

System iNtrusion Analysis & Reporting Environment

User Guide for Snare Server v7.0

© Intersect Alliance International 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 International 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 International 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.

Table of Contents

1. About this User Guide	3
2. Terms and Acronyms	4
3. Introduction to the Snare Server	5
4. Objectives - An Overview	20
5. Modular Objectives	27
6. Modular Objectives - Configuration & Output	31
7. Supporting Objectives - Status	54
8. Supporting Objectives - System	58
9. Output Modification Modules	77
10. Pre-Processed Tokens - FTokens	79
11. Snare Operational Checklists	81
12. Third Party Data Sources	85
13. Regulatory Reporting	89
14. Snare Agents	98
15. Modular Objective Templates	101
16. Collection subsystem	113
17. Expert configuration	116

1. About this User Guide

This guide aims to provide users of the Snare Server with the tools to be able to understand and work with it. Users expecting to work on the Snare Server should familiarize themselves with this document.



Important

This guide applies to the Snare Server v7 appliance only. If you are using an earlier, or later, version of this appliance please review the user guide for that specific version instead.



Other guides that may be useful to read

- Snare Server v7 Installation Guide
- Snare Server v7 Migration Guide
- Snare Server v7 Upgrade Guide
- Snare Enterprise Agent User Guides
- Snare Toolset White Paper



Third-party components

The Snare Server includes the following third-party components:

MaxMind – *This product includes GeoLite2 data created by MaxMind, available from <http://www.maxmind.com/>.*

The Snare Server includes the free GeoLite2 IP geolocation database for use with the IP address country identification functionality, and the Geographic IP Source-Destination mapping output module. Customers can replace the free version with their own licensed version of MaxMind's GeoIP2 database, if higher accuracy is required than the free GeoLite2 IP database provides.

2. Terms and Acronyms

Terms & Acronyms	Explanation
Categories	Navigation tools used to access particular objectives on the server.
Discriminators	Used to formulate an objective. Although there are many discriminators available each objective contains its own set of discriminators.
Event records	An event record contains information on when, what and where an activity has occurred on the host.
Objectives	An objective is a generic name for an interactive report, which performs a specific task or implements a set of analysis rules that are intended to derive useful information from event log data that is collected by the Snare Server.
Snare	The acronym for the System iNtrusion Analysis and Reporting Environment.
Snare Agent	A small program installed on clients (servers, desktops, etc) that gathers events from system logs and sends them to the Snare Server.
Snare Agent Management Console	Component within the Snare Server that provides basic remote configuration management of the Snare Agents within the Snare Server.
Snare Events	An event in Snare can be described as an occurrence in any specific or group of systems that, from an administrator's point of view, is important to note in the day to day running and security of the system.
Snare Server	The Snare Server is used to administer and monitor Snare objectives. It provides a simple web interface to all of the objectives and allows custom configuration of Snare's monitoring capabilities.
Snare System	This refers to both the Snare Server and Snare events.
Reflector	Component within the Snare Server that reflects all incoming events onto another collection server. The results being that both servers receive the same events, without the Agent needing to send to both.

3. Introduction to the Snare Server

The Snare Server allows a System Administrator to define, track and report on events occurring on the server. This section introduces the Snare System, and in particular, the Snare Server.

3.1. Logging on to the Snare Server

The Snare Server user interface can be accessed by any web browser by typing in the URL (server address) in the browser address bar. Any of the methods below work to access the Snare Server user interface:

- The Domain Name Service (DNS) name,
Example: snareserver.dept.gov
- Window domain name, or
Example: snareserver as <http://snareserver/>
- The IP address of the Snare Server.
Example: 192.168.1.1

Once the Snare Server login page loads, a user name and password need to be entered to authenticate access to the Snare Server pages.

It should be noted that both the Snare Server web interface, and the operating system login capabilities have been configured to implement password security and rotation controls. More details are available below.

What You Need

In order to log on to the Snare Server user interface you require:

- the server address
- a user name, and
- a password.

The default username is '*Administrator*' and uses the password defined during installation.

How To..

Log on to the Snare Server

- Open a web browser.
- In the Address bar, type:
<http://<server-address>>
- In **User Name**, type your user name.
- In **Password**, type your password.
- Select the **Log In** button. The Snare Server home page will open.



3.2. Initial Configuration

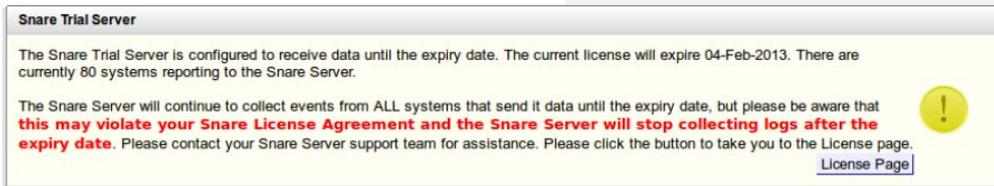
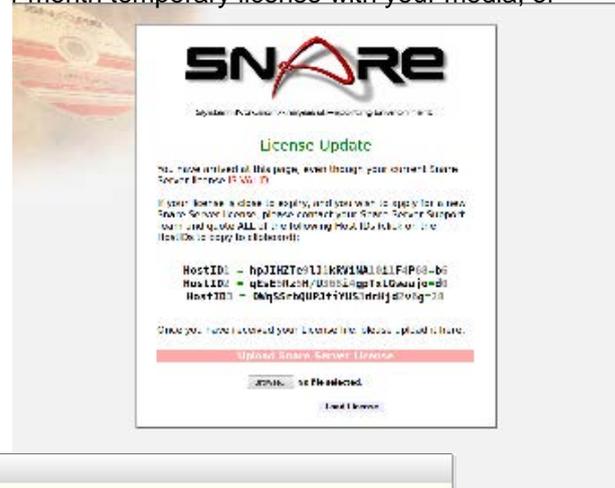
3.2.1. License

After the first logon to the web interface using the username 'administrator' and the password you selected during the installation process, you will be presented with the Host IDs of your new system. These IDs will be required to generate your final license. Please contact your Snare Server support team via the Snare Support online helpdesk service and provide these Host IDs.

Your Snare Server support team may also supply you with a 1 month temporary license with your media, or digital download - this temporary license can be installed at this point, using the License upload form shown in the graphic to the right.

If you need to return to the license page in the future, you can access it through the Snare Server Health Checker, or alternatively, you can point your browser at it directly by going to http://

snareserver.mycompany.com/Setup/License.php
 (where snareserver.mycompany.com is the DNS address or IP address of your Snare Server).



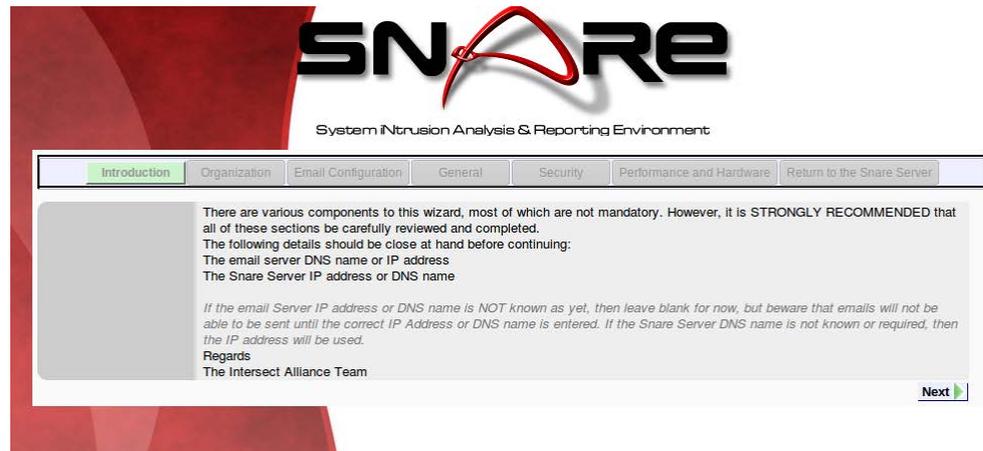
Once the license has been correctly loaded, the Snare Server Configuration Wizard will execute, allowing you to specify initial configuration settings.

3.2.2. Configuration Wizard

3.2.2.1. Introduction

The introduction screen will provide you with a general introduction to the Snare Server Wizard, and note any particular information that you should have on hand before continuing.

Click on the 'Next ->' button once the Snare Server returns control of the interface to you.



3.2.2.2. Organisation Settings

- Enter your organisation name.
- If you wish the login page of the Snare Server to display a notification/warning message, please enter the text you would like to see appear.
- Click on the 'Next ->' button.



- ✔ Although you can click on the navigation buttons at the top of the page in order to skip to a different section of the Snare Wizard, selecting one of these top navigation buttons will not save the changes that you have made to this page. Please use the 'Next' or 'Previous' buttons.

3.2.2.3. Email Configuration Settings

- Enter the DNS Name or IP address of a SMTP email server.
- If you set the default address to append for your organisation, Snare will add this on to any email addresses specified in the scheduled task settings associated with each objective.

- For instance, if you add 'dni.gov.au' here, you can specify 'fred.bloggs' in a scheduled task email configuration item, rather than 'fred.bloggs@dni.gov.au'.
- Enter the Reply-To address that the Snare Server should use to send emails from.
 - This will set any email 'reply to' addresses to this entry. If users hit their 'reply' button on a Snare email, this will be the address that email returns to. It is recommended you configure this to be your IT helpdesk, or a member of your security team.
- Select the preferred Snare Server Login mode.
 - Use HTTPS to force secure web access for Snare Server login.
- Select the preferred Email distribution mode
 - In general, it is recommended that each objective is configured to send out data independently of other objectives. If '*One email per user will go out.*' is selected, there may be a delay of up to 15 minutes after an individual objective completes, before the collection of generated objectives is sent to the destination user.
- If your organisation requires a classification header to be included within the electronic mail messages sent with an objective, add it here.
 - You may also choose to prepend, or append, the classification message to the subject line.
- If you are using an older mail client that cannot handle inline HTML formatted mail, this option gives you the chance to turn HTML content off. Objective output will still be included as an attachment to the electronic mail message.
- Click on the '**Next ->**' button.

3.2.2.4. General Settings

- Enter the DNS name for the Snare Server. This name is inserted into web addresses when electronic mails are sent out from the Snare Server. As such, users will be able to click a link within their electronic mail message, and can be taken to the Snare Server. You should ensure that the name input here matches the name added to your Domain Name Server.

The screenshot shows the 'General' configuration page of the SNARE System Intrusion Analysis & Reporting Environment. The page has a red header with the 'SNARE' logo and the text 'System Intrusion Analysis & Reporting Environment'. Below the header is a navigation bar with tabs: Introduction, Organization, Email Configuration, General (selected), Security, Performance and Hardware, Additional Objectives, and Return to the Snare Server. The main content area contains several configuration fields:

- The DNS Name for the Snare server:** A text input field containing '10.1.2.3' with a small 'x' icon to its right. Below it is a checkbox labeled 'DO NOT regenerate the SSL browser certificate even if the server name has changed' which is checked. A link 'Force a certificate regeneration now' is below the checkbox. A note states: 'This name will be used to identify your Snare Server SSL Certificate, and will also be used to provide absolute links to the Snare Server in electronic mails. Note that if you have requested that the existing SSL certificate be overwritten, a new self signed certificate will be created based on the server name you have entered. If you have modified the servername, or recreated your SSL certificate, it is recommended that you restart the Snare web server. The option to restart the server will be provided to you at the end of this wizard.'
- The DNS Name or IP address of an NTP Date Server:** A text input field containing 'pool.ntp.org'. A note below states: 'Snare will attempt to synchronise the local clock to this server, every 24 hours.'
- Time Zone in which the Snare Server is installed:** A dropdown menu showing 'Australia/Adelaide'.
- The Port on which to contact Snare Agents:** A text input field containing '6161'.
- The Password to use to talk to Snare Agents:** A password field with masked characters '*****'. A note below states: 'This password will be used by default, to contact Snare Agents. Alternatives can be specified in the Snare Agent management objective.'
- Snare Server authentication mode:** Radio buttons for 'Normal' (selected), 'LDAP', and 'Active Directory'. Below these are text input fields for 'Server:' and 'Domain:'. A note at the bottom states: 'User password authentication can be delegated to an external Active Directory or LDAP server. Note that the user must still have a Snare account to log in. If specified, the Domain will be added to the end of the username for authentication purposes (eg: A username of 'auser' and a domain of 'test.local' will imply an LDAP/AD authentication of auser@test.local).'

At the bottom right of the form are 'Previous' and 'Next' buttons.

- The domain name you enter, is used to generate a self-signed SSL certificate. If you have manually installed a certificate from a formal certificate registry, it is recommended that you choose the 'Do NOT regenerate' certificate option, or your existing certificate will be overwritten.
 - Installation of a custom certificate is covered in the section on 'Expert Configuration' within this guide.
 - Note that on the first run of the wizard after installation of the Snare Server, regardless of the state of the 'Do NOT regenerate' option, the wizard will upgrade the default 1024-bit certificate, with a more robust 2048 bit version at the conclusion of this step.
- It is recommended that Network Time Protocol be used on the Snare Server to provide a reasonable likelihood that the system date/time is less susceptible to hardware clock drift. The Snare Server will utilise the NTP server (IP address or DNS name) as a source for time information. If your organisation does not have an NTP server available on the local network, you may wish to choose a server from the list available from <http://www.pool.ntp.org/>
- Enter the Port your Snare Agents are listening on for their remote administration interface. This port will be used by the Agent Management Console to contact your agents. By default the port is 6161.
- Enter the Password set on the remote administration interface of your Snare Agents. It is used by the Agent Management Console, decrypting encrypted log messages, as well as retrieving such items as user and group retrieval from the agents.
- Snare is capable of delegating authentication to an external LDAP Directory or Active Directory server.
 - Note that the user must still have an account on the Snare Server with the same name as the LDAP/AD user, to log in.
 - Enter the IP address (or DNS name, as long as the Snare Server has been configured to use your local DNS) of your target LDAP or Active Directory server.
 - If specified, the Domain will be added to the end of the username for authentication purposes (eg: A username of 'auser' and a domain of 'test.local' will imply an LDAP/AD authentication of auser@test.local. Only 'auser' will be used locally on the Snare Server to determine access control settings).

3.2.2.5. Security Settings

- Enable or disable the SSH and FTP services.
- Enable or disable SMB (Windows Share) access to the main Snare data store.
 - The Snare data store is where the event logs are stored in compressed form. This area can be accessed as a read only

windows share, via userid/password authentication.

- Set a password that your remote Windows machine needs to use, to connect to the Snare Server. This will share out the Snare Archive directory in read-only format.
- The username is always 'snarearch'.
- The Windows share can then be accessed from your windows machine (or NAS box) as \\snare_server_IP_or_DNS\SnareArchive. For example, \\10.2.3.4\SnareArchive.
- Enable or disable the enhanced password security functionality for the operating-system-level accounts that are installed by default by the Snare Server.
 - By default, the Snare Server enables password complexity controls, account lockout (30 minutes after 5 failed password attempts), and password history checks. Normally, though the Snare Server system accounts are exempted from the more stringent requirements of an organisational security policy; particularly the requirement for password rotation. The accounts are generally used for either system administration or automated log transfers, and may not fit in with password rotation policies. Enhanced security and forced rotation can be enabled, or disabled via this setting, if required.
- Enable or disable password expiry in the Snare Server.
 - By default, the Snare Server does not expire passwords for user accounts. Forced expiry/rotation can be enabled, or disabled via this setting, if required.
- Enable or disable enhanced password expiry in the Snare Server.
 - PCI, and related regulatory compliance compatible password controls can be enforced by turning on this setting.
- Enable or disable the *Database Manager Interface* - unless instructed by your support team, this interface can be left disabled. It is not required for the normal operation of the Snare Server.
- Enable or disable the *Basic Snare Firewall*, which uses the UFW firewall to configure IPTables. For normal operation of the Snare Server, the firewall should be left enabled; it will only block those ports that do not have an associated snare-related service active.
 - The status of the firewall can be determined by logging into the server via SSH, and executing `sudo ufw status`.
 - More information on UFW can be found at: <https://help.ubuntu.com/community/UFW>
- Some security vulnerability scanners identify links to 'external sites' as reportable vulnerabilities. This setting turns off clickable links in the external link redirect page.
- Click on the 'Next ->' button.

3.2.2.6. Performance and Hardware Settings

- In situations where a workstation, or other client, has incorrect date/time settings, and is sending log data to the Snare Server significantly out of sync with the correct date/time, the collection subsystem can be configured to discard events that are older than a certain number of days.

- Note that date-based discard does introduce a small performance penalty for collection rates.
- Some syslog sources do not produce data that consistently includes readily identifiable information that will allow Snare to categorise the event as (for example), a PIX Firewall event, or an IBM SOCKS Server event. When the Snare Server definitively identifies an event from a particular IP address as belonging to a certain category, it can save off this information, and attempt to steer all future syslog messages from the same source IP address, to the same category. Some performance tweaks, such as query timeout, maximum audit events returned, and debug levels are available here. It is recommended that these values retain their default values, unless your Snare Server support team recommends a change.
 - By default, this capability introduces a small performance penalty for collection rates. It is highly recommended that this option be turned off, if a syslog client happens to send multiple types of syslog data to a destination Snare Server (e.g. PIX Firewall logs, and IPTables Firewall log data, and IBM SOCKS Server logs).
- For organisations with very specialised hardware, pre-caching generated log data can actually be a performance penalty. Although it is strongly recommended that this setting remain at it's default unchecked state, the configuration setting can be toggled here.
- If your server has an optical writer (CD / DVD) installed, you can select the preferred default device here. Click on the **Next ->** button. A final screen will be displayed, reminding you of the location of the Snare Server documentation.
 - This setting will be used by the automated data archive objective, if you choose to schedule it.
- If you have changed the server name, or have forced a regeneration of the Snare Server certificate, choose the 'Restart Apache, and return to the Snare Server' option, otherwise click on 'Return to the Snare Server'.

3.2.2.7. Additional Objectives

The Snare Server comes with baseline objectives suitable for a wide range of deployments. However, additional, special-purpose objectives can be downloaded from the InterSect Alliance web site, to supplement the defaults.

Several of the objective sets are also stored on your installation media, and are copied to the Snare Server during install. In situations where a direct path to the internet is unavailable, these cached objectives can be imported into the Snare Server.

In general, the objectives available from this page are either:

- Associated with the security and audit components of industry regulations such as PCI, NISPOM, or HIPAA, or
- Are newly developed, and have not yet been integrated directly into the default objectives distributed as part of a Snare Server release.

Objectives imported during this step, will be added to the 'Reports' area of the Snare Server, under a new folder

called 'Imported Objectives', and tagged with the date/time of import.

3.3. Snare Server Home Page

Once a user has been successfully logged into the Snare Server, the home page is loaded, as shown in the [Snare Server user interface - home page](#) image below. The [Section Overview](#) table below lists and explains each numbered component.



#	Name	Explanation	For more information...
1	Sections	These three sections provide access to the Reports, Status, and System areas of the Snare Server.	"Navigating the Snare Server User Interface"
2	Objective Actions	These icons allow you to select objective-related functions such as configuration, access controls, and scheduling. You can also regenerate the currently selected objective, access attachments, or log out of the Snare Server.	"Top Panel" "Working with Objectives"
3	Snare Logo / Home	This takes you to the Snare Server Home page.	
4	Objective Selection and Management	This area of the page displays the list of objectives available in each section. For objectives in the 'Reports' area, objectives can be created, cloned, rearranged, or moved into containers.	"Objective Navigation Panel" "Modular Objectives"
5	Main Document area	This is the main output area of the Snare Server where the information generated by objectives is actually displayed.	
6	Status Information	Informational status message may be displayed here by the Snare Server. You can also increase the size of the status area.	
7	Health Checker Notification	If the Snare Health Checker requires attention to resolve an issue, an animated icon will appear in this location (the bottom far right of the screen), which can be clicked to navigate to the Health Checker.	
8	Uptime Information	This text will display the current length of time that the Snare Server has been up and running, adjacent to the Health Checker Notification (if there are any alerts requiring attention).	

