

NAME

sadf – Display data collected by **sar** in multiple formats.

SYNOPSIS

```
sadf [ -C ] [ -c | -d | -g | -j | -l | -p | -r | -x ] [ -H ] [ -h ] [ -T | -t | -U ] [ -V ] [ -O opts [,...] ] [ -P { cpu_list | ALL } ] [ -s [ hh:mm[:ss] ] ] [ -e [ hh:mm[:ss] ] ] [ --dev= dev_list ] [ --fs= fs_list ] [ --iface= iface_list ] [ --sar_options ] [ interval [ count ] ] [ datafile | -[0-9]+ ]
```

DESCRIPTION

The **sadf** command is used for displaying the contents of data files created by the **sar**(1) command. But unlike **sar**, **sadf** can write its data in many different formats (CSV, XML, etc.) The default format is one that can easily be handled by pattern processing commands like **awk** (see option **-p**). The **sadf** command can also be used to draw graphs for the various activities collected by **sar** and display them as SVG (Scalable Vector Graphics) graphics in your web browser (see option **-g**).

The **sadf** command extracts and writes to standard output records saved in the *datafile* file. This file must have been created by a version of **sar** which is compatible with that of **sadf**. If *datafile* is omitted, **sadf** uses the standard system activity daily data file. It is also possible to enter **-1**, **-2** etc. as an argument to **sadf** to display data of that days ago. For example, **-1** will point at the standard system activity file of yesterday.

The standard system activity daily data file is named *saDD* or *saYYYYMMDD*, where *YYYY* stands for the current year, *MM* for the current month and *DD* for the current day. **sadf** will look for the most recent of *saDD* and *saYYYYMMDD*, and use it. By default it is located in the */var/log/sysstat* directory. Yet it is possible to specify an alternate location for it: If *datafile* is a directory (instead of a plain file) then it will be considered as the directory where the standard system activity daily data file is located.

The *interval* and *count* parameters are used to tell **sadf** to select *count* records at *interval* seconds apart. If the *count* parameter is not set, then all the records saved in the data file will be displayed.

All the activity flags of **sar** may be entered on the command line to indicate which activities are to be reported. Before specifying them, put a pair of dashes (**--**) on the command line in order not to confuse the flags with those of **sadf**. Not specifying any flags selects only CPU activity.

OPTIONS

- C** Tell **sadf** to display comments present in file.
- c** Convert an old system activity binary datafile (version 9.1.6 and later) to current up-to-date format. Use the following syntax:

```
sadf -c old_datafile > new_datafile
```

- d** Print the contents of the data file in a format that can easily be ingested by a relational database system. The output consists of fields separated by a semicolon. Each record contains the hostname of the host where the file was created, the interval value (or **-1** if not applicable), the timestamp in a form easily acceptable by most databases, and additional semicolon separated data fields as specified by *sar_options* command line options. Note that timestamp output can be controlled by options **-T**, **-t** and **-U**.

--dev=dev_list

Specify the block devices for which statistics are to be displayed by **sadf**. *dev_list* is a list of comma-separated device names. Useful with option **-d** from **sar**.

-e [*hh:mm[:ss]*]

Set the ending time of the report. The default ending time is 18:00:00. Hours must be given in 24-hour format.

--fs=fs_list
Specify the filesystems for which statistics are to be displayed by **sadf**. *fs_list* is a list of comma-separated filesystem names or mountpoints. Useful with option **-F** from **sar**.

-g Print the contents of the data file in SVG (Scalable Vector Graphics) format. This option enables you to display some fancy graphs in your web browser. Use the following syntax:

```
sadf -g your_datafile [ -- sar_options ] > output.svg
```

and open the resulting SVG file in your favorite web browser. Output can be controlled using option **-O** (see below).

-H Display only the header of the report (when applicable). If no format has been specified, then the header data (metadata) of the data file are displayed.

-h When used in conjunction with option **-d**, all activities will be displayed horizontally on a single line.

--iface=iface_list
Specify the network interfaces for which statistics are to be displayed by **sadf**. *iface_list* is a list of comma-separated interface names. Useful with options **-n DEV** and **-n EDEV** from **sar**.

-j Print the contents of the data file in JSON (JavaScript Object Notation) format. Timestamps can be controlled by options **-T** and **-t**.

-l Export the contents of the data file to a PCP (Performance Co-Pilot) archive. The name of the archive can be specified using the keyword *pcparchive=* with option **-O**.

-O opts [...]
Use the specified options to control the output of **sadf**. The following options are used to control SVG output displayed by **sadf -g**:

autoscale

Draw all the graphs of a given view as large as possible based on current view's scale. To do this, a factor (10, 100, 1000...) is used to enlarge the graph drawing. This option may be interesting when several graphs are drawn on the same view, some with only very small values, and others with high ones, the latter making the former hardly visible.

bwcol

Use a black and white palette to draw the graphs.

customcol

Use a customizable color palette instead of the default one to draw the graphs. See environment variable **S_COLORS_PALETTE** below to know how to customize that palette.

height=value

Set SVG canvas height to *value*.

oneday

Display graphs data over a period of 24 hours. Note that hours are still printed in UTC by default: You should use option **-T** to print them in local time and get a time window starting from midnight.

packed

Group all views from the same activity (and for the same device) on the same row.

showidle

Also display *%idle* state in graphs for CPU statistics.

showinfo

Display additional information (such as the date and the host name) on each view.

showtoc

Add a table of contents at the beginning of the SVG output, consisting of links pointing at the first graph of each activity.

skipempty

Do not display views where all graphs have only zero values.

