

# SNARE

System iNtrusion Analysis & Reporting Environment

## Guide to SNARE Epilog for UNIX, SnareApache and SnareSquid

INTERSECT  
ALLIANCE

---

## Documentation History

Version No.	Date	Edits	By whom
1.0	19 June 2006	First draft for the Guide to SNARE Epilog for UNIX documentation, in the updated format.	David Mohr
1.1	6 July 2006	Minor updates and introduced documentation on the new Epilog modules for SnareApache and SnareSquid	David Mohr

© 1999-2006 Intersect Alliance Pty Ltd. All rights reserved worldwide.

Intersect Alliance Pty Ltd shall not be liable for errors contained herein or for direct, or indirect damages in connection with the use of this material. No part of this work may be reproduced or transmitted in any form or by any means except as expressly permitted by Intersect Alliance Pty Ltd. This does not include those documents and software developed under the terms of the open source General Public Licence, which covers the Snare agents and some other software.

The Intersect Alliance logo and Snare logo are registered trademarks of Intersect Alliance Pty Ltd. Other trademarks and trade names are marks' and names of their owners as may or may not be indicated. All trademarks are the property of their respective owners and are used here in an editorial context without intent of infringement. Specifications and content are subject to change without notice.

## About this guide

This guide introduces you to the functionality of the SNARE Epilog for Text Files. SNARE Epilog for UNIX is currently released for the Linux and Solaris operating systems and facilitates objective-based filtering, and remote audit event delivery of text-based log files for Linux and Solaris based systems. SNARE Epilog for UNIX will also allow a security administrator to fully remote control the application through a standard web browser if so desired. SNARE Epilog also includes additional modules for specifically monitoring Apache and Squid logs.

Other guides that may be useful to read include:

- SNARE Server User's Guide.
- Installation Guide to the SNARE Server.
- SNARE Server Troubleshooting Guide.
- The SNARE Toolset - A White Paper.

### Table of contents:

<b>1 Introduction.....</b>	<b>5</b>
<b>2 Overview of Snare Epilog for UNIX.....</b>	<b>6</b>
<b>3 Installing and running Epilog.....</b>	<b>7</b>
3.1 Epilog installation.....	7
3.2 Running Epilog.....	8
3.3 SnareApache installation.....	9
3.4 Running SnareApache.....	10
3.5 SnareSquid installation.....	11
3.6 Running SnareSquid.....	12
<b>4 Setting the audit configuration.....</b>	<b>13</b>
4.1 Audit configuration.....	13
4.2 Configuration Control.....	13
4.3 Auditing control.....	14
4.4 Log configuration.....	15
<b>5 Remote control and management.....</b>	<b>17</b>
5.1 Remote control.....	17
5.2 Remote distribution.....	19
<b>6 SNARE Server.....</b>	<b>20</b>
<b>7 About InterSect Alliance.....</b>	<b>22</b>
<b>Appendix A - Event Output Format.....</b>	<b>23</b>
<b>Appendix B - SNARE Configuration File.....</b>	<b>24</b>

## 1 INTRODUCTION



The team at InterSect Alliance have experience with auditing and intrusion detection on a wide range of platforms - Solaris, Windows NT/2000/2003/XP, Netware, Tru64, Linux, IRIX, AIX even MVS (ACF2/RACF); and within a wide range of IT security in businesses such as - National Security and Defence Agencies, Financial Service firms, Government Departments and Service Providers.

This background gives us a unique insight into how to effectively deploy host and network intrusion detection systems that support and enhance an organisation's business goals.

The development of 'SNARE Epilog for UNIX' will now allow text-based log files to be collected from Linux and Solaris based operating systems, and forwarded to a remote audit event collection facility. SNARE Epilog for UNIX will also allow a security administrator to fully remote control the application through a standard web browser if so desired. SNARE has been designed in such a way as to allow the remote control functions to be readily effected manually, or by an automated process.

The overall project is called 'SNARE' - **System iNtrusion Analysis & Reporting Environment**. The '**SNARE Server**' is a commercial release of software beneficial to organisations that wish to collect from a wide variety of SNARE agents and appliances such as firewalls or routers.

InterSect Alliance welcomes and values your support, comments, and contributions. Our contact details are available from our contact page at [www.intersectalliance.com](http://www.intersectalliance.com).

## 2 OVERVIEW OF SNARE EPILOG FOR UNIX



SNARE operates through the actions of one key application; the '*Epilog*' process. SNARE will monitor the given log files and manage the generated events based on the objectives defined in the SNARE configuration files. Log files are filtered using the SNARE objectives, labeled according to the log type identified and then passed over the network, using the UDP or TCP protocol, to one or more remote servers for collection, analysis and archival. *The TCP protocol capability, and the ability to send events to multiple hosts is only available in the 'supported' versions of the agents, which are made available to Snare Server customers. See Chapter 7 of this document for further details.* The *Epilog* daemon is able to be remotely controlled using a standard web browser.

SnareApache and SnareSquid are modules of Epilog that allow for targeted logging of Apache and Squid logs. These modules can also be controlled using a web browser.

## 3 INSTALLING AND RUNNING EPILOG



### 3.1 EPILOG INSTALLATION

SNARE Epilog includes an installation script to allow for easy installation and configuration of all critical components. The Epilog installation file includes the following key components:

- The *Epilog* binary that performs all of the main SNARE functions.
- Configuration files necessary to allow *Epilog* to work with the supported log files.
- The installation script, *install.sh*, which installs all necessary components.
- A removal script - *uninstall.sh*, which removes Epilog components from your system.

