IOSTAT(1) Linux User's Manual IOSTAT(1)

NAME

iostat - Report Central Processing Unit (CPU) statistics and input/out? put statistics for devices and partitions.

SYNOPSIS

iostat [-c] [-d] [-h] [-k | -m] [-N] [-s] [-t] [-V] [-x] [-y] [-z] [--compact] [--dec={ 0 | 1 | 2 }] [{ -f | +f } directory] [-j { ID | LABEL | PATH | UUID | ... }] [-o JSON] [[-H] -g group_name] [--human] [--pretty] [-p [device[,...] | ALL]] [device [...] | ALL] [interval [count]]

DESCRIPTION

The iostat command is used for monitoring system input/output device loading by observing the time the devices are active in relation to their average transfer rates. The iostat command generates reports that can be used to change system configuration to better balance the in? put/output load between physical disks.

The first report generated by the iostat command provides statistics concerning the time since the system was booted, unless the -y option is used (in this case, this first report is omitted). Each subsequent report covers the time since the previous report. All statistics are

CPU header row followed by a row of CPU statistics. On multiprocessor systems, CPU statistics are calculated system-wide as averages among all processors. A device header row is displayed followed by a line of statistics for each device that is configured.

The interval parameter specifies the amount of time in seconds between each report. The count parameter can be specified in conjunction with the interval parameter. If the count parameter is specified, the value of count determines the number of reports generated at interval seconds apart. If the interval parameter is specified without the count parame? ter, the iostat command generates reports continuously.

REPORTS

The iostat command generates two types of reports, the CPU Utilization report and the Device Utilization report.

CPU Utilization Report

The first report generated by the iostat command is the CPU Uti? lization Report. For multiprocessor systems, the CPU values are global averages among all processors. The report has the fol? lowing format:

%user Show the percentage of CPU utilization that occurred while executing at the user level (application).

%nice Show the percentage of CPU utilization that occurred while executing at the user level with nice priority.

%system
Show the percentage of CPU utilization that occurred while executing at the system level (kernel).

%iowait
Show the percentage of time that the CPU or CPUs were idle during which the system had an outstanding disk I/O request.

%steal Show the percentage of time spent in involuntary wait by the virtual CPU or CPUs while the hypervisor was servic? ing another virtual processor.

%idle Show the percentage of time that the CPU or CPUs were idle and the system did not have an outstanding disk I/O request.

Device Utilization Report

The second report generated by the iostat command is the Device Utilization Report. The device report provides statistics on a per physical device or partition basis. Block devices and parti?

the command line. If no device nor partition is entered, then statistics are displayed for every device used by the system, and providing that the kernel maintains statistics for it. If the ALL keyword is given on the command line, then statistics are displayed for every device defined by the system, including those that have never been used. Transfer rates are shown in 1K blocks by default, unless the environment variable POSIXLY_COR? RECT is set, in which case 512-byte blocks are used. The report may show the following fields, depending on the flags used (e.g. -x, -s and -k or -m):

Device:

This column gives the device (or partition) name as listed in the /dev directory.

tps Indicate the number of transfers per second that were is? sued to the device. A transfer is an I/O request to the device. Multiple logical requests can be combined into a single I/O request to the device. A transfer is of inde? terminate size.

Blk_read/s (kB_read/s, MB_read/s)

Indicate the amount of data read from the device ex? pressed in a number of blocks (kilobytes, megabytes) per

have a size of 512 bytes.

Blk_wrtn/s (kB_wrtn/s, MB_wrtn/s)

Indicate the amount of data written to the device ex? pressed in a number of blocks (kilobytes, megabytes) per second.

Blk_dscd/s (kB_dscd/s, MB_dscd/s)

Indicate the amount of data discarded for the device ex? pressed in a number of blocks (kilobytes, megabytes) per second.

Blk_w+d/s (kB_w+d/s, MB_w+d/s)

Indicate the amount of data written to or discarded for the device expressed in a number of blocks (kilobytes, megabytes) per second.

Blk_read (kB_read, MB_read)

The total number of blocks (kilobytes, megabytes) read.

Blk_wrtn (kB_wrtn, MB_wrtn)

The total number of blocks (kilobytes, megabytes) writ? ten.

The total number of blocks (kilobytes, megabytes) dis? carded.

