

Table of Contents

What is awstats.....	1
How to setup awstats.....	1
Setting up awstats for virtual hosting.....	1
Master awstats.....	3
Installation.....	3
Configuration.....	3
Awstats maintenance.....	4
Adding and removing of hosts, log files, etc.....	4
Having separate log files.....	4
Having several databases/config files.....	4
Adding statistics per virtual host only.....	4
XML database format.....	4

What is awstats

Awstats stands for "Advanced web statistics". In dashboard infrastructure it is used as an httpd logs parser.

The official site is at <http://awstats.sourceforge.net/>, the official documentation is at <http://awstats.sourceforge.net/docs/>.

How to setup awstats

The awstats on dashboard machines is installed as a awstats-dashb rpm from dashboardexternals repo:

```
# yum install awstats-dashb
```

There are 3 types of installation, depending on your needs. If your host is a simple httpd without virtualhosting in its configuration, then the package will configure itself upon installation. Thus, you can access the statistics using the following link :

```
http://<your_host_running_httpd>/awstats/awstats.pl?config=<your_host_running_httpd.cern.ch>
```

Note: if you changed log location with CustomLog directive in httpd configuration, you should also change LogFile directive in awstats configuration file at

```
/etc/awstats/awstats.<your_host_running_httpd.cern.ch>.conf
```

Note2: if statistics show "Never updated", run `# /usr/share/awstats/wwwroot/cgi-bin/awstats.pl -update -config=<your_host_running_httpd.cern.ch>` and check that there aren't any errors.

Setting up awstats for virtual hosting

This installation scheme should be used if there's a need to to monitor several or all virtual hosts running on single httpd.

1. To differentiate between different virtual hosts entries in log file, a SiteName of virtual host should be written to them. This is accomplished by inserting "%v" in front of the first parameter of the LogFormat directive. E.g., modify `/etc/httpd/conf/httpd.conf` according to this:

```
LogFormat "%v %h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-Agent}i\"" combinedvhost
CustomLog logs/access_log combinedvhost
```

Be sure to have only one CustomLog directive globally and only one or no CustomLog directives for each virtual host.

You also may remove or rename old httpd log files (`/var/log/httpd/access_log*`), but it is not necessary - awstats will treat old records as "broken" and skip them.

2. Copy `/etc/awstats/awstats.<your_hostname>.conf` to `/etc/awstats/awstats.vhosts.conf`:

```
# cp -p /etc/awstats/awstats.`hostname -s`.cern.ch.conf /etc/awstats/awstats.vhosts.conf
```

3. Modify `/etc/awstats/awstats.vhosts.conf`:

3.1. Change LogFormat so awstats can parse logs with included virtual host hostname:

```
LogFormat = "%virtualname %host %other %logname %time1 %methodurl %code %bytesd %refererquot %uaq
```

3.2. Modify "HostAliases" so it includes all hostnames of virtual hosts. E.g. the following directives may be used if the server hosts three virtual servers with hostnames vhost1, vhost2 and vhost3:

```
# My virtual hosts:
HostAliases="REGEX[^\.*vhost1(\.cern.ch|)\$] REGEX[^\.*vhost2(\.cern.ch|)\$]"
# The line above is too long, thus specifying vhost3 in new line:
HostAliases="REGEX[^\.*vhost3(\.cern.ch|)\$]"
```

3.3. Uncomment and setup the LoadPlugin directive for dashb_apps_vhosts plugin:

```
LoadPlugin="dashb_apps_vhosts vhost1, _:vhost1.cern.ch vhost2, _:vhost2.cern.ch vhost3, _:vhost3.cern.ch"
```

The LoadPlugin directive contains the name of plugin (dashb_apps_vhosts) and its parameters, delimited with spaces. A parameter is defined as:

PARAMETER ::= [VHOSTNAMESLIST ":"] MAPPINGACCESSSTR

the VHOSTNAMESLIST is a list of possible addresses that may be used to access one virtual host instance. For example, if virtual server dashb-atlas-ssb may be accessed with either dashb-atlas-ssb, dashb-atlas-ssb.cern.ch, dashboard13 and dashboard13.cern.ch, all of them should be included in VHOSTNAMESLIST:

```
LoadPlugin="dashb_apps_vhosts dashb-atlas-ssb, _, dashboard13, dashboard13.cern.ch: dashb-atlas-ssb.cern.ch, dashboard13.cern.ch"
```

During log parsing, all hostname specified in VHOSTNAMESLIST is treated as a bunch of aliases for a single virtual host.

The definition of VHOSTNAMESLIST is below:

```
VHOSTNAMESLIST ::= VHOSTNAME [ ", " VHOSTNAME ]*
```

VHOSTNAME is a hostname or underscore ("_"), which is substituted with the hostname deduced from MAPPINGACCESSSTR. If there isn't VHOSTNAMESLIST present, the awstats would insert a "_" in front of MAPPINGACCESSSTR.

MAPPINGACCESSSTR is a string that is either url MAPPINGURL that is used to get mapping via http, or a hostname MAPPINGHOSTNAME of some machine, in which case it is substituted with **http://<==MAPPINGHOSTNAME&==gt;/dashboard/request.py.**

Note: the LoadPlugin directive cannot be broken in several lines, and "/" before EOL does not work as expected.

