

NAME

systat - display system statistics

SYNOPSIS

systat [*-display*] [*display-commands*] [*refresh-interval*]

DESCRIPTION

The **systat** utility displays various system statistics in a screen oriented fashion using the curses screen display library, ncurses(3).

While **systat** is running the screen is usually divided into two windows (an exception is the vmstat display which uses the entire screen). The upper window depicts the current system load average. The information displayed in the lower window may vary, depending on user commands. The last line on the screen is reserved for user input and error messages.

By default **systat** displays the processes getting the largest percentage of the processor in the lower window. Other displays show swap space usage, disk I/O statistics (a la iostat(8)), virtual memory statistics (a la vmstat(8)), TCP/IP statistics, and network connections (a la netstat(1)).

Input is interpreted at two different levels. A “global” command interpreter processes all keyboard input. If this command interpreter fails to recognize a command, the input line is passed to a per-display command interpreter. This allows each display to have certain display-specific commands.

Command line options:

-display The *-* flag expects *display* to be one of: **icmp**, **icmp6**, **ifstat**, **iolat**, **iostat**, **ip**, **ip6**, **netstat**, **pigs**, **sctp**, **swap**, **tcp**, **vmstat**, or **zarc**. These displays can also be requested interactively (without the *-*) and are described in full detail below.

refresh-interval The *refresh-value* specifies the screen refresh time interval in seconds. Time interval can be fractional.

display-commands

A list of commands specific to this display. These commands can also be entered interactively and are described for each display separately below. If the command requires arguments, they can be specified as separate command line arguments. A command line argument *--* will finish display commands. For example:

```
systat -ifstat -match bge0,em1 -pps
```

This will display statistics of packets per second for network interfaces named as bge0 and em1.

systat -iostat -numbers -- 2.1

This will display all IO statistics in a numeric format and the information will be refreshed each 2.1 seconds.

Certain characters cause immediate action by **systat**. These are

- ^L** Refresh the screen.
- ^G** Print the name of the current “display” being shown in the lower window and the refresh interval.
- :** Move the cursor to the command line and interpret the input line typed as a command. While entering a command the current character erase, word erase, and line kill characters may be used.

The following commands are interpreted by the “global” command interpreter.

- help** Print the names of the available displays on the command line.
- load** Print the load average over the past 1, 5, and 15 minutes on the command line.
- stop** Stop refreshing the screen.

[start] [number]

Start (continue) refreshing the screen. If a second, numeric, argument is provided it is interpreted as a refresh interval (in seconds). Supplying only a number will set the refresh interval to this value.

- quit** Exit **systat**. (This may be abbreviated to **q**.)

The available displays are:

- pigs** Display, in the lower window, those processes resident in main memory and getting the largest portion of the processor (the default display). When less than 100% of the processor is scheduled to user processes, the remaining time is accounted to the “idle” process.

icmp Display, in the lower window, statistics about messages received and transmitted by the Internet Control Message Protocol ("ICMP"). The left half of the screen displays information about received packets, and the right half displays information regarding transmitted packets.

The **icmp** display understands two commands: **mode** and **reset**. The **mode** command is used to select one of four display modes, given as its argument:

rate: show the rate of change of each value in packets (the default) per second
delta: show the rate of change of each value in packets per refresh interval
since: show the total change of each value since the display was last reset
absolute: show the absolute value of each statistic

The **reset** command resets the baseline for **since** mode. The **mode** command with no argument will display the current mode in the command line.

icmp6 This display is like the **icmp** display, but displays statistics for IPv6 ICMP.

ip Otherwise identical to the **icmp** display, except that it displays IP and UDP statistics.

ip6 Like the **ip** display, except that it displays IPv6 statistics. It does not display UDP statistics.

sctp Like **icmp**, but with SCTP statistics.

tcp Like **icmp**, but with TCP statistics.

