# Easy way to look for irregularities and problems in log files / Facilitate reading log files on GNU / Linux and FreeBSD

**Author :** admin


As a System Administrator I need to check daily the log files produced on various GNU / Linux distributions or FreeBSD. This can sometimes take too much time if the old fashioned way using the normal system tools **cat, less and tail etc.** is used.

Reading logs one by one eats too much of my time and often as logs are reviewed in a hurry some crucial system irregularities, failed ssh or POP3 / Imap logins, filling disk spaces etc. are missed.

Therefore I decided to implement **automated log parsing programs** which will summary and give me an overview (helicopter view) on what were the system activities from the previous day (24h) until the moment I logged the system and issued the log analyzer program.
There are plenty of programs available out there that does **"wide scale" log analysis**, however there are two applications which on most GNU / Linux and BSD systems had become a de-facto standard programs to **scan system log files for interesting lines**.

These are:

- **1. logwatch** - *system log analyzer and reporter*
- **2. logcheck** - *program to scan system log files for interesting lines*

**1. logwatch** is by default installed on most of the Redhat based Linux systems (Fedora, RHEL, CentOS etc.). On Debian distributions and as far as I know (Ubuntu) and the other deb based distros *logwatch* is not installed by default. Most of the servers I manage these days are running Debian GNU / Linux so, to use *logwatch* I needed to install it from the available repository package, e.g.:

```
debian:~# apt-get install logwatch
...
```

**logwatch** is written in perl and with some big files to analyze, parsing them might take hell a lot of time.

It does use a bunch of configuration scripts which defines how *logwatch* should read and parse the various services logwatch support by default. These conf scripts are also easily extensible, so if one has to analyze some undefined service in the conf files he can easily come up with a new conf script that will support the service/daemon of choice.Using logwatch is very easy, to get an overview about server system activity invoke the *logwatch* command:

```
debian:~# logwatch
 ##################### Logwatch 7.3.6+cvs20080702-debian (07/02/08) ####################
Processing Initiated: Thu Nov 24 05:22:07 2011
Date Range Processed: yesterday
( 2011-Nov-23 )
Period is day.
Detail Level of Output: 0
Type of Output/Format: stdout / text
Logfiles for Host: debian
 ##################################################

---------- dpkg status changes Begin ----------

  Upgraded:
libfreetype6 2.3.7-2+lenny7 => 2.3.7-2+lenny8
libfreetype6-dev 2.3.7-2+lenny7 => 2.3.7-2+lenny8

---------- dpkg status changes End ----------

  ---------- httpd Begin ----------

  Requests with error response codes
400 Bad Request
HTTP/1.1: 2 Time(s)
admin/scripts/setup.php: 2 Time(s)
401 Unauthorized
...
...
---------- vpopmail End ----------

  ---------- Disk Space Begin ----------

Filesystem Size Used Avail Use% Mounted on
/dev/md0 222G 58G 154G 28% /

  ---------- Disk Space End ----------

  ########## Logwatch End ##########
```

The execution might take up from 10 to 20 seconds up to 10 or 20 minutes depending on the log files size and the CPU / RAM hardware on the machine where **/var/log/...** logs will be analyzed.

*logwatch* output can be easily mailed to a custom mail address using a crontab if the server runs a properly configured SMTP server. Using a cron like:

00 5 * * * /usr/sbin/logwatch | mail -s "$(hostname) log files for $(date)"

Here is time to make a note that **logwatch** is ported also to FreeBSD and is available from BSD's port tree, from a port with path:

**/usr/ports/security/logcheck**

**2. logcheck** is another handy program, which does very similar job to *logwatch* . The "interesting" information it returns is a bit less than compared to *logwatch*

The good thing about **logcheck** is that by default it is made to mail every 1 hour a brief data summary which might be of an interest to the sys admin.
*Logcheck* is available for install on RedHat distros via **yum** and has existing package for Debian as well as a port for FreeBSD under the port location **/usr/ports/security/logcheck**

To **install on logcheck on Debian**:

debian:~# apt-get install logcheck
...

After installation I found it wise to change the default mailing time from each and every hour to just once per day to prevent my email from overfilling with "useless" mails.

This is done by editting the default cron tab installed by the package located in **/etc/cron.d/logcheck**

The default file looks like so:

# /etc/cron.d/logcheck: crontab entries for the logcheck package
PATH=/usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin
MAILTO=root
@reboot logcheck if [ -x /usr/sbin/logcheck ]; then nice -n10 /usr/sbin/logcheck -R; fi
2 * * * * logcheck if [ -x /usr/sbin/logcheck ]; then nice -n10 /usr/sbin/logcheck; fi
# EOF

To change it run only once per day its content should looks something like:

# /etc/cron.d/logcheck: crontab entries for the logcheck package

```
PATH=/usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin
MAILTO=root
@reboot logcheck if [ -x /usr/sbin/logcheck ]; then nice -n10 /usr/sbin/logcheck -R; fi
2 5 * * * logcheck if [ -x /usr/sbin/logcheck ]; then nice -n10 /usr/sbin/logcheck; fi
# EOF
```

Altering it that way the log summary interesting info analysis will be sent on mail every day in 05:02 a.m. Changing the default email *logcheck* will ship its log analyzer report emails on deb based distros is done via editting the file:

/etc/logcheck/logcheck.conf

And changing the **SENDMAILTO=""** variable to point to the appropriate admin email email addr.