

[howto](#), [linux](#), [windows](#), [aix](#), [diagnosis](#), [monitoring](#), [tools](#), [top](#), [htop](#), [atop](#), [netatop](#), [apache](#), [apachetop](#), [modstatus](#), [proftpd](#), [ftptop](#), [mytop](#), [mysql](#), [mariadb](#), [powertop](#), [iotop](#), [ntopng](#), [iftop](#), [jnettop](#), [bandwidthd](#), [etherape](#), [ethtool](#), [nethogs](#), [iptraf](#), [ngrep](#), [mrtg](#), [bmon](#), [traceroute](#), [traceroute6](#), [tracert6](#), [iptstate](#), [darkstat](#), [vnstat](#), [netstat](#), [ss](#), [nmap](#), [mtr](#), [tcpdump](#), [justniffer](#), [serverdensity](#), [opennms](#), [sysusage](#), [pcp](#), [ksysguard](#), [munin](#), [nagios](#), [icinga](#), [zenoss](#), [cacti](#), [zabbix](#), [nmon](#), [conky](#), [glances](#), [squid](#), [sarg](#), [saidar](#), [libstatgrab](#), [rrdtool](#), [procexp](#), [linuxprocessexplorer](#), [df](#), [dstat](#), [netsnmp](#), [incron](#), [monitorix](#), [vmstat](#), [mpstat](#), [pmap](#), [uptime](#), [ps](#), [sar](#), [sysstat](#), [iostat](#), [collectl](#), [free](#), [goaccess](#), [logwatch](#), [swatchdog](#), [multitail](#)



# Linux Diagnosis and Monitoring Tools

## Sources

- <https://www.serverdensity.com/monitor/linux/how-to/>
- <https://www.tecmint.com/linux-performance-monitoring-tools/>
- <https://www.tecmint.com/command-line-tools-to-monitor-linux-performance/>

## top

- [https://en.wikipedia.org/wiki/Top\\_\(software\)](https://en.wikipedia.org/wiki/Top_(software))
- <https://www.tecmint.com/12-top-command-examples-in-linux/>

top (table of processes) is a task manager program, found in many Unix-like operating systems, that displays information about CPU and memory utilization. The program produces an ordered list of running processes selected by user-specified criteria, and updates it periodically. Default ordering is by CPU usage, and only the top CPU consumers are shown. top shows how much processing power and memory are being used, as well as other information about the running processes. Some versions of top allow extensive customization of the display, such as choice of columns or sorting method. top is useful for system administrators, as it shows which users and processes are consuming the most system resources at any given time.

## htop

- <https://htop.dev/>
- <https://man7.org/linux/man-pages/man1/htop.1.html>

htop is an interactive cross-platform process viewer. It is a text-mode application (for console or X terminals) and requires ncurses.

## atop

- <https://linux.die.net/man/1/atop>
- <https://www.atoptool.nl/>

The program atop is an interactive monitor to view the load on a Linux system. It shows the occupation of the most critical hardware resources (from a performance point of view) on system level, i.e. cpu, memory, disk and network. It also shows which processes are responsible for the indicated load with respect to cpu- and memory load on process level. Disk load is shown if per process "storage accounting" is active in the kernel or if the kernel patch 'cnt' has been installed. Network load is only shown per process if the kernel patch 'cnt' has been installed.

Atop is an ASCII full-screen performance monitor for Linux that is capable of reporting the activity of all processes (even if processes have finished during the interval), daily logging of system and process activity for long-term analysis, highlighting overloaded system resources by using colors, etc. At regular intervals, it shows system-level activity related to the CPU, memory, swap, disks (including LVM) and network layers, and for every process (and thread) it shows e.g. the CPU utilization, memory growth, disk utilization, priority, username, state, and exit code. In combination with the optional kernel module netatop, it even shows network activity per process/thread.

## apachetop

- <https://github.com/JeremyJones/Apachetop>
- <https://github.com/JeremyJones/Apachetop/wiki>

apachetop is a console-based (non-gui) tool for monitoring the threads and overall performance of a set of Apache web servers. It runs on Unix systems which have Perl, LWP, and Term::ReadKey installed. It is based largely on the excellent mytop tool written by Jeremy

Zawodny.

## Apache mod\_status

- [https://httpd.apache.org/docs/2.4/mod/mod\\_status.html](https://httpd.apache.org/docs/2.4/mod/mod_status.html)

The Status module allows a server administrator to find out how well their server is performing. A HTML page is presented that gives the current server statistics in an easily readable form. If required this page can be made to automatically refresh (given a compatible browser). Another page gives a simple machine-readable list of the current server state.

The details given are:

- The number of workers serving requests
- The number of idle workers
- The status of each worker, the number of requests that worker has performed and the total number of bytes served by the worker (\*)
- A total number of accesses and byte count served (\*)
- The time the server was started/restarted and the time it has been running for
- Averages giving the number of requests per second, the number of bytes served per second and the average number of bytes per request (\*)
- The current percentage CPU used by each worker and in total by all workers combined (\*)
- The current hosts and requests being processed (\*)

The lines marked “(\*)” are only available if ExtendedStatus is On. In version 2.3.6, loading mod\_status will toggle ExtendedStatus On by default.

## ftptop

- <https://linux.die.net/man/1/ftptop>

ftptop - display running status on proftpd server connections

## mytop

- <https://linux.die.net/man/1/mytop>
- <https://jeremy.zawodny.com/mysql/mytop/>
- <https://jeremy.zawodny.com/mysql/mytop/mytop.html>

mytop is a console-based (non-gui) tool for monitoring the threads and overall performance of a MySQL 3.22.x, 3.23.x, and 4.x server. It runs on most Unix systems (including Mac OS X) which have Perl, DBI, and Term::ReadKey installed. And with Term::ANSIColor installed you even get color. If you install Time::HiRes, you'll get good real-time queries/second stats. As of version 0.7, it even runs on Windows (somewhat).

## powertop

- <https://github.com/fenrus75/powertop>

PowerTOP is a Linux\* tool used to diagnose issues with power consumption and power management. In addition to being a diagnostic tool, PowerTOP also has an interactive mode you can use to experiment with various power management settings, for cases where the Linux distribution has not enabled those settings.

## iotop

- <https://linux.die.net/man/1/iotop>
- <http://guichaz.free.fr/iotop/>

Linux has always been able to show how much I/O was going on (the bi and bo columns of the vmstat 1 command). Iotop is a Python program with a top like UI used to show of behalf of which process is the I/O going on. It requires Python ≥ 2.7 and a Linux kernel ≥ 2.6.20 with the TASK\_DELAY\_ACCT CONFIG\_TASKSTATS, TASK\_IO\_ACCOUNTING and CONFIG\_VM\_EVENT\_COUNTERS options on.

## ntopng

- <https://www.ntop.org/products/traffic-analysis/ntop/>

ntopng - High-Speed Web-based Traffic Analysis and Flow Collection. ntopng is the next generation version of the original ntop, a network traffic probe that monitors network usage. ntopng is based on libpcap/PF\_RING and it has been written in a portable way in order to virtually run on every Unix platform, MacOS and on Windows as well. ntopng – yes, it's all lowercase – provides a intuitive, encrypted web user interface for the exploration of realtime and historical traffic information.