There are three general components to the Epilog installation package:

- **Epilog binary**  
The *Epilog* daemon is contained in the '*Epilog*' binary. This binary provides the capability to read event log records from text files, filter the events according to the 'objectives' defined by the user, provide a web based remote control interface, and specify the log files to monitor.
- **install.sh/uninstall.sh**  
These two scripts undertake the installation and uninstallation functions for Epilog for UNIX. The scripts may prompt the user for confirmation, or for specific configuration options during the installation and uninstallation processes.
- **Configuration File**  
A single configuration file is required to correctly run *Epilog*. The file *epilog.conf* is copied to the */etc/snare/epilog* directory on your local filesystem during the installation process.

#### ▶ HOW TO... Install the SNARE Epilog package for UNIX

1. Download the 'Epilog' file from the Intersect Alliance website to the target server.
2. Change directory to the folder containing Epilog and, as the 'root' user, type:  

```
# gzip -d Epilog-<version>.tar.gz
# tar xvf Epilog-<version>.tar
```

*(where <version> is the version of SNARE Epilog for UNIX you wish to install)*
3. A directory called Epilog-<version> will be created. Enter this directory:  

```
# cd Epilog-<version>
```
4. In order to commence the installation, type in './install.sh'. A prompt will only be displayed if an existing configuration file is found, otherwise a basic configuration file is used by default.
5. Once the installation process has completed, the *Epilog* daemon will start automatically (although no log monitors will be configured) and the daemon will also be integrated into your normal boot process.

### 3.2 RUNNING EPILOG

Upon installation of SNARE Epilog for UNIX, the *Epilog* binary will be installed in the */usr/bin* directory. The *Epilog* process will be controlled by the */etc/init.d/epilogd* daemon control script, so there is no need to start or stop *Epilog* directly.

The **Epilog** daemon must be running, if the events are to be passed to a remote host. The **Epilog** daemon may be stopped, started or restarted by issuing the commands: `/etc/init.d/epilogd stop`, `/etc/init.d/epilogd start` or `/etc/init.d/epilogd restart`, respectively.

**▶ HOW TO...** Run the **Epilog** Daemon:

1. Login as root.
2. Execute the command `/etc/init.d/epilogd start`.
3. Execute the command `ps -ef | grep epilog`, and check that there is one (or two if the micro-web server is active) process called `/usr/bin/epilog`.

**▶ HOW TO...** Enable Remote Audit Control

If the **Epilog** daemon is run on a system that has remote control enabled in the `epilog.conf` file, then the audit subsystem may be remotely controlled using a standard web browser. Note that for this to work, the remote control facility should be set (see the following section of the documentation for specific instructions on remote control settings), and the `/etc/snare/epilog/epilog.conf` MUST have AT LEAST the 'allow=1' line under the [Remote] configuration category specified (NB: Epilog must also have a different `listen_port` if it is operating on the same system as a Snare operating system audit daemon):

```
[Remote]
  allow=1
  listen_port=6162
  restrict_ip=10.0.0.1
  accesskey=SnYlb.gT4Gk2k
```

If the 'restrict\_ip' line is in the `epilog.conf` file, then the only machine that can access the remote control feature, is the system that is listed on that line. If the 'accesskey' line is specified, then a password is required to access the remote control function (the username for remote control is always **snare**). The password in the SNARE configuration file, is 'encrypted' using the standard UNIX 'crypt' function. Using a web browser type in the following on the URL bar:

```
http://<ip address or DNS hostname>:6162
```

(NOTE that '6162' will be the port number specified in the '`listen_port`' of the `/etc/snare/epilog/epilog.conf` file).

### 3.3 SNAREAPACHE INSTALLATION

SnareApache includes an installation script to allow for easy installation and configuration of all critical components. The SnareApache archive includes the following key components:

- Configuration files necessary to allow SnareApache to work with the given log files.
- The installation script, *snareapache\_install.sh*, which enables the easy installation of all necessary components.
- A removal script - *snareapache\_uninstall.sh*, which removes SnareApache components from your system.

The event processing is handled by the Epilog binary. There are two general components to the SnareApache installation package:

- **snareapache\_install.sh/snareapache\_uninstall.sh**  
These two scripts undertake the installation and uninstallation functions for SnareApache. The scripts may prompt the user for confirmation, or for specific configuration information, if required (discussed in detail below).
- **Configuration File**  
A single configuration file is required to correctly run *SnareApache*. The file *apache.conf* is copied to the */etc/snare/epilog* directory of your local filesystem during the installation process.

#### ▶ HOW TO... Install the SnareApache package for UNIX

1. Ensure that Snare Epilog is correctly installed.
2. Download the 'SnareApache' file from the Intersect Alliance website to the target server.
3. Change directory to the folder containing SnareApache and, as the 'root' user, type:  

```
# gzip -d SnareApache-<version>.tar.gz
# tar xvf SnareApache-<version>.tar
```

*(where <version> is the version of SnareApache you wish to install)*
4. A directory called SnareApache-<version> will be created. Enter this directory:  

```
# cd SnareApache-<version>
```
5. In order to commence the installation, type in './snareapache\_install.sh'. A prompt will only be displayed if an existing configuration file is found, otherwise a basic configuration file is used by default.
6. Once the installation process has completed, the *SnareApache* daemon will start automatically (although no destination server will be configured) and the daemon will also be integrated into your normal boot process. The default configuration will monitor the file */var/log/httpd/access\_log*.