3.4. Overview of Sections

Version 7 of the Snare Server provides a dynamic, application-style point-and-click user interface, directly from your web browser. Each of the three section buttons displayed in the top-right corner of your browser (Reports, Status, System) corresponds to a range of capabilities, as discussed below.

3.4.1. Reports



Additional objectives can be:

- Created by Snare Server users, either from scratch, or by cloning and modifying an existing objective.
- Downloaded from the InterSect Alliance web site.

By default, the reports area will contain objectives relating to:

Active Scanning

- Example: Scan the local network, and conduct a network vulnerability analysis on hosts that are found.
- Example: Connect to the organisational border router and download the current configuration settings. Compare these settings to an authorised baseline configuration, and highlight any changes that have been made.

Application Audit

- Example: Display a list of inappropriate material that has been accessed through the organisational proxy server.
- Example: List users who have utilised the UNIX 'SUDO' command.

Network

- Example: Display a geographic map of IP addresses that have been denied access by the organisational Checkpoint Firewall.
- Example: Report on the top ten hosts that have initiated a port scan against the organisation, as reported by the gateway network intrusion detection system.

Operating Systems

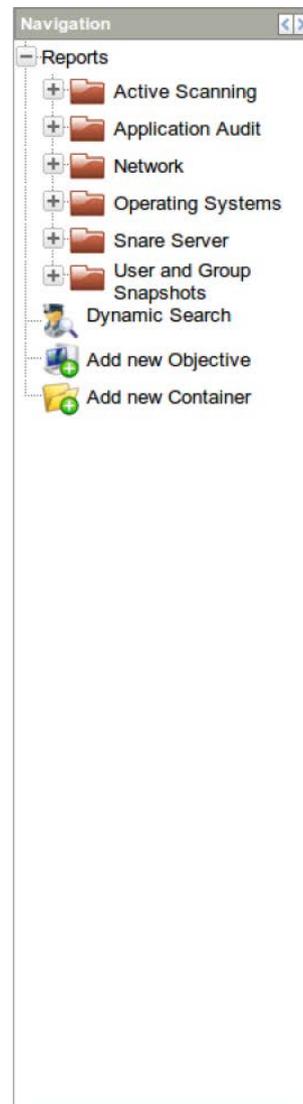
- Example: Generate a real-time alert when a user outside an authorised list, attempts to access a sensitive file on a Windows file server.
- Example: Send a daily email to security administrators, if the list of users in the Domain Administrators group changes.

Snare Server

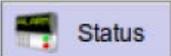
- Example: Display a report that shows users who have modified the configuration of any Snare Server objectives.

User and Group Snapshots

- Example: Based on the information provided by the Snare Agent for Solaris, produce a report showing any unauthorised members of the 'sensitivedata' UNIX group.



3.4.2. Status



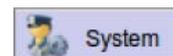
This section allows you to access information relating to the status of the Snare Server, including:

- General statistics on the type, distribution and volume of log data that currently resides on the Snare Server.
- An overview of the data that is currently coming into the Snare Server, in order to determine whether a newly installed agent is reporting to the server.
- General system information, relating to the hardware on which the Snare Server resides.
- Potential problems that the Snare Server has detected, and wishes to inform you of.

More details on the objectives available in "Status" are available further down in this document.



3.4.3. System



This section allows you to access functions that manage and maintain the Snare Server and its users, and also manage the configuration of Snare Agents that report to the Snare Server.

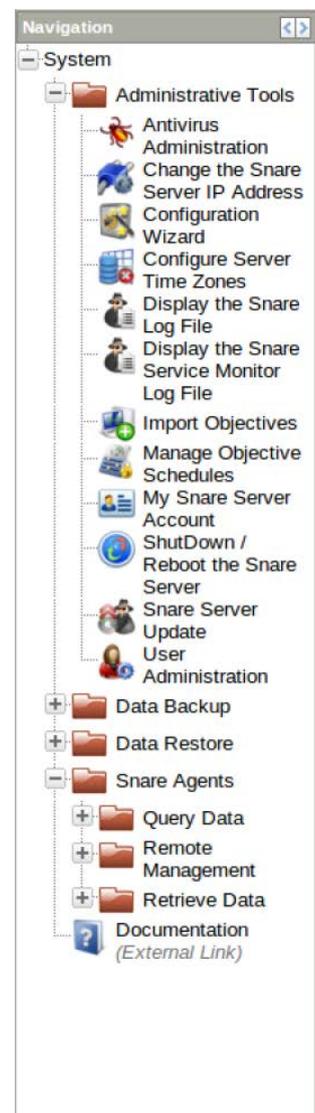
In the Administrative Tools section, you will find functions to:

- Manage the anti-virus installation, including the capability to update to the latest signatures.
- Modify system configuration settings such as IP address, DNS servers, and time zones.
- Display log files that may help the Snare Server support team to provide you with assistance.
- Update the Snare Server with new software, patch existing applications, or install new objectives.
- Manage users, and internal Snare Server settings.

Data Backup and Restoration allows you to archive Snare Server log data and objectives to optical media, or synchronise the data store to externally attached USB drives.

Snare Agents provides the ability to retrieve system data, such as Users and Groups, from the Snare Agents reporting to the Snare Server. It also provides a remote management interface for viewing and managing configuration on your agents in a central place.

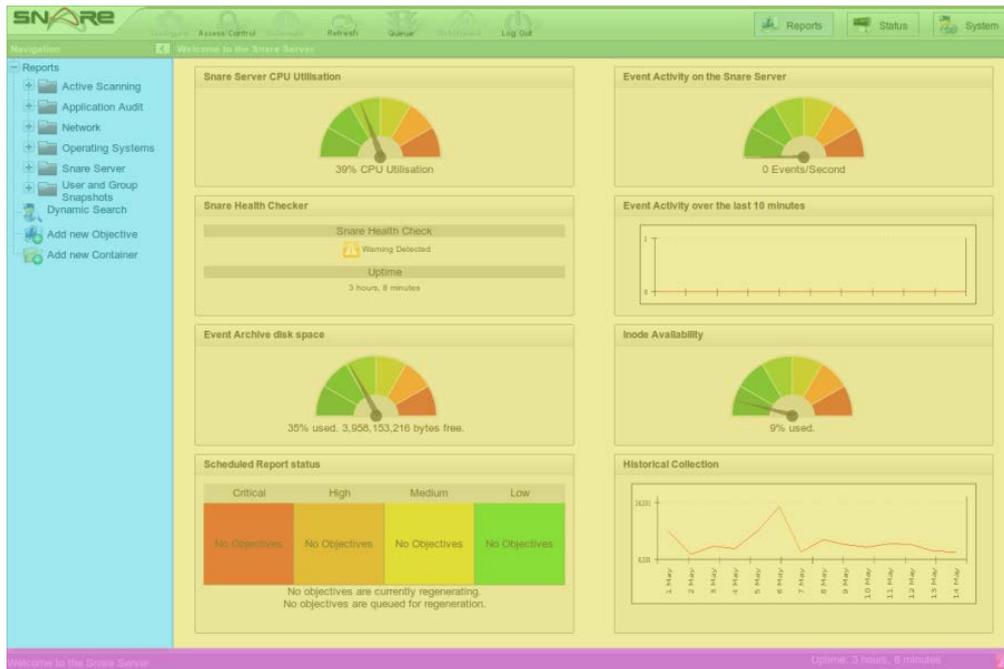
More details on the objectives available in the [System section](#) of this document.



3.5. Navigating the Snare Server User Interface

Snare Server employs a 'Web 2.0' style interactive application interface, employing drag and drop, popup dialogs, and dynamically updating data.

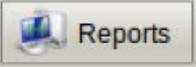
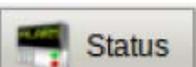
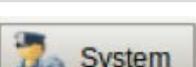
The Snare Server interface is generally divided into four 'panels', as shown in false-colour, in the image below.



3.5.1. Top Panel

The 'green' panel provides buttons for performing common functions and switching between the different navigation menus. Some of these buttons are greyed out when they are not available for use for the current objective.

- These buttons are, in order left to right:

	The Snare Logo, which takes you back to the dashboard.
 Configure	Modify the configuration of the currently open objective.
 Access Control	Change who can access, or modify the configuration of the currently open objective.
 Schedule	Configure the objective to regenerate on schedule, and modify the email distribution list.
 Regenerate	Add the currently displayed objective to the regeneration queue.
 Queue	Displays the currently queued and regenerating objectives.
 Attachment	Provides the option to download attachments generated by the currently open objective.
 Log Out	Logs you out of your current session on the Server Server.
 Reports	Switch to the Reports navigation menu mode.
 Status	Switch to the Status navigation menu mode.
 System	Switch to the System navigation menu mode.

3.5.2. Objective Navigation Panel

The 'blue' panel, also known as the 'Objective Navigation Panel' provides the ability to select individual objectives to be displayed. The behaviour of this section is slightly different, depending on whether you have chosen the 'Reports', the 'Status' or 'System' areas.

All three areas support a 'tree' style interface, where containers can be expanded or contracted to hide or reveal objectives. All three areas display objectives in natural alphabetic order, with containers prioritised before objectives. All three also provide the ability to expand, contract or collapse the navigation panel by using the left-right arrow icons that are displayed to the right of the navigation panel title. The 'Reports' area, however, also offers the ability to:

- Drag and drop objectives from one container to another.
- Create new containers, and objectives using the 'Add New Objective' and 'Add New Container' links.
- Clone, rename, delete, or modify the icons of objectives, by right-clicking on an objective. A pop-up menu will appear providing these options.

- Rename, recursively delete, or export the contents of a container, by right-clicking on a container. A pop-up menu will appear providing these options.
- Search for events using a search-engine style interface across multiple log sources, with 'Dynamic Search'.

A range of default objectives will be installed in the 'Reports' area for you by the Snare Server Installation process.

Dynamic Search



Dynamic Search may be used to quickly sift through information across multiple log sources, at the expense of completeness. The following filters are available for this tool:

- Find Events that contain: enter a string or event id
- Within the following date range: select from a date range or time period e.g. This Month
- Data Sources to Search: potential data sources which may be sending log data to the Snare Server e.g. WinSecurity, GenericSyslog
- Query Timeout (seconds): defaults to 60 seconds, but may be increased if searching on a larger subset of data sources or time range.

The Objective Navigation Panel can be partially hidden from view, by clicking the left-pointing arrow at the top-right of the navigation panel. The panel will be 'folded' into the side of the window. A small, right-pointing arrow will remain in place, to restore the panel to normal size.

The panel can also be expanded and contracted by smaller increments, using the left/right pointing arrows.

3.5.3. Status Panel

The 'purple' panel will display information relating to the currently selected objective, such as a summary of the objective configuration settings that have been modified, or the current progress towards completion while the objective is regenerating.

Over to the right-hand side of the area, the amount of time that the server has been running without reboot will also be displayed, and if the Snare Server Health Checker needs to inform you of an issue that requires your attention, an animated notification icon may also be displayed. This icon may be clicked on for further information.

3.5.4. Objective Panel

The 'yellow' panel is where objectives are actually displayed. When you select an objective from the 'blue' panel, this panel updates to show you the objective.

Many objectives display portions of the objective results in 'tabs' at the top of the page. These can be individually clicked to scan through the results. The type and function of these components is objective dependent, but will often include:

- A 'pattern map', which shows volumes of events, divided up into 15 minute segments for the reporting period.
- Tabular details, which displays a configurable proportion of the results.
- Various line graphs, bar graphs, port-maps, geolocation maps, or pie graphs.



Many objectives will also have interactive components that can be clicked to:

- Drill down for more information
- Page between results
- Sort data within a table

✔ In order to provide a modern, interactive user interface, the Snare Server utilises some features available only in more modern browsers. Users of Internet Explorer version 8 or prior, or Firefox 3 or prior, may experience slow response from JavaScript engines, poor quality graphics, or other degraded capabilities.

✔ You may notice a slight reduction in user interface performance once per hour, just after the new hour turns. Snare takes the opportunity to pre-cache results from your currently defined objectives on a regular basis, which can reduce interface response on single processor systems for a short time.

This feature can be turned off in the performance section of the Snare Server wizard.

Pre-caching can provide significant performance benefits for objectives that are generated weekly or at greater time periods. Objectives that are generated daily, may notice slight performance benefits. If your objectives are consistently being regenerated more frequently, having pre-cache turned on, may actually negatively impact the overall performance of your Snare Server.

4. Objectives - An Overview

4.1. What is an Objective?

An objective is a generic name for an interactive report, which performs a specific task or implements a set of analysis rules that are intended to derive useful information from event log data that is collected by the Snare Server.

In most circumstances, the term 'Objective' refers to the set of clickable items found in the 'Reports' section of the Snare Server - these are generally known as 'Modular Objectives'. However, the term 'objective' is also used interchangeably for items in the 'Status', and 'System' sections.

4.1.1. Modular Objectives

The objectives that are found within the 'Reports' section of the Snare Server user interface are known as 'Modular Objectives'. A modular objective is highly configurable, and generally includes:

- A query builder that allows you to create very complex search criteria, incorporating precedence, logical operations, and advanced matching capabilities.
- A 'Token' definition system that can pull fields contained within particular consistent patterns, out of a larger string.
- A range of potential output modules, such as 15-minute pattern maps, tabular event data, graphs, and so on.
- The ability to be scheduled to run on a regular, defined basis, and the potential to send output via electronic mail to data owners, system administrators, network administrators, and security administrators.
- Real-time reporting capabilities for events that match the search criteria.

Modular objectives are discussed in more detail below.

4.1.2. System and Status Objectives

The objectives found in the Status area of the Snare Server, generally provide overview information on total collection volumes and speeds, and checks associated with the health of the Snare Server and its associated agents. The System area provides access to objectives that perform general system administrative tasks, or facilitate agent management activities.

4.2. What is an Event?

Snare objectives create reports based on events that are generated by servers, workstations, applications or appliances. An event in Snare is a significant occurrence that is used to track or benchmark an organisation's performance or security. All events used in Snare will express the following properties:

The first characteristic of an event is that it occurs in the time scale. Time scale is very important as events are time dependent. All events in Snare are time tracked. In many objectives the organisation of events is wholly based on the time of the occurrence.

Example

A user login may be a normal occurrence during the day, however that user logging in at 12am in the morning may be considered unnecessary and thus is important to keep track of.

Events can be grouped, or referred to, based on the type of event.

Example

An event may relate to security, an application, a user or access to a certain resource. Each of these events can be grouped and organised for simple tracking. Snare uses these ideals to group Objectives.

Within an event group, the actual occurrence is given an event name using data from the actual event.

Example

A user may access a system and we need to know why or what happened during the access. This information is stored as a value, such as 'allowed access' or 'denied access' to a particular resource. This is important because we can then watch events based on being denied access to a security critical system.

Any system information at the time of the event can also be stored, and is Objective/event dependent. Shown below is a very generalized example of an event, and does not reflect the nature or structure of data stored by Snare.

Example

```
12:35am 19/09/03 # Snare Server Security : User login # Denied User  
'joe.bloggs' bad password.
```

4.3. Working with Objectives

Objectives on the Snare Server, whether modular, or otherwise, generally share two common features - the ability to set access controls, and the ability to configure the objective. There are some exceptions; for example, the 'System Status' objective within the 'Status' area, does not offer any configuration options.

4.3.1. Accessing Objectives

The objective navigation panel provides a tree-like view of the set of objectives that are available to you. Access controls, set by the Snare Server Administrator, may limit your view of objectives to a subset of those available.

To access an objective, single-left-click on either the text, or the icon associated with the objective you wish to display. The objective panel will update with the output from the objective as at its most recent regeneration point.



If an Administrator has added an objective to your access list, after you have logged into the Snare Server, hitting the 'Refresh Page' button on your web browser will make the new objective appear in the Objective Navigation area.

4.3.2. Configuring Objectives



Once you have accessed an objective by clicking on its icon or title within the objective navigation panel, the 'Configuration' button in the top panel can be clicked to modify the settings associated with the objective.

Once clicked, a new dialog will appear in the objective navigation panel. Configuration settings for modular objectives will be covered in more detail later in this document, but configuration dialogs will generally share the following common components:

- A title that tells you which objective you are currently changing.
- A series of form elements that will allow you to change settings associated with the objective.
- A "Set" button, to confirm the actions you have undertaken.
- A "Cancel" button, which will revert the objective configuration back to its previously saved state.



Hovering your mouse over the description field (the grey part to the left of the configuration setting options) will often reveal more information about the configuration settings on offer.

Configure Objective: General Statistics

Total Number of Agents to Display. Display all agents
 Display top 10 agents
 Display top 50 agents

Amount of time to display. Display all time
 Display last 12 weeks
 Display last 12 days

Set Cancel

4.3.3. Scheduling an Objective



Objectives can be configured to automatically regenerate on a periodic basis. Click on the 'Schedule' button in the top panel, in order to modify both the regeneration schedule for an objective, and also the users and/or groups who should receive an electronic mail message in the event that the objective produces data.

Objectives can be configured to regenerate:

- Hourly
- Daily
- Weekly
- Monthly
- Quarterly
- Yearly, or
- Once only, at a specified year, month, date and time (5 minute granularity).

In addition, each schedule configuration option has some additional flexibility available; for example, the 'Hourly' setting can be modified so that the objective always regenerates at 40 minutes past the hour. A 'Weekly' objective can be forced to regenerate every Tuesday, at 3:05 PM.

Each objective can have its own email distribution list. It is also possible to specify that emails are only sent out if there is something to report for that objective. Electronic mail can be sent either to all members of a Snare Server group, or individual recipients can be specified.

Set Objective Schedule

Schedule Never
 Once
 Hourly
 Daily
 Weekly
 Monthly
 Quarterly
 Yearly

Enable custom start time: 2012 May 21 Mon 15 10

Enable email distribution of report? On/Off
 No email when there is nothing to report

Email Distribution List

Default
Administrators
SuperUsers
leigh.purdie@intersectalliance.com

Tony.Ferris@dni.gov.au

Set Cancel

-  You can view, and change the schedules of all installed objectives by browsing to the 'Manage Objective Schedules' objective within the Administrative Tools container of the 'System' section.

4.3.4. Regenerating an Objective



In addition to scheduling an objective to be regenerated, a user can interactively submit the currently displayed objective for immediate regeneration by clicking on the Regenerate / Refresh icon. This will add the objective to the regeneration queue, and display the Queue dialog to track progress.

4.3.5. Objective Queue



The Snare Server objective queue dialog can be accessed by the 'Queue' icon in the top panel. It will also automatically appear when the currently displayed objective, is manually added to the regeneration queue by clicking the 'Regenerate' button as discussed above.

When an objective is in the regeneration queue, the Queue icon will also turn red to signify that one or more objectives are currently regenerating.