Blk_w+d (kB_w+d, MB_w+d)

The total number of blocks (kilobytes, megabytes) written or discarded.

r/s The number (after merges) of read requests completed per second for the device.

w/s The number (after merges) of write requests completed per second for the device.

d/s The number (after merges) of discard requests completed per second for the device.

f/s The number (after merges) of flush requests completed per second for the device. This counts flush requests exe? cuted by disks. Flush requests are not tracked for parti? tions. Before being merged, flush operations are counted as writes.

sec/s (kB/s, MB/s)

The number of sectors (kilobytes, megabytes) read from,

rsec/s (rkB/s, rMB/s)

The number of sectors (kilobytes, megabytes) read from the device per second.

wsec/s (wkB/s, wMB/s)

The number of sectors (kilobytes, megabytes) written to the device per second.

dsec/s (dkB/s, dMB/s)

The number of sectors (kilobytes, megabytes) discarded for the device per second.

rqm/s The number of I/O requests merged per second that were queued to the device.

rrqm/s The number of read requests merged per second that were queued to the device.

wrqm/s The number of write requests merged per second that were queued to the device.

drqm/s The number of discard requests merged per second that were queued to the device.

%rrqm The percentage of read requests merged together before being sent to the device.

%wrqm The percentage of write requests merged together before being sent to the device.

%drqm The percentage of discard requests merged together before being sent to the device.

areq-sz

The average size (in kilobytes) of the I/O requests that were issued to the device.

Note: In previous versions, this field was known as av? grq-sz and was expressed in sectors.

rareq-sz

The average size (in kilobytes) of the read requests that were issued to the device.

wareq-sz

The average size (in kilobytes) of the write requests that were issued to the device.

dareq-sz

that were issued to the device.

await The average time (in milliseconds) for I/O requests is? sued to the device to be served. This includes the time spent by the requests in queue and the time spent servic? ing them.

r_await

The average time (in milliseconds) for read requests is? sued to the device to be served. This includes the time spent by the requests in queue and the time spent servic? ing them.

w_await

The average time (in milliseconds) for write requests is? sued to the device to be served. This includes the time spent by the requests in queue and the time spent servic? ing them.

d_await

The average time (in milliseconds) for discard requests issued to the device to be served. This includes the time spent by the requests in queue and the time spent servic? ing them.

f_await

The average time (in milliseconds) for flush requests is? sued to the device to be served. The block layer com? bines flush requests and executes at most one at a time. Thus flush operations could be twice as long: Wait for current flush request, then execute it, then wait for the next one.

aqu-sz The average queue length of the requests that were issued to the device.

Note: In previous versions, this field was known as avgqu-sz.

%util Percentage of elapsed time during which I/O requests were issued to the device (bandwidth utilization for the de? vice). Device saturation occurs when this value is close to 100% for devices serving requests serially. But for devices serving requests in parallel, such as RAID arrays and modern SSDs, this number does not reflect their per? formance limits.

OPTIONS

-c Display the CPU utilization report.

Don't break the Device Utilization Report into sub-reports so that all the metrics get displayed on a single line.

-d Display the device utilization report.

--dec={ 0 | 1 | 2 }

Specify the number of decimal places to use (0 to 2, default value is 2).

-f directory

+f directory

Specify an alternative directory for iostat to read devices sta? tistics. Option -f tells iostat to use only the files located in the alternative directory, whereas option +f tells it to use both the standard kernel files and the files located in the al? ternative directory to read device statistics.

directory is a directory containing files with statistics for devices managed in userspace. It may contain:

- a "diskstats" file whose format is compliant with that located in "/proc",

- statistics for individual devices contained in files whose format is compliant with that of files located in "/sys".

In particular, the following files located in directory may be used by iostat:

directory/block/device/stat
directory/block/device/partition/stat

partition files must have an entry in directory/dev/block/ di? rectory, e.g.:

directory/dev/block/major:minor --> ../../block/device/partition

-g group_name { device [...] | ALL }
Display statistics for a group of devices. The iostat command reports statistics for each individual device in the list then a line of global statistics for the group displayed as group_name and made up of all the devices in the list. The ALL keyword means that all the block devices defined by the system shall be included in the group.

-H This option must be used with option -g and indicates that only global statistics for the group are to be displayed, and not statistics for individual devices in the group.

-h This option is equivalent to specifying --human --pretty.

--human

Print sizes in human readable format (e.g. 1.0k, 1.2M, etc.)