The following option may be used when data are exported to a PCP archive:

pcparchive=name

Specify the name of the PCP archive to create.

The following option is used to control raw output displayed by **sadf -r**:

debug

Display additional information, mainly useful for debugging purpose.

-P { cpu_list | ALL }

Tell **sadf** that processor dependent statistics are to be reported only for the specified processor or processors. *cpu_list* is a list of comma-separated values or range of values (e.g., **0,2,4-7,12-**). Note that processor 0 is the first processor, and processor **all** is the global average among all processors. Specifying the **ALL** keyword reports statistics for each individual processor, and globally for all processors.

-p Print the contents of the data file in a format that can easily be handled by pattern processing commands like `awk`. The output consists of fields separated by a tab. Each record contains the host-name of the host where the file was created, the interval value (or `-1` if not applicable), the timestamp, the device name (or `-` if not applicable), the field name and its value. Note that timestamp output can be controlled by options `-T`, `-t` and `-U`.

-r Print the raw contents of the data file. With this format, the values for all the counters are displayed as read from the kernel, which means e.g., that no average values are calculated over the elapsed time interval.

-s [hh:mm[:ss]]

Set the starting time of the data, causing the **sadf** command to extract records time-tagged at, or following, the time specified. The default starting time is 08:00:00. Hours must be given in 24-hour format.

-T Display timestamp in local time instead of UTC (Coordinated Universal Time).

-t Display timestamp in the original local time of the data file creator instead of UTC (Coordinated Universal Time).

-U Display timestamp (UTC - Coordinated Universal Time) in seconds from the epoch.

-V Print version number then exit.

-x Print the contents of the data file in XML format. Timestamps can be controlled by options `-T` and `-t`. The corresponding DTD (Document Type Definition) and XML Schema are included in the `sysstat` source package. They are also available at <http://pagesperso-orange.fr/sebastien.godard/download.html>

ENVIRONMENT

The **sadf** command takes into account the following environment variables:

S_COLORS_PALETTE

Specify the colors used by **sadf -g** to render the SVG output. This environment variable is taken into account only when the custom color palette has been selected with the option *customcol* (see option **-O**). Its value is a colon-separated list of capabilities associated with six-digit, three-byte hexadecimal numbers (hex triplets) representing colors that defaults to

**0=000000:1=1a1aff:2=1affb2:3=b21aff:4=1ab2ff:5=ff1a1a:6=ffb31a:7=b2ff1a:
8=efefef:9=000000:A=1a1aff:B=1affb2:C=b21aff:D=1ab2ff:E=ff1a1a:F=ffb31a:
G=bebebe:H=000000:I=000000:K=ffffff:L=000000:T=000000:W=000000:X=000000**

Capabilities consisting of an hexadecimal digit (0 through F) are used to specify the first sixteen colors in the palette (these colors are used to draw the graphs), e.g., 3=ffffff would indicate that the third color in the palette is white (0xffffff). Other capabilities are:

- G=** Specify the color used to draw the grid lines.
- H=** Specify the color used to display the report header.
- I=** Specify the color used to display additional information (e.g., date, hostname...)
- K=** Specify the color used for the graphs background.
- L=** Specify the default color (which is for example used to display the table of contents).
- T=** Specify the color used to display the graphs title.
- W=** Specify the color used to display warning and error messages.
- X=** Specify the color used to draw the axes and display the graduations.

S_TIME_DEF_TIME

If this variable exists and its value is **UTC** then **sadf** will use UTC time instead of local time to determine the current daily data file located in the */var/log/sysstat* directory.

EXAMPLES

sadf -d /var/log/sysstat/sa21 -- -r -n DEV

Extract memory and network statistics from system activity file 'sa21', and display them in a format that can be ingested by a database.

sadf -p -P 1

Extract CPU statistics for processor 1 (the second processor) from current daily data file, and display them in a format that can easily be handled by a pattern processing command.

BUGS

SVG output (as created by option **-g**) is fully compliant with SVG 1.1 standard. Graphics have been successfully displayed in various web browsers, including Firefox, Chrome and Opera. Yet SVG rendering is broken on Microsoft browsers (tested on Internet Explorer 11 and Edge 13.1): So please don't use them.

FILES

/var/log/sysstat/saDD
/var/log/sysstat/saYYYYMMDD

The standard system activity daily data files and their default location. YYYY stands for the current year, MM for the current month and DD for the current day.

AUTHOR

Sebastien Godard (sysstat <at> orange.fr)

SEE ALSO

sar(1), **sadc(8)**, **sa1(8)**, **sa2(8)**, **sysstat(5)**

<https://github.com/sysstat/sysstat>

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