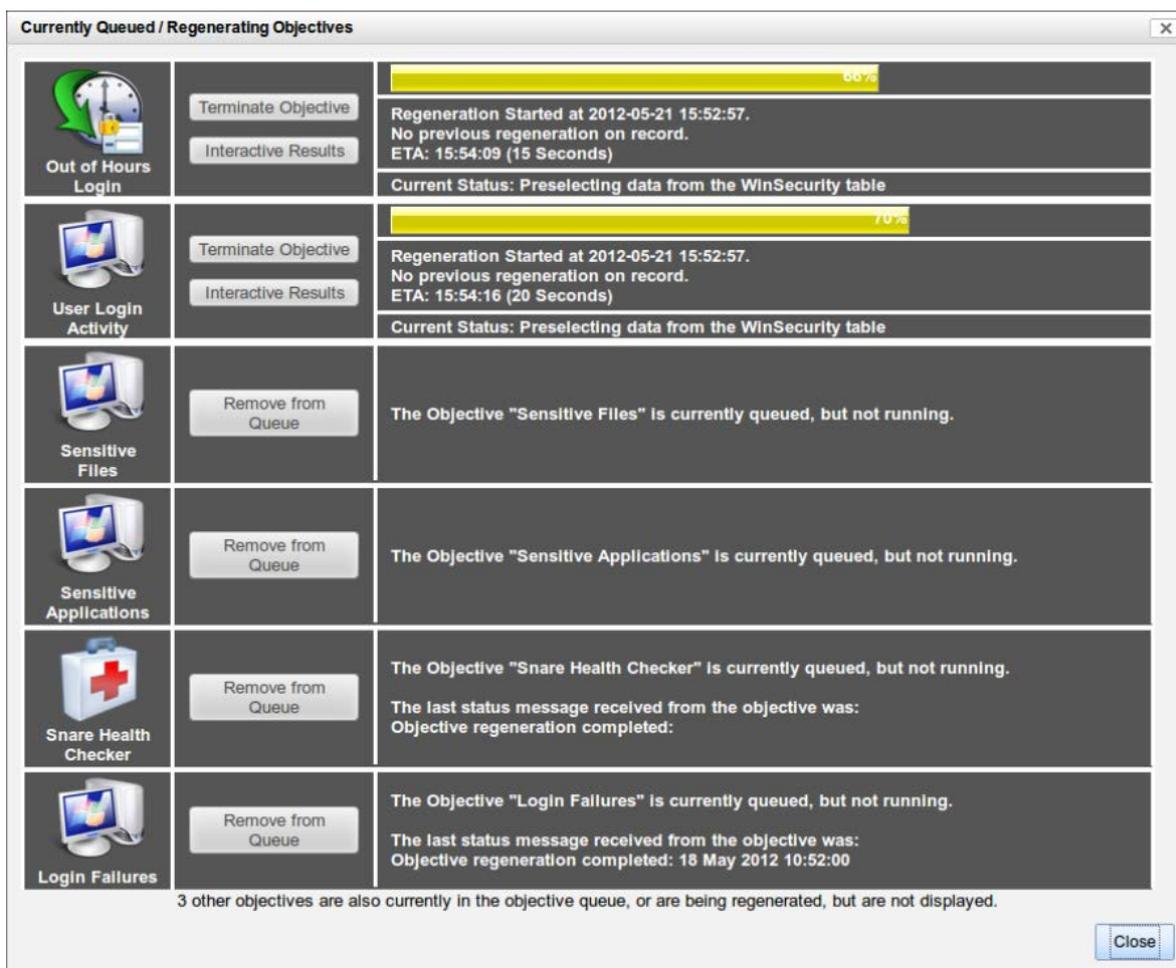
The Queue dialog lists the six objectives that are currently highest in the regeneration queue. For those objectives that are actively regenerating, the following information, and options, are presented:

- The objective title, and icon.
- A 'Terminate Objective' button, which will halt regeneration of the objective.
- For modular objectives, an 'Interactive Results' button, which will attempt to display a snapshot of the results that are currently being retrieved from the Snare Server data store.
- A progress bar that shows the approximate completion state of the objective.
- Information on:
 - When the regeneration process was started.
 - The time taken for the previous regeneration of this objective.
 - An estimated completion time (absolute, and elapsed).
- Any status updates delivered by the objective.

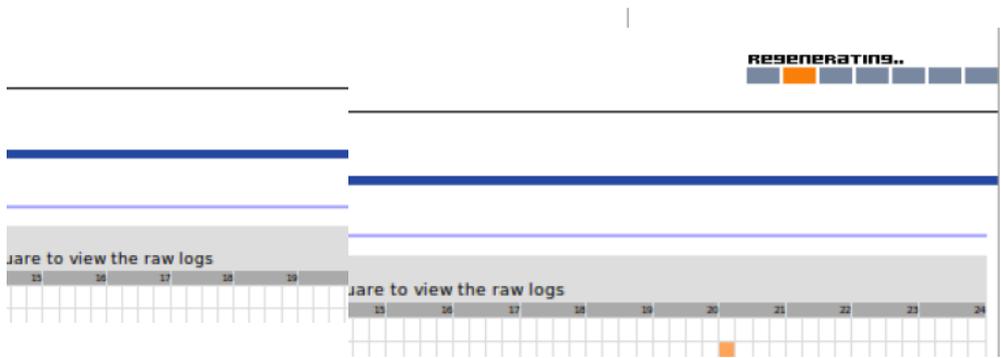
For objectives that are not yet regenerating, the following information, and options, are presented:

- The objective title, and icon.
- A 'Remove from Queue' button, which will delete the objective from the regeneration queue.

Information on how many objectives are currently in the regeneration queue, but are not displayed in the current dialog, is also available at the bottom of the dialog window.



- ✔ The objective queue dialog can be closed without interrupting the regenerating, or queued objectives.
- ✔ The Snare Server will attempt to concurrently regenerate a number of objectives commensurate with the number of CPU cores your system has available. A single-core CPU will generally only run a single objective at a time. A dual-core machine will run two objectives concurrently.
- ✔ When the currently displayed objective is either in the regeneration queue, or actively regenerating, a notification will also appear in the top-right-hand corner of the objective panel.





4.3.6. Access Control

Every objective created on the Snare Server can be individually secured so that only



authorized staff have access to it. Access is granted at group level; therefore, a user must be attached to a group in order to view or change an objective.

One of two levels can be granted:

- Write access. This provides a user with the ability to change the configuration settings for the objective.
- Read access. This provides a user with access to view the output of this objective, and also regenerate the objective.

In addition, users who create, or clone an objective, are identified as the owner of the objective. Both the owner, and Snare Server administrators have the ability to:

- Delete the objective, and
- Add new users to the objective.

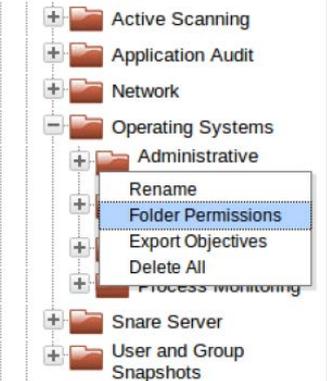
✔ Users and Groups are created and managed in the System -> Administrative Tools -> User Administration objective.

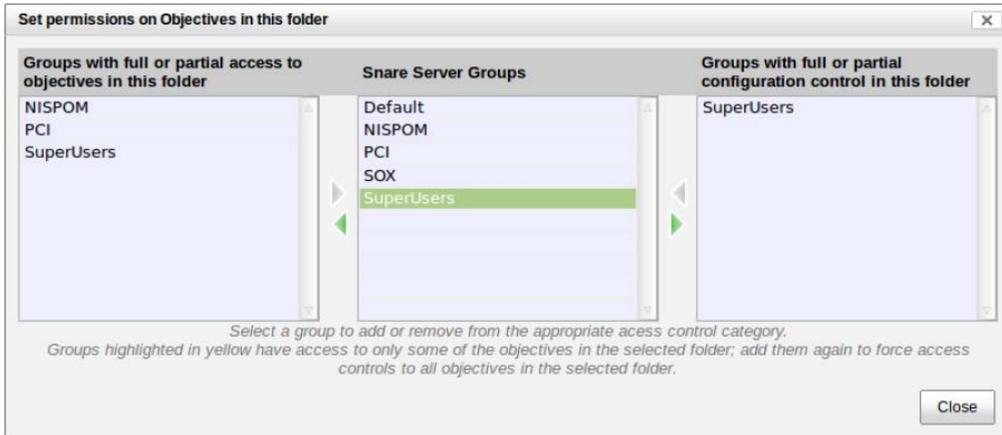
✔ **How To..**
How to change a group's access rights.

- In **Select a Snare group**, select the group. The information in 'Access Control Settings' will update.
- Select the appropriate access level check box.
- Click the **Set** button. This makes the new settings to take effect immediately.

In situations where access controls need to be applied to an entire folder of objectives, recursively, the 'Reports' navigation panel offers a 'Folder Permissions' menu option when you right click on a folder.

Selecting the "Folder Permissions" option will generate a dialog box that lists the Groups that are currently defined on the Snare Server, and provides the opportunity to add or remove groups from the 'Read' or 'Configure' capabilities.





4.3.7. Objective Documentation



Monitor access and changes to the Snare Server.

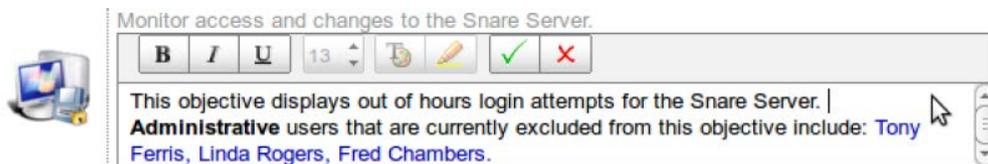
Generated: 28 May 2012 11:54:07

Objective documentation

is available at the top of the main objective output panel. By default, the objective will display text that has been either:

- Hard coded in the Snare Server, for the particular log source from which the objective derives its data, or
- Encoded with the actual objective, in the case of objectives that have been imported from the InterSect Alliance objective download area.

However, objective documentation can be added to, or modified, by those that have the ability to configure an objective. Double-clicking on the text of the documentation, will bring an editable field, that provides you with basic word-processor style functionality, such as font sizes, colours, and weights.



Clicking on the green 'Tick', will save the current documentation. The red 'cross' will cancel the current edits.

5. Modular Objectives

5.1. Overview

Modular objectives are the core of Snare's analysis capabilities. They are found within the 'Reports' section of the Snare Server user interface, and are highly configurable. They will generally include the following components:

- A query builder that allows you to create very complex search criteria, incorporating precedence, logical operations, and advanced matching capabilities.
- A 'Token' definition system that can pull fields contained within particular consistent patterns, out of an event of interest.
- A range of potential output modules, such as 15-minute pattern maps, tabular event data, graphs, and so on.
- The ability to be scheduled to run on a regular, defined basis, and the potential to send output via electronic mail to data owners, system administrators, network administrators, and security administrators.
- Real-time reporting capabilities for events that match the search criteria.

5.1.1. Objective Templates

Snare includes a range of 'templates' (often referred to as an 'Objective Type' in the Snare Server user interface) to make the job of a security administrator easier when crafting a new objective.

These templates are hard-coded in the Snare Server, may pre-define custom search criteria for you, will sometimes include custom code to perform tasks, and may be updated and expanded on each release of the Snare Server. More information on Objective Templates is available below.

5.2. Arranging

The Reports objective navigation panel provides an interactive tree, allowing you to not only view the objectives that are available, but also to rearrange objectives in a custom structure.

Although the Snare Server presents the objectives in alphabetical order (containers first, then objectives), you can:

- Create new containers in which to store objectives.
- Move objectives from one container to another.
- Move containers to another container, or back to the root of the tree.

✔ Rearranging the location of an objective, or container, will change the location for all users of Snare - not just your account.

✔ When you expand or contract a particular container within the Reports area, Snare will save this information off, so that the same settings will be applied next time you log in.

5.3. Creating

5.3.1. Creating a New Container

At the base of the Reports objective navigation panel, is the "Add new Container" link (item 1, in the navigation graphic to the right). Clicking this link will create a container called "New Container", which will be inserted into the navigation tree in the appropriate alphabetised position (item 2, in the navigation graphic to the right).

✔ A new container is a temporary item that only exists for two hours, and will not be visible to other users of the Snare Server. It will not become permanent, or visible to other users, until you add an objective to the container.

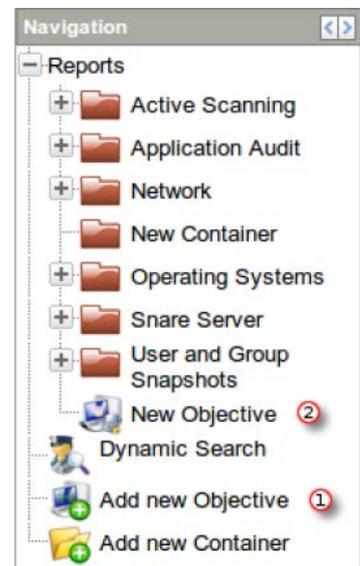


5.3.2. Creating a New Objective

Near the base of the Reports objective navigation panel, is the "Add new Objective" link (item 1, in the navigation graphic to the right). Clicking this link will create a new objective (called 'New Objective'), which will be inserted into the navigation tree in the appropriate alphabetised position (item 2, in the navigation graphic to the right).

By default, the new objective will be configured with very simple settings.

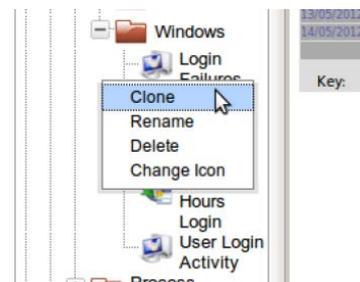
Once the objective is visible on your navigation panel, you can select it using your left mouse button, and change the configuration, access controls, or schedule settings to your requirements.



5.4. Cloning

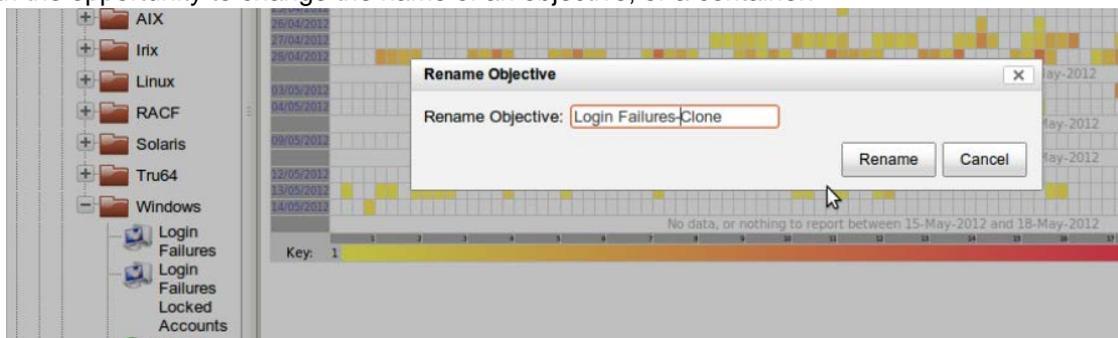
Right-clicking on an existing modular objective will raise a pop-up menu (otherwise known as a 'context menu'). From the menu, you can select the 'Clone' option in order to make a functional copy of the objective you selected.

Once you have clicked the clone option, the new objective will be added to the Reports navigation panel, with the name of the original objective appended with '-Clone' (eg: "Test Objective" will become "Test Objective-Clone"). A new dialog will appear in the main objective display panel, giving you the opportunity to rename the objective.



5.5. Renaming

From the pop-up menu that appears when you click your right-mouse-button, the 'Rename' option will provide you with the opportunity to change the name of an objective, or a container.



Enter a new name for the objective or folder, and click the 'Rename' button to complete the process.

✔ Objectives each have a unique 'Objective ID'. Since it is the objective ID that is used by the Snare Server to differentiate objectives, you can potentially have two objectives with exactly the same name, that have different configurations, access controls, and scheduling. Although Snare will be happy to allow this, in order to limit confusion, it may be worth avoiding this practice.

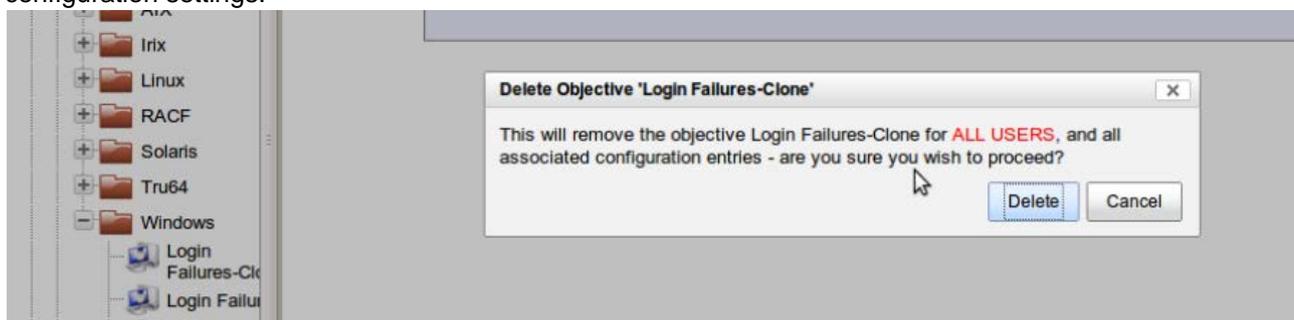
✔ Renaming a container takes a little longer than renaming an objective, since Snare has to recursively search through the contents of the container, and modify the path of each objective contained therein.

5.6. Removing

5.6.1. Individual Objectives

When you choose the 'Delete' option from the context menu, a dialog will appear, notifying you that the objective will be removed for ALL USERS of the Snare Server, and will ask for confirmation before proceeding.

Selecting the 'Delete' button from the dialog, will remove the objective, and associated objective configuration settings.



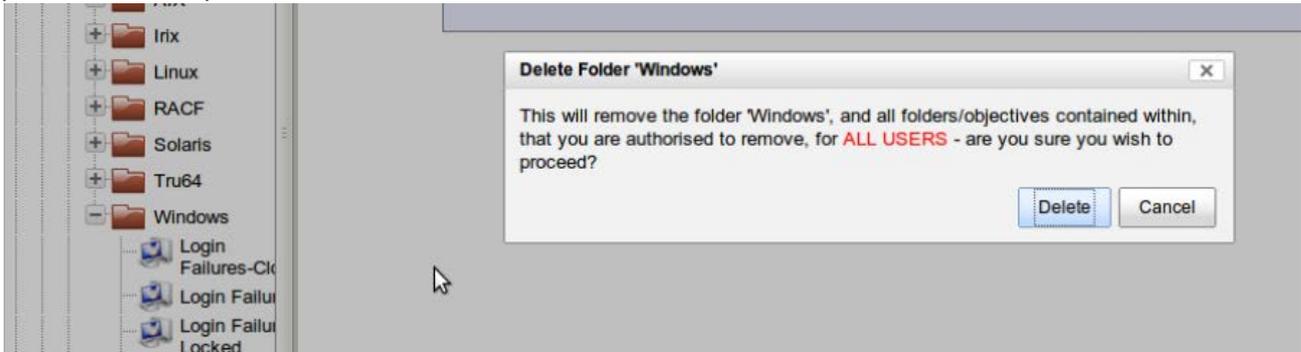
5.6.2. Containers

Right-clicking on a container, will allow you to remove all objectives within the container that your Snare Server user account has permission to remove.

A dialog will appear, notifying you that the objectives will be removed for ALL USERS of the Snare Server, and will ask for confirmation before proceeding.

In a situation where you have chosen to remove a container, but you do not have permission to remove

some or all of the underlying objectives, the Snare Server will check each objective for authorisation, and only remove those that you are authorised to delete. In this case, the original container will remain after the process has completed.

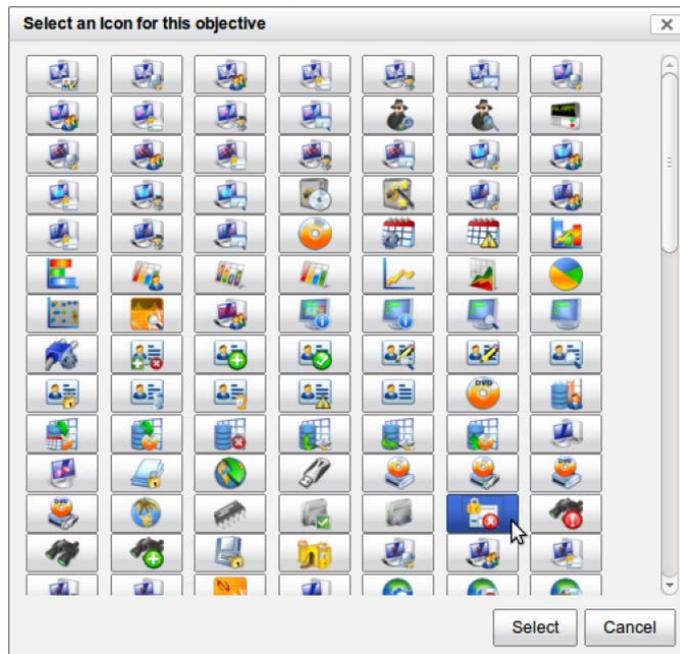


5.7. Icon

Snare generally selects an icon for an objective by examining (in descending priority order):

- The icon associated with the objective from which the current objective has been cloned.
- The icon for the 'Objective Type' (eg: Windows logins) from which the objective is descended.
- The icon for the 'log type' (eg: Windows Security) that the objective scans.