### 3.4 RUNNING SNAREAPACHE

Upon installation of SnareApache, a symlink will be created in `/usr/bin` pointing to the *Epilog* binary. The *SnareApache* process will be controlled by the `/etc/init.d/snareapached` daemon control script, so there is no need to start or stop *SnareApache* directly.

The *SnareApache* daemon must be running, if the events are to be passed to a remote host. The *SnareApache* daemon may be stopped, started or restarted by issuing the commands: `/etc/init.d/snareapached stop`, `/etc/init.d/snareapached start` or `/etc/init.d/snareapached restart`, respectively.

#### ▶ HOW TO... Run the *SnareApache* Daemon:

1. Login as root.
2. Execute the command `/etc/init.d/snareapached start`.
3. Execute the command `ps -ef | grep snareapache`, and check that there is one process called `/usr/bin/snareapache` (or two if the snareapache micro-web server is active, NB: this is not the default setting)..

#### ▶ HOW TO... Enable Remote Audit Control

By default, the SnareApache daemon can be controlled through the Epilog web interface, however, individual remote control can still be provided using the procedure below.

If the *SnareApache* daemon is run on a system that has remote control enabled in the `apache.conf` file, then the audit subsystem may be remotely controlled using a standard web browser. Note that for this to work, the remote control facility should be set (see the following section of the documentation for specific instructions on remote control settings), and the `/etc/snare/epilog/apache.conf` MUST have AT LEAST the 'allow=1' line under the [Remote] configuration category specified (NB: SnareApache must also have a different `listen_port` if it is operating on the same system as a SnareCore or Epilog daemon):

```
[Remote]
allow=1
listen_port=6163
restrict_ip=10.0.0.1
accesskey=SnYlb.gT4Gk2k
```

If the 'restrict\_ip' line is in the `apache.conf` file, then the only machines that are able to remote control the agent are those listed on that line. If the 'accesskey' line is specified, then a password is required to access the remote control function (the username for remote control is always *snare*). The password in the SNARE configuration file, is 'encrypted' using the standard UNIX 'crypt' function. Using a web browser type in the following on the URL bar:

```
http://<ip address or DNS hostname>:6163
```

(NOTE that '6163' will be the port number specified in the '`listen_port`' of the `/etc/snare/epilog/apache.conf` file).

### 3.5 SNARESQUID INSTALLATION

SnareSquid includes an installation script to allow for easy installation and configuration of all critical components. The SnareSquid archive includes the following key components:

- Configuration files necessary to allow SnareSquid to work with the given log files.
- The installation script, *snare squid\_install.sh*, which enables the easy installation of all necessary components.
- A removal script - *snare squid\_uninstall.sh*, which removes SnareSquid components from your system.

The event processing is handled by the Epilog binary. There are two general components to the SnareSquid installation package:

- **snare squid\_install.sh/snare squid\_uninstall.sh**  
These two scripts undertake the installation and uninstallation functions required to ensure SnareSquid works as required. The scripts prompt the user on the steps that need to be undertaken and the choices to be made (discussed in detail below).
- **Configuration File**  
A single configuration file is required to correctly run *SnareSquid*. The file *squid.conf* is copied to the */etc/snare/epilog* directory of your local filesystem during the installation process.

#### **HOW TO...** Install the SnareSquid package for UNIX

1. Ensure that Snare Epilog is correctly installed.
2. Download the 'SnareSquid' file from the Intersect Alliance website to the target server.
3. Change directory to the folder containing SnareSquid and, as the 'root' user, type:  

```
# gzip -d SnareSquid-<version>.tar.gz
# tar xvf SnareSquid-<version>.tar
```

**(where <version> is the version of SnareSquid you wish to install)**
4. A directory called SnareSquid-<version> will be created. Enter this directory:  

```
# cd SnareSquid-<version>
```
5. In order to commence the installation, type in './snare squid\_install.sh'. A prompt will only be displayed if an existing configuration file is found, otherwise a basic configuration file is used by default.
6. Once the installation process has completed, the *SnareSquid* daemon will start automatically (although no destination server will be configured) and the daemon will also be integrated into your normal boot process. The default configuration will monitor the file */var/log/squid/access.log*.

### 3.6 RUNNING SNARESQUID

Upon installation of SnareSquid, a symlink will be created in `/usr/bin` pointing to the *Epilog* binary. The *SnareSquid* process will be controlled by the `/etc/init.d/snaresquidd` daemon control script, so there is no need to start or stop *SnareSquid* directly.

The *SnareSquid* daemon must be running, if the events are to be passed to a remote host. The *SnareSquid* daemon may be stopped, started or restarted by issuing the commands: `/etc/init.d/snaresquidd stop`, `/etc/init.d/snaresquidd start` or `/etc/init.d/snaresquidd restart`, respectively.

#### ▶ HOW TO... Run the *SnareSquid* Daemon:

1. Login as root.
2. Execute the command `/etc/init.d/snaresquidd start`.
3. Execute the command `ps -ef | grep snaresquid`, and check that there is one process called `/usr/bin/snaresquid` (or two if the snaresquid micro-web server is active, NB: this is not the default setting).

#### ▶ HOW TO... Enable Remote Audit Control

By default, the SnareSquid daemon can be controlled through the Epilog web interface, however, individual remote control can still be provided using the procedure below.