## iftop

- <https://linux.die.net/man/8/iftop>
- <https://code.blinkace.com/pdw/iftop>
- <https://www.ex-parrot.com/pdw/iftop/>

iftop is a free software command-line system monitor tool that produces a frequently updated list of network connections. By default, the connections are ordered by bandwidth usage, with only the “top” bandwidth consumers shown. The iftop website gives the following description: “iftop does for network usage what top(1) does for CPU usage. It listens to network traffic on a named interface and displays a table of current bandwidth usage by pairs of hosts. Handy for answering the question 'why is our ADSL link so slow?'”. iftop monitors network traffic and displays a table of current bandwidth usage. An interface may be specified or, if not, it will listen on the first interface it finds which looks like an external interface (with libpcap and libncurses). iftop must be run with sufficient permissions to monitor all network traffic; on most systems this means that it must be run as a root user, see sudo. By default, iftop will look up hostnames associated with addresses and counts all IP packets that pass through the filter. Hostname look-up can add substantial traffic, in and of itself, and may result in an inaccurate display of network traffic. You may wish to suppress display of DNS traffic by using filter code such as “not port domain”, or switch it off entirely, by using the -n option or by pressing “n” when the program is running. Using the -F option makes it possible to show packets entering and leaving a given network.

## jnettop

- <https://linux.die.net/man/8/jnettop>

jnettop - View hosts/ports taking up the most network traffic. jnettop captures traffic coming across the host it is running on and displays streams sorted by bandwidth they use. Result is a nice listing of communication on network by host and port, how many bytes went through this transport and the bandwidth it is consuming.

## bandwidthd

- <http://bandwidthd.sourceforge.net/>

BandwidthD tracks usage of TCP/IP network subnets and builds html files with graphs to display utilization. Charts are built by individual IPs, and by default display utilization over 2 day, 8 day, 40 day, and 400 day periods. Furthermore, each ip address's utilization can be logged out at intervals of 3.3 minutes, 10 minutes, 1 hour or 12 hours in cdf format, or to a backend database server. HTTP, TCP, UDP, ICMP, VPN, and P2P traffic are color coded. BandwidthD runs on most platforms including windows. Required libraries for unix are only: libpcap, libgl and libpng. Bandwidthd now produces output in 2 ways. The first is as a standalone application that produces static html and png output every 200 seconds. The second is as a sensor that transmits it's data to a backend database which is then reported on by dynamic php pages. The visual output of both is similar, but the database driven system allows for searching, filtering, multiple sensors and custom reports.

## EtherApe

- <https://etherape.sourceforge.io/>

EtherApe is a graphical network monitor for Unix modeled after etherman. Featuring link layer, IP and TCP modes, it displays network activity graphically. Hosts and links change in size with traffic. Color coded protocols display. It supports Ethernet, FDDI, Token Ring, ISDN, PPP, SLIP and WLAN devices, plus several encapsulation formats. It can filter traffic to be shown, and can read packets from a file as well as live from the network. Node statistics can be exported.

## ethtool

- <https://linux.die.net/man/8/ethtool>
- <https://git.kernel.org/pub/scm/network/ethtool/ethtool.git>
- <https://mirrors.edge.kernel.org/pub/software/network/ethtool/>

ethtool is the standard Linux utility for controlling network drivers and hardware, particularly for wired Ethernet devices. It can be used to: Get identification and diagnostic information. Get extended device statistics. Control speed, duplex, autonegotiation and flow control for Ethernet devices. Control checksum offload and other hardware offload features. Control DMA ring sizes and interrupt moderation. Control receive queue selection for multiqueue devices. Upgrade firmware in flash memory. Most features are dependent on support in the specific driver. See the manual page for full information.

## NetHogs

- <https://linux.die.net/man/8/nethogs>
- <https://github.com/raboof/nethogs>

NetHogs is a small 'net top' tool. Instead of breaking the traffic down per protocol or per subnet, like most tools do, it groups bandwidth by process. NetHogs does not rely on a special kernel module to be loaded. If there's suddenly a lot of network traffic, you can fire up NetHogs and immediately see which PID is causing this. This makes it easy to identify programs that have gone wild and are suddenly taking up your bandwidth. Since NetHogs heavily relies on /proc, most features are only available on Linux. NetHogs can be built on Mac OS X and FreeBSD, but it will only show connections, not processes.

## iptraf

- <http://iptraf.seul.org/>

IPTraf is a console-based network statistics utility for Linux. It gathers a variety of figures such as TCP connection packet and byte counts, interface statistics and activity indicators, TCP/UDP traffic breakdowns, and LAN station packet and byte counts. An IP traffic monitor that shows information on the IP traffic passing over your network. Includes TCP flag information, packet and byte counts, ICMP details, OSPF packet types. General and detailed interface statistics showing IP, TCP, UDP, ICMP, non-IP and other IP packet counts, IP checksum errors, interface activity, packet size counts. A TCP and UDP service monitor showing counts of incoming and outgoing packets for common TCP and UDP application ports. A LAN statistics module that discovers active hosts and shows statistics showing the data activity on them. TCP, UDP, and other protocol display filters, allowing you to view only traffic you're interested in. Logging. Supports Ethernet, FDDI, ISDN, SLIP, PPP, and loopback interface types. Utilizes the built-in raw socket interface of the Linux kernel, allowing it to be used over a wide range of supported network cards. Full-screen, menu-driven operation.

## ngrep

- <https://github.com/jpr5/ngrep>

ngrep is like GNU grep applied to the network layer. It's a PCAP-based tool that allows you to specify an extended regular or hexadecimal expression to match against data payloads of packets. It understands many kinds of protocols, including IPv4/6, TCP, UDP, ICMPv4/6, IGMP and Raw, across a wide variety of interface types, and understands BPF filter logic in the same fashion as more common packet sniffing tools, such as tcpdump and snoop.

## MRTG

- <https://oss.oetiker.ch/mrtg/>

Tobi Oetiker's MRTG - The Multi Router Traffic Grapher. You have a router, you want to know what it does all day long? Then MRTG is for you. It will monitor SNMP network devices and draw pretty pictures showing how much traffic has passed through each interface. Routers are only the beginning. MRTG is being used to graph all sorts of network devices as well as everything else from weather data to vending machines. MRTG is written in perl and works on Unix/Linux as well as Windows and even Netware systems. MRTG is free software licensed under the Gnu GPL.

## bmon

- <https://github.com/tgraf/bmon>

bmon - Bandwidth Monitor. bmon is a monitoring and debugging tool to capture networking related statistics and prepare them visually in a human friendly way. It features various output methods including an interactive curses user interface and a programmable text output for scripting.

## traceroute

- <https://linux.die.net/man/8/traceroute>

In computing, traceroute and tracert are computer network diagnostic commands for displaying possible routes (paths) and measuring transit delays of packets across an Internet Protocol (IP) network. The history of the route is recorded as the round-trip times of the packets received from each successive host (remote node) in the route (path); the sum of the mean times in each hop is a measure of the total time spent to establish the connection. Traceroute proceeds unless all (usually three) sent packets are lost more than twice; then the connection is lost and the route cannot be evaluated. Ping, on the other hand, only computes the final round-trip times from the destination point. For Internet Protocol Version 6 (IPv6) the tool sometimes has the name traceroute6 and tracert6.

## IPTState

- <https://www.phildev.net/iptables/>

IP Tables State. IPTState is a top-like interface to your netfilter connection-tracking table. Using iptstate you interactively watch where traffic crossing your netfilter/iptables firewall is going, sort by various criteria, limit the view by various criteria. But it doesn't stop there: as of version 2.2.0 you can even delete states from the table! The only requirements are a curses library (usually ncurses), and libnetfilter\_conntrack version 0.0.50 or later. IPTState is now in the Debian, Redhat, Fedora Core, Mandrake, Gentoo, FloppyFW, and many other distributions you can find a list of.

## darkstat

- <https://linux.die.net/man/8/darkstat>
- <https://unix4lyfe.org/darkstat/>
- <https://github.com/emikulic/darkstat>

darkstat captures network traffic, calculates statistics about usage, and serves reports over HTTP. Traffic graphs, reports per host, shows ports for each host. Embedded web-server with deflate compression. Asynchronous reverse DNS resolution using a child process. Small. Portable. Single-threaded. Efficient. Supports IPv6.