However, you can set a specific custom icon for an objective by choosing the 'Change Icon' option from the objective context menu. A dialog will appear on the main objective panel, that provides a selection of icons. Choose the 'Select' button to finalise the selection.

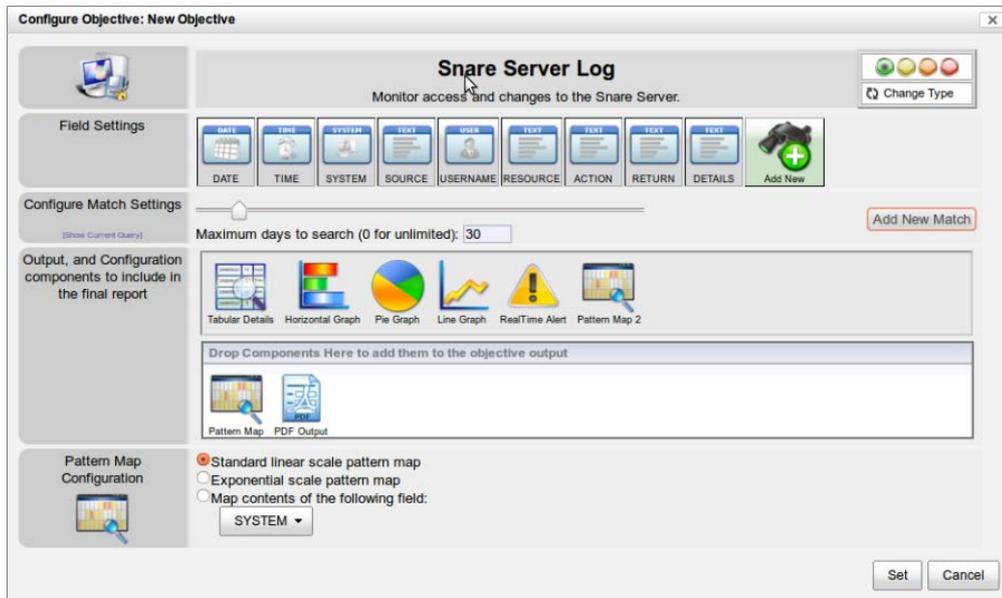


6. Modular Objectives - Configuration & Output

Snare modular objectives begin life extremely simply. As you add more components, and more complex match settings, Snare will enable greater flexibility, and more configuration options.

i Example

The simple configuration dialog shown below scans the "Snare Server Log" data source, for any events produced over the course of the last 30 days, and displays a 15 minute 'Pattern Map' of the resulting data. A PDF has also been added to the output component list.

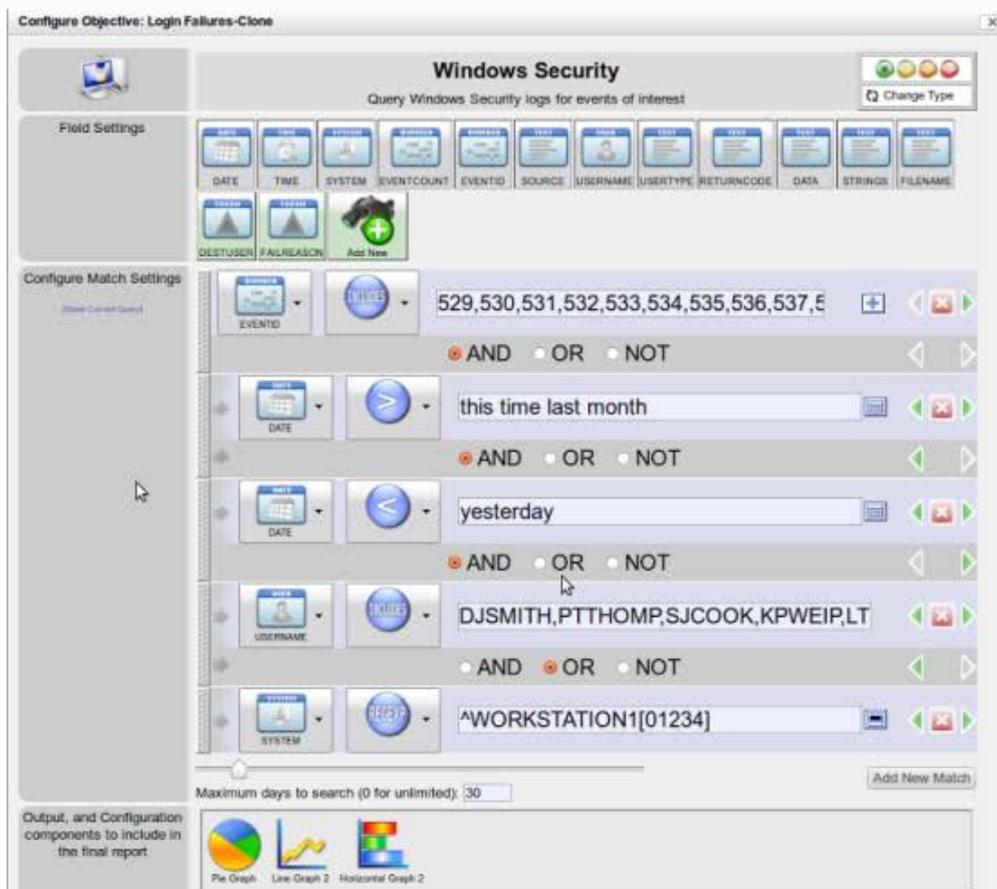


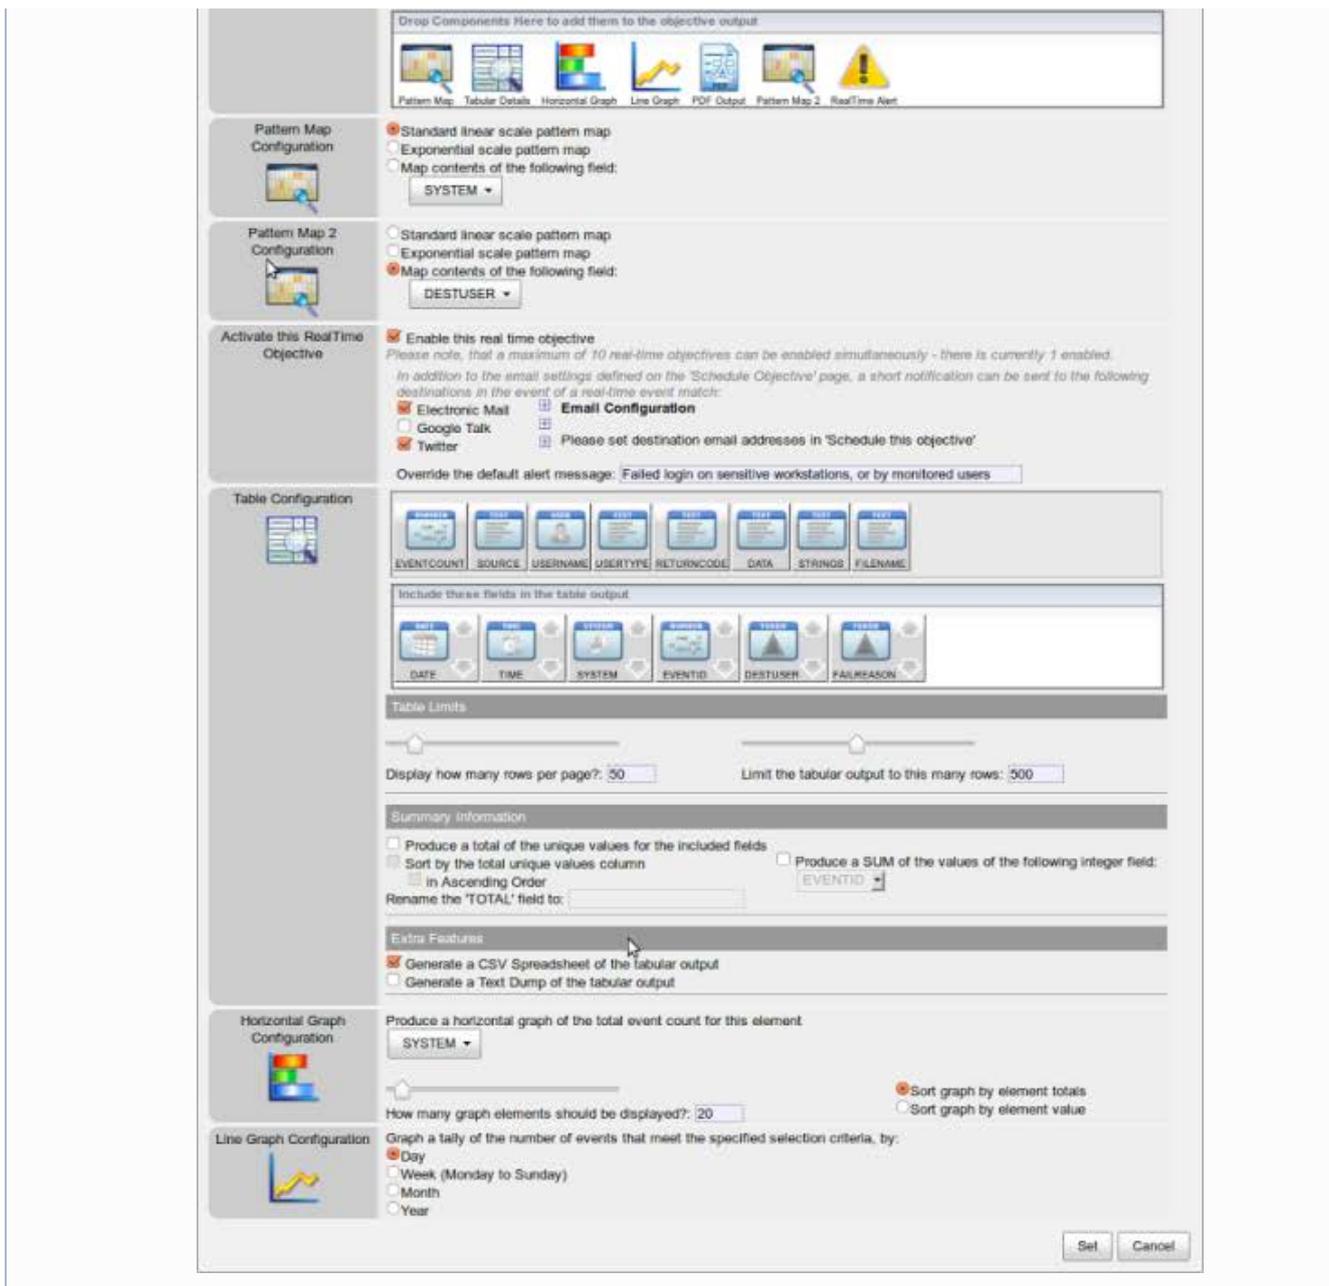
i Example

A more complex objective is introduced below, as an indication of how flexible and comprehensive the Snare Modular objective query and output builder can be. The objective:

- Defines search criteria that looks for any events that match:
 - Dates between this time last month, and yesterday, non-inclusive.
 - Events related to login failures.
 - For EITHER the users DJSMITH, PTTHOMP, SJCOOK, KPWEIP, or LTROMA) OR from any source system starting with the word 'WORKSTATION', followed by a number that starts with 10,11,12,13 or 14.
- Defines two additional 'tokens':
 - 'DESTUSER' scans the 'STRINGS' field of Windows Security logs, for a target account name, which is pulled out of the event using the following regular expression search:
 - (?<=logon to account:| User Name:| Target Account Name:|Account For Which Logon Failed: Security ID: Account Name:)(.?) *(?=by:|Domain:|Target Account ID:|Account Domain:)
 - 'FAILREASON' also scans the 'STRINGS' field of Windows Security logs, for information on why a login failure occurred, using the following regular expression match:
 - (?<=Reason:|User Account)(.) *(?=Status:|User Name:|Domain:| *Target Account Name:)
- Outputs a standard 15 minute pattern map, using yellow and red to show parts of the day where low or high volumes (respectively) of event data have been discovered.
- Outputs a pattern map highlighting occurrences of each destination user account name.
- Sets the objective up as a real-time objective, sending the output to the email list defined in the 'Objective Schedule' area of the objective. The title of the email has been configured to be "Failed login on sensitive workstations, or by monitored users".
- Outputs a Table of the first 500 results that match the criteria outlined above, 50 results per page, displaying only the fields Date, Time, System, EventID, DestUser, and FailReason. A CSV dump of the table will also be produced.
- Outputs a horizontal graph of the top 20 systems that match the search criteria for the objective.
- Outputs a line graph of events-per-day that match the search criteria for an objective.

Although the information above, and the image below, are likely to be quite overwhelming when first encountered, this document will explain each section in more detail.





6.1. Objective Header

The objective header displays:



- The icon that is currently associated with the objective.
- The objective title (Failed User Logins), the data source it currently interrogates (Windows Security), and the documentation assigned to the objective.
- A modular objective configuration management panel, described below.

6.1.1. Criticality

An objective can be assigned a criticality level by clicking on the green, yellow, orange or red radio buttons. If the objective has any information to report in any of the modular output components, the objective will be tagged with the appropriate colour in the objective navigation panel.

✔ An objective tagged with a 'green' criticality will retain the default 'black' writing when it is displayed in the objective navigation panel.

✔ The navigation panel will not refresh immediately in response to the change in criticality status for an objective. Generally, the updated status can be seen on next login, but it may be sooner if you, or another Snare Server user, modifies an objective or container name, or position, in the objective navigation panel.

6.1.2. Objective Type

Snare includes a range of 'templates' (often referred to as an 'Objective Type' in the Snare Server user interface) to make the job of a Snare administrator easier when crafting a new objective. These templates are hard-coded in the Snare Server, may pre-define custom search criteria for you, will sometimes include custom code to perform tasks, and may be updated and expanded on each release of the Snare Server.

A list of the templates included in the Snare Server is available in the '*Modular Objective Templates*' chapter, but here are some representative samples:

- Report whenever a user attempts to access a sensitive file on a Windows file server.
- Notify administrators when a particular Solaris user attempts to run a command.
- Show modifications to permission flags, for ACF2 accounts.
- Show all attempts to gain access to the root account on AIX systems.
- Compare the current CISCO PIX or Router configuration to an authorised version.
- Display events related to electronic mail delivery, for Gauntlet firewalls.
- Search syslog data for attempts to use the 'sudo' or 'su' commands to escalate privileges.
- Search IPTables firewall logs for dropped packets that have a source address of a non-routable IP block.
- Highlight attempts to port-scan a NetScreen firewall.
- Show attempts to change the configuration of a Nortel VPN Router.
- Monitor attempts to access RACF resources.
- Highlight failed authentication attempts on a SOCKS server.
- Display results from the Snort network intrusion detection system.
- Report on inappropriate material accessed through the corporate proxy server.
- Show out-of-hours login access for Windows systems.

i Example

A Windows failed login template, will pre-define a match setting that looks for events that contain an EventID of **529, 530, 531, 532, 533, 534, 535, 536, 537, 539, 644, 681 or 4625** - all of which indicate a failed login event. If Microsoft adds a new failed login event to Windows, a future version of the Snare Server will update the windows failed login template so that existing objectives also pick up the new information.

i Example

A Windows successful login template, also defines two new modular features: "User Flags", and "Event Threshold". When the "User Flags" module is added to the objective, it allows users to be included or excluded from the final report, based on whether the account is disabled, locked, or has the "Don't expire password" or "Password cannot change" flags set. The objective will derive this information from scans performed by appropriate Snare Agents.

Once the "Objective Type" button is selected, a new dialog will appear in the objective window.



General objective template categories are displayed in the tree-menu to the left. Once an objective category is selected, a list of available objective types will be displayed in the right-hand section of the dialog. Click on the appropriate 'Select' button to choose a template.

A checkbox is available at the bottom left hand side of the dialog window, which will hide categories for which there is no event data on your Snare Server.

6.1.3. Unlock Objective

Objectives that are based on a pre-defined objective template can be 'unlocked', in order to change the pre-defined match settings, but once they are unlocked, they may no longer include the custom tokens, or custom modular components, and will not have their components upgraded by the InterSect Alliance team with new releases.

Objectives that are 'locked' can still have additional match settings and tokens, added to the mix. Match settings and tokens are explored further below.

6.2. Field Settings

The field settings section displays a list of the 'fields' that are available to use in your search criteria, and also as input fields for modular output components.



Snare also allows you to 'break apart' an existing field, and place the resulting sub-string into a field with a new name; this is known as a 'Token'.

6.2.1. Tokens

If a small part of an existing field needs to be captured for further analysis, or reporting reasons, a new token can be defined by clicking on the green 'Add New' button.

A new dialog window will appear, which will allow you to configure your new token.

Add/Modify Field X

Field Name

Configure the Field What field contains the information you are interested in extracting?

Search Criteria

This functionality uses perl-compatible regular expressions, to specify a pattern that represents this field.
 If, for example, you were trying to extract a username from the following string:
 Alter Details: User: fred@mycompany.com.au Computer Machine1234
 The following search & selection criteria would search for the term 'User:' followed by 0 or more whitespace characters, and will then pull out any alphanumeric characters (including dots, dashes and '@' symbols) that occur in sequence, until a non-matching character is found:
 User:\s*([A-Za-z0-9@!-]+)
[Regular Expression Tester](#)

i "Field Name" defines the name that you wish to assign to the new field.

"Configure the Field" asks you to select the source field that contains the information you are looking for.

"Search Criteria" asks you to define the regular expression that will be used to pull the substring out of the field content.

✔ A regular expression is a complex, but extremely powerful tool, that will facilitate flexible matching, and extraction of substrings. We will cover regular expressions in more detail below, but in Snare, they take the general form:

- **Text to match before substring** (Substring match) **Text to match after substring**

So, for example, assume the DETAILS field of an event includes the following string:

```
Safend Protector File Logging Alter Details: User: george Computer: Machine1234 Client GMT: 8/10/2011 4:13:23 AM Client Local Time: 8/10/2011 3:13:23 PM Server Time: 8/10/2011 4:13:30 AM Group: Policy: policy Device Description: Disk drive Device Info: Kingston DataTraveler 2.0 USB Device Port: USB Device Type: Removable Storage Devices Vendor: 13FE Model: 1D00 Distinct ID: ID001 Details: File Name: file.pdf File Type: pdf File Size: 50945 Created: 8/10/2011 3:13:23 PM Modified: 8/10/2011 2:23:44 PM Action: Write
```

In order to capture the user (highlighted in bold and red) from the above string, the regular expression would need to look for the word after the "User: " sub-string, that is composed of alphanumeric characters (with the addition of the '@' symbol).

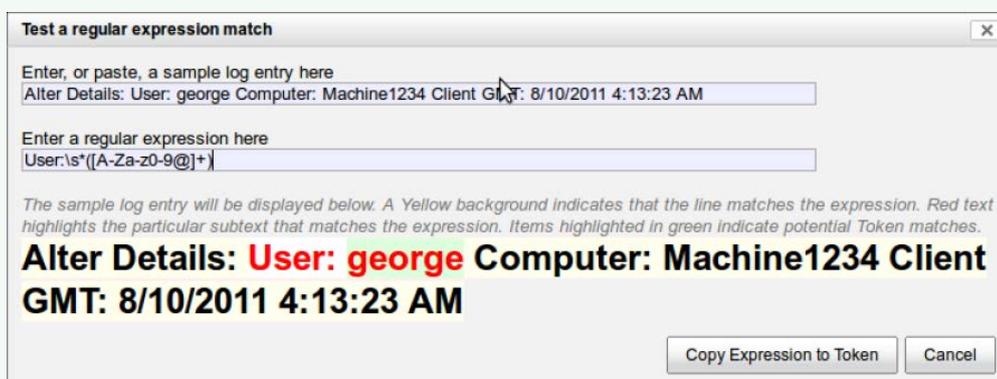
The token required to achieve this looks like:

Field Name: USER
Source Field: STRINGS
Search Criteria: User:\s*([A-Za-z0-9@]+)*

This translates as: look for a "User:" sub-string, then 0 or more white spaces, then anything after that which contains 1 or more letters, numbers, or an @ symbol. This is then a valid token according to our search criteria.