The units displayed with this option supersede any other default units (e.g. kilobytes, sectors...) associated with the metrics.

-j { ID | LABEL | PATH | UUID | ... } [device [...] | ALL]

Display persistent device names. Keywords ID, LABEL, etc. spec? ify the type of the persistent name. These keywords are not lim? ited, only prerequisite is that directory with required persis? tent names is present in /dev/disk. Optionally, multiple de? vices can be specified in the chosen persistent name type. Be? cause persistent device names are usually long, option --pretty is implicitly set with this option.

-k Display statistics in kilobytes per second.

-m Display statistics in megabytes per second.

-N Display the registered device mapper names for any device mapper devices. Useful for viewing LVM2 statistics.

-o JSON

Display the statistics in JSON (JavaScript Object Notation) for? mat. JSON output field order is undefined, and new fields may

-p [{ device[,...] | ALL }]

Display statistics for block devices and all their partitions that are used by the system. If a device name is entered on the command line, then statistics for it and all its partitions are displayed. Last, the ALL keyword indicates that statistics have to be displayed for all the block devices and partitions defined by the system, including those that have never been used. If op? tion -j is defined before this option, devices entered on the command line can be specified with the chosen persistent name type.

--pretty

Make the Device Utilization Report easier to read by a human. The device name will be printed on the right side. The report may also be broken into sub-reports if there are many metrics to display (use --compact option to prevent this).

-s Display a short (narrow) version of the report that should fit in 80 characters wide screens.

-t Print the time for each report displayed. The timestamp format may depend on the value of the S_TIME_FORMAT environment vari? able (see below).

-V Print version number then exit.

-x Display extended statistics.

-y Omit first report with statistics since system boot, if display? ing multiple records at given interval.

-z Tell iostat to omit output for any devices for which there was no activity during the sample period.

ENVIRONMENT

The iostat command takes into account the following environment vari? ables:

POSIXLY_CORRECT

When this variable is set, transfer rates are shown in 512-byte blocks instead of the default 1K blocks.

S_COLORS

By default statistics are displayed in color when the output is connected to a terminal. Use this variable to change the set? tings. Possible values for this variable are never, always or auto (the latter is equivalent to the default settings).

Please note that the color (being red, yellow, or some other

issue simply because of the color. It only indicates different ranges of values.

S_COLORS_SGR

Specify the colors and other attributes used to display statis? tics on the terminal. Its value is a colon-separated list of capabilities that defaults to H=31;1:I=32;22:M=35;1:N=34;1:Z=34;22. Supported capabilities are:

H= SGR (Select Graphic Rendition) substring for percentage values greater than or equal to 75%.

I= SGR substring for device names.

M= SGR substring for percentage values in the range from 50% to 75%.

N= SGR substring for non-zero statistics values.

Z= SGR substring for zero values.

S_TIME_FORMAT

If this variable exists and its value is ISO then the current

header. The iostat command will use the ISO 8601 format (YYYY-MM-DD) instead. The timestamp displayed with option -t will also be compliant with ISO 8601 format.

EXAMPLES

iostat Display a single history since boot report for all CPU and De? vices.

iostat -d 2

Display a continuous device report at two second intervals.

iostat -d 2 6

Display six reports at two second intervals for all devices.

iostat -x sda sdb 2 6

Display six reports of extended statistics at two second inter? vals for devices sda and sdb.

iostat -p sda 2 6

Display six reports at two second intervals for device sda and all its partitions (sda1, etc.)

BUGS

/proc filesystem must be mounted for iostat to work.

Kernels older than 2.6.x are no longer supported.

Although iostat speaks of kilobytes (kB), megabytes (MB)..., it actu? ally uses kibibytes (kiB), mebibytes (MiB)... A kibibyte is equal to 1024 bytes, and a mebibyte is equal to 1024 kibibytes.

FILES

/proc/stat contains system statistics.

/proc/uptime contains system uptime.

/proc/diskstats contains disks statistics.

/sys contains statistics for block devices.

/proc/self/mountstats contains statistics for network filesystems.

/dev/disk contains persistent device names.

AUTHOR

Sebastien Godard (sysstat <at> orange.fr)

SEE ALSO

sar(1), pidstat(1), mpstat(1), vmstat(8), tapestat(1), nfsiostat(1), cifsiostat(1)

https://github.com/sysstat/sysstat

http://pagesperso-orange.fr/sebastien.godard/