iolat Display statistics describing the hardware latencies of I/O operations as computed by the *CAM_IOSCHED_DYNAMIC* option. This option must be in the kernel config file of the running kernel for this display to work. All devices are displayed as there is currently no way to filter them. The statistics displayed for the I/O latencies are the percentiles with sufficient data during the polling interval to compute. If a value cannot be estimated “-” is displayed. The P50 (also known as the median), P90, P99 and P99.9 values are computed if more than 2, 10, 100 or 1000 operations occurred during the polling interval. The latency is the hardware latency values, and does not include any software queuing time. The latencies are estimated based on histogram data computed by the CAM I/O scheduler and represent estimates of the actual value that are only good to two or three significant digits. The display of latency changes based on the scale of the latency to reflect the precision of the estimates and to fit on the available screen space. All latencies are reported in milliseconds. When color is enabled

- Values below the medium latency threshold are displayed in green.
- Values between the minimum latency and high latency thresholds are displayed in magenta.
- Values above the high latency thresholds are displayed in red.

When color is disabled, the default foreground and background colors are always used.

The following commands are specific to the **iolat** display; the minimum unambiguous prefix may be supplied.

color	Toggle the use of color in the display. The default is on.
hi=XXX	Set the high latency threshold to XXX milliseconds.
med=XXX	Set the medium latency threshold to XXX milliseconds.
read	Toggle the display of statistics about read operations. The default is on.
write	Toggle the display of statistics about write operations. The default is on.
trim	Toggle the display of statistics about trim operations. The default is on.

iostat

Display, in the lower window, statistics about processor use and disk throughput. Statistics on processor use appear as bar graphs of the amount of time executing in user mode (“user”), in user mode running low priority processes (“nice”), in system mode (“system”), in interrupt mode (“interrupt”), and idle (“idle”). Statistics on disk throughput show, for each drive, megabytes per second, average number of disk transactions per second, and average kilobytes of data per transaction. This information may be displayed as bar graphs or as rows of numbers which scroll downward. Bar graphs are shown by default.

The following commands are specific to the **iostat** display; the minimum unambiguous prefix may be supplied.

numbers	Show the disk I/O statistics in numeric form. Values are displayed in numeric columns which scroll downward.
bars	Show the disk I/O statistics in bar graph form (default).
kbpt	Toggle the display of kilobytes per transaction. (the default is to not display kilobytes per transaction).

swap

Show information about swap space usage on all the swap areas compiled into the kernel and processes that are swapped out as well as a summary of disk activity.

The swap areas are displayed first with their name, sizes and usage percentage. The *Used* column indicates the total blocks used so far; the graph shows the percentage of space in use on each partition. If there are more than one swap partition in use, a total line is also shown. Areas known to the kernel, but not in use are shown as not available.

Below the swap space statistics, processes are listed in order of higher swap area usage. Pid, username, a part of command line, the total use of swap space in bytes, the size of process, as well as per-process swap usage percentage and per-system swap space percentage are shown per process.

At the bottom left is the disk usage display. It reports the number of kilobytes per transaction, transactions per second, megabytes per second and the percentage of the time the disk was busy averaged over the refresh period of the display (by default, five seconds). The system keeps statistics on most every storage device. In general, up to seven devices are displayed. The devices displayed by default are the first devices in the kernel's device list. See `devstat(3)` and `devstat(9)` for details on the `devstat` system.

vmstat Take over the entire display and show a (rather crowded) compendium of statistics related to virtual memory usage, process scheduling, device interrupts, system name translation caching, disk I/O etc.

The upper left quadrant of the screen shows the number of users logged in and the load average over the last one, five, and fifteen minute intervals. Below this line are statistics on memory utilization. The first row of the table reports memory usage only among active processes, that is processes that have run in the previous twenty seconds. The second row reports on memory usage of all processes. The first column reports on the number of kilobytes in physical pages claimed by processes. The second column reports the number of kilobytes in physical pages that are devoted to read only text pages. The third and fourth columns report the same two figures for virtual pages, that is the number of kilobytes in pages that would be needed if all processes had all of their pages. Finally the last column shows the number of kilobytes in physical pages on the free list.

Below the memory display is a list of the average number of threads (over the last refresh interval) that are runnable ('r'), in page wait ('p'), in disk wait other than paging ('d'), sleeping ('s'), and swapped out but desiring to run ('w'). The row also shows the average number of context switches ('Csw'), traps ('Trp'; includes page faults), system calls ('Sys'), interrupts ('Int'), network software interrupts ('Sof'), and page faults ('Flt').