Tokens, once created, are then treated as if they were a normal field, and can be filtered, grouped, sorted, or used as a target field in any modular output component that uses fields (eg: Graphs or Tables). This creates a powerful mechanism to effectively query sub-strings which are contained within a much larger string. Any number of tokens can be created which allows for a variety of choices when querying strings within strings.

A regular expression tester is also available, which can assist you with the process of creating a token; it can be accessed by clicking the 'Regular expression tester' link near the base of the token definition dialog.



If the expression you are using has a match somewhere in the sample log entry, it will be highlighted in yellow. Red text indicates the area of the sample that exactly matches your token expression, and a section highlighted in green shows the actual substring that will be pulled out by the token.

Once you are happy that the regular expression meets your requirements, you can copy the expression back to your token with a click of a button, rather than copying/pasting the information from the dialog.

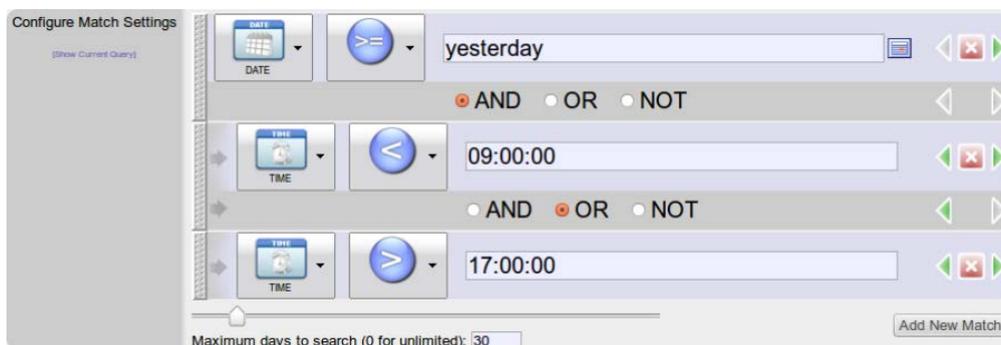
i Example

Regular expression samples:

- `from=<([>]+)`
 - search for the word 'from' followed by an equal sign, and a less than symbol. Retrieve any characters after that, until you encounter a greater-than symbol (>)
 - `from=<john@somewhere.com>`
- `(?<=Account Name:)(.*?)(?= New Account| New Domain| Target Domain| Account| Domain| Caller)`
 - Retrieve any text between the strings "Account Name: " and one of either " New Account", " New Domain", " Target Domain", " Account", " Domain" or " Caller"
 - **Account Name:** DNIUserLP **Target Domain:** DNI
- `open(([^)]+)`
 - Search for the string "open(", and retrieve any characters after that until you encounter a close bracket.
 - `open(O_RDONLY|O_RDWR)`
- `^.....(.....)`
 - Ignore the first 13 characters of this field, and retrieve the next 10.
 - **INFORMATION:** SUCCESS USER: JABLOGGS
- `^. {12}(. {9})`
 - Same as the above example, but using the regular expression 'repeat' function: Ignore the first 13 characters of this field, and retrieve the next 10.
 - **INFORMATION:** SUCCESS USER: JABLOGGS
- `^. {12}([a-zA-Z0-9]{9})`
 - Same as the above example, but cuts out the extra spaces after "SUCCESS", by limiting the valid retrieved characters to alphanumeric only.
 - **INFORMATION:** SUCCESS USER: JABLOGGS

✔ Tokens that you have created, or can modify, are highlighted in green. Tokens that are part of an underlying objective template, and are therefore locked, will be highlighted in red.

6.3. Configure Match Settings



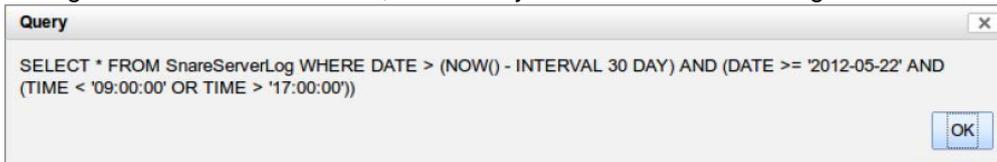
Snare's query builder is a flexible tool that allows you to create very complex search criteria, incorporating precedence, logical operations, and advanced matching capabilities.

6.3.1. Show Current Query

Although Snare does not utilise a database back-end for data storage, queries created with the Snare query builder are translated into SQL syntax, and passed through a database translation layer.

Selecting the '[Show Current Query]' link in the title panel, will pop up a new dialog that displays the SQL query

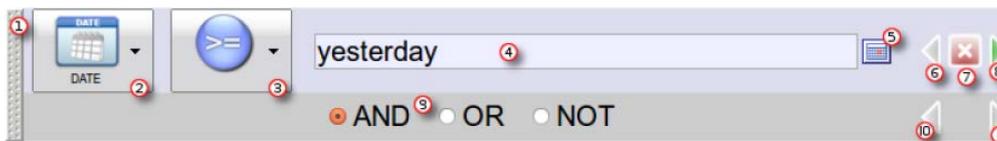
that would be run against the Snare datastore, based on your current match settings.



6.3.2. Adding a New Match

Selecting the 'Add New Match' button will append a new row to your existing match settings. By default, the new row will use 'Date' as the target field, ">" as the comparison operator, and the input field will be initially blank.

6.3.3. Match Row Components



6.3.3.1. Drag and Drop grab bar

Each match row can be moved up or down, and positioned before or after other match criteria. Click and hold the grab bar, and drag the match row, to rearrange. Snare evaluates matches from top to bottom.

6.3.3.2. Field to use

Select the field to use for your search criteria, by clicking on this button. A drop-down menu will appear, that will detail the fields that you can choose from.

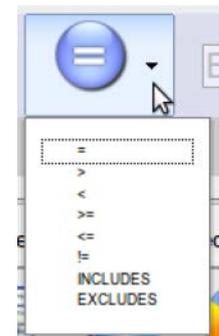


- ✓ Snare breaks up event logs into a series of fields for you, when the event arrives at the Snare Server. As described in the section on 'Tokens' above, you can also choose to create meta-fields that represent a predictable portion of a larger field. These tokens will appear in the drop-down menu after you create them.

6.3.3.3. Comparison operator

The comparison operators available for selection depend on the field you have chosen. Numeric, date, and time values will have the following comparison operators available:

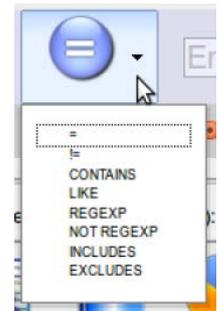
- Equals (=)
- Greater than (>)
- Less than (<)
- Greater than or equal to (>=)
- Less than or equal to (<=)
- Not equal to (!=)
- Includes
 - You may include several comma-separated values in the input field - eg: 1,2,3,7,9
- Excludes
 - You may include several comma-separated values in the input field - eg: 1,2,3,7,9



String values will have the following comparison operators available:

- Equals (=)
- Not equal to (!=)
- Contains
 - This will search for a simple case insensitive substring

- Like
 - Implements a SQL LIKE operator. LIKE uses the 'percent' sign for wildcards - so for example, a search for "%login%failed%" will match the string "attempted login for user 'fred' failed at 17:23:01"
- Regexp
 - Implements a perl-compatible regular expression search. As highlighted [above](#), regular expressions are complex, but extremely powerful and flexible string search functions.
 - **Tip:** Snare co-opts the "start of string" and "end of string" characters ("^" and "\$" respectively) to refer to the start of the contents of the field you are currently operating on, and the end of the field, rather than referencing the entire line.
- Not Regexp
 - Excludes all fields that match the supplied regular expression.
- Includes
 - You may include several comma-separated values in the input field - eg: fred,jim,tony
- Excludes
 - You may include several comma-separated values in the input field - eg: fred,jim,tony



6.3.3.4. Input field

A flexible input field that allows you to specify search criteria based on your field and comparison operator.

Note that Snare also allows you to compare two fields, rather than entering an actual value here. If, for example, you had two username fields (*USERNAME*, and *TARGETUSER*) in your data source, and you wished to return events in the query where the two were not equal, you could:

- Set **USERNAME** as your 'Field to Use'
- Set **!=** as your 'Comparison Operator'
- Enter **TARGETUSER** in the Input field, surrounded by two '@' characters:
 - **@TARGETUSER@**

These two '@' symbols, indicate to Snare that the contents of the input field refers to a "Field to use" as highlighted above, rather than a static comparison value. The '@' symbols will be removed, and processed by Snare. Tokens are supported, and the following comparison operations are valid:

- =
- !=
- >
- <
- >=
- <=

✓ Some fields allow you to specify indirect values. The 'Date' field, for example, generally takes arguments of the format "YYYY-MM-DD", but values such as the following are also valid, and will be reinterpreted each time the objective runs:

- last week
- next thursday
- -7 weeks
- first day of next month
- last day of april 2012
- first day of last month
- 3rd april this year + 20 days

The 'Time' field generally prefers times in the format "HH:MM:DD", however it also accepts standard integer values. A number in the 'Time' field will be interpreted as "Now minus 'x' minutes", where 'x' is the value you entered, and 'Now' is the time at which the objective runs. For example:

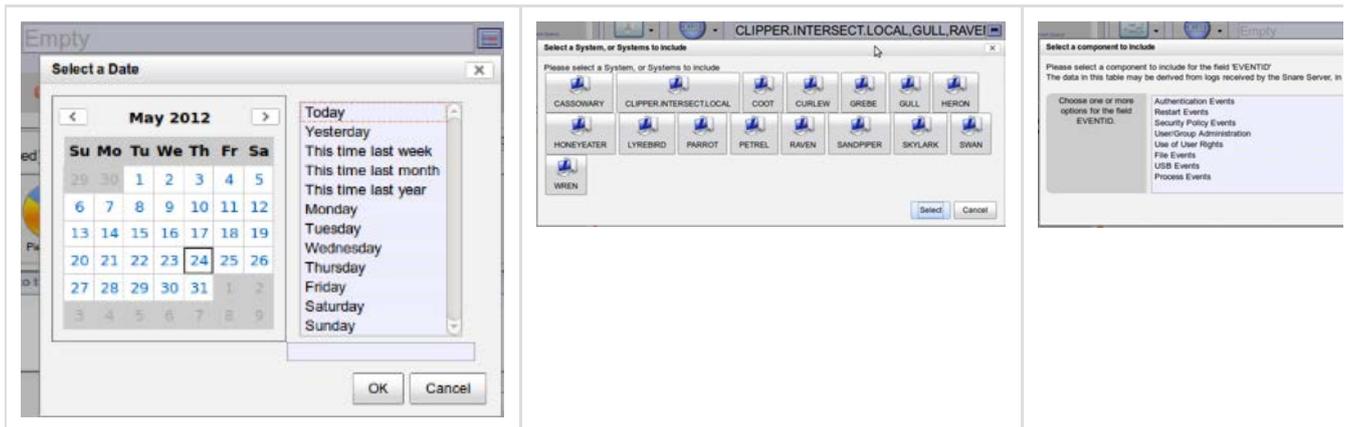
- **TIME >= 120**
 - This will be interpreted by Snare as: "**TIME is greater or equal to NOW minus 120 minutes**".
 - If the objective runs at 10:05am, Snare will calculate the value to be 08:05:00
 - The match term will then be extrapolated to "**TIME IS greater or equal to 08:05:00**"
- Note that TIME >= 120, does **NOT NECESSARILY** equate to "in the last two hours". Since Snare extrapolates the value to TIME >= 08:05:00, then if you have set your objective to report on data from the last 30 days, it will mean your objective will display data from ANY DAY that matches your other match terms, between 08:05:00 and 23:59:59. Data between 00:00:00 and 08:04:59 will not be displayed by the objective. If you specifically need to ONLY show data within the last two hours, you would need to add a second match term: DATE = today, or alternatively, set your "Maximum Days to Search" slider to "1 day".
- Note also that the logical sign direction is somewhat counter-intuitive if you do not understand how the Snare Server extrapolates the time value. It could be argued that "<= 120 (minutes)" would be better syntax for "in the last two hours"; however, that would assume that a value in the time field also automatically controls the date field; this is not the case. Please be cautious with your logical sign direction.

6.3.3.5. Contextual selection button

This button appears for some fields, and provides you with the ability to quickly select either:

- A range of values that have been captured from the last 30 days' worth of log data, or
- A range of common, or recommended, values.

✓ Selecting multiple values will generally turn on the 'INCLUDES' comparison operator, if you have not already selected 'INCLUDES' or 'EXCLUDES'.



6.3.3.6. Decrease match row indentation

The Snare Server query builder is able to implement explicit operation precedence for your match terms, by using indentation of match and logic rows. In this way, groups of match terms can be joined using a variety of logical operations such as AND and OR.

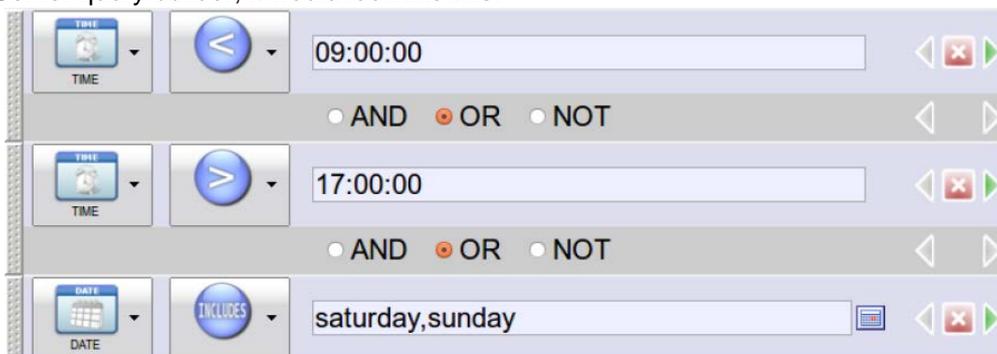
If you are familiar with the use of brackets in mathematical operations, an increase in indentation is analogous to opening a bracket. A decrease of indentation is analogous to closing a bracket.

Many queries do not require row indentation. For example, if we were trying to implement a simple 'outside of work hours' query, we would want to look for events that occur outside of 9am to 5pm (for example), or that occur on a weekend.

If we were writing this in mathematical, or SQL-like notation, we would use the following:

```
IF TIME < '09:00:00' OR TIME > '17:00:00' OR DATE = 'saturday' OR DATE = 'sunday'
```

In the Snare Server query builder, it would look like this:

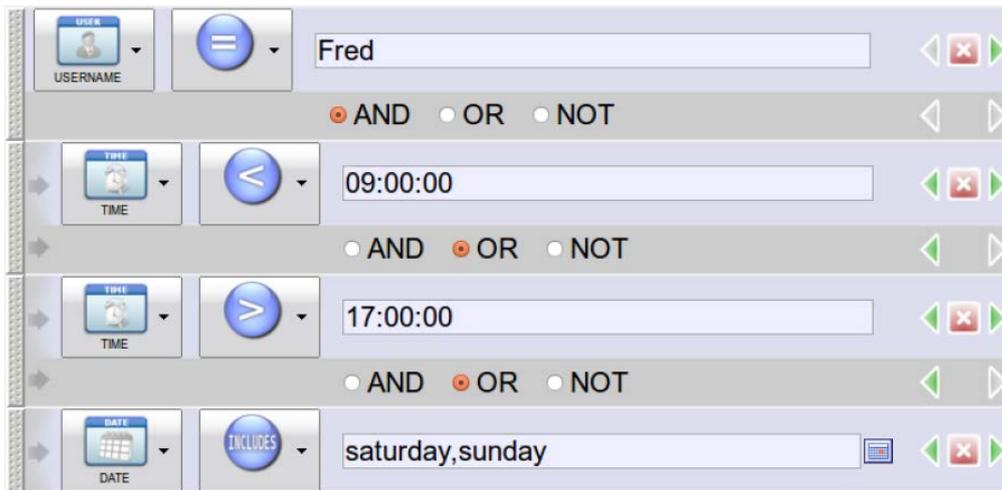


However, if we wanted to ONLY report on out-of-hours events that are tagged with the username "Fred", then we would need to do something a little more complex, and add in operational precedence / brackets. Our new string would look like this:

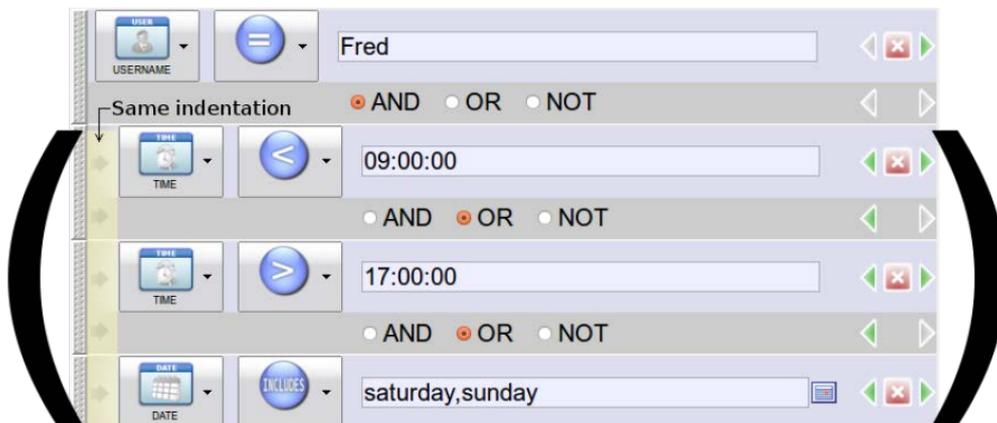
```
IF USERNAME = "Fred" AND (TIME < '09:00:00' OR TIME > '17:00:00' OR DATE = 'saturday' OR DATE = 'sunday')
```

Note the addition of the **AND**, and the **brackets**. This would mean that both the date/time component AND the UserName component had to be matched, for an event to be reported. If an event occurred on a weekend, but it was not for the user Fred, it would not meet the criteria we specified.

Recall that increasing our indentation factor for a particular row, is equivalent to opening a bracket. In this case, we would need to add the 'UserName = Fred' match row at the top of the match rows, and then increase the indentation of every row after the logic element associated with the UserName match:



In effect, because all of the rows with the same indentation (as indicated by the number of exposed arrows to the left of the row) are 'grouped', they are enclosed within the same bracket group, as illustrated by the following image.



By utilising the indentation of match and logic rows, complex logical precedence operations can be designed.

✔ It is sometimes easy to overlook an un-indented logic row. Remember that an increase, or decrease in indentation for both match AND logic rows, will affect your bracket placement.

✔ The 'decrease row indentation' arrow will be enabled (green) or disabled (grey) depending on the current row indentation, and whether it is possible to decrease the indentation of the row any more.

6.3.3.7. Remove match and logic row

Removes the current match row, and the associated logic row, from the Snare query builder. Snare will ask you for confirmation that you wish to remove the row, and then remove the appropriate match from the configuration settings.

6.3.3.8. Increase match row indentation

Increases the indentation for the associated match row. An increase of indentation is analogous to opening a bracket.

6.3.3.9. Logical operators

Choose from "AND", "OR", or "NOT" (which translates to "AND NOT").

6.3.3.10. Decrease logic row indentation

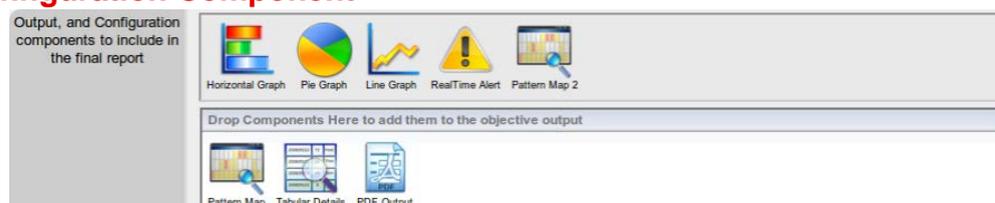
A decrease of indentation in the logic row, like the match row, is analogous to closing a bracket.