If the *SnareSquid* daemon is run on a system that has remote control enabled in the `squid.conf` file, then the audit subsystem may be remotely controlled using a standard web browser. Note that for this to work, the remote control facility should be set (see the following section of the documentation for specific instructions on remote control settings), and the `/etc/snare/epilog/squid.conf` MUST have AT LEAST the 'allow=1' line under the [Remote] configuration category specified (NB: SnareSquid must also have a different `listen_port` if it is operating on the same system as a SnareCore or Epilog daemon):

```
[Remote]
allow=1
listen_port=6163
restrict_ip=10.0.0.1
accesskey=SnYlb.gT4Gk2k
```

If the 'restrict\_ip' line is in the `squid.conf` file, then the only machines that are able to remote control the agent are those listed on that line. If the 'accesskey' line is specified, then a password is required to access the remote control function (the username for remote control is always *snare*). The password in the SNARE configuration file, is 'encrypted' using the standard UNIX 'crypt' function. Using a web browser type in the following on the URL bar:

```
http://<ip address or DNS hostname>:6163
```

(NOTE that '6163' will be the port number specified in the '`listen_port`' of the `/etc/snare/epilog/squid.conf` file).

## 4 SETTING THE AUDIT CONFIGURATION

### 4.1 CONFIGURATION

The configuration files are stored in */etc/snare/epilog*. This directory contains necessary configuration files with all the details required by the audit daemons to successfully execute. Failure to have a correct configuration file available in this location will not 'crash' the daemons, but will result in logs not being processed, or forwarded to your central log server.

**Tip:** Manual editing of the configuration files is possible, but care should be taken to ensure that it conforms to the required format for the audit daemon. Also, any use of the Remote Control SNARE capability to modify security objectives or selected events, will result in any manual configuration file changes being overwritten. Details on the configuration file format can be viewed in Appendix B - SNARE Configuration File.

The most effective and simplest way to configure the Epilog audit daemons (including SnareApache and SnareSquid) is to use the remote control capability. Additionally, the installation script will ensure that a correctly formatted configuration file is generated, based on broad user requirements when initially installing the Epilog tools.

The installation scripts may request the choice of 2 installation profiles. There is only one starting configuration for each agent which will be automatically selected unless the agent is being reinstalled. Where the agent is being reinstalled, there is the second option to preserve the existing configuration file.

### 4.2 CONFIGURATION CONTROL

Due the modular nature of Epilog, a single remote control interface is available to access and modify all of the available configuration files, as shown in Figure 1. Once the chosen configuration file is selected (indicated by the **bold** name), all of the functions discussed below will operate on the selected configuration file. By default, all functions will operate on the *epilog.conf* file.

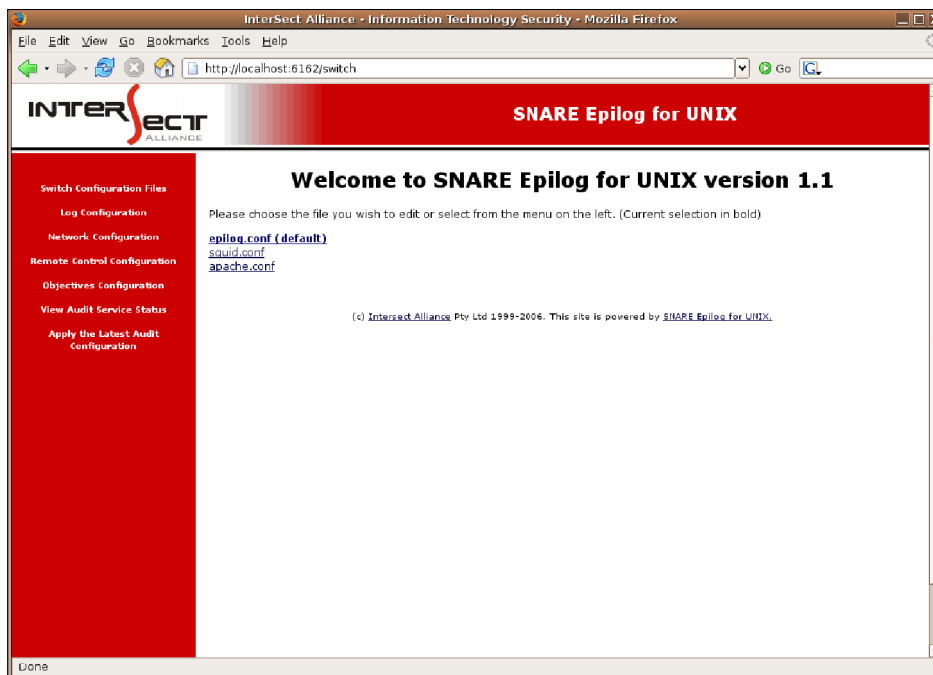


Figure 1: Configuration File Selection

### 4.3 LOGGING CONTROL

The initial log configuration parameters to consider are:

- The hostname, IP address and UDP or TCP port of the remote collection servers,

These three parameters are shown in the '*Network Configuration*' menu, shown in Figure 2 below. Note that the figure below shows all the menus for other features, such as 'Log Configuration', 'Objectives Configuration' and 'Remote Control Configuration'. This and other functions are discussed later in this documentation.