Below the process queue length listing is a numerical listing and a bar graph showing the amount of system (shown as '='), interrupt (shown as '+'), user (shown as '>'), nice

(shown as '-'), and idle time (shown as ' ').

Below the process display are statistics on name translations. It lists the number of names translated in the previous interval, the number and percentage of the translations that were handled by the system wide name translation cache, and the number and percentage of the translations that were handled by the per process name translation cache.

To the right of the name translations display are lines showing the number of dirty buffers in the buffer cache ('dtbuf'), desired maximum size of vnode cache ('desvn'), number of vnodes actually allocated ('numvn'), and number of allocated vnodes that are free ('frevn').

At the bottom left is the disk usage display. It reports the number of kilobytes per transaction, transactions per second, megabytes per second and the percentage of the time the disk was busy averaged over the refresh period of the display (by default, five seconds). The system keeps statistics on most every storage device. In general, up to seven devices are displayed. The devices displayed by default are the first devices in the kernel's device list. See devstat(3) and devstat(9) for details on the devstat system.

Under the date in the upper right hand quadrant are statistics on paging and swapping activity. The first two columns report the average number of pages brought in and out per second over the last refresh interval due to page faults and the paging daemon. The third and fourth columns report the average number of pages brought in and out per second over the last refresh interval due to swap requests initiated by the scheduler. The first row of the display shows the average number of disk transfers per second over the last refresh interval; the second row of the display shows the average number of pages transferred per second over the last refresh interval.

Below the paging statistics is a column of lines regarding the virtual memory system. The first few lines describe, in units (except as noted below) of pages per second averaged over the sampling interval, pages copied on write ('cow'), pages zero filled on demand ('zfod'), pages optimally zero filled on demand ('ozfod'), the ratio of the (average) ozfod / zfod as a percentage ('%ozfod'), pages freed by the page daemon ('daefr'), pages freed by exiting processes ('prcfr'), total pages freed ('totfr'), pages reactivated from the free list ('react'), the average number of times per second that the page daemon was awakened ('pdwak'), pages analyzed by the page daemon ('pdpgs'), and in-transit blocking page faults ('intrn'). Note that the units are special for '%ozfod' and 'pdwak'. The next few lines describe, as amounts of memory in kilobytes, pages wired down ('wire'), active pages ('act'), inactive pages ('inact'), dirty pages queued for laundering ('laund'), and free pages ('free'). Note that the values displayed are the current transient ones; they are not averages.

At the bottom of this column is a line showing the amount of virtual memory, in kilobytes, mapped into the buffer cache ('buf'). This statistic is not useful. It exists only as a placeholder for the corresponding useful statistic (the amount of real memory used to cache disks). The most important component of the latter (the amount of real memory used by the vm system to cache disks) is not available, but can be guessed from the 'inact' amount under some system loads.

Running down the right hand side of the display is a breakdown of the interrupts being handled by the system. At the top of the list is the total interrupts per second over the time interval. The rest of the column breaks down the total on a device by device basis. Only devices that have interrupted at least once since boot time are shown.

The following commands are specific to the **vmstat** display; the minimum unambiguous prefix may be supplied.

boot	Display cumulative statistics since the system was booted.
run	Display statistics as a running total from the point this command is given.
time	Display statistics averaged over the refresh interval (the default).
zero	Reset running statistics to zero.

zarc display arc cache usage and hit/miss statistics.

netstat Display, in the lower window, network connections. By default, network servers awaiting requests are not displayed. Each address is displayed in the format "host.port", with each shown symbolically, when possible. It is possible to have addresses displayed numerically, limit the display to a set of ports, hosts, and/or protocols (the minimum unambiguous prefix may be supplied):

all	Toggle the displaying of server processes awaiting requests (this is the equivalent of the -a flag to netstat(1)).
numbers	Display network addresses numerically.
names	Display network addresses symbolically.
proto protocol	Display only network connections using the indicated <i>protocol</i> . Supported protocols are "tcp", "udp", and "all".
ignore [items]	Do not display information about connections associated with the specified hosts or ports. Hosts and ports may be specified by name ("vangogh", "ftp"), or numerically. Host addresses use the Internet dot notation ("128.32.0.9"). Multiple items may be specified with a single command by separating them with spaces.
display [items]	Display information about the connections associated with the specified

hosts or ports. As for *ignore*, [*items*] may be names or numbers.