6.3.3.11. Increase logic row indentation

An increase in indentation in the logic row, like the match row, is analogous to opening a bracket.

6.3.4. Output and Configuration Component

Output and configuration components can be dragged from the top half of this section, into the bottom half, titled "*Drop Components Here to add them to the objective output*". This will result in either:



- A new output component being added to the report (for example, a line graph, or a table of log data), or
- Additional configuration settings displayed, that can be applied to the objective (for example, adding the option to only display users who have an expired account, on the Windows domain).

The number, and type of output components, depends significantly on the data source that is being interrogated.

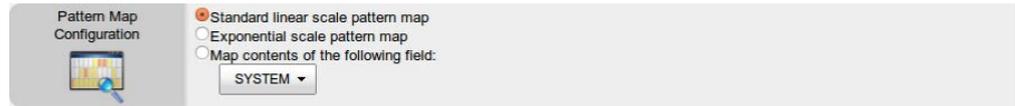
Example

- Objectives that scan firewall or router related event log data may include a 'geolocation map' that attempts to draw a line from the approximate geographic location of a source IP address, to destination IP addresses, on a map of the world.
- Objectives that scan proxy log data may include special output components that measure bandwidth by target site or user.
- UNIX systems that use the concept of a 'home directory' may make a special output component available that highlights attempts by users other than the owner, to access a home directory.
- Objectives focused on network intrusion detection system log data may introduce a 'target port map' output component, that can visually map attempts to port scan a corporate network.

Most components, when added to the objective, will also create a 'configuration panel' that allows you to control the output of each component. A '15 minute pattern map', for example, will provide the option of using a standard linear colour scale for the output, an exponential colour scale that highlights different ranges of data, or even a visual map of a particular target output field. Some of the more common output components are highlighted below.

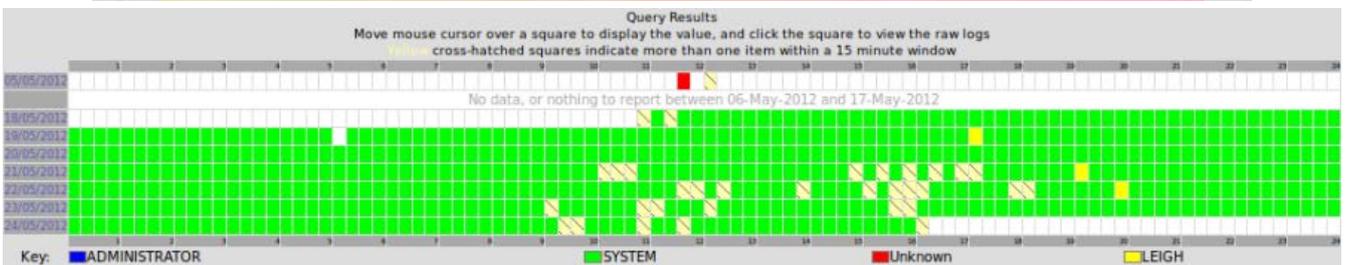
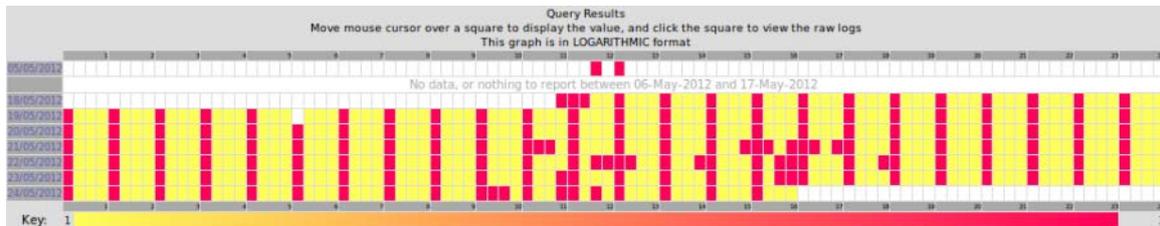
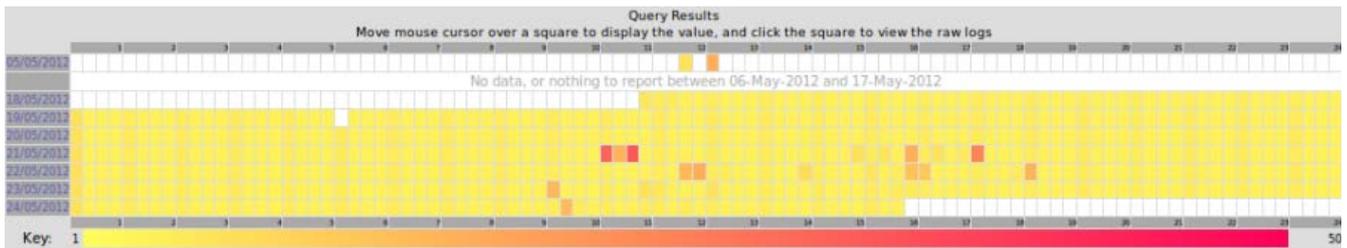
-  Some components, when dragged to the drop area, will reveal a second version of the same component in the drag section (eg: Pattern Map, and Pattern Map 2). As such, an objective can have two copies of many components, with slightly different configuration settings applied to each.

6.3.5. Pattern Map



The 15 minute pattern map provides a visual overview of event log data, displaying an indication of the volume, or contents of each separate 15 minute segment within the reporting period, as a colour selected from an appropriate area of graduated scale.

The pattern map can be configured to use a standard scale, an exponential scale, or to map the contents of a particular field. Exponential mode can highlight particular patterns that are difficult to see in the standard colour mode.



Each element of the pattern map can be clicked on, with your left mouse button, to search for the data that comprises that particular 15 minute segment. A new dialog will appear in the objective panel that shows the underlying data. The data can be sorted by clicking on a column header.

✔ Sorting on 'Date' will sort on both Date and Time. Selecting 'Time' will only sort on the Time column.

Event Details

<< first < prev 1 next > last >>

DATE	TIME	SYSTEM	ACTION	PROTO	SRCINT	SRCADDR	SRCPORT	DSTINT	DSTADDR	DSTPORT
2012-05-12	04:03:11	BUDGERIGAR	deny	TCP	dec0	194.94.252.3 Germany	5861	lo0	129.165.93.173 United States	139 (NETBIOS Session Service)
2012-05-12	04:07:03	BUDGERIGAR	deny	UDP	dec0	129.15.87.28 United States	8724	lo0	129.127.76.34 Australia	49 (Login Host Protocol (TACACS))
2012-05-12	04:08:17	BUDGERIGAR	deny	TCP	dec0	141.149.176.224 United States	3754	lo0	129.127.76.34 Australia	53 (Domain Name Server)
2012-05-12	04:11:45	BUDGERIGAR	deny	UDP	dec0	129.15.31.170 United States	1287	lo0	129.15.6.81 United States	68 (Bootstrap Protocol Client)
2012-05-12	04:12:37	BUDGERIGAR	deny	UDP	dec0	129.15.87.28 United States	8989	lo0	129.127.76.34 Australia	135 (DCE endpoint resolution)
2012-05-12	04:13:04	BUDGERIGAR	deny	TCP	dec0	153.96.180.2 European Union	1116	lo0	129.15.6.127 United States	80 (Web Server)
2012-05-12	04:13:10	BUDGERIGAR	deny	TCP	dec0	153.96.180.2 European Union	7511	lo0	129.63.44.1 United States	25 (Simple Mail Transfer)
2012-05-12	04:13:48	BUDGERIGAR	deny	TCP	dec0	68.90.156.118 United States	4505	lo0	129.15.95.244 United States	23 (Telnet)

<< first < prev 1 next > last >>

Close

Clicking on a date, to the left of the pattern map, will attempt to generate a table listing all events for that particular day, that match the objective search criteria.

✔ For high volume sites, this process may take a long time to complete.

6.3.6. Table

Table Configuration

SOURCE DETAILS

Include these fields in the table output

DATE TIME SYSTEM RESOURCE ACTION USERNAME RETURN

Table Limits

Display how many rows per page?: 50 Limit the tabular output to this many rows: 500

Summary Information

Produce a total of the unique values for the included fields
 Sort by the total unique values column
 in Ascending Order
 Produce a SUM of the values of the following integer field:
 No Values Available

Rename the 'TOTAL' field to:

Extra Features

Generate a CSV Spreadsheet of the tabular output
 Generate a Text Dump of the tabular output

To include a dump of event data that matches the search criteria specified for an objective, the 'Tabular Details' modular component can be dragged into the inclusion list.

Fields that should be included within the table output can be dragged from the top half of this section, into the panel titled "Include these fields in the table output".

When a field is dragged into the inclusion list, ascending and descending sort buttons will appear next to the field. Sort criteria is evaluated left to right. Fields can be reordered within the inclusion list in order to modify the sort output. The order of the fields in the inclusion list, will also define the order that they appear in the tabular output component.

DATE	TIME	SYSTEM	ACTION	PROTO	SRCADDR	DSTADDR	DSTPORT
2012-04-25	00:14:47	BUDGERIGAR	deny	TCP	137.226.144.3 European Union	129.15.6.127 United States	80 (Web Server)
2012-04-25	00:12:54	BUDGERIGAR	deny	TCP	129.15.87.28 United States	129.15.95.244 United States	2600 (zebrasrv)
2012-04-25	00:07:54	BUDGERIGAR	deny	TCP	194.94.253.3 Germany	129.15.6.127 United States	67 (Bootstrap Protocol Server)
2012-04-25	00:04:36	BUDGERIGAR	deny	TCP	129.15.114.131 United States	129.113.241.247 United States	369 (rpc2portmap)
2012-04-25	00:00:37	BUDGERIGAR	deny	TCP	137.226.144.3 European Union	129.15.95.244 United States	139 (NETBIOS Session Service)
2012-04-25	00:27:46	BUDGERIGAR	deny	TCP	129.15.114.131 United States	129.205.199.153	6346 (gnutella-svc)
2012-04-25	00:24:50	BUDGERIGAR	deny	TCP	129.15.114.131 United States	129.63.44.1 United States	80 (Web Server)
2012-04-25	00:22:12	BUDGERIGAR	deny	TCP	129.15.114.131 United States	129.239.2.107 United States	25 (Simple Mail Transfer)
2012-04-25	00:17:25	BUDGERIGAR	deny	UDP	194.94.253.3 Germany	129.113.241.247 United States	68 (Bootstrap Protocol Client)
2012-04-25	00:42:42	BUDGERIGAR	deny	TCP	194.94.252.3 Germany	129.15.6.127 United States	194 (Internet Relay Chat Protocol)
2012-04-25	00:38:19	BUDGERIGAR	deny	TCP	194.94.253.3 Germany	129.165.93.173 United States	80 (Web Server)

By default, the table will display a subset of the data that matches the objective search criteria. The default settings are 500 rows, at 50 rows per page.

- ✓ The table width will be set to the size of your browser window, minus a small space around the border of the table. For small screens, this can mean that long lines 'squash up' into very narrow columns, and you see very few lines per page.

You can make your entire table bigger by scrolling up to the top-right corner of the table, grabbing the very top-right edge (click and hold your left mouse button), and dragging your mouse off to the right hand side, beyond the boundaries of your page (ie: to the right hand limit of your browser window, or beyond). This will increase the size of the table beyond the visible area of your browser window, and a new scroll-bar will appear at the bottom of your browser window. You can then rearrange the width of each column as appropriate, and each line will take up less vertical screen real-estate.

Results can be 'grouped' to produce a tally of events that contain common field values. In order to activate this, choose the fields that should participate in the 'group', and add them to the table field inclusion list. Select the checkbox next to "Produce a total of the unique values for the included fields", in the "Summary Information" section of the table configuration component. You may also choose to sort by the total unique values column, and potentially rename the column from the default "TOTAL" to something that better represents your data.

For example, based on the table output screenshot above, if you wanted to analyse the most common destination ports by date, you could add fields 'Date', 'Proto', 'Action' and 'DstPort' to the field inclusion list.

Table Configuration

ADDRESS ADDRESS SYSTEM TIME TEXT PORT TEXT
DSTADDR SRCADDR SYSTEM TIME SRCINT SRCPORT DSTINT

Include these fields in the table output

DATE PROTOCOL TEXT PORT
DATE PROTO ACTION DSTPORT

Table Limits

Display how many rows per page?: 50 Limit the tabular output to this many rows: 500

Summary Information

Produce a total of the unique values for the included fields Produce a SUM of the values of the following integer field:
 Sort by the total unique values column in Ascending Order DSTPORT
Rename the 'TOTAL' field to: Port Count

Extra Features

Generate a CSV Spreadsheet of the tabular output
 Generate a Text Dump of the tabular output

After the objective regenerates, your table displays a new column, which shows how many events share the same Date, Protocol, Action and Destination port, out of all events that match the search criteria specified in the objective.

Pattern Map Tabular Details Line Graph Destination Port Map Geolocation Map

<< first < prev 1 2 3 4 5 6 7 8 9 10 next > last >>

DATE	PROTO	ACTION	DSTPORT	PORTCOUNT
2012-04-29	TCP	deny	80 (Web Server)	368
2012-04-29	UDP	deny	80 (Web Server)	205
2012-05-14	TCP	deny	80 (Web Server)	136
2012-05-09	TCP	deny	80 (Web Server)	135
2012-05-16	TCP	deny	80 (Web Server)	130
2012-05-05	TCP	deny	80 (Web Server)	127
2012-05-02	TCP	deny	80 (Web Server)	125
2012-05-17	TCP	deny	80 (Web Server)	123
2012-05-18	TCP	deny	80 (Web Server)	123
2012-05-01	TCP	deny	80 (Web Server)	122
2012-05-08	TCP	deny	80 (Web Server)	121
2012-04-26	TCP	deny	80 (Web Server)	120

Fields that are numeric, or tokens that are derived from a numeric field, will also be included as an option under the 'Produce a SUM of the values of the following integer field'. This feature is useful in situations where you wish to know information like:

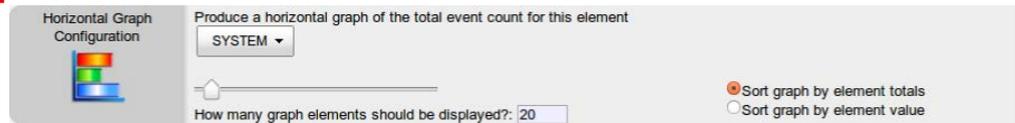
- Who are the top 10 users of bandwidth, through our corporate proxy server or firewall? (ie: Produce a sum of 'Bytes' per-user or per-IP)

✔ SUMMED column values will respect the sort criteria you have attached to the original field. If you ask Snare to produce a SUM of the 'Bytes' field, for example, and have chosen to sort Bytes in descending order, the SUMMED values will be sorted in descending order.

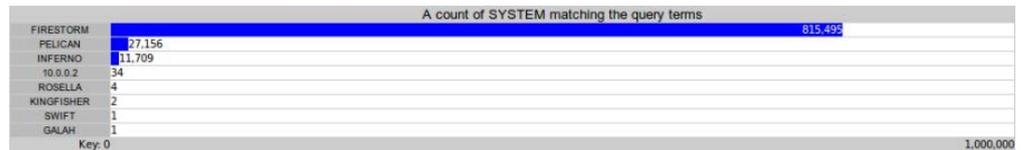
CSV (Tab delimited) and text dumps of the table data can also be produced. These will be available as an attachment to the objective.

6.3.7. Horizontal Graph

Horizontal graphs can be created by adding this element to your modular inclusion list. Select a field to use as a basis for the graph by choosing an option under "Produce a horizontal graph of the total event count for this element".

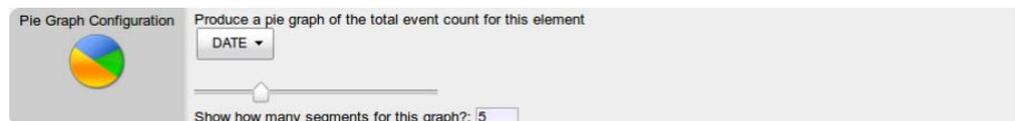


The number of graph 'rows' to be included can be defined, and you can also sort the graph by either the total count (descending), or by the actual field value (alphanumerically).



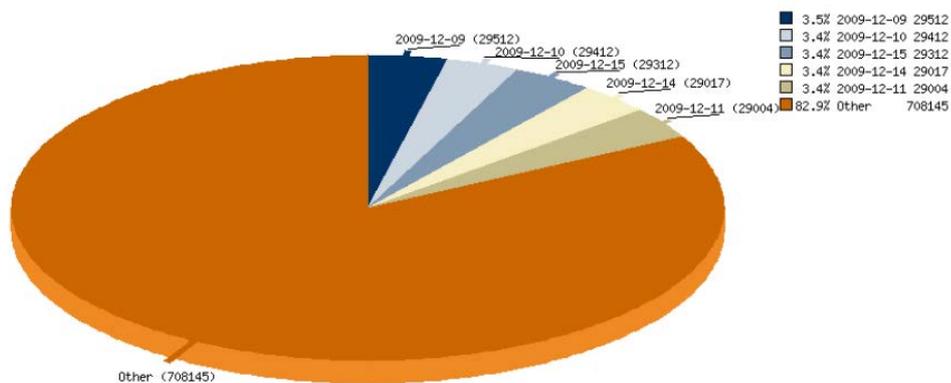
6.3.8. Pie Graph

Pie graphs can be created by adding this element to your modular inclusion list. Select a field to use as a basis for the graph by choosing an option under "Produce a pie graph of the total event count for this element".



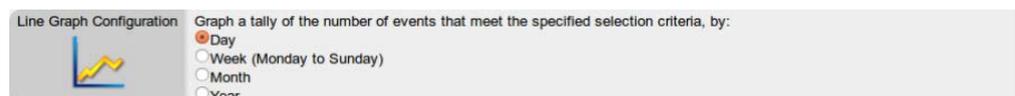
You can specify the preferred number of segments to be shown in the pie graph. If these segments do not represent 100% of the returned results, an additional 'Other' segments will be displayed on the pie graph.

Pie Graph of the event field 'DATE'

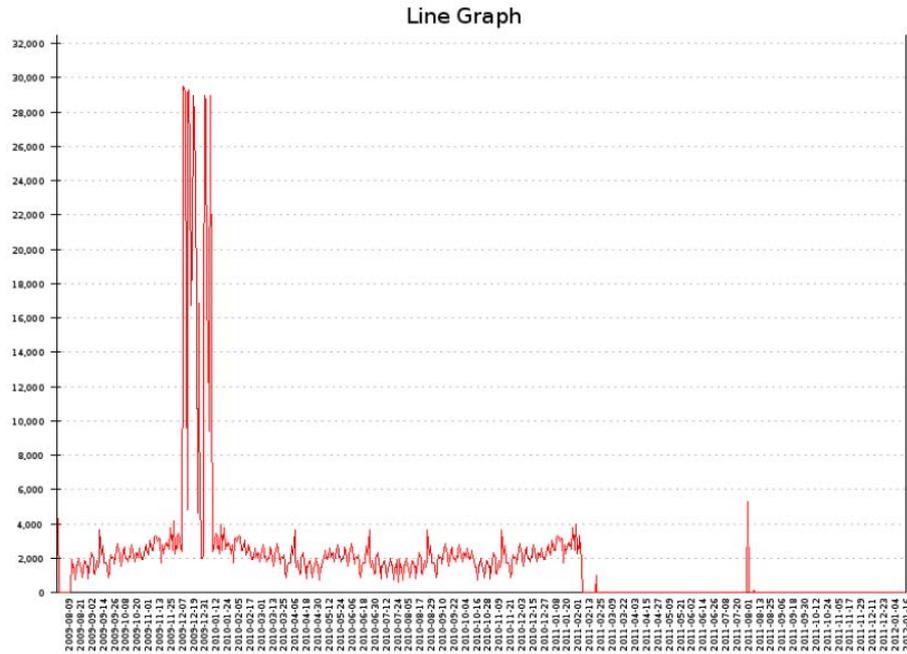


6.3.9. Line Graph

A line graph of total events can be created by adding this element to your modular inclusion list. You may specify that events should be graphed by Day, Week, Month or Year.