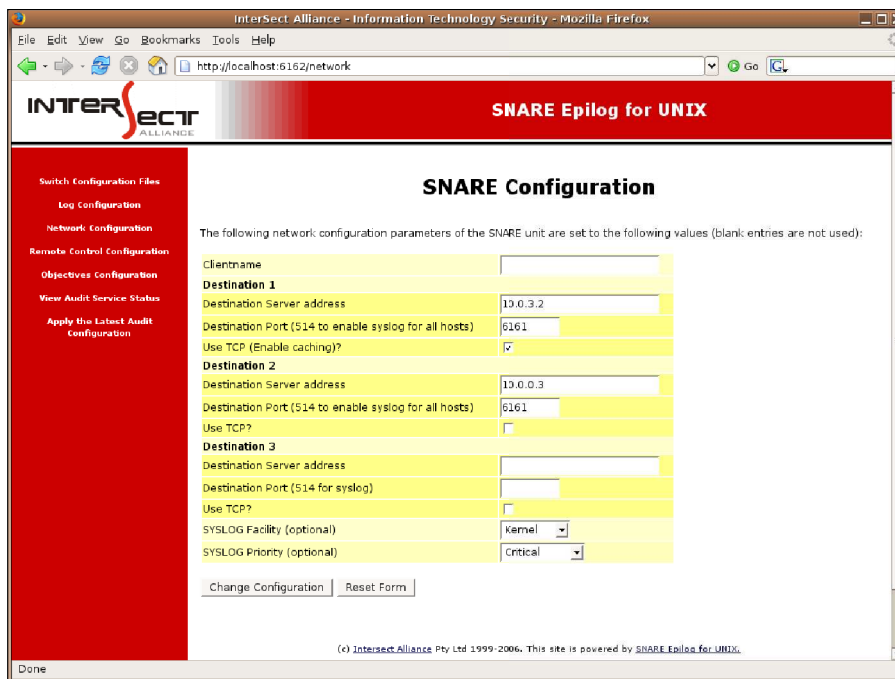


Figure 2: Network Configuration

The hostname field can be used to override the fully qualified domain name of the host system, which will be used by Epilog if this field is blank. Note that executing the command 'hostname' on a command prompt will display the current host name allocated to the host.

The 'port' shown in Figure 2 is the SNARE Server's port that will be used to collect the events. If, for example, the Intersect Alliance SNARE Server is used, then this should be the default port of 6161. Supported agents will have an additional options to enable TCP (and optional caching for one server) and configure multiple hosts. See Chapter 7 below for more details on the supported versions of the SNARE agents. Additional hosts can be added one at a time by clicking "Change Configuration" after each addition. To remove a host, delete the "Destination server address" and click "Change Configuration". The caching feature will store unsent messages in memory until the destination server is once again contactable. The cache is limited to 320000 messages or the available memory of the host system (whichever comes first). Restarting the agent will purge this cache, freeing all the memory used by the cache.

A major function of SNARE Epilog for UNIX is the capability to filter events. This is accomplished via the advanced auditing 'objectives' function. Any number of objectives may be specified, and are displayed within the 'Objectives Configuration' menu on the remote control browser page, as shown in Figure 3 below.

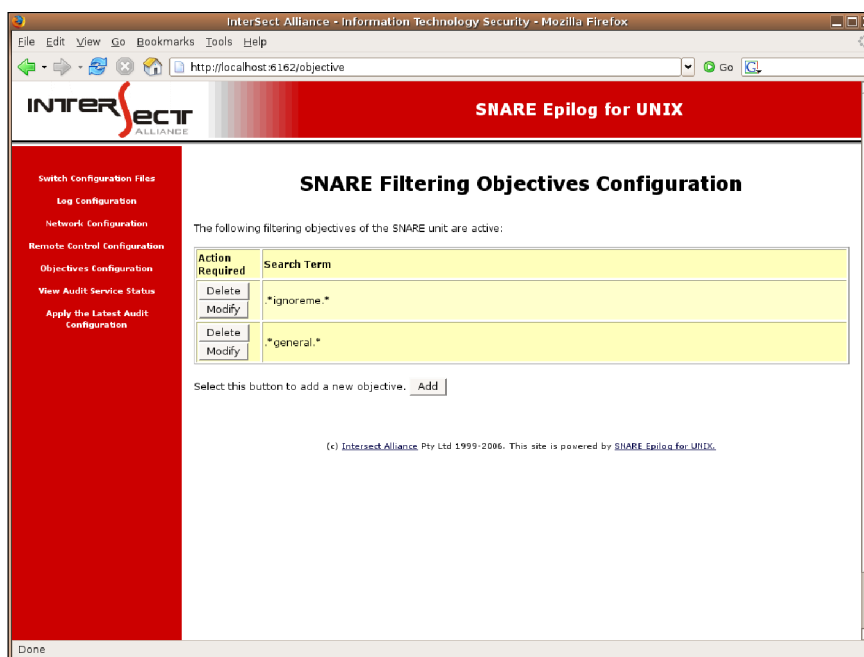


Figure 3: Objectives Configuration

Each of the objectives provides a high level of control over which events are selected and reported. Events are selected using specific filters called 'Objectives'. Due to the generic nature of SNARE Epilog for UNIX, no default objectives are defined and subsequently, all events will be passed directly to the configured network destination.

The 'Search Term' allows a 'regular expression' match term to check against the event-specific matchable item. Regular expressions are an advanced form search filter. For example, the term `'.*[Pp]ass(word|wd).*` would match the following:

- /etc/passwd
- /tmp/PasswordFile

but would not match

- /etc/PASSWD/
- /home/red/PaSsWoRd .txt

The search term will be used to search the entire string for any matches against the given expression. So for example, this means that an included search term of `'.*pwd.*'` would apply to any single line with the term 'pwd' contained in it. If the objective is set to exclude, then lines matching the search term will be discarded. All events are included by default.

#### 4.4 LOG CONFIGURATION

The *Epilog* daemon's main focus is the ability to monitor any text-based log file. The initial log configuration parameters to consider are:

- The location of the log files to be monitored, and
- The type of log files being monitored.

These parameters are shown in the '*Log Configuration*' menu, shown in Figure 4 below.