## vnStat

- <https://humdi.net/vnstat/>
- <https://github.com/Hulxv/vnstat-client>

vnStat is a console-based network traffic monitor for Linux and BSD that keeps a log of network traffic for the selected interface(s). It uses the network interface statistics provided by the kernel as information source. This means that vnStat won't actually be sniffing any traffic and also ensures light use of system resources regardless of network traffic rate. This program is open source / GPL'ed and can be installed either as root or as a single user. Better instructions are included in the README.

## netstat

- <https://linux.die.net/man/8/netstat>

In computing, netstat (network statistics) is a command-line network utility that displays network connections for Transmission Control Protocol (both incoming and outgoing), routing tables, and a number of network interface (network interface controller or software-defined network interface) and network protocol statistics. It is available on Unix, Plan 9, Inferno, and Unix-like operating systems including macOS, Linux, Solaris and BSD. It is also available on IBM OS/2 and on Microsoft Windows NT-based operating systems including Windows XP, Windows Vista, Windows 7, Windows 8 and Windows 10. It is used for finding problems in the network and to determine the amount of traffic on the network as a performance measurement. On Linux this program is mostly obsolete, although still included in many distributions. On Linux, netstat (part of "net-tools") is superseded by ss (part of iproute2). The replacement for netstat -r is ip route, the replacement for netstat -i is ip -s link, and the replacement for netstat -g is ip maddr, all of which are recommended instead.

## SS

- <https://man7.org/linux/man-pages/man8/ss.8.html>
- <https://www.redhat.com/sysadmin/ss-command>

ss - another utility to investigate sockets. ss is used to dump socket statistics. It allows showing information similar to netstat. It can display more TCP and state information than other tools.

## nmap

- <https://nmap.org/>
- <https://github.com/nmap/nmap>
- <https://nmap.org/book/man.html>

Nmap (Network Mapper) is a network scanner created by Gordon Lyon (also known by his pseudonym Fyodor Vaskovich).[4] Nmap is used to discover hosts and services on a computer network by sending packets and analyzing the responses. Nmap provides a number of features for probing computer networks, including host discovery and service and operating system detection. These features are extensible by scripts that provide more advanced service detection,[6] vulnerability detection,[6] and other features. Nmap can adapt to network conditions including latency and congestion during a scan. Nmap started as a Linux utility and was ported to other systems including Windows, macOS, and BSD. It is most popular on Linux, followed by Windows.

## mtr

- <https://www.redhat.com/sysadmin/linux-mtr-command>
- <https://linux.die.net/man/8/mtr>
- <https://www.bitwizard.nl/mtr/>
- <https://github.com/traviscross/mtr>

mtr combines the functionality of the 'traceroute' and 'ping' programs in a single network diagnostic tool. As mtr starts, it investigates the network connection between the host mtr runs on and a user-specified destination host. After it determines the address of each network hop between the machines, it sends a sequence ICMP ECHO requests to each one to determine the quality of the link to each machine. As it does this, it prints running statistics about each machine. For a preview take a look at the screenshots.

## tcpdump

- <https://www.tcpdump.org/index.html>
- <https://www.tcpdump.org/index.html#documentation>
- <http://www.alexonlinux.com/tcpdump-for-dummies>
- <https://danielmiessler.com/study/tcpdump/>
- <https://packetlife.net/media/library/12/tcpdump.pdf>
- <https://blog.wains.be/2007/2007-10-01-tcpdump-advanced-filters/>

tcpdump is a data-network packet analyzer computer program that runs under a command line interface. It allows the user to display TCP/IP and other packets being transmitted or received over a network to which the computer is attached.[3] Distributed under the BSD license,[4] tcpdump is free software. Tcpdump works on most Unix-like operating systems: Linux, Solaris, FreeBSD, DragonFly BSD, NetBSD, OpenBSD, OpenWrt, macOS, HP-UX 11i, and AIX. In those systems, tcpdump uses the libpcap library to capture packets. The port of tcpdump for Windows is called WinDump; it uses WinPcap, the Windows version of libpcap.

## Justniffer

- <https://onotelli.github.io/justniffer/>
- <https://github.com/onotelli/justniffer>

Justniffer - Network TCP Packet Sniffer. Justniffer is a network protocol analyzer that captures network traffic and produces logs in a customized way, can emulate Apache web server log files, track response times and extract all "intercepted" files from the HTTP traffic. It lets you interactively trace tcp traffic from a live network or from a previously saved capture file. Justniffer's native capture file format is libpcap format, which is also the format used by tcpdump and various other tools.

Reliable TCP Flow Rebuilding - The main Justniffer's feature is the ability to handle all those complex low level protocol issues and retrieve the correct flow of the TCP/IP traffic: IP fragmentation, TCP retransmission, reordering. etc. It uses portions of Linux kernel source code for handling all TCP/IP stuff. Precisely, it uses a slightly modified version of the libnids libraries that already include a modified version of Linux code in a more reusable way. Optimized for "Request / Response" protocols. It is able to track server response time

Justniffer was born as tool for helping in analyzing performance problem in complex network environment when it becomes impractical to analyze network captures solely using wireshark. It will help you to quickly identify the most significant bottlenecks analyzing the performance at "application" protocol level.

In very complex and distributed systems is often useful to understand how communication takes place between different components, and when this is implemented as a network protocol based on TCP/IP (HTTP, JDBC, RTSP, SIP, SMTP, IMAP, POP, LDAP, REST, XML-RPC, IIOP, SOAP, etc.), justniffer becomes very useful. Often the logging level and monitoring systems of these systems does not report important information to determine performance issues such as the response time of each network request. Because they are in a "production" environment and cannot be too much verbose or they are in-house developed applications and do not provide such logging.

Other times it is desirable to collect access logs from web services implemented on different environments (various web servers, application servers, python web frameworks, etc.) or web services that are not accessible and therefore traceable only on client side.

Justniffer can capture traffic in promiscuous mode so it can be installed on dedicated and independent station within the same network "collision domain" of the gateway of the systems that must be analyzed, collecting all traffic without affecting the system performances and requiring invasive installation of new software in production environments.

## Server Density

- <https://support.serverdensity.com/hc/en-us/articles/360001066243>

Server Density: SaaS Monitoring - Proactive infrastructure monitoring for cloud, servers, containers & websites. Cross platform monitoring for Linux, Windows, Docker, Kubernetes, FreeBSD and Mac. Automatic agent installation using our API and integration with Chef, Puppet, Ansible and SaltStack.

## OpenNMS

- <https://www.opennms.com/>



Serious remote code execution (RCE) and denial of service (DOS) vulnerabilities in Apache Log4j could affect customers running some OpenNMS products.

OpenNMS is a free and open-source enterprise grade network monitoring and network management platform. It is developed and supported by a community of users and developers and by the OpenNMS Group, offering commercial services, training and support. The goal is for OpenNMS to be a truly distributed, scalable management application platform for all aspects of the FCAPS network management model while remaining 100% free and open source. Currently the focus is on Fault and Performance Management. All code associated with the project is available under the Affero General Public License. The OpenNMS Project is maintained by The Order of the Green Polo.

## SysUsage

- <https://github.com/darold/sysusage>
- <https://sysusage.darold.net/>

SysUsage: the sysstat and sar grapher. SysUsage is a tool used to continuously monitor a system and generate daily/weekly/monthly/yearly graphical report using rrdtool and sar. SysUsage generate graphical reports on all system activity information. His periodical reports allow you to keep track of the machine activity during his life and will be a great help for performance analysis and resources management. SysUsage can be run periodically from 10 seconds cycle in daemon mode to 1 minute or more using crond. SysUsage can be run from a central server to call a ssh remote execution of the sysusage perl script so that collected data will be stored in this central place. You also will have just one place where rrdtool and related Perl modules need to be installed as well as just one place where sysusagegraph or sysusagejqgraph need to be executed.

