command will allow the requesting party to have the report dprinted and directed to his location.

6.0 REFERENCES

1. Multics HEALS II
Product Functional Specification
Spec. No. 58004068

2. HEALS II Manual
Field Engineering Division May, 1975

3. Honeywill Error Analysis and Logging System (HEALS)
Series 60(level 66)/6000 GCOS
Manual DD44, Rev. 0

4. Summary Edit Program
Series 60(level 66)/6000 GCOS
Manual DD24, Rev. 0

5. System Tables
Series 60(level 66)/6000 GCOS
Manual DD14, Rev. 0

6. HEALS 2/H
TDM-TOS-18
J.P. Brown, May 1, 1975
Computer Engineering Operations, Phoenix

7. Engineering Product Specification, Part 1
6000 Processor
No. 43A219601

8. Engineering Product Specification (EPS-1)
6100 CPU
No. 58001068

9. Engineering Product Specification-1
6000 System Controller
No. 43A219602

10. Mass Storage/Tape MPC Basic Firmware Specification
No. 43A177755

11. Unit Record MPC Controller/Subsystem
Engineering Product Specification
No. 58001108

12. Engineering Product Specification-1
DSC181/DSC190 Controller
No. 43A232230

13. Engineering Product Specification-1
MTC 500/MTP0600 Controller
No.43A232215

14. Engineering Product Specification, Part 1
    6000 Semiconductor Storage Unit
    No. 58001116