Line Graph of Events, by Day



6.3.10. PDF Output

To include a PDF of the objective output, add this component to the modular objective inclusion list. The PDF will be available from the 'Attachments' button in the top panel, and will be included with any electronic mail messages that are sent as a result of this objective being regenerated.

6.3.11. Real-time Alert

Activating real-time alerts for any objective activates a module in the collection subsystem,

Activate this RealTime Objective Enable this real time objective
Please note, that a maximum of 10 real-time objectives can be enabled simultaneously - there is currently 1 enabled.
 Send Email Alerts.
Alerts are sent to the User Account(s) or Email address(es) specified in the Objective Schedule options.
Override the default alert message:

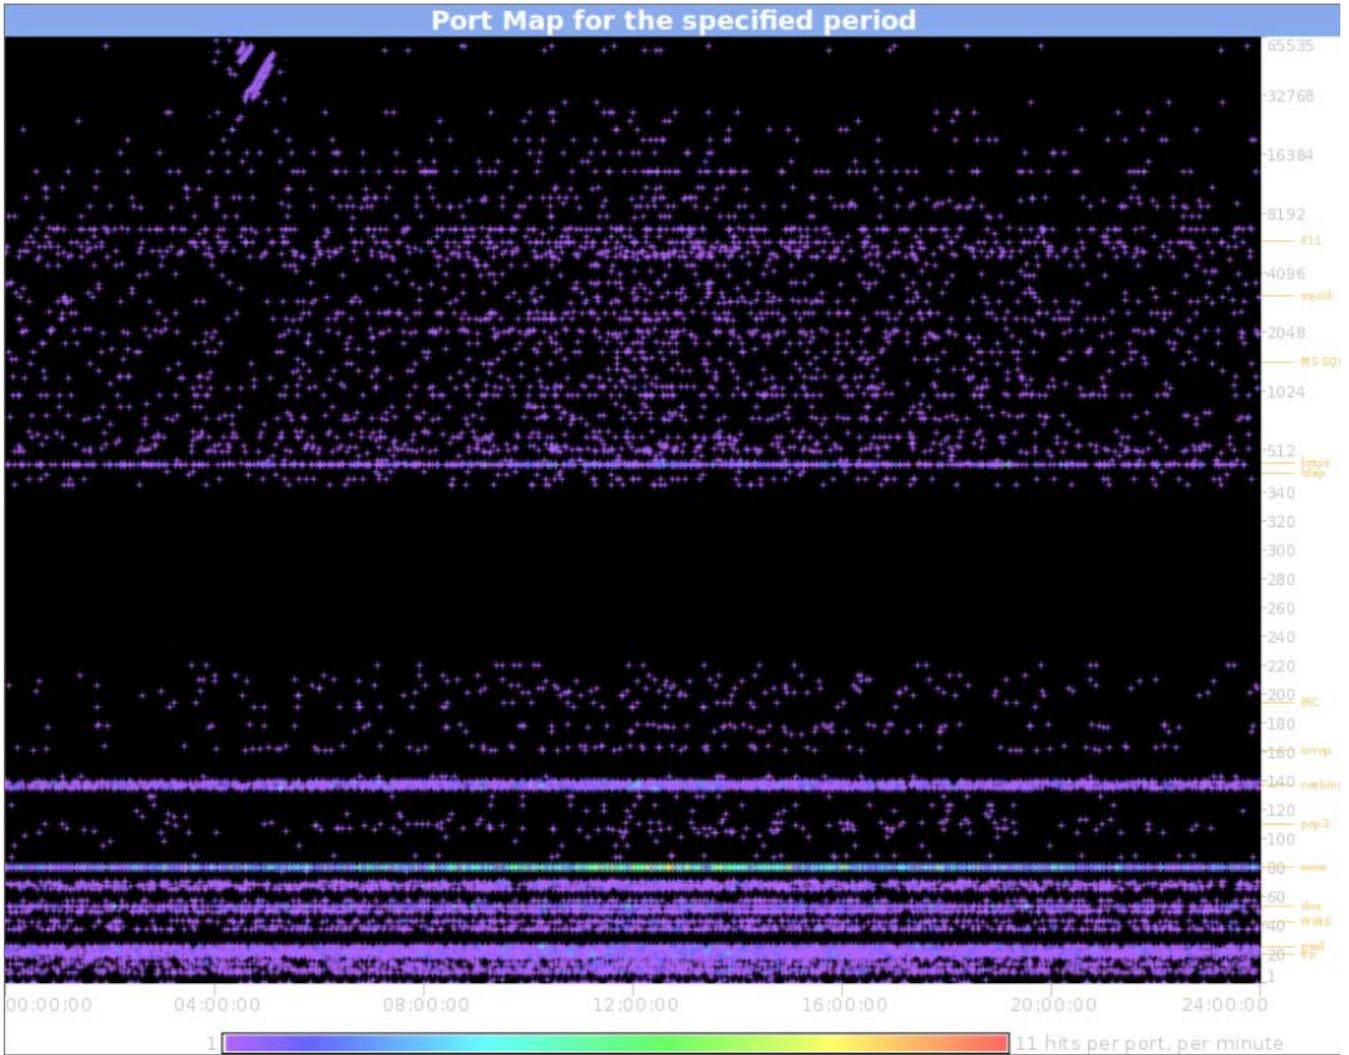
that scans incoming data for events that match your query terms. Real-time alerts can be sent out via email.

- ✔ Activating real-time alerts will reduce your maximum potential event collection speeds. Each additional real-time alert that is activated, will also increase the amount of processing that your server needs to do, per-event, and will slightly decrease your maximum potential event collection speed.

6.3.12. Destination Port Map

This output component appears for data sources that include a destination IP address, and destination port - such as firewalls, or network intrusion detection systems.

The destination port map shows destination ports hit during the period specified in the objective match settings, as a clickable exponentially-scaled dot-map. Areas of higher activity are represented as colours towards the top end of the colour spectrum.



6.3.13. Geolocation Map

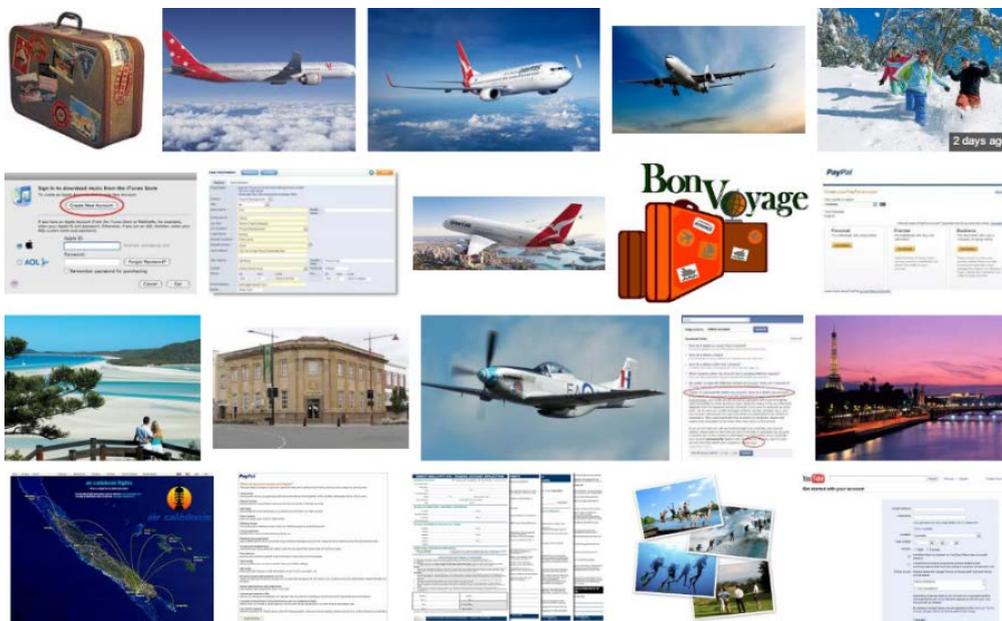
This output component draws lines that represent the country of origin for source and destination IP addresses, for firewall/NIDS related data sources.



6.3.14. Random Image Selection

For proxy-server related objectives, a random selection of images can be displayed to provide a general overview of image-related browsing habits.

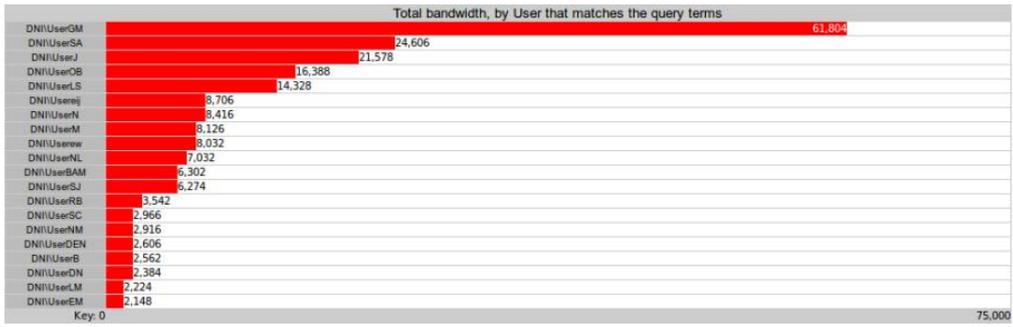
Random Image Count Display how many random images?:



6.3.15. Bandwidth by User / Site

For proxy-server related objectives, the top sites by bandwidth, and/or top authenticated users by bandwidth utilisation can be displayed.

Bandwidth by Site Configuration How many graph elements should be displayed?:



7. Supporting Objectives - Status

The Status category contains objectives used to monitoring the status and performance of the Snare Server. This includes information on user access to the Snare Server, current scripts and processes that are running or queued to run, summaries of the data in the data store and general health check information.

The key sub-categories are:

- General Statistics
- Monitor Live Data
- Retrieve Integrity Check of the Data Store
- Snare Health Checker
- System Status
- Total events Plotted per 15 minutes



7.1. General Statistics

This objective provides a number of graphical displays, summarising the data currently held in the Snare Server data store.

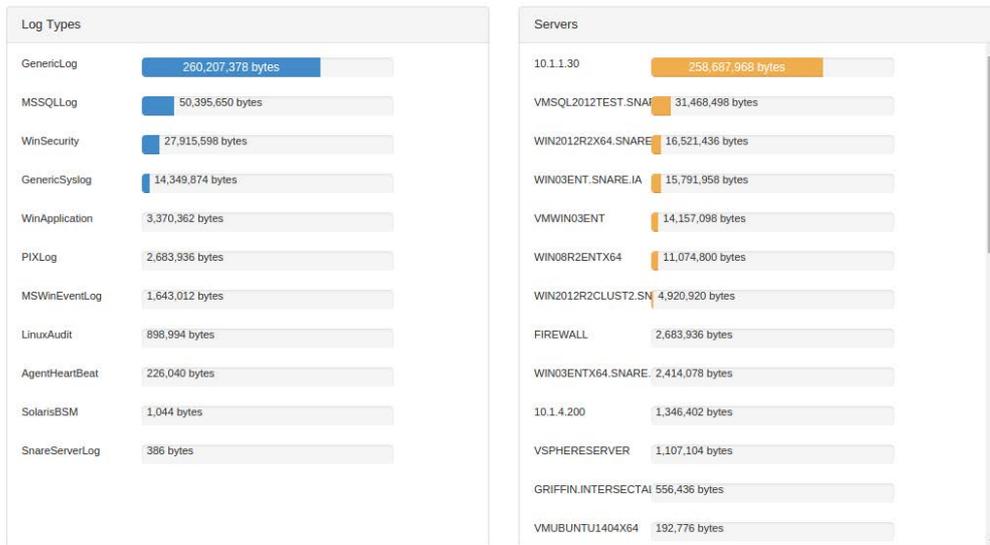
Tabs include:

- A stacked horizontal bar graph of events per month.
- A vertical bar graph of total events for the current year.
- A vertical bar graph of events per second, per day, for the last 12 weeks.
- A collective clickable graph that displays total number of events, compressed storage size, and average compressed bytes per event for each log type, and each agent within the log type.
- A pattern map of events per system over the last 12 weeks.
- A horizontal graph of total events per system, sorted by system.

LogType	Total Number of Events	Compressed Storage Size	Average Compressed Bytes per Event
AC2Log	274 events	95,236 bytes	347 bytes
AIXAudit	18,349 events	286,250 bytes	15 bytes
CISCORouterLog	27,395 events	1,034,028 bytes	37 bytes
CKEFPoSLog	18,124 events	1,947,722 bytes	107 bytes
CyberGuardFirewallLog	22,124 events	1,859,022 bytes	84 bytes
FirewallLog	22,348 events	1,925,801 bytes	86 bytes
GauntletFirewallLog	27,233 events	1,800,134 bytes	66 bytes
GenericLog	425 events	2,792 bytes	6 bytes
GenericSysLog	285 events	87,283 bytes	306 bytes
IPTablesFirewall	27,118 events	1,215,442 bytes	45 bytes
ISAFWSLog	27,243 events	1,217,293 bytes	45 bytes
IrixSAT	1,558 events	40,772 bytes	26 bytes
LinuxAudit	130 events	23,028 bytes	177 bytes
DARTER	12 events	136 bytes	11 bytes
LAPWING	12 events	1,443 bytes	120 bytes
QUAL	485 events	11,792 bytes	24 bytes
SILVEREYE	2 events	618 bytes	310 bytes
MSWinEventLog	10 events	5,887 bytes	589 bytes
MailLog	250 events	8,189 bytes	33 bytes
NetScreenFirewall	24 events	3,523 bytes	147 bytes
NetgearFirewallLog	27,118 events	1,301,808 bytes	48 bytes
NetgearRouterLog	27,277 events	897,249 bytes	33 bytes
NortelVPNRouter	282 events	25,445 bytes	90 bytes
ObjectAccess	273 events	1,482,772 bytes	54 bytes
PIXLog	22,255 events	1,109,973 bytes	49 bytes
RACFLog	11,236 events	308,290 bytes	27 bytes
SOCKSLog	210 events	18,395 bytes	88 bytes
SnareServerLog	27,382 events	2,182,024 bytes	79 bytes
Snort	10,274 events	172,888 bytes	17 bytes
SolarisBSM	2 events	200 bytes	100 bytes
TopicLog	280 events	2,382 bytes	8 bytes
Tru64Audit	24 events	5,928 bytes	247 bytes
UNKNOWN	24 events	2,228 bytes	93 bytes
UniversalLog	2,024 events	245,880 bytes	121 bytes
WebLog	1,174 events	200,818 bytes	171 bytes
WinApplication	10,287 events	2,403,238 bytes	234 bytes
WinSecurity	10,287 events	2,403,238 bytes	234 bytes
WinSystem	10,287 events	2,403,238 bytes	234 bytes

7.2. Monitor Live Data

This objective provides a way to preview the events that are being received by the Snare Server live. It is designed for debugging and event collection health checking, rather than for auditing the exact events received by the server.



The box on the left lists all of the **Log Types** for the incoming Events, and the number of bytes received for each *Log Type*. Clicking on a specific *Log Type* filters the other displays to make it easier to drill down and see specific events coming into the server.

The box on the right lists all of the **Servers** or hosts that are sending events to the Snare Server. Like the *Log Types* list, it shows the number of bytes received. Clicking on a *Log Type* will filter the *Servers* listed in this box to only those that have sent events of that specific type.

Snapshot of the last 10 Generic Syslog events from VMWIN03ENT

DATE	TIME	SYSTEM	TABLE	CRITICALITY	SOURCE	EVENT
2014-12-12	00:34:36	VMWIN03ENT	GenericSyslog	13	MSSQLSERVER/master	MSSQLLog 2014-12-12 00:34:35.427 09.00.5057 15 1 55 MSSQLSERVER/master SNARE\administrator TextData, Success.1 SessionLoginName,SNARE\administrator NTUserName,administrator HostName,WIN03ENT ApplicationName,SQLCMD
2014-12-12	00:34:36	VMWIN03ENT	GenericSyslog	13	MSSQLSERVER/master	MSSQLLog 2014-12-12 00:34:35.610 09.00.5057 14 1 55 MSSQLSERVER/master SNARE\administrator TextData,-- network protocol, LPC set quoted_identifier on set arithabort off set numeric_roundabort off set ansi_warnings on set ansi_padding on set ansi_nulls on set concat_null_yields_null on set cursor_close_on_commit off set implicit_transactions off set language us_english set dateformat mdy set datefirst 7 set transaction isolation level read committed Success.1 NTUserName,administrator HostName,WIN03ENT ApplicationName,SQLCMD
2014-12-12	00:34:36	VMWIN03ENT	GenericSyslog	13	MSSQLSERVER/master	MSSQLLog 2014-12-12 00:34:35.610 09.00.5057 33 0 55 MSSQLSERVER/master SNARE\administrator TextData,Error: 208, Severity: 16, State: 1 Success.0 SessionLoginName,SNARE\administrator NTUserName,administrator HostName,WIN03ENT ApplicationName,SQLCMD Error:208 ErrorString,Invalid object name '%*s'. TransID,92647966
2014-12-12	00:34:36	VMWIN03ENT	GenericSyslog	13	MSSQLSERVER/master	MSSQLLog 2014-12-12 00:34:35.610 09.00.5057 15 1 55 MSSQLSERVER/master SNARE\administrator TextData, Success.1 SessionLoginName,SNARE\administrator NTUserName,administrator HostName,WIN03ENT ApplicationName,SQLCMD
2014-12-12	00:34:36	VMWIN03ENT	GenericSyslog	13	MSSQLSERVER/master	MSSQLLog 2014-12-12 00:34:35.820 09.00.5057 14 1 55 MSSQLSERVER/master SNARE\administrator TextData,-- network protocol, LPC set quoted_identifier on set arithabort off set numeric_roundabort off set ansi_warnings on set ansi_padding on set ansi_nulls on set concat_null_yields_null on set cursor_close_on_commit off set implicit_transactions off set language us_english set dateformat mdy set datefirst 7 set transaction isolation level read committed Success.1 NTUserName,administrator HostName,WIN03ENT ApplicationName,SQLCMD

The bottom box shows the last 10 events received, to provide a preview of the events coming in for the selected *Log Type* and *Server*.

 This objective consumes system resources whilst active. It may have a small negative effect on event collection rates if left open for long periods of time.

7.3. Retrieve Integrity Check of the Data Store

This objective scans the current data store, and generates an MD5 checksum for each file found. The results can be downloaded from the objective, and compared against previous runs, in order to verify that data has not been tampered with since the last run.

Sample output from the integrity check process

```
/data/SnareArchive/2012-04-20/COCKATIEL/CISCORouterLog-21-0-68513.822876.log.gz  
  
d7287c9c3efd7a07c998f06f0314403f  
  
/data/SnareArchive/2012-04-20/COCKATIEL/CISCORouterLog-06-3-68513.61309.log.gz  
  
aa517d816dfe305323b5ebcc59b27b40  
  
/data/SnareArchive/2012-04-20/COCKATIEL/CISCORouterLog-13-0-68513.708384.log.gz  
  
38c4b225ade619ddd8def2dcad4ed4a  
  
/data/SnareArchive/2012-04-20/COCKATIEL/CISCORouterLog-02-3-68513.566432.log.gz  
  
f8958812a174c441c8f807d31a671457  
  
/data/SnareArchive/2012-04-20/COCKATIEL/CISCORouterLog-07-3-68513.6255701.log.gz  
  
bf63b8bcd9f97342dc31a2fa85efde64  
  
/data/SnareArchive/2012-04-20/COCKATIEL/CISCORouterLog-08-3-68513.6394141.log.gz  
  
5744823f6f57be8be1934c382bfb5c13
```

7.3.1. Snare Health Checker

This objective provides a 'health check' for the Snare Server by querying the status of key functions of the Snare Server, including, but not limited to:

- licensing,
- whether the key services are still functioning and,
- the amount of disk space available.