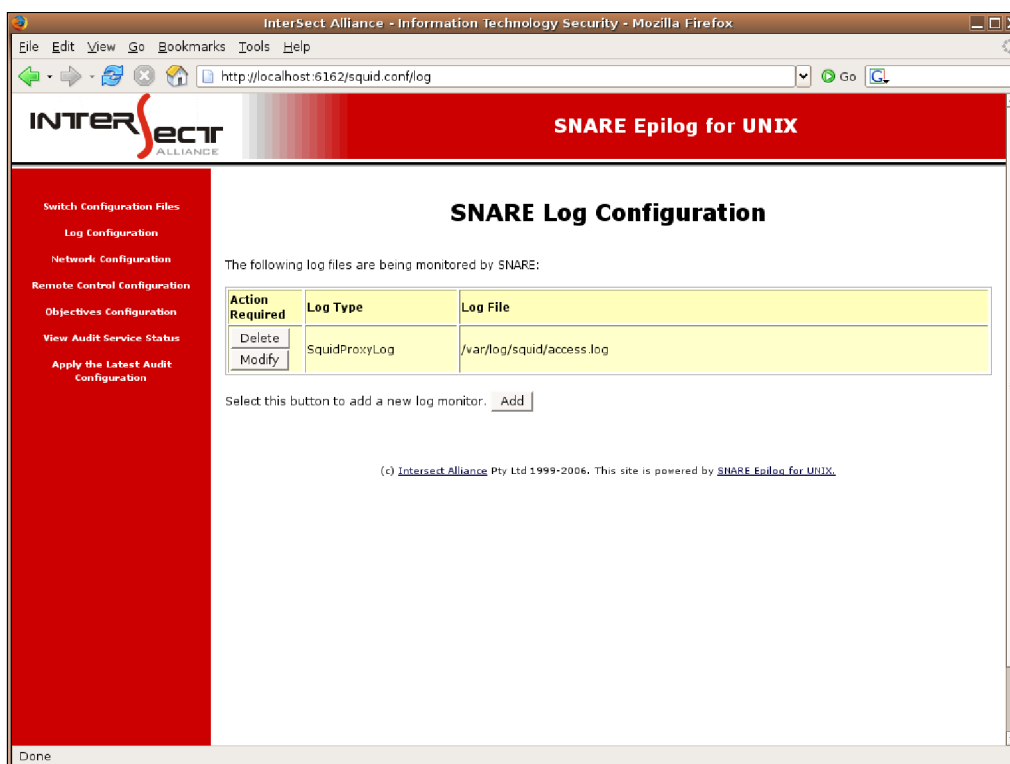


Figure 4: SnareSquid Log Configuration

From this page, log monitors can be added, deleted and modified. The 'Log File' field must be defined as the fully qualified path to the desired log file (and may include spaces). SNARE Epilog for UNIX will then continuously monitor this file for any changes, immediately reporting them to the identified SNARE servers. SNARE Epilog for UNIX will follow the exact name of the file even if it is rotated, truncated, replaced or deleted. In the event that the file is removed, the Epilog daemon will wait until the file is recreated and then resume normal monitoring. The log type of a file will tell the SNARE server how to handle the incoming data stream and in which database table the processed information should be stored. The currently available log types are:

- GenericLog - Generic log format (default)
- ApacheLog - Apache web logs
- ISAWebLog - Microsoft ISA web logs
- MSProxySvr - Microsoft proxy server logs
- SMTPSvcLog - Microsoft SMTP logs
- SquidProxyLog - Squid proxy logs

Once the above settings have been finalised, clicking 'Change Configuration' on the remote control page will save the configuration to the designated configuration file (as defined by the 'Switch Configuration' page). However, to ensure the designated daemon has received the new configuration, the daemon MUST be restarted via the 'Apply the Latest Audit Configuration' menu item, or alternatively, by issuing the restart command to the associated daemon control script.

## 5 REMOTE CONTROL AND MANAGEMENT



The *Epilog* service is a separate standalone component of the SNARE system, as described in *2 Overview of Snare Epilog for UNIX on page 5*. The audit configuration can be developed and set using the remote control web browser, set via the installation script, or configured manually as per the exact requirements detailed in Appendix B: SNARE Configuration File.

The *Epilog* daemon (along with the SnareApache and SnareSquid daemons) can be restarted remotely from the menu item *Apply the Latest Audit Configuration*. This will instruct the audit daemon to re-read the configuration file, clear the buffers and restart. This function is useful when changes to the audit configuration have simply been saved to the configuration file, without being 'applied'. The user can therefore select when to activate a new configuration by selecting this link. This restart process can also be executed via the command line as follows:

- As root, execute the command: `ps -ef | grep epilog`
- It should return something like:

```
# ps -ef | grep epilog
root 1271 1270 0 13:17:27 pts/4    0:00 /usr/bin/epilog
root 1270      1 0 13:17:27 pts/4    0:00 /usr/bin/epilog
root 1273 1059 0 13:17:33 pts/4    0:00 grep epilog
```

- As root, execute the command: `/etc/init.d/epilogd restart`
- As root, execute the command: `ps -ef | grep epilog`, and check that the processes have been restarted by ensuring the '*Epilog*' processes have new process IDs.
- SnareApache and SnareSquid can be restarted using the same process with "snareapache" or "snaresquid" used in place of "epilog".

### 5.1 REMOTE CONTROL

Another function of the *Epilog* services are their ability to be remote controlled. This facility has been incorporated to allow all the functions normally available through the configuration file, to also be available through a standard web browser. The *Epilog* daemon employs a custom micro-web server to facilitate configuration through a browser, or via an automated custom designed tool such as the SNARE Server. The *Epilog* micro-web server also has the added ability to control the configuration files for SnareApache and SnareSquid (as discussed in Section 4.2), so the remote control of the individual agents is not necessary. Figure 5 below shows a web browser connecting to a SNARE agent.

