

Anemone 1.0 Documentation

Erik Lotspeich <erik@electricc.com>

July 20, 2000

Contents

1	License	2
2	Introduction	2
3	Packing list	3
4	Installation	4
4.1	Automated installation	4
4.2	Manual package installation	4
4.3	Rebuilding the package from source	5
5	Using Anemone	5
5.1	File locations	5
5.1.1	Configuration Files	6
5.1.2	Web content	6
5.1.3	Administration files	6
5.2	Accessing the MRTG graphs	6
5.3	Using the Administration interface	6
5.3.1	Accessing the Administration interface	6
5.3.2	Using the MRTG Status	6
5.3.3	Using the Log Rolling section	7
5.3.4	Changing the Administration interface password	7

6	Technical Features	7
6.1	Installer	8
6.1.1	Squid configuration file changes (/etc/squid/squid.conf)	8
6.1.2	Crontab file changes (/etc/crontab)	8
6.2	Administration interface	8
7	Questions or Comments	8

1 License

Copyright (C) 2000 The Electric Connection (<http://www.electricc.com>)

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.

2 Introduction

Webmasters frequently analyze their web server access logs with logfile-analysis software that generates graphs of server usage, most commonly-accessed URLs, what is the most common web browser, and more. Though logfile-analysis software generates detailed output and a wealth of information regarding activity of the website, it does not provide an instantaneous outlook on the current activity of the website *now*. It is often useful to know how busy a website is at the current moment instead of waiting until the logs are rolled nightly and viewing them the day after.

Anemone provides an easy-to-use interface to the well-known and powerful Multi-Router Traffic Grapher (MRTG) software package (<http://ee-staff.ethz.ch/~oetiker/webtools/mrtg/mrtg.html>). As the name would imply, MRTG was originally intended to graph the traffic of Internet routers in real-time. MRTG uses SNMP to obtain the data from the device in question and uses the innovative

“gd” graphics library (which generates PNG files from a simple scripting language) to accomplish this.

Squid provides SNMP data-reporting capabilities which allows a properly-configured MRTG to collect data from Squid and report it in graph form. Unfortunately the difficulty of setting up Squid, MRTG, gd, and properly-configuring the SNMP can seem a daunting task for a novice and a difficult one for even an expert user.

Anemone simplifies this process dramatically. Anemone works with Tomahawk (the Squid/Apache web-server product) and enables the SNMP capability of Squid, sets up the MRTG configuration to request data from Squid, and properly installs the gd library. Anemone also provides a number of sample graphs to provide a wealth of real-time server information. In addition, Anemone provides an easy-to-use administration interface to control the basic operation of MRTG and to view the graphs that have been created.

3 Packing list

Anemone is distributed as a Unix tar archive compressed with gzip. Virtually every modern Linux system provides a version of GNU tar that makes extracting such an archive trivial. To extract the archive on such as system, type the following command in the directory of the archive:

```
tar zxvf anemone-1.0.tar.gz
```

On systems that do not include GNU tar, the following command will work equally well:

```
gzip -dc anemone-1.0.tar.gz | tar xvf -
```

The archive will create a directory called “anemone-1.0” and place the following files into it.

COPYING The GNU Public License (GPL).

COPYRIGHT A file containing copyright information.

INSTALL A file containing basic installation instructions.

README A file that should be read first.

ChangeLog A file that describes recent changes and version history of the package.

Credits A file that describe contributions by those outside of The Electric Connection.

anemone-documentation.ps Documentation in Adobe PostScript format.

anemone-documentation.pdf Documentation in Adobe PDF format.

anemone-documentation.txt Documentation as a text file.

anemone-1.0-1.i386.rpm The RH Linux 6.x RPM for the Intel x86 architecture.

anemone-1.0-1.src.rpm The SRPM which contains the source code so the package can be built on different architectures.

install.sh An install script that will automate the package installation process on systems of Intel x86 architecture.

4 Installation

The Anemone software is delivered in RPM (RedHat Package Manager) format. When installed, it will configure a working MRTG and GD installation on your workstation. As such, the package has built-in RPM dependencies to protect against installing over an existing MRTG or GD packages. The most common way to resolve these dependencies is to remove your existing MRTG or GD. RPM will leave the existing MRTG configuration files and will rename these existing files when you install Anemone. Your previous configuration should be merged into the configuration files that Anemone will install. This is done to maintain consistency among configuration files and differing versions of software. It also fits with the RPM package model.

4.1 Automated installation

Anemone provides an automated package installation program that will aid in installation on systems of Intel x86 architecture. To install Anemone, simply change to the directory “anemone-1.0” after extracting the tar gzip archive and type “. /install.sh”. Depending on your machine’s configuration and the software you may have installed on your machine, you may be asked a series of questions.

4.2 Manual package installation

Many experienced Linux users are wary of automated installation programs and desire more control over the package installation process. Such a user will typically prefer to manage installations and removals of RPM package files directly. Typically RPM installs use a command similar to the following:

```
rpm -Uvh anemone-1.0-1.i386.rpm
```

Removals of conflicting packages can be accomplished with a command similar to the following:

```
rpm -e gd
rpm -e mrtg
```

It is likely that RPM will not let you remove a package with the above command due to existing dependencies. Since Anemone will install Apache in the same locations as the original RedHat install, a command such as:

```
rpm -e gd --nodeps
rpm -e mrtg --nodeps
```

Followed by a Anemone install will serve the desired purpose.

4.3 Rebuilding the package from source

If you would like to port Anemone to a different architecture, or if you are a developer who would like to see or change how Anemone is put together, you will want to recompile the package from the source RPM (SRPM). Most Linux distributions ship with a directory “/usr/src/redhat” in place.

