NAME

atopgpud - GPU statistics daemon

SYNOPSIS

atopgpud [-v]

DESCRIPTION

The *atopgpud* daemon gathers statistical information from all Nvidia GPUs in the current system. With a sampling rate of one second, it maintains the statistics of every GPU, globally (system level) and per process. When *atopgpud* is active on the target system, *atop* connects to this daemon via a TCP socket and obtains all GPU statistics with every interval.

The approach to gather all GPU statistics in a separate daemon is required, because the Nvidia driver only offers the GPU busy percentage of the last second. Suppose that *atop* runs with a 10-minute interval and would fetch the GPU busy percentage directly from the Nvidia driver, it would reflect the busy percentage of the last second instead of the average busy percentage during 600 seconds. Therefore, the *atopgpud* daemon fetches the GPU busy percentage every second and accumulates this into a counter that can be retrieved by *atop* regularly. The same approach applies to other GPU statistics.

When the *atopgpud* daemon runs with root privileges, more process level counters (i.e. GPU busy and GPU memory busy per process) are provided that are otherwise not applicable.

Notice that certain GPU statistics are only delivered for specific GPU types. For older or less sophisticated GPUs, the value -1 is returned for counters that are not maintained. In the output of *atop* these counters are shown as 'N/A'.

When no (Nvidia) GPUs can be found in the target system, *atopgpud* immediately terminates with exit code 0.

Log messages are written via the *rsyslogd* daemon with facility 'daemon'. With the -v flag (verbose), *atopgpud* also logs debug messages.

INSTALLATION

The *atopgpud* daemon is written in Python, so a Python interpreter should be installed on the target system. This can either be Python version 2 or Python version 3 (the code of *atopgpud* is written in a generic way). Take care that the first line of the *atopgpud* script contains the proper command name to activate a Python interpreter that is installed on the target system!

The *atopgpud* daemon depends on the Python module *pynvml* to interface with the Nvidia driver. This module can be installed by the *pip* or *pip3* command and is usually packaged under the name *nvidia-ml-py* Finally, the *pynvml* module is a Python wrapper around the *libnvidia-ml* shared library that needs to be installed as well.

After installing the *atop* package, the *atopgpud* is not automatically started, nor will the service be enabed by default. When you want to activate this service (permanently), enter the following commands (as root):

systemctl enable atopgpu
systemctl start atopgpu

INTERFACE DESCRIPTION

Client processes can connect to the *atopgpud* daemon on TCP port 59123. Subsequently, such client can send a request of two bytes, consisting of one byte request code followed by one byte integer being the API version number.
The request code in the first byte can be 'T' to obtain information about the GPU types installed in this system (usually only requested once).
The request code can be 'S' to obtain all statistical counter values (requested for every interval).

The response of the daemon starts with a 4-byte integer. The first byte is the API version number that determines the response format while the subsequent three bytes indicate the length (big endian order) of the response string that follows.
In the response strings the character '@' introduces system level information of one specific GPU and the character '#' introduces process level information related to that GPU.

For further details about the meaning of the counters in a response string, please consult the source code.

SEE ALSO

atop(1), atopsar(1), atoprc(5), netatop(4), netatopd(8), atopacctd(8)
https://www.atoptool.nl

AUTHOR

Gerlof Langeveld (gerlof.langeveld@atoptool.nl)