Note2: much of this can be automated by parsing http -s output. It should be in the ToDo list.

4. Now, test the new configuration interactively with the following command:

```
# /usr/share/awstats/wwwroot/cgi-bin/awstats.pl -update -config=vhosts
```

Check that there aren't any errors. Then, try to access the web view with this URL:

```
http://<your_host_running_httpd>/awstats/awstats.pl?config=vhosts
```

If everything is ok, the awstats with virtual hosting will be updated automatically every hour by cron.

Master awstats

Master awstats is a dashboard service that is built on top of awstats for monitoring multiple "real" hosts. It periodically downloads necessary log files from remote machines and parses them using awstats.pl. Also it automatically modifies awstats .conf files as needed.

Installation

To be described. The awstats-master should be available in dashboard rpm repository.

The awstats-master requires **awstats-dashb** to be installed on the master host, but other hosts may run without awstats-dashb installed.

Configuration

Configuration is partially automated (if compared to awstats with virtual hosting).

1. Create directories and set read and write permissions for needed user:

```
# mkdir -p /opt/dashboard/var/lib/mawstats/  
# mkdir -p /opt/dashboard/var/lib/mawstats/logs  
# mkdir -p /opt/dashboard/var/lib/mawstats/conf  
# mkdir -p /opt/dashboard/var/log/mawstats/
```

2. Copy `/etc/awstats/awstats.model.conf` to

`/opt/dashboard/var/lib/mawstats/conf/awstats.master.conf` and edit it as needed. Be sure that the config contains one and only one line starting with `LoadPlugin="dashb_apps_vhosts`, either commented or uncommented.

3. Add an entry to `service-config.xml`:

```
<service name="awstats.master"  
  module="dashboard.awstats.master.MasterAWStatsService"  
  class="MasterAWStatsService">  
  <config>  
    <param name="interval">3600</param> <!-- default value is 3600 -->  
    <param name="remote.hosts">REMOTE_HOST_1 REMOTE_HOST_2 ...</param> <!-- mandatory -->  
    <param name="awstats.conf.path">/opt/dashboard/var/lib/mawstats/conf/awstats.master.conf</param>  
    <param name="awstats.debug.level">0</param> <!-- default value is 0 -->  
    <!--param name="local.logs.dir">/opt/dashboard/var/lib/mawstats/logs</param-->  
    <!--param name="awstats.db.dir">/opt/dashboard/var/lib/mawstats</param-->  
    <!--param name="remote.user">root</param-->  
    <param name="awstats.stdout.file">/opt/dashboard/var/log/mawstats/awstats.stdout.log</param>  
  </config>  
</service>
```

Note: More on dashboard services and service groups at

<http://dashb-build.cern.ch/build/nightly/doc/guides/common/html/dev/serviceConfigSection.html> .

4. Start service with `# dashb-agent-start <service.group.name> .`

Check dashboard log file and `/opt/dashboard/var/log/mawstats/awstats.stdout.log` for errors.

Awstats maintenance

Unfortunately, sometimes there is a need to interfere with awstats internals. Here are some hints upon it.

Adding and removing of hosts, log files, etc

If you add or remove hosts or log files to/from awstats.XXX.conf or awstats-master, keep in mind that old log entries are going to be skipped and will **not** influence statistics.

But, if you're setting up awstats for the first time for this host, you can delete the database files: all collected statistics are going to be lost, but all log entries will be parsed completely. Also check the log rotation policy: check `/etc/logrotate.d/httpd` and `/etc/logrotate.conf`. With default configuration, all logs for 4 weeks are kept, plus log file for the current week.

The database files are kept in `/var/lib/awstats` by default, and in `/opt/dashboard/var/lib/mawstats` for master awstats. They have names in the following format:

```
==awstats<MONTH><YEAR20xx>.<CONFIG_NAME>.txt==
```

Having separate log files

Sometimes it is needed to have separate log files for several virtual hosts. No problem: they can be merged during the parsing process with `logresolvemerge.pl` utility that comes with awstats. Put this directive in `awstats.conf`:

```
LogFile="/usr/share/awstats/tools/logresolvemerge.pl <LIST OF FILES> |"
```

Keep in mind that using single plain log file is faster as awstats can jump to the place in stopped parsing the last time, while with `logresolvemerge` it creates a pipe and thus have to parse and skip all old log records.

Having several databases/config files

Sometimes you may want to have different awstats configurations to work simultaneously. No problem: copy config file and edit it as you wish. Awstats will use different database files for different configs. Also there's `awstats.model.conf` that is used to create `awstats.<hostname>.conf` upon rpm installation.

The config file name should follow the pattern `/etc/awstats/awstats.<config_name>.conf`. The database will be created with `awstats_updateall.pl` utility that is run by cron with `/etc/cron.hourly/awstats` configuration. The results can be viewed by following the

```
http://<hostname>/awstats/awstats.pl?config=<config_name>
```

Adding statistics per virtual host only

See http://awstats.sourceforge.net/docs/awstats_extra.html#domainaliases.

XML database format

At this moment XML database format is not supported by `dashb_apps_vhosts` plugin.

-- SergeyMitsyn - 03-Dec-2010

This topic: Sandbox > SergeyMitsynSandbox
Topic revision: r1 - 2010-12-03 - SergeyMitsynExCern



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)