

NAME

`iostat` — report I/O and system statistics

SYNOPSIS

`iostat [-atpisbecgqmzfdhI] [lrate] [interval [count]]`

DESCRIPTION

Iostat reports statistics kept about various activities within the system.

For each disk, the system counts IO completions and the number of words transferred. Also, each sixtieth of a second, the state of each disk is examined and a tally is made if the disk is active. This tally goes into one of four categories, depending on whether the system is executing in user mode, in 'nice' (background) user mode, in system mode, or idle. From all these numbers and from the known transfer rates of the devices it is possible to determine information such as the degree of IO overlap and average seek times for each device.

The optional *interval* argument causes *iostat* to report once each *interval* seconds. The first report is for all time since a reboot and each subsequent report is for the last interval only.

The optional *count* argument restricts the number of reports.

The interval and count arguments may be repeated in pairs to provide varied lengths and numbers of reporting intervals.

With no flag argument *iostat* reports the elapsed time since boot and gives the percentage of time the system has spent in each of the four categories mentioned above.

There are zillions of options:

- a** Turns on all options except for **-p**, **-s**, **-z**, **-m**, **-f**, **-d**, **-h** and **-I**.
- t** For typewriters collectively, the input and output character rate is reported.
- p** Report the percentage of time spent in system mode, in nice (low priority) user mode, in user mode, and in idle mode.
- i** Report, in addition to the percentage of time spent in each of the four categories for the **-p** option, the percentage of time each disk was active (seeking or transferring), the percentage of time any disk was active, the percentage of time spent in 'IO wait' (processor idle, but with a disk active,) and the percentage of time spent in interrupt routines with a priority less than that of the system clock.
- s** Report the raw timing information: 24 numbers indicating the percentage of time spent in each of the possible configurations of 4 system states and 8 IO states (3 disks each active or not).
- b** For each disk, the number of transfers per minute, the number of transfers per second, the milliseconds per average seek, and the milliseconds per data transfer exclusive of seek time is reported.
- e** For this invocation of the program report the elapsed time in minutes since the last report. For the first report, this is the time since boot.
- o** Print out the number of times the process table, the text table, the inode table, and the file table have overflowed.
- c** Print out the rate and count of interrupts (excluding the system clock,) the average number of milliseconds per interrupt, the rate and count of traps caused by kernel switchable text, calls to the swap routine, the number of forks, the number of executes, the number of disk reads, the number of disk writes, and the number of process switches. The interrupt timing information is valid only if there is no hardware interrupting at or above the priority of the system clock. If there is some other high priority device which interrupts at a constant rate this information can be supplied to

iostat using the `-I` option.

- `-g` Report the percentage of time a buffer was needed and none was on the free list, the percentage of time a block was requested and it was found to already be in core, the rate per second of calls to get block (as well as the total number of calls), the percentage of time a physical IO buffer header was needed and marked busy, and the rate per second of successful calls to the physio subroutine (as well as the total number of calls).
- `-q` Report the average queue size and percent occupancy for each of three queue: the queue of jobs on the swap device but marked runnable (the swap queue); the queue of jobs in core and marked runnable (the run queue); and the queue of jobs waiting for disk IO other than swap IO to complete (the disk queue). Note that the average reported is an average of only those samples with a non-zero occupancy: to compute the true average queue size the size given must be multiplied by the occupancy.
- `-m` From analyzing the memory map report the number of free 64 byte segments; the average, maximum, and minimum size of the segments; and the total number of segments. Notice that this is the one measurement which is *not* a time average, but an instantaneous report.
- `-z` Suppress the first report in a series so that the cumulative statistics since boot are not given and only interval statistics are given.
- `-f` Provide a form feed between reports.
- `-d -h` The `-d` and `-h` options are used to obtain disk drive cylinder usage and seek distance profiles. Counts are kept for a single drive (defined by variable `dk_unit` in operating system) of the cylinder desired and the required seek distance in 8 cylinder increments. The `-d` option dumps the data in tabular form. The `-h` option produces a histogram on the printer of the data.
- `-I` Requires the user to supply the optional *Irate* argument which specifies the number of interrupts per second at a priority higher than or equal to the system clock. This rate should not include the system clock.

FILES

/dev/mem, /unix