SysUsage continuously monitor your systems informations and generate periodical graph reports using rrdtool or javascript jqplot library. All reports are shown throught a web interface.

SysUsage grabs all system activities using Sar and system commands allowing you to keep tracks of your computer or server activity during his life. It is a great help for performance analysis and resources management. The threshold notification can alarm you when the system capabilities are reached by sending SMTP messages or throught Nagios reports.

By default it will monitor all you need to know on your server activity (See Features), it is written in Perl and should works on all Unix like plateforms. It doesn't require a Database system like MySQL or PostgreSQL but lie on rrdtool. In addition you can embeded your own plugins written in any proگرامing language.

Since release 5.0 SysUsage can be run from a centralized place where collected statistics will be stored and where graphics will be rendered. Unless other monitoring tools with lot of administration work, SysUsage is design to have the lesspossible things to configure and a high level of admin system knowledge. Each server can also be self monitored and you just have to connect your browser to the web interface to know his health level.

SysUsage is design with simplicity in mind. I want all relevant statistics from my servers within an intuitive web interface and without spending too much time to configure it, if you know Nagios, you know what I mean. You will especially like SysUsage for that.

## PCP

- <http://www.pcp.io/>
- <https://github.com/performancecopilot/pcp>

PCP - Performance Co-Pilot. Performance Co-Pilot is a system performance analysis toolkit. Lightweight : Collect performance metrics from your systems efficiently. Distributed : Collate metrics from multiple hosts and a variety of operating systems. Included : Everything you need is already included in the major distributions: Fedora, RHEL, Debian, SUSE, Ubuntu, Gentoo. Analyze systems' performance metrics in real-time or using historical data. Compare performance metrics between different hosts and different intervals. Observe trends and identify abnormal patterns. Extend the collected performance metrics in a simple way. PCP offers a multitude of APIs and libraries to extract and make use of performance metrics from your own application.

## KDE system guard

- <https://apps.kde.org/ksysguard/>

System Guard allows you to monitor information and statistics about your system. In addition to monitoring the local system, it can connect to remote systems running the System Guard Daemon.

## Munin

- <https://munin-monitoring.org/>
- <https://guide.munin-monitoring.org/en/latest/>

Munin is a networked resource monitoring tool that can help analyze resource trends and “what just happened to kill our performance?” problems. It is designed to be very plug and play. A default installation provides a lot of graphs with almost no work. In Norse mythology Hugin and Munin are the ravens of the god king Odin. They flew all over Midgard for him, seeing and remembering, and later telling him. “Munin” means “memory”.

## Nagios

- <https://www.nagios.org/>
- <https://github.com/NagiosEnterprises/nagioscore>

Nagios Core /'nɑ:gi:ɔ:ʊs/, formerly known as Nagios, is a free and open-source computer-software application that monitors systems, networks and infrastructure. Nagios offers monitoring and alerting services for servers, switches, applications and services. It alerts users when things go wrong and alerts them a second time when the problem has been resolved.

Ethan Galstad and a group of developers originally wrote Nagios as NetSaint. As of 2015 they actively maintain both the official and unofficial plugins. Nagios is a recursive acronym: “Nagios Ain't Gonna Insist On Sainthood” - “sainthood” makes reference to the original name NetSaint, which changed in response to a legal challenge by owners of a similar trademark. “Agios” (or “hagios”) also transliterates the Greek word ἅγιος, which means “saint”.

Nagios was originally designed to run under Linux, but it also runs on other Unix variants. It is free software licensed under the terms of the GNU General Public License version 2 as published by the Free Software Foundation. Network Monitoring: When it comes to open source network monitoring tools, the World's largest organizations turn to Nagios. Nagios monitors the network for problems caused by overloaded data links or network connections, as well as monitoring routers, switches and more. Easily able to monitor availability, uptime and response time of every node on the network, Nagios can deliver the results in a variety of visual representations and reports. Network Monitoring Software. Network Traffic Monitoring. Network Analyzer.

Nagios is known for being the best server monitoring software on the market. Server monitoring is made easy in Nagios because of the flexibility to monitor your servers with both agent-based and agentless monitoring. With over 5000 different addons available to monitor your servers, the community at the Nagios Exchange have left no stone unturned. Server Monitoring Software. Windows Server Monitoring. Linux Server Monitoring.

Implementing effective application monitoring with Nagios allows your organization to quickly detect application, service, or process problems, and take action to eliminate downtime for your application users. Nagios provides tools for monitoring of applications and application state - including Windows applications, Linux applications, UNIX applications, and Web applications. Application Monitoring Tools. Web Application Monitoring. Application Log Monitoring.



## Icinga

- <https://github.com/Icinga/icinga2>
- <https://icinga.com/docs/icinga-2/latest/doc/02-installation/>
- <https://icinga.com/>

Icinga is a monitoring system which checks the availability of your network resources, notifies users of outages, and generates performance data for reporting. Scalable and extensible, Icinga can monitor large, complex environments across multiple locations. Icinga 2 is the monitoring server and requires Icinga Web 2 on top in your Icinga Stack. The configuration can be easily managed with either the Icinga Director, config management tools or plain text within the Icinga DSL.

## Zenoss

- <https://community.zenoss.com/>

Zenoss Community Edition is a free and open-source application, server, and network management platform based on the Zope application server. Released under the GNU General Public License version 2, Zenoss Community Edition provides a web interface that allows system administrators to monitor availability, inventory/configuration, performance, and events. Zenoss Community Edition is by Zenoss Inc., which was founded in 2005 and is headquartered in Austin, Texas. The company develops hybrid IT monitoring and analytics software.

## Cacti

- <https://www.cacti.net/>
- <https://github.com/Cacti/cacti>
- <https://docs.cacti.net/>

Cacti is an open-source, web-based network monitoring and graphing tool designed as a front-end application for the open-source, industry-standard data logging tool RRDtool. Cacti allows a user to poll services at predetermined intervals and graph the resulting data. It is generally used to graph time-series data of metrics such as CPU load and network bandwidth utilization. A common usage is to monitor network traffic by polling a network switch or router interface via Simple Network Management Protocol (SNMP).

The front end can handle multiple users, each with their own graph sets, so it is sometimes used by web hosting providers (especially dedicated server, virtual private server, and colocation providers) to display bandwidth statistics for their customers. It can be used to configure the data collection itself, allowing certain setups to be monitored without any manual configuration of RRDtool. Cacti can be extended to monitor any source via shell scripts and executables.

Cacti can use one of two back ends: "cmd.php", a PHP script suitable for smaller installations, or "Spine" (formerly Cactid), a C-based poller which can scale to thousands of hosts.

## Zabbix

- <https://git.zabbix.com/projects/zbx/repos/zabbix/browse>
- <https://www.zabbix.com/>

Zabbix is an open-source monitoring software tool for diverse IT components, including networks, servers, virtual machines (VMs) and cloud services. Zabbix provides monitoring metrics, among others network utilization, CPU load and disk space consumption. Zabbix monitoring configuration can be done using XML based templates which contain elements to monitor.[3] The software monitors operations[vague] on Linux, Hewlett Packard Unix (HP-UX), Mac OS X, Solaris and other operating systems (OSes); however, Windows monitoring is only possible through agents. Zabbix can use MySQL, MariaDB, PostgreSQL, SQLite, Oracle or IBM DB2 to store data.[4]

Backend of Zabbix is written in C and the web frontend is written in PHP.

Zabbix offers several monitoring features:

- Simple checks can verify the availability and responsiveness of standard services such as SMTP or HTTP without installing any software on the monitored host.
- A Zabbix agent can also be installed on UNIX and Windows hosts to monitor statistics such as CPU load, network utilization, disk space, etc.
- As an alternative to installing an agent on hosts, Zabbix includes support for monitoring via SNMP, TCP and ICMP checks, as well as over IPMI, JMX, SSH, Telnet and using custom parameters. Zabbix supports a variety of near-real-time notification mechanisms, including XMPP.

