THE BELL SYSTEM TECHNICAL JOURNAL
SAFEGUARD SUPPLEMENT
Printed in U.S.A.

SAFEGUARD Data-Processing System:

A Means to Effective Computer Resource Utilization

By J. P. KUONI

(Manuscript received January 3, 1975)

"Bird-dogging" is the process of tracking down computer center users who are either having problems and therefore are not getting their job done or who are using a disproportionate share of the computer's resources. Analysis of utilization data for the SAFEGUARD support computer centers has shown that the problems caused by these users can be of alarming magnitude, leading some observers to believe that bird-dogging is the single most effective system performance tuning activity that can be performed. Bird-dogging is an integral component in reliable project scheduling and effective cost control. This paper discusses the methods now used to identify problem users and some experiences gained from the effort.

I. INTRODUCTION

This paper describes the function of bird-dogging as the main tool for achieving the most efficient use of the computer. Specifically, through analysis of computer utilization data (which may be sampled on a daily, weekly, or monthly basis), the use of computer center resources and the problems of its users are monitored in detail. This is followed as needed with a program of counseling. The purpose of counseling is to better educate computer users to employ effectively the computing resources available to them (hardware, operating system, and application software). Counseling also provides feedback to the designers of application software to allow implementation of designs that would permit better utilization of the hardware and operating system features.

Some segments of the bird-dogging campaign are conducted on a daily basis for short-term gains, and other segments take the form of more extensive investigations yielding long-range gains. The latter activity more closely approximates the traditional system tuning.

Bird-dogging has been actively supported at several project support computer centers since the fall of 1971. Manpower allotment during this period is estimated to be two or three full-time technical staff members at each location. This total includes manpower employed to develop programs for automated report generation.

II. WHY BIRD-DOG?

Although many installations are committed to ongoing efforts in the traditional areas of systems performance analysis, few are engaged in bird-dogging campaigns. Why, then, are the SAFEGUARD project centers actively supporting this activity? There are two main reasons: schedule reliability and cost control.

First, schedule reliability. During the years of developing the system's software, timely completion of the hundreds of interlocking software modules has been critical for project delivery. It has been imperative, therefore, that everyone, even the below-average programmer, complete his or her responsibilities on time and successfully. To increase confidence in meeting project schedules, those who are unable to make it on their own must be helped.

Second, cost control. Bird-dogging helps reduce costs through short-term immediate benefits and long-range improvements. For example, bird-dogging usually produces immediate benefits by reducing the resubmittal rates of "problem" programs, which increases the turnaround potential of other programs competing for the limited computing resources.

In the long run, for example, many users having similar problems may reveal that the documentation of how to use a particular feature is inadequate. Following through on individual problems to gain insight into underlying causes is often worthwhile and carries considerable long-range benefit.

III. UTILIZATION DATA

To permit monitoring the center's users, several types of utilization data are obtained from a series of automated reports and other sources.

3.1 Automated reports

The bulk of bird-dogging data is generated by several special-purpose report programs developed by project personnel. Most of these programs use the System Management Facilities (SMF)* data as input. A brief description of each report and its use follows.

---

* SMF is an optional feature of the Operating System (OS) (Ref. 1), which collects system, job-management, and data-management information.

The stat (statistics) card report shows detailed accounting information about each job run on the computer, sorted by supervisory group and department. Information such as CPU time, lines printed, region size, disc and tape setups, read-in time, and purge time are shown for each job processed. This report is produced and examined daily and gives indications of overall throughput, average turnaround, distribution of work among departments, and unusual jobs. It also provides a reference for the day's activities.

The abnormal end (ABEND) report provides data about each job that aborts. Information such as failure code, programmer name, job name, and CPU time is provided. These data are also printed and examined daily to give indications of particular users who consistently have problems, specific programs that frequently fail, and repeated ABEND codes that may be symptomatic of system problems.

The usage report provides detailed characteristics of the high-usage programs executed by each department. It also shows a rank order list of these high-usage programs. These data are used to pinpoint programs to be considered for performance analysis and improvement, as well as to pinpoint possible inefficient or unusual use of a program by a particular department.

The high-resource report and the exception report highlight users whose jobs exhibited certain high-resource characteristics such as exceptionally long turnaround time, extended use of central processor time, great volume of printed output, very large use of core memory, and utilization of several setup devices, or those jobs that experience a job control language error after significant expense of resources.

3.2 Other sources of data

In addition to the various automated report programs that provide utilization data, there are several other important sources of bird-dogging data. Direct problem program monitoring and feedback from operations personnel are the two most significant sources.

Program monitoring is achieved through use of a proprietary software monitor that provides valuable execution profiles of user programs. Several monitors are on the market; the project centers are using Boole and Babbage's Program Evaluator (PPE).[2] Experience to date indicates PPE is easy to use, well documented, and consistently helpful in providing areas for program performance improvements. PPE indicates where and how the monitored program spends its time and how compute-limited or input/output-limited the program is. The effects of subsequent improvements to the program are readily apparent by remonitoring.