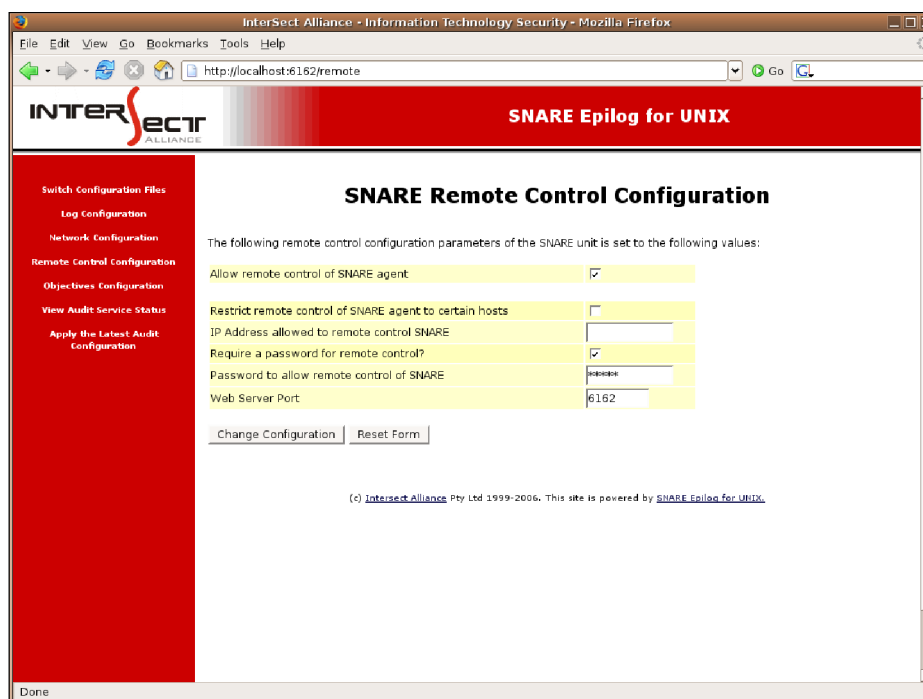


Figure 5: Remote Control Configuration

The functions available through the web browser are identical to those available on the SNARE configuration file. The parameters which may be set for remote control operation are shown in Figure 5, and are discussed in detail below:

- **Allow remote control of SNARE agent.** Selecting this checkbox will allow the SNARE agent to be remotely controlled from a web browser. This host may be independent from the central audit collection server. If the remote control function is disabled, and you wish to enable the facility, follow the instructions detailed in 'Enable Remote Audit Control' in Section 3.2 of this document.
- **IP Address allowed to remote control SNARE.** Remote control actions may be limited to a given host. This host, entered as an IP address in this field, will only allow remote connections to be effected from the stated IP address. Note that access control based on source IP address is prone to spoofing, and should be considered as a security measure to be used in conjunction with other countermeasures (such as ensuring your organisational firewall does not allow external connections to the SNARE micro-web server port).
- **Password to allow remote control of SNARE.** A password may be set so that only authorised individuals may access the remote control functions. If accessing the remote control functions through a browser or batch-mode tool (such as 'curl' or 'wget'), note that the UserID is always 'snare', and the password is whatever has been defined by the user. This password is not encrypted when being transmitted via the http session, but is encrypted when stored in the respective configuration files.
- **Web Server Port.** Normally, a traditional web server operates on port 80. If this is the case, then a user need only type the address into the browser to access the site. If however, a web server is operating on a port other than 80 (eg. 6162), then the user needs to type **http://mysite.gov:6162** to reach the web server. The default *Epilog* web server port may be changed using this setting, if it conflicts with an established web server or SNARE agent. However, care should be taken to note the new server port, as it will need to be placed in the URL needed to access the SNARE agent.

---

## 5.2 REMOTE DISTRIBUTION

Epilog provides the facility to send events to a remote host, using UDP or TCP. The purpose of the SNARE agents is to filter the local events, and send them in real time to a remote server. No local log file is available for this agent.

The remote server may be a SNARE Server (discussed in the next section), or a custom tool that listens on port 6161 for SNARE events.

## 6 SNARE SERVER



The SNARE Server collects events and logs from a variety of operating systems, applications and appliances including, but not limited to: Windows NT/2000/XP/2003, Solaris, AIX, Irix, Linux, Tru64, ACF2, RACF, CISCO Routers, CISCO PIX Firewall, CyberGuard Firewall, Checkpoint Firewall1, Gauntlet Firewall, Netgear Firewall, IPTables Firewall, Microsoft ISA Server, Microsoft IIS Server, Lotus Notes, Microsoft Proxy Server, Apache, Squid, Snort Network Intrusion Detection Sensors, IBM SOCKS Server, and Generic Syslog Data of any variety.