Released under the terms of GNU General Public License version 2, Zabbix is free software.

## nmon

- <http://nmon.sourceforge.net/pmwiki.php>
- <http://nmon.sourceforge.net/pmwiki.php?n=Site.Documentation>

nmon (short hand for Nigel's Monitor) is a computer performance system monitor tool for the AIX and Linux operating systems. The nmon tool has two modes a) displays the performance stats on-screen in a condensed format or b) the same stats are saved to a comma-separated values (CSV) data file for later graphing and analysis to aid the understanding of computer resource use, tuning options and bottlenecks. nmon for Linux is open source and available under GNU General Public License while the nmon for AIX is a proprietary software integrated into AIX.

## conky

- <https://github.com/brndnmtthws/conky>
- <https://github.com/brndnmtthws/conky/wiki>

Conky is a free software desktop system monitor for the X Window System. It is available for Linux, FreeBSD, and OpenBSD. Conky is highly configurable and is able to monitor many system variables including the status of the CPU, memory, swap space, disk storage, temperatures, processes, network interfaces, battery power, system messages, e-mail inboxes, Arch Linux updates, many popular music players (MPD, XMMS2, BMPx, Audacious, etc.), weather updates, breaking news, and much more. Unlike system monitors that use high-level widget toolkits to render their information, Conky is drawn directly in an X window. This allows it to be configured such that it consumes relatively few system resources.

Conky was hailed as “one of the best maintained, and definitely one of the most useful, programs in the world of open source” in Linux Magazine.

It comes pre-configured in the default installation of Pinguy OS and CrunchBang Linux.

## Glances

- <https://nicolargo.github.io/glances/>
- <https://github.com/nicolargo/glances/wiki>

Glances is a cross-platform monitoring tool which aims to present a large amount of monitoring information through a curses or Web based interface. The information dynamically adapts depending on the size of the user interface. It can also work in client/server mode. Remote monitoring could be done via terminal, Web interface or API (XML-RPC and RESTful). Stats can also be exported to files or external time/value databases. Glances is written in Python and uses libraries to grab information from your system. It is based on an open architecture where developers can add new plugins or exports modules.

## sarg

- <https://sourceforge.net/projects/sarg/>
- <https://linux.die.net/man/1/sarg>

sarg - Squid Analysis Report Generator is a tool that allow you to view “where” your users are going to on the Internet. sarg generate HTML reports, with informations about users, IP Addresses, bytes, sites and times. These HTML files can appear in your web server's directory for browsing by users or administrators. You may also have sarg email the reports to the Squid Cache administrator.

sarg is a logfile parser and analyzer for the Squid Web Proxy Cache, which can be found at <http://www.squid-cache.org/>. More information is available at <http://sarg.sourceforge.net/>. sarg can take the squid access log as its input, or optionally the squidGuard <http://www.squidguard.org/> Squid filter/redirector logfile format.

## libstatgrab / saidar

- <https://libstatgrab.org/>
- <https://libstatgrab.org/#documentation>
- <https://github.com/libstatgrab/libstatgrab>

libstatgrab is a library that provides cross platform access to statistics about the system on which it's run. It's written in C and presents a selection of useful interfaces which can be used to access key system statistics. The current list of statistics includes CPU usage, memory utilisation, disk usage, process counts, network traffic, disk I/O, and more.

The current list of supported and tested platforms includes FreeBSD, Linux, NetBSD, OpenBSD, Solaris, DragonFly BSD, HP-UX and AIX. We are always happy to accept patches to extend support to other operating systems.

The package also includes a couple of useful tools. The first, `saidar`, provides a curses-based interface to viewing the current state of the system. The second, `statgrab`, gives a `sysctl`-style interface to the statistics gathered by `libstatgrab`. This extends the use of `libstatgrab` to people writing scripts or anything else that can't easily make C function calls. Included with `statgrab` is a script to generate an MRTG configuration file to use `statgrab`.

The library part of `libstatgrab` is licensed under the GNU LGPL, but the tools and examples are licensed under the GNU GPL.

## RRDtool

- <https://oss.oetiker.ch/rrdtool/>
- <https://oss.oetiker.ch/rrdtool/doc/index.en.html>

RRDtool is the OpenSource industry standard, high performance data logging and graphing system for time series data. RRDtool can be easily integrated in shell scripts, perl, python, ruby, lua or tcl applications.

## monit / M/Monit

- <https://mmonit.com/monit/>
- <https://mmonit.com/>
- <https://bitbucket.org/tildeslash/monit/src/master/>
- <https://mmonit.com/monit/documentation/monit.html>

Monit is a free, open-source process supervision tool for Unix and Linux. With Monit, system status can be viewed directly from the command line, or via the native HTTP(S) web server. Monit is able to do automatic maintenance, repair, and run meaningful causal actions in error situations.[3] Monit rose to popularity with Ruby on Rails and the Mongrel web server, because a tool was needed that could manage the many identical Mongrel processes that needed to be run to support a scalable Ruby on Rails site, and Monit was fairly uniquely suited for the needs of the Ruby on Rails community. Many popular Rails sites have used Monit, including Twitter and scribd.

Monit can restart a process automatically if process dies or monitor process characteristics, such as memory or cpu cycles and alert by email or execute and action.

Additionally M/Monit can monitor and manage distributed computer systems, M/Monit uses Monit as an agent and can manage and monitor. M/Monit is licensed software.

## Linux process explorer

- <https://github.com/wolfc01/procexp/blob/master/README.md>

Graphical process explorer for Linux. Shows process information: process tree, TCP IP connections and graphical performance figures for processes. Aims to mimic Windows `procxp` from `sysinternals`, and aims to be more usable than `top` and `ps`, especially for advanced users.

## df

- [https://en.wikipedia.org/wiki/Df\\_\(Unix\)](https://en.wikipedia.org/wiki/Df_(Unix))
- <https://pubs.opengroup.org/onlinepubs/9699919799/utilities/df.html>

`df` (abbreviation for disk free) is a standard Unix command used to display the amount of available disk space for file systems on which the invoking user has appropriate read access. `df` is typically implemented using the `statfs` or `statvfs` system calls.

## discus

- <https://manpages.ubuntu.com/manpages/bionic/man1/discus.1.html>

`discus` aims to make `df` prettier. Features include color, bar graphs, and smart formatting of numbers (automatically choosing the most suitable size from kilobytes, megabytes, gigabytes, or terabytes). Or choose your own size, along with specifying the number of decimal places you'd like to see. You may also copy `/etc/discusrc` to `$HOME/.discusrc` and customize things to your preference.

## dstat

- <https://linux.die.net/man/1/dstat>
- <http://dag.wiee.rs/home-made/dstat/>

dstat - versatile tool for generating system resource statistics. Dstat is a versatile replacement for vmstat, iostat and ifstat. Dstat overcomes some of the limitations and adds some extra features. Dstat allows you to view all of your system resources instantly, you can eg. compare disk usage in combination with interrupts from your IDE controller, or compare the network bandwidth numbers directly with the disk throughput (in the same interval). Dstat also cleverly gives you the most detailed information in columns and clearly indicates in what magnitude and unit the output is displayed. Less confusion, less mistakes, more efficient.

Dstat is unique in letting you aggregate block device throughput for a certain diskset or network bandwidth for a group of interfaces, ie. you can see the throughput for all the block devices that make up a single filesystem or storage system.

Dstat allows its data to be directly written to a CSV file to be imported and used by OpenOffice, Gnumeric or Excel to create graphs.