Operations personnel can provide valuable bird-dogging data. In many cases, user problems may not appear in the automated reports, or problems do appear but their magnitude is hidden.

IV. CASE STUDIES

This section presents several cases that typify many of the long-range studies undertaken as a result of the analysis of weekly and monthly computer utilization data.

4.1 Study 1

For a period of several months, the types and frequencies of ABENDs at the computer centers were investigated. It was found that 15 to 20 percent of all jobs submitted eventually ABENDed and 25 to 30 percent of the total Central Processing Unit (CPU) time was spent executing these jobs. The ABENDs were grouped into four categories:

(*i*) Those that were a result of insufficient estimates of the computer resources required by the job (resources include CPU time, memory, and I/O estimates).
(*ii*) Those that reflected problems of a data base nature.
(*iii*) Those that resulted from a program check condition.
(*iv*) Those that were symptomatic of a hardware malfunction.

The most striking observation from this study was that the inability of users to correctly estimate the computer resources required for their job appeared to be by far the biggest obstacle to successful job execution. As a result of this and other related studies:

(*i*) The support software user manuals were revised to include algorithms for estimating required computer resources.
(*ii*) Modifications were implemented to OS that allowed selected critical modules to complete execution even though the actual CPU time consumed has exceeded the programmer's estimate.

As a corollary to the problem of insufficient estimates, system performance was often degraded by serious overestimation. An educational campaign was initiated by distributing to all project programmers an informational bulletin that clarified the specification of job and of job step region parameters.

Because of the changing nature of the project and its computation requirements and the scattered implementation of study recommendations, objective measurements of subsequent improvements have not as yet been attempted.

4.2 Study 2

The usage report indicated heavy use by one department of a "home-grown" data reduction routine. By revising the program only slightly, CPU time was dropped from 110 to 8 seconds per execution.

4.3 Study 3

Analysis of the execution profile for the CENTRAN compiler demonstrated that a much higher than average number of accesses to the CENTRAN symbol tables were required during the compilation of large programs with certain characteristics. By specifying additional core memory in the region size over the default, overall resource requirements were reduced (and, hence, cost to process the job was reduced).

Detailed data for each CENTRAN compilation were available through the automated reports. The 75 programmers who were responsible for programs with exceptional characteristics were contacted over a period of several months and were requested to allocate additional memory for their compiles. Most individuals complied and experienced a decrease of turnaround time (by reduced elapsed time), with an attendant system cost reduction.

4.4 Study 4

The exception report provided a list of jobs requiring high resource use. With the cooperation of the users, these jobs were scheduled for evening or weekend shifts. Rescheduling of these jobs eliminated them from competition with other jobs for limited prime-shift computing resources.

4.5 Study 5

It was observed by operations personnel, and later confirmed by examination of reports that correlated turnaround time and resource usage, that certain users were taking advantage of a loophole in the computer centers' job-scheduling algorithm. The slightly higher priority assigned by the algorithm to jobs requiring a setup led to the submittal of jobs with unneeded setups. A job-scheduling adjustment corrected the problem.

4.6 Study 6

The usage report indicated that the SAFEGUARD Data Reduction SYSTEM (SDRS) was the largest single user of CPU resources, consuming 20 to 30 percent of all CPU time. Analysis of the facility with PPE indicated that much of this time was spent communicating with the operating system. Interrupt recovery capabilities were provided for each type of input data. These required many different recovery routines that necessitated specifying different interrupt exit addresses to OS many times. The same capabilities were preserved by some minor restructuring of the program and the addition of logic to determine the appropriate interrupt recovery. Post-modification benchmarking revealed an average 60-percent savings of CPU time for this program.

V. CONCLUSIONS

It is the belief of the project centers that bird-dogging is the single most effective tuning activity that can be performed. Bird-dogging is an integral component in reliable project scheduling and effective cost control. As in other areas of system tuning, although the fruits of individual events and incidents seem indisputable, the successes (or failures) of bird-dogging can seldom be proven objectively by quantitative measure. Justification, therefore, remains mostly in the subjective domain.

The bird-dogging effort has been hindered by design errors and limitations in the SMF portion of the operating system and by the lack of commercially available SMF data reduction systems suitable for project needs.* Hence, considerable manpower was expended in developing a series of automated report programs.

The computer centers have found a software monitor, in this case Boole and Babbage's PPE, helpful in providing data for program performance improvement. Every bird-dogger should have something of this sort available.

The ultimate success of any bird-dogging program depends heavily upon the degree of cooperation received from the user community and its management. Care should be taken from the outset to present suggestions and criticism in a positive manner. Helping users to help themselves will contribute to improved confidence in meeting schedules and to lower computer center costs.

REFERENCES

1. IBM System 360 Operating System, System Management Facilities, Order Number GC28-6712-5, International Business Machines Corporation.
2. Systems Measurement Software (SMS/360), User's Guide for PPE, Boole and Babbage, Inc., Cupertino, California.

* To this author's knowledge, there is only one "off-the-shelf" SMF reduction system available, the SMF Selective Analyzer, FDP-5798-AAR, IBM Corporation.