Copy the source RPM into the appropriate location with the following command:

```
cp anemone-1.0-1.src.rpm /usr/src/redhat/SRPMS
```

Change to the SRPMS directory and execute the following command:

```
rpm --recompile anemone-1.0-1.src.rpm
```

You can also use the following command to accomplish a similar purpose:

```
rpm --rebuild anemone-1.0-1.src.rpm
```

Using the “–rebuild” command is recommended if you are simply rebuilding the package and you do not need to change any of the source code. This is because the “–rebuild” option will remove the source code from your disk after the package is built. The “–recompile” option leaves the source code on your disk so that you can see and change it.

5 Using Anemone

5.1 File locations

This section describes the location of files in the Anemone package. The file locations have been carefully chosen to complement the Anemone organization that is based on the RedHat Apache organization.

5.1.1 Configuration Files

/etc/httpd/mrtg/run MRTG configuration file. Includes both “mrtg.conf”, the main MRTG configuration file.

5.1.2 Web content

/home/httpd/html/mrtg The default location for web-based content is in this directory. Graph files and images are placed in this directory by MRTG.

5.1.3 Administration files

/etc/httpd/admin Both the wrapper scripts and the administration script files are located in this directory. The wrapper scripts are protected by Apache access controls so that they are inaccessible by all except the scripts themselves.

5.2 Accessing the MRTG graphs

Immediately after the MRTG package is installed, it will generate a set of graphs based on current data obtained from Squid. It is likely that you will have to wait the full five minutes to see useful results on the graphs. The easiest way to access the graphs is by the URL path “/viewgraph”. Simply type “http://your.domain.com/viewgraph” into your web browser to see a list of graphs.

5.3 Using the Administration interface

5.3.1 Accessing the Administration interface

The Administration interface can be accessed by entering the URL path “/mrtg-admin” after the hostname in your browser (i.e. `http://your.domain.com/mrtg-admin`). The Administration interface is password protected with the same username/password combination as the Squid Cache Manager which requires a login of “admin” and a password of “admin”.

5.3.2 Using the MRTG Status

The MRTG Status section of the Administration interface provides at-a-glance status of Squid and Apache including version information and whether or not MRTG is running or not. It also provides access to the three time settings that relate to MRTG and log creation. The first number listed is a

the number of seconds specified in the MRTG configuration file that the browser will wait before refreshing the MRTG graph page with updated graphs. The number is converted to minutes by the MRTG Administration interface for easy readability. The second number is the value that MRTG uses to accurately calculate statistics. It should indicate the frequency that MRTG is run. The third number is the frequency at which MRTG is run by cron.

These values should be the same in normal operation. The numbers reflect the actual values – since the values can be different, the differences will appear here. To adjust the frequency at which MRTG graphs are produced, simply use the MRTG interval pulldown to select the desired time and click the “Save Settings” button. Note that changing the settings this way will update all three time values (this is the recommended behavior).

5.3.3 Using the Log Rolling section

The Log Rolling section of the Administration interface displays the current frequency at which the logs will be rolled. It also lets the log rolling frequency be modified. Time log rolling frequency can be adjusted by using the pulldowns which will adjust the frequency of log rolling of the Apache Error Log and the Squid Access Log accordingly. Since Squid uses cron to initiate its log rolling, there is an additional option for rolling the Access Log that allows the time of day that the logs will be rolled to be adjusted. The time-of-day field is not applicable to a log rolling frequency of less than a day. Changes to the settings can be saved by adjusting the values and clicking the “Save Settings” button to make the change.

5.3.4 Changing the Administration interface password

The administration interface uses “DB” files to record the username and password. The username for use in the Administration interface is “admin”. To change the password, you must first be the “root” user (either by logging in as root or with su). Change to the “/etc/httpd/auth” directory and issue the following command:

```
dbmmanage users update admin
```

You will be prompted to enter the new password twice. After this, the password will be updated and it will work immediately upon reloading the administration pages in your web browser.

6 Technical Features

This section attempts to explain the features that Anemone adds to a typical Squid/Apache web server setup and the modifications made to the various configuration files, startup scripts, and how

the installer works. For additional information, please see the comments in the source code of the programs and shell scripts included with the package.

6.1 Installer

6.1.1 Squid configuration file changes (/etc/squid/squid.conf)

The RPM post-install script makes modifications to the Squid configuration file, `squid.conf`, in order to enable the SNMP abilities of Squid:

1. Makes changes to the “`acl`” and “`snmp_access`” directives.

6.1.2 Crontab file changes (/etc/crontab)

The RPM post-install script makes modifications to the “`/etc/crontab`” file so that MRTG will run every five minutes.

6.2 Administration interface

The Anemone Administration interface consists of a second Apache server that is configured by a separate Apache configuration file called “`httpd-admin.conf`”. This file is located in the “`/etc/httpd/conf`” directory with the main server configuration file. A modified version of the “`/usr/sbin/apachectl`” program is provided. This script is called “`apachectl-admin`” and is located in the “`/usr/sbin`” directory. This script takes the same options as the standard “`apachectl`” program and can be used to stop, start, and restart the administration interface.

Since both the Administration Apache server and the main server run as user “`nobody`”, additional work is needed to allow the administration interface to read and write system configuration files and execute the necessary programs. This is accomplished by use of C programs that run “`suid root`”. Using Apache access controls, only the Administration script can be executed by web requests – the wrapper scripts located in the “`/etc/httpd/admin`” directory are protected.

7 Questions or Comments

If you have any questions or comments about this software or about other services provided by The Electric Connection (<http://www.electricc.com>), please contact us at info@electricc.com.