## Net-SNMP

- <http://www.net-snmp.org/>
- <http://www.net-snmp.org/docs/man/>

Simple Network Management Protocol (SNMP) is a widely used protocol for monitoring the health and welfare of network equipment (eg. routers), computer equipment and even devices like UPSs. Net-SNMP is a suite of applications used to implement SNMP v1, SNMP v2c and SNMP v3 using both IPv4 and IPv6. The suite includes command-line applications to:

- retrieve information from an SNMP-capable device, either using single requests (snmpget, snmpgetnext), or multiple requests (snmpwalk, snmptable, snmpdelta).
- manipulate configuration information on an SNMP-capable device (snmpset).
- retrieve a fixed collection of information from an SNMP-capable device (snmpdf, snmpnetstat, snmpstatus).
- convert between numerical and textual forms of MIB OIDs, and display MIB content and structure (snmptranslate).
- A graphical MIB browser (tkmib), using Tk/perl.
- A daemon application for receiving SNMP notifications (snmptrapd). Selected notifications can be logged (to syslog, the NT Event Log, or a plain text file), forwarded to another SNMP management system, or passed to an external application.
- An extensible agent for responding to SNMP queries for management information (snmpd). This includes built-in support for a wide range of MIB information modules, and can be extended using dynamically loaded modules, external scripts and commands, and both the SNMP multiplexing (SMUX) and Agent Extensibility (AgentX) protocols.
- A library for developing new SNMP applications, with both C and perl APIs.

Net-SNMP is available for many Unix and Unix-like operating systems and also for Microsoft Windows. Note: Functionality can vary depending on the operating system. Please see the README files for information specific to your platform.

## incron

- <https://github.com/ar-/incron>
- <http://inotify.aiken.cz/?section=incron&page=about&lang=en>

This program is the “inotify cron” system. It consist of a daemon and a table manipulator. You can use it a similar way as the regular cron. The difference is that the inotify cron handles filesystem events rather than time periods.

This project was kicked off by Lukas Jelinek in 2006 and then unfortunately abandoned in 2012. Upstream development and bug-tracking/fixing continued in 2014 on GitHub: <https://github.com/ar-/incron> .

inotify is an inode-based filesystem notification technology. It provides possibility to simply monitor various events on files in filesystems. It is a very much powerful replacement of (obsolete) dnotify. inotify brings a comfortable way how to manage files used in your applications.

inotify can be used for such tasks:

- detecting changes in files and directories (e.g. configuration files, mail directories)
- guarding critical files and their eventual automatic recovery
- file usage statistics and similar purposes
- automatic upload handling
- monitoring installations outside of packaging systems
- automatic on-change backup and/or versioning
- reflecting changes to search databases

There is much more that can be done using inotify. Switch your fantasy on and incorporate inotify to your application the best way.

## monitorix

- <https://github.com/mikaku/Monitorix>
- <https://www.monitorix.org/>
- <https://www.monitorix.org/documentation.html>
- <https://wiki.archlinux.org/title/monitorix>

Monitorix is a free, open source, lightweight system monitoring tool designed to monitor as many services and system resources as possible. It has been created to be used under production Linux/UNIX servers, but due to its simplicity and small size can be used on embedded devices as well.

It consists mainly of two programs: a collector, called monitorix, which is a Perl daemon that is started automatically like any other system service, and a CGI script called monitorix.cgi. Monitorix includes its own HTTP server built in (which is listening by default on port 8080/TCP) to see the statistics graphs, so you aren't forced to install a third-party web server to use it. Just point your browser at <http://localhost:8080/monitorix>.

All of its development was initially created for monitoring Red Hat, Fedora and CentOS Linux systems, so this project was made keeping in mind these type of distributions. Today it runs on different GNU/Linux distributions and even in other UNIX systems like FreeBSD, OpenBSD and NetBSD.

It is currently in active development adding new features, new graphs and correcting bugs in the attempt to offer a great tool for daily systems administration.

## vmstat

- <https://linux.die.net/man/8/vmstat>
- <https://en.wikipedia.org/wiki/Vmstat>

vmstat (virtual memory statistics) is a computer system monitoring tool that collects and displays summary information about operating system memory, processes, interrupts, paging and block I/O. Users of vmstat can specify a sampling interval which permits observing system activity in near-real time.

The vmstat tool is available on most Unix and Unix-like operating systems, such as FreeBSD, Linux or Solaris.

## Web VMStat

- <https://github.com/joewalnes/web-vmstats>

Display live Linux system stats (memory, CPU, IO, etc) in a pretty web-page, with charts and everything.

## uptime

- <https://man7.org/linux/man-pages/man1/uptime.1.html>
- [https://en.wikipedia.org/wiki/Uptime#Using\\_uptime](https://en.wikipedia.org/wiki/Uptime#Using_uptime)

uptime gives a one line display of the following information. The current time, how long the system has been running, how many users are currently logged on, and the system load averages for the past 1, 5, and 15 minutes.

## mpstat

- <https://linux.die.net/man/1/mpstat>
- <http://sebastien.godard.pagesperso-orange.fr/>

mpstat - Report processors related statistics. The mpstat command writes to standard output activities for each available processor, processor 0 being the first one. Global average activities among all processors are also reported. The mpstat command can be used both on SMP and UP machines, but in the latter, only global average activities will be printed. If no activity has been selected, then the default report is the CPU utilization report.

The interval parameter specifies the amount of time in seconds between each report. A value of 0 (or no parameters at all) indicates that processors statistics are to be reported for the time since system startup (boot). The count parameter can be specified in conjunction with the interval parameter if this one is not set to zero. The value of count determines the number of reports generated at interval seconds apart. If the interval parameter is specified without the count parameter, the mpstat command generates reports continuously.

## pmap

- <https://linux.die.net/man/1/pmap>

pmap - report memory map of a process. The pmap command reports the memory map of a process or processes.

## ps

- <https://man7.org/linux/man-pages/man1/ps.1.html>
- [https://en.wikipedia.org/wiki/Ps\\_\(Unix\)](https://en.wikipedia.org/wiki/Ps_(Unix))

ps displays information about a selection of the active processes. If you want a repetitive update of the selection and the displayed information, use top instead. In most Unix and Unix-like operating systems, the ps program (short for “process status”) displays the currently-running processes. A related Unix utility named top provides a real-time view of the running processes.

## sar

- <https://linux.die.net/man/1/sar>
- <http://sebastien.godard.pagesperso-orange.fr/>
- [http://sebastien.godard.pagesperso-orange.fr/man\\_sar.html](http://sebastien.godard.pagesperso-orange.fr/man_sar.html)

sar - Collect, report, or save system activity information. The sar command writes to standard output the contents of selected cumulative activity counters in the operating system. The accounting system, based on the values in the count and interval parameters, writes information the specified number of times spaced at the specified intervals in seconds. If the interval parameter is set to zero, the sar command displays the average statistics for the time since the system was started. If the interval parameter is specified without the count parameter, then reports are generated continuously. The collected data can also be saved in the file specified by the -o filename flag, in addition to being displayed onto the screen. If filename is omitted, sar uses the standard system activity daily data file, the /var/log/sa/sadd file, where the dd parameter indicates the current day. By default all the data available from the kernel are saved in the data file.

## Sysstat

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

The sysstat utilities are a collection of performance monitoring tools for Linux. These include sar, sady, mpstat, iostat, tapestat, pidstat, cifsostat and sa tools. Go to the Features page to display a list of sysstat's features, or see the Documentation page to learn some more about them. Sysstat is Open Source / Free Software, and is freely available under the GNU General Public License, version 2.

