NAME

netatop - per-process network statistics gathering

SYNOPSIS

modprobe netatop

DESCRIPTION

This manual page documents the *netatop* kernel module. This module uses the netfilter interface to gather statistics about TCP and UDP traffic per task, on level of process (thread group) and individual thread.

Once the module is loaded via modprobe, it is called for every packet that is transmitted by a local process and every packet that is received from an interface. For each packet it tries to determine the related process and thread to maintain statistical counters about the number of packets transmitted/received and the number of bytes transmitted/received. Separate counters are maintained for TCP and UDP packets. It does not only view the packets that contain the user data but also the protocol related packets (like SYN, ACK, ... for the TCP protocol).
Received packets can only be identified as related to a process after that process has transmitted at least one packet in the process' context.

DETAILS

When the module discovers a packet for a new connection quintuple (TCP) or for a new local UDP port, it creates a so-called sockinfo structure. As soon as the *netatop* module knows to what process or thread the sockinfo struct is related, a reference is made from the sockinfo struct to the taskinfo struct that represents the proces or thread within the module. However, the related task can only be determined when a packet is transmitted, i.e. the module is called in the context of the transmitting process. At such moment the tgid (process) and pid (thread) can be obtained from the process administration to be stored in the module's own taskinfo structs (one for the process, one for the thread).
For the time that the sockinfo struct can not be related to a taskinfo struct (e.g. when only packets are received so far), counters are maintained temporarily in the sockinfo struct. After a related taskinfo struct has been discovered (i.e. the task transmits), counters will be maintained in the taskinfo struct itself. When packets are only **received** for a socket (e.g. another machine is sending UDP packets to the local machine) while the local task never responds, no match to a process can be made and the packets remain unidentified by the netatop module. At least one packet should have been sent by a local process to be able to identify packets for such process.

The module uses a garbage collector to cleanup the unused sockinfo structs if connections do not exist any more (TCP) or have not been used for some time (TCP/UDP). Furthermore, the garbage collector checks if taskinfo structs still represent existing processes or threads. If not, the taskinfo struct is destroyed (in case of a thread) or it is moved to a separate list of finished processes (in case of a process). Analysis programs can read the taskinfo of such finished process. When the taskinfo struct of a finished process is not read within 15 seconds, the taskinfo struct will be destroyed from the exitlist.
The garbage collector can be acivated by issueing a special getsockopt call from an analysis program (e.g. atop). Apart from that, a time-based garbage collector activation is issued anyhow every 15 seconds.

SUPPORTED IOCTLS

Programs can open an IP socket and use the getsockopt() system call to issue commands to this module. With this system call the following commands can be issued:

ATOP_GETCNT_TGID

Obtain the current counters for a specific process (thread group) in a netpertask structure. When the required process does not exist, errno ESRCH is given.

ATOP_GETCNT_PID

Obtain the current counters for a specific thread in a netpertask structure. When the required thread does not exist, errno ESRCH is given.

ATOP_GETCNT_EXIT

Obtain the counters of an exited process. This command has to be issued within 15 seconds after a process has been declared 'finished' by the garbage collector. Whenever this command is issued while there is no exited process in the exitlist, the requesting process is blocked until an exited process is

available.

NETATOP_FORCE_GC
Activate the garbage collector of the *netatop* module to determine if sockinfo structs of old connections/ports can be destroyed and to determine if taskinfo structs of exited processes can be moved to the exitlist.

NETATOP_EMPTY_EXIT
Block the calling process until the exitlist with the taskinfo structs of exited processes is empty.

FILES

/proc/netatop
In this file, counters can be found that show the total number of packets sent/received and the number of packets that were unidentified (i.e. not accounted to a process/thread).
Furthermore, counters can be found about the current number of sockinfo structs, active taskinfo structs and taskinfo structs of exited processes.

SEE ALSO

netatopd(8), atop(1), atopsar(1), atoprc(5)
http://www.atoptool.nl

AUTHOR

Gerlof Langeveld (gerlof.langeveld@atoptool.nl)