show [*ports/hosts*]

Show, on the command line, the currently selected protocols, hosts, and ports. Hosts and ports which are being ignored are prefixed with a '!'. If *ports* or *hosts* is supplied as an argument to **show**, then only the requested information will be displayed.

reset Reset the port, host, and protocol matching mechanisms to the default (any protocol, port, or host).

ifstat Display the network traffic going through active interfaces on the system. Idle interfaces will not be displayed until they receive some traffic.

For each interface being displayed, the current, peak and total statistics are displayed for incoming and outgoing traffic. By default, the **ifstat** display will automatically scale the units being used so that they are in a human-readable format. The scaling units used for the current and peak traffic columns can be altered by the **scale** command.

scale [*units*] Modify the scale used to display the current and peak traffic over all interfaces. The following units are recognised: kbit, kbyte, mbit, mbyte, gbit, gbyte and auto.

pps Show statistics in packets per second instead of bytes/bits per second. A subsequent call of **pps** switches this mode off.

match [*patterns*] Display only interfaces that match pattern provided as an argument. Patterns should be in shell syntax separated by whitespaces or commas. If this command is called without arguments then all interfaces are displayed. For example:

```
match em0, bge1
```

This will display em0 and bge1 interfaces.

```
match em*, bge*, lo0
```

This will display all **em** interfaces, all **bge** interfaces and the loopback interface.

Commands to switch between displays may be abbreviated to the minimum unambiguous prefix; for example, “io” for “iostat”. Certain information may be discarded when the screen size is insufficient

for display. For example, on a machine with 10 drives the **iostat** bar graph displays only 3 drives on a 24 line terminal. When a bar graph would overflow the allotted screen space it is truncated and the actual value is printed “over top” of the bar.

The following commands are common to each display which shows information about disk drives. These commands are used to select a set of drives to report on, should your system have more drives configured than can normally be displayed on the screen.

ignore [*drives*] Do not display information about the drives indicated. Multiple drives may be specified, separated by spaces.

display [*drives*]

Display information about the drives indicated. Multiple drives may be specified, separated by spaces.

only [*drives*] Display only the specified drives. Multiple drives may be specified, separated by spaces.

drives Display a list of available devices.

match *type,if,pass* [| ...]

Display devices matching the given pattern. The basic matching expressions are the same as those used in `iostat(8)` with one difference. Instead of specifying multiple **-t** arguments which are then ORed together, the user instead specifies multiple matching expressions joined by the pipe (‘|’) character. The comma separated arguments within each matching expression are ANDed together, and then the pipe separated matching expressions are ORed together. Any device matching the combined expression will be displayed, if there is room to display it. For example:

```
match da,scsi | cd,ide
```

This will display all SCSI Direct Access devices and all IDE CDROM devices.

```
match da | sa | cd,pass
```

This will display all Direct Access devices, all Sequential Access devices, and all passthrough devices that provide access to CDROM drives.

FILES

/boot/kernel/kernel For the namelist.
/dev/kmem For information in main memory.
/etc/hosts For host names.
/etc/networks For network names.
/etc/services For port names.

SEE ALSO

netstat(1), kvm(3), icmp(4), icmp6(4), ip(4), ip6(4), tcp(4), udp(4), gstat(8), iostat(8), vmstat(8)

HISTORY

The **systat** program appeared in 4.3BSD. The **icmp**, **ip**, and **tcp** displays appeared in FreeBSD 3.0; the notion of having different display modes for the ICMP, IP, TCP, and UDP statistics was stolen from the **-C** option to netstat(1) in Silicon Graphics' IRIX system.

BUGS

Certain displays presume a minimum of 80 characters per line. Ifstat does not detect new interfaces. The **vmstat** display looks out of place because it is (it was added in as a separate display rather than created as a new program). The **iolat** command does not implement the common device commands including filtering, as it does not use the devstat(3) mechanism to obtain its statistics.