Sysstat's main features:

- Includes four groups of monitoring tools (sar / sady / sady, iostat / tapestat / cifsostat, mpstat, pidstat) for global system performance analysis.
- Can monitor a huge number of different metrics:
  - Input / Output and transfer rate statistics (global, per device, per partition, per network filesystem and per Linux task / PID).
  - CPU statistics (global, per CPU, per NUMA nodes and per Linux task / PID), including support for virtualization architectures.
  - Memory, hugepages and swap space utilization statistics.
  - Virtual memory, paging and fault statistics.
  - Per-task (per-PID) memory and page fault statistics.
  - Global CPU and page fault statistics for tasks and all their children.
  - Process creation activity.
  - Interrupt statistics (global, per CPU and per interrupt, including potential APIC interrupt sources, hardware and software interrupts).
  - Extensive network statistics: network interface activity (number of packets and kB received and transmitted per second, etc.) including failures from network devices; network traffic statistics for IP, TCP, ICMP and UDP protocols based on SNMPv2 standards; support for IPv6-related protocols.
  - Fibre Channel traffic statistics.
  - Software-based network processing (softnet) statistics.
  - NFS server and client activity.
  - Sockets statistics.
  - Run queue and system load statistics.
  - Kernel internal tables utilization statistics.
  - System and per Linux task switching activity.
  - Swapping statistics.
  - TTY devices activity.

- Power management statistics (instantaneous and average CPU clock frequency, fans speed, devices temperature, voltage inputs)
- USB devices plugged into the system.
- Filesystems utilization (inodes and blocks).
- Tape drives statistics.
- Pressure-Stall Information statistics.
- Can generate graphs (SVG format - Scalable Vector Graphics) that can be displayed in your favorite web browser!
- Most system statistics can be saved in a file for future inspection.
- Allows to configure the length of data history to keep.
- On the fly detection of new devices (disks, network interfaces, etc.) that are created or registered dynamically.
- Support for UP and SMP machines, including machines with hyperthreaded or multi-core processors.
- Support for hotplug CPUs (it detects automatically processors that are disabled or enabled on the fly) and tickless CPUs.
- Works on many different architectures, whether 32- or 64-bit.
- Needs very little CPU time to run (written in C).
- System statistics collected by sar/sadc can be exported in various different formats (CSV, XML, JSON, SVG, etc.). DTD and XML Schema documents are included in sysstat package. JSON output format is also available for mpstat and iostat commands.
- sar data can also be exported by sadf to PCP (Performance Co-Pilot) archive.
- Smart color output for easier statistics reading.
- Internationalization support (sysstat has been translated into numerous different languages). Sysstat is now part of the Translation Project.
- Sysstat commands can automatically select the unit used to display sizes for easier reading (see option -human).
- Many programs available on the internet to use sysstat's data to make graphs (one of them, isag, is included in sysstat).
- iostat has support for devices managed by drivers in userspace like spdk.

## iostat

- <http://sebastien.godard.pagesperso-orange.fr/>
- [http://sebastien.godard.pagesperso-orange.fr/man\\_iostat.html](http://sebastien.godard.pagesperso-orange.fr/man_iostat.html)

iostat is also part of the sysstat package. This command is used for monitoring system input/output. The reports themselves can be used to change system configurations to better balance input/output load between hard drives in your machine.

## collectl

- <https://linux.die.net/man/1/collectl>
- <http://collectl.sourceforge.net/>
- <http://collectl.sourceforge.net/Documentation.html>

There are a number of times in which you find yourself needing performance data. These can include benchmarking, monitoring a system's general health or trying to determine what your system was doing at some time in the past. Sometimes you just want to know what the system is doing right now. Depending on what you're doing, you often end up using different tools, each designed to for that specific situation.

Unlike most monitoring tools that either focus on a small set of statistics, format their output in only one way, run either interactively or as a daemon but not both, collectl tries to do it all. You can choose to monitor any of a broad set of subsystems which currently include buddyinfo, cpu, disk, inodes, infiniband, lustre, memory, network, nfs, processes, quadrics, slabs, sockets and tcp.

## free

- <https://man7.org/linux/man-pages/man1/free.1.html>

free displays the total amount of free and used physical and swap memory in the system, as well as the buffers and caches used by the kernel. The information is gathered by parsing /proc/meminfo.

## /proc file system

- <https://www.kernel.org/doc/html/latest/filesystems/proc.html>
- <https://tldp.org/LDP/Linux-Filesystem-Hierarchy/html/proc.html>

/proc is very special in that it is also a virtual filesystem. It's sometimes referred to as a process information pseudo-file system. It doesn't contain 'real' files but runtime system information (e.g. system memory, devices mounted, hardware configuration, etc). For this reason it can be regarded as a control and information centre for the kernel. In fact, quite a lot of system utilities are simply calls to files in this directory. For example, 'lsmod' is the same as 'cat /proc/modules' while 'lspci' is a synonym for 'cat /proc/pci'. By altering files located in this directory you can even read/change kernel parameters (sysctl) while the system is running.

## GKrellM

- <http://gkrellm.srcbox.net/>
- <https://git.srcbox.net/gkrellm/gkrellm>

GKrellM is a single process stack of system monitors which supports applying themes to match its appearance to your window manager, Gtk, or any other theme.

## Gnome system monitor

- <https://wiki.gnome.org/Apps/SystemMonitor>
- <https://help.gnome.org/users/gnome-system-monitor/stable/>

System Monitor is a tool to manage running processes and monitor system resources.

## GoAccess

- <https://github.com/allinurl/goaccess>
- <https://goaccess.io/>
- <https://goaccess.io/man>

GoAccess is an open-source web analytics application for Unix-like operating systems. It has both a text-based and a web application user interface. It can provide real-time analytics by continuously monitoring web server logs. GoAccess is an open source real-time web log analyzer and interactive viewer that runs in a terminal in \*nix systems or through your browser.

It provides fast and valuable HTTP statistics for system administrators that require a visual server report on the fly.

## Logwatch

- <https://linux.die.net/man/8/logwatch>
- <https://sourceforge.net/projects/logwatch/>

Logwatch is a customizable, pluggable log-monitoring system. It will go through your logs for a given period of time and make a report in the areas that you wish with the detail that you wish. Logwatch is being used for Linux and many types of UNIX.

## Swatchdog

- <http://sourceforge.net/projects/swatch/>
- <http://swatch.sourceforge.net/>

Swatch: the active log file monitoring tool. Swatch started out as the "simple watchdog" for actively monitoring log files produced by UNIX's syslog facility. It has since been evolving into a utility that can monitor just about any type of log.

## MultiTail

- <https://github.com/halturin/multitail>
- <http://www.vanheusden.com/multitail/>
- <https://www.vanheusden.com/multitail/examples.html>
- <https://linux.die.net/man/1/multitail>

MultiTail allows you to monitor logfiles and command output in multiple windows in a terminal, colorize, filter, and merge.

acct or psacct - Monitor User Activity

acct or psacct (depending on if you use apt-get or yum) allows you to monitor all the commands a users executes inside the system including CPU and memory time. Once installed you get that summary with the command 'sa'. \* <http://www.gnu.org/software/acct/>

psacct or acct tools are very useful for monitoring each users activity on the system. Both daemons runs in the background and keeps a close watch on the overall activity of each user on the system and also what resources are being consumed by them.

These tools are very useful for system administrators to track each users activity like what they are doing, what commands they issued, how



much resources are used by them, how long they are active on the system etc.