In addition to the above, the benefits of purchasing the SNARE Server include:

- Official support mechanism for the SNARE open source agents. Note that official SNARE agent support is not offered through *any* other channels.
- All future SNARE Server versions and upgrades included as part of an annual maintenance fee.
- Ability to collect any arbitrary log data, either via UDP or TCP protocols.
- Proven technology that works seamlessly with the SNARE agents.
- SNARE reflector technology that allows for all collected events to be sent, in real time, to a standby/backup SNARE Server.
- Ability to continuously collect large numbers of events. SNARE Server collection rates exceed 60,000 events per minute using a low end, workstation class, Intel based PC on a 100Mbps network.
- Automatic archiving of events to compressed text format after a configurable event time period. This is to prevent the database from slowing down due to storage of old events.
- Ability to drill down from top level reports. This reduces the amount of data “clutter” and allows a system administrator to fine tune the reporting objectives.
- Ability to create “cloned” objectives that allow very specific reporting against any collection profile. These reports, along with all SNARE Server objectives, may be scheduled and emailed to designated staff.
- The SNARE Server uses extensive discriminators for each objective, allowing system administrators to finely tune reporting based on inclusion or exclusion of certain parameters.
- Very simple, single CD installation for those users not requiring a hardware based appliance.

The SNARE Server uses a hardened version of the Linux operating system base for stability and its ability to use a myriad of stable and functional open source tools. A SNARE Server user, however need not be concerned with managing a Linux server. The SNARE Server, once installed, is a fully contained appliance, and does not require any system administrator level maintenance. The SNARE Server will operate on commonly available Intel based PCs, with hardware specifications shown on the next page.

There are supported versions of the SNARE agents which are only available through the purchase of a SNARE Server. Functionality includes, but is not limited to, ability to send events via TCP as well as UDP, and the ability to send events to many destinations, not just one host.




Figure 6 Screen shot from the SNARE Server

## 7 ABOUT INTERSECT ALLIANCE



InterSect Alliance is a team of leading information technology security specialists in both the 'technical' and 'policy' areas. In particular, Intersect Alliance are noted leaders in key aspects of IT Security, including host intrusion detection. Our solutions have and continue to be used in the most sensitive areas of Government and business sectors. Intersect Alliance consult and contract to number of agencies in Australia and the Asia Pacific, for both the business and Government sectors.

The Intersect Alliance business strategy includes demonstrating our commitment and expertise in IT security by releasing open source products such as SNARE, and the proprietary SNARE Server. Intersect Alliance intend to continue releasing tools that enable users, administrators and clients worldwide to achieve a greater level of productivity and effectiveness in the area of IT Security, by simplifying, abstracting and/or solving complex security problems.

Visit the Intersect Alliance website for more information at [www.intersectalliance.com](http://www.intersectalliance.com).

---

## APPENDIX A - EVENT OUTPUT FORMAT

The *Epilog* daemon collects data from the identified logs files and passes it unaltered to the identified network destination. Whitespace is the primary element used separate elements within the data. An audit event may look something like this:

```
soll0dev      SquidProxyLog      0      1152134522.688      857 127.0.0.1 TCP_MISS/200 12149 GET  
http://www.intersectalliance.com/ - DIRECT/150.101.115.22 text/html
```

The information in blue, as shown in the above record, is information added by the *Epilog* daemon. The format of this information is as follows:

<hostname> <log\_type> <unused> <log\_event>

## APPENDIX B - SNARE CONFIGURATION FILE

Details on the audit configuration were discussed previously. The purpose of this section is to discuss the makeup of the configuration file. The Epilog, SnareApache and SnareSquid configuration files are located in */etc/snare/epilog*, and their locations may not be changed. If the configuration file does not exist, the audit daemon will execute, but will not actively audit events until a correctly formatted configuration file is present, or unless specific instructions are passed to the audit module at load time.

SNARE can be configured in several different ways, namely:

- Via the installation script (*Recommended*), or
- Via the web server (*Recommended*), or
- By manually editing the configuration file.

The format of the audit configuration file is discussed below.

[HostID]	This item stores the hostname, if it is different from the assigned hostname.
name=<hostname>	This is the name of the host.
[Output]	By default, if no output section exists within the configuration file, the audit daemon will <b>NOT</b> send any audit data out. Note that audit events will be sent to all valid network destinations specified in the Output section.
network=hostname:port:tcp network=hostname:port	Audit data can be sent to a remote system using the UDP (default) or TCP protocol. Data will be sent to the remote host, and network port specified here. Each additional host must be specified on a new line. Caching will be enabled for the first host only if TCP is enabled.
[Input]	This section identifies the log files to be monitored.
log=LogType:/fully/qualified/file /name log=/fully/qualified/file/name	The audit daemon will continuously monitor the identified files by name and send data to the network destinations specified within the [Output] section. Spaces are valid characters. Note that if the audit daemon is not running as root, the file must be readable by the user under which the audit daemon is running. The LogType is optional and is used to inform the SNARE server how to process the data stream. A list of valid log types can be found in Section 4.3.
[Objectives]	This section describes the format of the objectives. Objectives are composed of: <ul style="list-style-type: none"> <li>1. The match term : a filter expression, and is defined in extended regular expression format.</li> </ul> <p>Note that whitespace will be trimmed from the start and end of items, but will be assumed to be valid when bracketed by other characters.</p>
match=.*more.* match!=.*less.*	Include any lines that contains the word "more" Exclude any lines that contains the word "less"
[Remote]	This subkey stores all the remote control parameters.
allow=1	"Allow" is an integer, and set to either 0 or 1 to allow remote control. 1= allow remote, 0=do not allow.
listen_port=6162	This value is the web server port. A missing "listen_port" will default the web server to port 80.
restrict_ip=10.0.0.1	This is an IP address, that will be used so that this address will be the only host that is allowed to connect to the web server. If this item does not exist, then the web server will not restrict by IP address.
accesskey=snare	This value is the password that is used to log into the SNARE web server. If this item does not exist, then a password will not be requested when connecting to the web server. The password is encrypted when stored in the snare.conf, using the standard UNIX "crypt" facility, with salt.