

## Accounting Logs

- Advantages:
  - Built in
  - The data reflects real-system usage.
  - Use them before developing a new monitor
- Disadvantages:
  - No analysis programs. Only utilities to convert them in a readable format are supplied.
  - Data may not be at the desired level of granularity. For example: distribution of I/O sizes
  - Only charged resources may be recorded
  - Low accuracy
  - No system-level information, such as queue lengths or device utilizations.

- Not possible to separate the effects of queueing from user service demands.

## Data in Accounting Logs

- Resource usage:
  - CPU time
  - Elapsed time
  - Number and total size of disk I/O
  - Paging I/Os
  - Terminal I/Os
  - Network I/Os, ...
- Granularity:
  - System-wide resource usage
  - Resources used by each program
  - Resources used by each user session.
- Abnormal conditions such as:
  - Security violations
  - Abnormal job terminations
  - System restarts
  - Down times

– Device errors

# Typical Data Recorded by Accounting Utilities

---

Name of the program

Program start time: date and time of the day

Program end time: date and time of the day

CPU time used by the program

Number of disk writes

Total disk write bytes

Number of disk reads

Total disk read bytes

Number of terminal writes

Total terminal write bytes

Number of terminal reads

Total terminal read bytes

Number of page-read I/Os to the paging device

Number of pages read from the paging device

Number of page faults

---

## Derived Quantities

- Per activation: Useful for
  - Modeling program behavior
  - Constructing synthetic workloads.
- Percentage of total:  
Resource consumed by all activations of a particular program.
  - Programs with a high percent have a high impact on the resource
  - These programs should be optimized.
  - These programs should be included in test workloads.
- Per second resource-consumption rates:  
Divide the resource consumption by the elapsed time.
  - Indicate the intensity of resource usage (utilization)

- Help calculate the number of users that can be supported
- Help determine the number of resources required.
- Help find programs that can cause the resource to become bottleneck.
- Per CPU second resource consumption:  
Divide the resource consumption by the CPU time consumed.
  - Less variable than per second consumption.
  - Gives resource demand per instruction of the program.
  - Represents the ratio of the resource demand to CPU demand.
  - Programs with high disk I/O per CPU second are disk bound;



## Commonly Asked Questions

- *Which programs should the programmers be trained to use more efficiently? Or, which programs provide the highest opportunity for better human interface?*

The programs with high percentage of total elapsed time provide the best opportunity for training as well as better human interface.

- *Which programs are good candidates for code optimization? Or, which programs should be included in a workload to be used in analyzing a new CPUs performance?*

The programs with high percentage of total CPU time.

- *Which programs offer the highest opportunities for code restructuring to*

*minimize the page faults?*

The programs with a high percentage of total page faults, a high percentage of total page-read operations, or a high percentage of total pages read

- *Which programs have a poor locality of reference?*

The programs with high “page faults per CPU second” have a poor locality of reference.

## Commonly Asked Questions (Continued)

- *How many jobs can run simultaneously without undue performance degradation?*

The CPU rate (CPU time per second) for CPU bound programs. Or, the device with the highest utilization.

- *Which programs are I/O-bound?*

The programs with high “disk I/Os per CPU second”

- *What is the average think time?*

Approximately = reciprocal of terminal reads per second.

- *What workload characteristics should be used in a simulation or analytical model?*

To analyze a new device, the programs that

use a high percentage of that resource should be included in the workload.

To compare two different systems the overall average data per workload unit should be used.

## Case Study: Accounting Logs

Data collected at six different universities for a period of six months.

- There are more disk reads than writes. Similarly, there are more terminal writes than reads. The ratio of reads to writes were as follows:

For disk I/O:

Number of writes : Number of reads :: 1:3  
Size of writes : Size of reads :: 1:2.5  
Total write bytes : Total read bytes :: 1:7.5

For terminal I/O:

Number of writes : Number of reads :: 2:1  
Size of writes : Size of reads :: 4:1  
Total write bytes : Total read bytes :: 8:1

Although the ratios may be different for

other environments, the following inequalities should generally hold:

Number of disk writes	<	Number of disk reads
Size of disk writes	<	Size of disk reads
Number of terminal reads	<	Number of terminal writes
Size of terminal reads	<	Size of terminal writes

law of slow creation  $\Rightarrow$  More terminal write than reads.

## Case Study: Accounting Logs (Continued)

- In an educational environment, system programs are used very extensively.
  - 50% of all resources consumed
  - 83 out of every 100 programs
  - 57% of the user time
- A large portion of the system resources was used on compilers.
  - Only 6% of all activations
  - One-fourth of the CPU consumption
  - One-third of all page faults
  - One-fifth of all page reads.
- About one-third of user time was spent on editors.
- LOGINOUT was the most commonly executed program. 18% of all activations

were LOGINOUTs.  $\Rightarrow$  Shortage of terminals

DELETE (file deletion utility) constituting about 10% of all activations  $\Rightarrow$  Shortage of disk space

- Average session lengths of some of the interactive programs were:

BASIC 7 minutes

SYSGEN 5 minutes (To Change Sys. Parameters)

SOS 5 minutes (Text Editor)

EDT 4 minutes (Text Editor)

TECO 4 minutes (Text Editor)

CREATE 2 minutes (File-Creation Utility)

DISPLAY 2 minutes (System-Monitoring Utility)



## Exercises

- 8.1 Using accounting log data, how would you answer the following questions:
- a. Which programs should be used to analyze a new terminal concentrator?
  - b. How many program X users can be supported with a given disk device?
  - c. Is the RTE script of a benchmark representative?
  - d. Which programs should be chosen for I/O optimization?