For installation and example usage of commands read the article on Monitor User Activity with psacct or acct

#### whowatch

Similar to acct this tool monitors users on your system and allows you to see in real time what commands and processes they are using. It gives you a tree structure of all the processes and so you can see exactly what's happening. \* <http://whowatch.sourceforge.net/>

#### strace

~\ownCloud\zim\Bilder\DEF\artikel-linux-monitoring-tools\strace.jpg

strace is used to diagnose, debug and monitor interactions between processes. The most common thing to do is making strace print a list of system calls made by the program which is useful if the program does not behave as expected. \* <http://sourceforge.net/projects/strace/>

#### DTrace

~\ownCloud\zim\Bilder\DEF\artikel-linux-monitoring-tools\dtrace.jpg

DTrace is the big brother of strace. It dynamically patches live running instructions with instrumentation code. This allows you to do in-depth performance analysis and troubleshooting. However, it's not for the weak of heart as there is a 1200 book written on the topic. \* <http://dtrace.org/blogs/about/>

#### webmin

~\ownCloud\zim\Bilder\DEF\artikel-linux-monitoring-tools\webmin.jpg

Webmin is a web-based system administration tool. It removes the need to manually edit unix configuration files and lets you manage the system remotely if need be. It has a couple of monitoring modules that you can attach to it. \* <http://www.webmin.com/>

#### stat

~\ownCloud\zim\Bilder\DEF\artikel-linux-monitoring-tools\stat.jpg

Stat is a built-in tool for displaying status information of files and file systems. It will give you information such as when the file was modified, accessed or changed.

#### ifconfig

~\ownCloud\zim\Bilder\DEF\artikel-linux-monitoring-tools\ifconfig.jpg

ifconfig is a built-in tool used to configure the network interfaces. Behind the scenes network monitor tools use ifconfig to set it into promiscuous mode to capture all packets. You can do it yourself with `ifconfig eth0 promisc` and return to normal mode with ``ifconfig eth0 -promisc``.

#### ulimit

~\ownCloud\zim\Bilder\DEF\artikel-linux-monitoring-tools\unlimit.jpg

ulimit is a built-in tool that monitors system resources and keeps a limit so any of the monitored resources don't go overboard. For instance making a fork bomb where a properly configured ulimit is in place would be totally fine. \* <http://ss64.com/bash/ulimit.html>

#### cpulimit

CPUlimit is a small tool that monitors and then limits the CPU usage of a process. It's particularly useful to make batch jobs not eat up too many CPU cycles. \* <https://github.com/opsengine/cpulimit>

#### lshw

~\ownCloud\zim\Bilder\DEF\artikel-linux-monitoring-tools\lshw.jpg

lshw is a small built-in tool extract detailed information about the hardware configuration of the machine. It can output everything from CPU version and speed to mainboard configuration. \* <https://github.com/opsengine/cpulimit>

#### w

W is a built-in command that displays information about the users currently using the machine and their processes.

#### lsof

Lsof command used in many Linux/Unix like system that is used to display list of all the open files and the processes. The open files included are disk files, network sockets, pipes, devices and processes. One of the main reason for using this command is when a disk cannot be unmounted and displays the error that files are being used or opened. With this command you can easily identify which files are in use. The most common format for this command is.

~\ownCloud\zim\Bilder\DEF\artikel-linux-monitoring-tools\lsof.jpg

lsof is a built-in tool that gives you a list of all open files and network connections. From there you can narrow it down to files opened by processes, based on the process name, by a specific user or perhaps kill all processes that belongs to a specific user.

More lsof command usage and examples : [10 lsof Command Examples in Linux](#)

collectd

Collectd is a Unix daemon that collects all your monitoring statistics. It uses a modular design and plugins to fill in any niche monitoring. This way collectd stays as lightweight and customizable as possible. \* <https://collectd.org/>

Observium

Observium is also a network monitoring tool, it was designed to help you manage your network of servers easily, there are 2 versions from it; Community Edition which is free & open-source and Commercial version which costs £150/year.

~\ownCloud\zim\Bilder\DEF\artikel-linux-monitoring-tools\observium.png

~\ownCloud\zim\Bilder\DEF\artikel-linux-monitoring-tools\Observium2.jpeg

Observium is an auto-discovering network monitoring platform supporting a wide range of hardware platforms and operating systems. Observium focuses on providing a beautiful and powerful yet simple and intuitive interface to the health and status of your network. \* <http://www.observium.org/>

Features of Observium \* Written in PHP with MySQL database support. \* Has a nice web interface to output information and data. \* Ability to manage and monitor hundreds of hosts worldwide. \* The community version from it is licensed under QPL license. \* Works on Windows, Linux, FreeBSD and more.

Read More: [Observium - Network Management and Monitoring Tool for RHEL/CentOS](#)

PHP Server Monitoring

Unlike the other tools on this list, PHP Server Monitoring is a web script written in PHP that helps you to manage you websites and hosts easily, it supports MySQL database and is released under GPL 3 or later. PHP Server Monitor

~\ownCloud\zim\Bilder\DEF\artikel-linux-monitoring-tools\PHP-Server-Monitor.png

Features \* A nice web interface. \* Ability to send notifications to you via Email & SMS. \* Ability to view the most important information about CPU and RAM. \* A very modern logging system to log connection errors and emails that are sent. \* Support for cronjob services to help you monitor your servers and websites automatically.

Read More: [Install PHP Server Monitoring Tool in Arch Linux](#)

nload

It's a command line tool that monitors network throughput. It's neat because it visualizes the in and and outgoing traffic using two graphs and some additional useful data like total amount of transferred data. You can install it with

```
yum install nload
```

or

```
sudo apt-get install nload
```

SmokePing

SmokePing keeps track of the network latencies of your network and it visualises them too. There are a wide range of latency measurement plugins developed for SmokePing. If a GUI is important to you it's there is an ongoing development to make that happen. \* <http://oss.oetiker.ch/smokeping/>

Shinken monitoring

~\ownCloud\zim\Bilder\DEF\artikel-linux-monitoring-tools\shinken.jpg

Shinken is a monitoring framework which is a total rewrite of Nagios in python. It aims to enhance flexibility and managing a large environment. While still keeping all your nagios configuration and plugins. \* <http://www.shinken-monitoring.org/>

Linux Dash - Linux Server Performance Monitoring

From its name, "Linux Dash" is a web dashboard that shows you the most important information about your Linux systems such as RAM, CPU, file-system, running processes, users, bandwidth usage in real time, it has a nice GUI and it's free & open-source.

~\ownCloud\zim\Bilder\DEF\artikel-linux-monitoring-tools\linux-dash.jpeg

Read More: [Install Linux Dash \(Linux Performance Monitoring\) Tool in Linux](#)

Wireshark - Network Protocol Analyzer

Also, unlike all the other tools on our list, Wireshark is an analyzer desktop program which is used to analyze network packets and to monitor network connections. It's written in C with the GTK+ library and released under the GPL license. [Wireshark Network Analyzer](#)

~\ownCloud\zim\Bilder\DEF\artikel-linux-monitoring-tools\wireshark.jpg

Features \* Cross-platform: it works on Linux, BSD, Mac OS X and Windows. \* Command line support: there's a command line based version from Wireshark to analyze data. \* Ability to capture VoIP calls, USB traffic, network data easily to analyze it. \* Available in most Linux distributions repositories.

Read More: [Install Wireshark - Network Protocol Analyzer Tool in Linux](#)

Arpwatch - Ethernet Activity Monitor

Arpwatch is a kind of program that is designed to monitor Address Resolution (MAC and IP address changes) of Ethernet network traffic on a Linux network. It continuously keeps watch on Ethernet traffic and produces a log of IP and MAC address pair changes along with a timestamps on a network. It also has a feature to send an email alerts to administrator, when a pairing added or changes. It is very useful in detecting ARP spoofing on a network.

Read More : [Arpwatch to Monitor Ethernet Activity](#)

Suricata - Network Security Monitoring

Suricata is an high performance open source Network Security and Intrusion Detection and Prevention Monitoring System for Linux, FreeBSD and Windows. It was designed and owned by a non-profit foundation OISF (Open Information Security Foundation).

Read More : [Suricata - A Network Intrusion Detection and Prevention System](#)

---

~~DISCUSSION~~

From:

<https://wiki.nanoscopic.de/> - **nanoscopic wiki**

Permanent link:

<https://wiki.nanoscopic.de/doku.php/pages/howtos/diagnose/linux-diagnosis-and-monitoring-tools>

Last update: **2023/06/15 16:42**