These functions, and others, are configurable via the "Configure" tab.

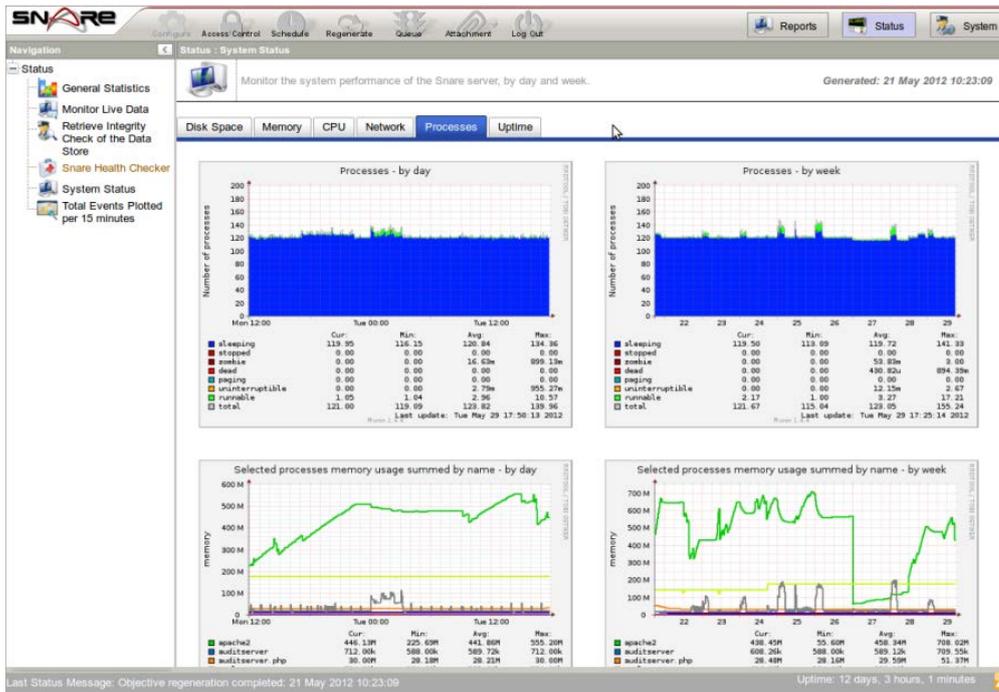
It is recommended that any (red) problem indications are reported and resolved immediately.

Warning messages (in orange) should be investigated when time permits.

 Unlike most other Snare Server objectives, it is not necessary to 'regenerate' this objective. The results are calculated 'on the fly' every time it is loaded.

7.4. System Status

This objective provides the details of the Snare Server. It includes hardware description, operating system distribution, uptime and information and graphs on CPU, network, memory, swap and mounted file system usage.



7.4.1. Total Events Plotted per 15 Minutes

This objective displays the total number of events received for every host over a 35 day period. The coloured rectangles indicate the number of events received during the 15 minute period relative to the scale shown at the bottom of the graph. Further details on the number of events received can be ascertained by placing the mouse cursor over the coloured rectangles.

Furthermore, the raw logs for this 15 minute period can be viewed by 'clicking' the rectangle. Note that for the current month, the details will obviously be for the period since the report was generated. Collection details for a specific host can be viewed by selecting the list of hosts shown at the bottom of the graph.

8. Supporting Objectives - System

This section allows you to access functions that manage and maintain the Snare Server and its users, and also manage the configuration of Snare Agents that report to the Snare Server.

8.1. Administrative Tools

8.1.1. Antivirus Administration

The Snare Server is based on a custom distribution of Linux, and is therefore potentially susceptible to (significantly) less than 1% of all viruses currently in the wild. The Snare Server does not provide desktop-level functionality, and the risk profile for virus infection on the Snare Server is extremely low. However, the Snare Server integrates the ClamAV virus checker, which is an open source (GPL) antivirus engine designed for detecting Trojans, viruses, malware and other malicious threats. It includes a high performance multi-threaded scanning daemon that provides numerous file format detection mechanisms, file unpacking support, archive support, and multiple signature languages for detecting threats.

The anti-virus scan can be run on a scheduled basis, and can be configured to perform:

- a complete system scan,
- exclude the Snare Data Store, and results cache from the scan (recommended), or
- only scan the home directories of Snare Server user accounts.

The reason that it is recommended that the Data Store and results cache be excluded from the scan, is that there is a significant risk that the virus scanner will pick up false-positives in those directories, due to the nature and volume of data stored therein.

It is the customers responsibility to ensure the antivirus software is kept up to date and is scheduled to run in accordance with your corporate security policy.

8.1.2. Change the Snare Server IP address

The Snare Server IP address, netmask, default gateway, and DNS servers can be modified using this objective. IP, netmask and default gateway values can be modified on a per-ethernet-card basis.

It should be noted that once the IP address has changed, the server will no longer be contactable via the old IP address, so if you were connecting to the old IP address with your web browser, your browser may become unresponsive after the address change.

8.1.3. Configuration Wizard

The configuration wizard is covered earlier in this documentation.

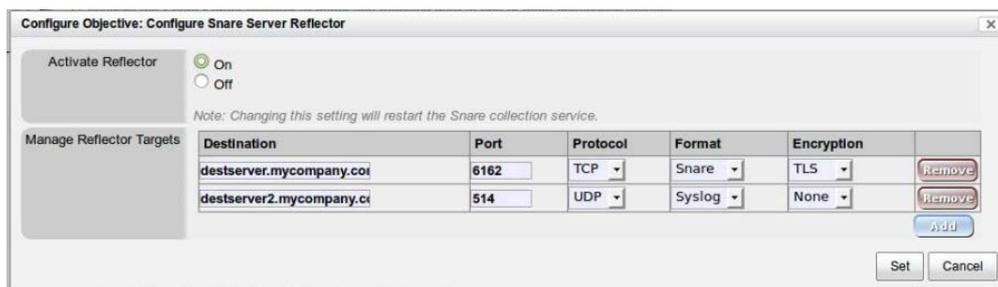
8.1.4. Configure Server Time Zones

The Snare Server has the ability, on a per-source basis, to time-shift the data at query time. In general, agents will report data back to the Snare Server using their local time and time-zone. For objectives such as "Tell me whenever someone logs in before business hours", this strategy works perfectly well. However, if you have a reporting agent in Paris, another in London, and your Snare server was based in New York, and your reports predominantly needed to be based around the time in New York (EST), then you may wish to turn time zone manipulation on.

This will allow you to construct a report such as "Tell me any events that occurred between 08:30 and 09:30 US Eastern Standard Time, regardless of what the local time on the destination server was reporting".

Please note that enabling time zone manipulation, will slow down Snare Server queries by an appreciable amount, and whilst enabled, will affect all data from the configured source system. Note also that the original data will not be modified - turning off time zone manipulation will return reporting to normal.

8.1.5. Configure Snare Server Reflector



The Snare Server reflector is capable of sending data to arbitrary ports, in either 'Snare' traditional, or syslog encapsulated formats. TLS/SSL encryption is available, if the destination server supports it.

Each additional reflector destination will have some impact upon the maximum potential collection rates of the Snare Server, where the amount depends significantly on your choice of hardware and network bandwidth availability.

As a general guide, low single-figure percentage differences have been noted on high end workstation-equivalent hardware, when comparing an unreflected server, with a server reflecting to two destination points.

- ✔ The Snare Server will reflect all incoming data to your destination points, regardless of original format. For example, if the source and destination formats are both syslog, the event will be pushed through unchanged. If the source format was 'Snare' or 'SNMP', and the destination format was syslog, then a syslog header will be prepended to the data before pushing it to the remote server.

8.1.6. Display the Snare Log File

In situations where you request assistance from your Snare Server support team, you may be asked to email a copy of the Snare debug log file. This file contains generic information on what objectives run, and what scheduled tasks are currently implemented. Increasing the Snare Server debug level (see the section above on "Configuration Wizard" for more information), will significantly increase the amount of data that is written to this file.

8.1.7. Display the Snare Service Monitor Log File

Collection is the process that the Snare Server is most anxious to ensure is robust and reliable. If something causes the collection subsystem to fail, it will be restarted as soon as possible, and the server will attempt to collect as many useful statistics relating to memory usage, disk usage, and process information, as it can, in order to support debugging efforts by your Snare Server support team.

8.1.8. Import Objectives

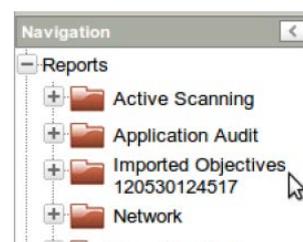


The team at InterSect Alliance have come up with a quantity of default objectives that suit a diverse range of organisations, and security-related regulatory requirements. However, there may be situations where additional specialised objectives are made available to users of the Snare Server. The **'Import from the InterSect Alliance Objective Store'** button will allow you to select, and import, objectives.

Objectives	Description	Modified	
 NISPOM	Objectives that meet NISPOM Chapter 3 Requirements	30 May 2012	<input type="button" value="Import"/>
 PCI	Objectives that meet Payment Card Industry regulatory requirements	30 May 2012	<input type="button" value="Import"/>
 ReleaseObjectives	Snare Server default objectives as of 2012-05-30	30 May 2012	<input type="button" value="Import"/>
 Test	A single test objective to test the Objective import capability of Snare Server version 6.0	21 Jul 2011	<input type="button" value="Import"/>

Objectives will be imported into a new container, called "Imported Objectives YYMMDDHHMMSS" (where YYMMDDHHMMSS represents the date/time of import).

In addition to importing objectives from the InterSect Alliance web site, there is also an option to upload objectives from a file stored on your local workstation. In situations where you have previously used the 'Objective Export' capability by right-clicking on a container, the objectives will be exported to either a local file, or via email, to a selected destination user.



8.1.9. Manage Objective Schedules

This objective provides summary information on current objective scheduling, target email addresses, and access controls. A link to each objective also enables you to modify the associated configuration settings.

8.1.10. Prepare Server for Upgrade

This objective runs a number of checks on your Snare Server to ensure it is ready to be upgraded to the next major version using the 'over-the-top' upgrade method. Note that until a new major version of the Snare Server is available, this objective will not provide any significant functionality.

8.1.11. My Snare Server Account

Your Snare Server password can be changed in this objective. Last login date/time information is also available. Note that the Snare Server implements several password security policies, including:

- 90 Day Rotation
- Password reuse protection
- Last password similarity checks
- Password complexity requirements
- Dictionary word exceptions

8.1.12. Shutdown / Reboot the Snare Server

Users with administrative-level access to the Snare Server will be able to shut down, or reboot the Snare Server from this objective.

8.1.13. Snare Server Update

The team at InterSect Alliance will release updates to:

- Add features to the Snare Server
- Fix issues that have been reported
- Update operating system components in response to security issues that specifically affect Snare
- Update virus checker signatures.

The update will be made available in the form of a PGP signed compressed archive. This objective will accept such an update file, verify that the PGP signature is valid, and apply the update to your Snare Server installation.

8.1.13.1. Troubleshooting Updates



Troubleshooting Updates

Blank navigation/screen after upgrade process.

It is unlikely, but possible, that after an upgrade the navigation section, or the entire page, may end up on a blank white screen. This is caused by your web browser caching some of the old page components and preventing the server from using the upgraded components. While we have put checks in place within Snare to try and prevent this, it is possible that some browsers may bypass these checks. To resolve the issue, you can (in most browsers) hold down the *Shift* key while pressing *Refresh* on the browser. If this doesn't work, try clearing the browser cache and restarting the browser. If this still does not work, try using a different browser.

8.1.14. User Administration

It is recommended that a number of users be created after the Snare Server has been installed, so that:

- The Administrator username and password do not have to be shared and
- It will be possible to identify which user is accessing and configuring Snare.

This objective allows you to create users and groups.



This objective details the users and groups that are allowed to access the Snare Server.

Snare Server Users (LDAP / Active Directory authentication is not active)					
User	Name	Last Known Logon	Group Membership	Email Address	
	Administrator	Administrator Account	2012-05-30 11:34:50	Default, Administrators, SuperUsers	
	leigh	Leigh	2012-05-21 10:35:06	Default	

[Create a new Snare Server User](#)

Snare Server Groups	
Group	Group Comments
 Accounting	Accounting
Administrators	SNARE Administrators group
Default	Default SNARE group
Finance	Finance
NetworkAdministrators	Network Administrators
SuperUsers	SNARE Super Users group
SystemAdministrators	System Administrators

[Create a new Snare Server Group](#)

There are three groups that are built into the Snare Server: Administrators, SuperUsers, and 'Default'.

All users are automatically included in the 'Default' group. The 'Administrators' group has the same access as the 'administrator' userid with the exception of a number of functions that are restricted to the 'administrator' (eg: Changing the password of the Administrator account). The 'SuperUser' group has no particular privileges, but can be used to group accounts with significant privileges to objectives, if you wish to take advantage of it.

You may define as many additional Groups as you wish.

The Snare Server implements several password security policies, including:

- 90 Day Rotation
- Password reuse protection
- Last password similarity checks
- Password complexity requirements
- Account locking on multiple failed login attempts
- Dictionary word exceptions

If a password does not meet the requirements identified above, an error message will be displayed during password definition.

UserID	<input type="text" value="SnareUser"/>
User Name	<input type="text" value="Snare User"/>
Password	<input type="password" value="●●●●●●●●"/>
Password (repeat)	<input type="password" value="●●●●●●●●"/> Poor password: it is based on a dictionary word
Email Address	<input type="text"/>
Group Membership	<input type="text" value="Administrators"/> <input type="text" value="SuperUsers"/> <input type="text" value="NISPOM"/>

In situations where an account is locked due to several failed login attempts, an additional configuration setting on the user management screen will offer the administrator the capability to unlock a Snare Server user account. If an account is not unlocked, the account will automatically unlock after 30 minutes.

If a users account exceeds the 90 day password validity limit, the Snare Server will request a password update.



8.1.14.1. Operating System Password Controls

The operating system password controls are managed by the Pluggable Authentication Modules (PAM) in Linux. The configuration files are located in /etc/pam.d directory. The password controls for the Snare Server are detailed in the /etc/pam.d/common-password file. The file can be updated to reflect your corporations security policy.

The default settings are as follows and enforces a password retry of 3 attempts before failure, length of 10 characters, a difference of three characters from previous password, one uppercase letter, one numeric, one special character, and one lowercase letter:

- `password requisite pam_cracklib.so retry=3 minlen=10 difok=3 ucredit=-1 dcredit=-1 ocredit=-1 lcredit=-1`

The configuration will enforce the password policy rules for the following operating system accounts root, snare and snarexfer. For additional information on the values of each setting refer to the manual pages for pam.d and pam_cracklib.

8.2. Snare Agents

This section assumes that you have Snare Agents installed and are using the agents as part of your logging environment. It provides tools to perform the following tasks:

- Gather user and group information to support other Snare Server objectives.
- Query User and Group information gathered from your Snare Agents.
- Audit and manage the configuration of Snare Agents within your environment.

8.2.1. Query Data

This is a simple objective that scans the user and group details retrieved from various Snare Agents as part of the "Retrieve Data" objectives within "Snare Agents".

For the database table **AuthGroupMembers** select your query conditions

Field	Type	Function	Value	Sort	Group By
GroupName	VARCHAR	LIKE		<input checked="" type="radio"/>	<input type="checkbox"/>
ID	VARCHAR	LIKE		<input type="radio"/>	<input type="checkbox"/>
Members	TEXT	LIKE		<input type="radio"/>	<input type="checkbox"/>
Domain	VARCHAR	LIKE		<input type="radio"/>	<input type="checkbox"/>
Comment	TEXT	LIKE		<input type="radio"/>	<input type="checkbox"/>
Retrieved	DATETIME	=		<input type="radio"/>	<input type="checkbox"/>
(AND)	DATETIME	=		<input type="radio"/>	<input type="checkbox"/>
Source	VARCHAR	LIKE		<input type="radio"/>	<input type="checkbox"/>

Limit number of results to queries per page Reverse Sort Order

[Run Query](#) | [Export CSV Spreadsheet](#) | [Export Text Dump](#)

Utilise the search functions to scan for particular users or groups of interest. The search function provides a very basic query builder. Results are returned in tabular form.

AuthGroupMembers						
GroupName	ID	Members	Domain	Comment	Retrieved	Source
Administrators	-	LUNA\Administrator	10.0.0.9	Administrators have complete and unrestricted access to the computer/domain	2009-05-11 12:03:59	Windows
Backup Operators	-		10.0.0.9	Backup Operators can override security restrictions for the sole purpose of backing up or restoring files	2009-05-11 12:03:59	Windows
Guests	-	LUNA\Guest	10.0.0.9	Guests have the same access as members of the Users group by default, except for the Guest account which is further restricted	2009-05-11 12:03:59	Windows
HelpServicesGroup	-	LUNA\SUPPORT_388945a0	10.0.0.9	Group for the Help and Support Center	2009-05-11 12:03:59	Windows
Network Configuration Operators	-		10.0.0.9	Members in this group can have some administrative privileges to manage configuration of networking features	2009-05-11 12:03:59	Windows
Power Users	-	S-1-5-21-448539723-1417014170-725345543-1113	10.0.0.9	Power Users possess most administrative powers with some restrictions. Thus, Power Users can run legacy applications in addition to certified applications	2009-05-11 12:03:59	Windows
Remote Desktop Users	-		10.0.0.9	Members in this group are granted the right to logon remotely	2009-05-11 12:03:59	Windows
Replicator	-		10.0.0.9	Supports file replication in a domain	2009-05-11 12:03:59	Windows
Users	-	NT AUTHORITY\INTERACTIVE,NT AUTHORITY\Authenticated Users,LUNA\Fred Bloggs	10.0.0.9	Users are prevented from making accidental or intentional system-wide changes. Thus, Users can run certified applications, but not most legacy applications	2009-05-11 12:03:59	Windows

8.2.2. Retrieve Data

8.2.2.1. AIX Users and Groups

Retrieve users and groups by connecting to all, or specific, Snare for AIX Agents that have sent data to the Snare Server, and requesting a dump of the user and group data.

User and group information will be used by AIX objectives to convert numeric user and group ID information into user/group names, and to implement user/group snapshot objectives.

In order to run this objective successfully, you should have at least one 'Snare for AIX' agent installed on a server that has full YP visibility, with 'remote control' activated, and a password set that matches either the 'override' password explicitly configured for this objective, or the password set under the 'Configuration Wizard'. In addition, the system in question should be reachable by the Snare Server from a network perspective (eg: firewalls between the Snare Server and the YP master should allow TCP connections from the Snare Server to the remote system on TCP port 6161).

8.2.2.2. Cognos Users and Groups

Retrieve users and groups by connecting to a Cognos-specific LDAP server that has been configured to allow the Snare Server IP address to download Cognos user and group information.

User and group information will be used by Cognos objectives to convert numeric user and group ID information

into user/group names, and to implement user/group snapshot objectives.

8.2.2.3. Irix Users and Groups

Retrieve users and groups by connecting to all, or specific, Snare for Irix Agents that have sent data to the Snare Server, and requesting a dump of the user and group data.

User and group information will be used by Irix objectives to convert numeric user and group ID information into user/group names, and to implement user/group snapshot objectives.

8.2.2.4. LDAP Users and Groups

Retrieve users and groups by connecting to a generic LDAP server that has been configured to allow the Snare Server IP address to scan for user and group information.

8.2.2.5. Linux Users and Groups

Retrieve users and groups by connecting to all, or specific, Snare for Linux Agents that have sent data to the Snare Server, and requesting a dump of the user and group data.

User and group information will be used by Linux objectives to convert numeric user and group ID information into user/group names, and to implement user/group snapshot objectives.

8.2.2.6. OS400 Users and Groups

Search for files generated with the AS/400 DSPUSRPRF tool, that have been transferred to the /data/SnareCollect/OS400Users directory on the Snare Server, and retrieve user account information, and related user flags from the file.

8.2.2.7. Lotus Notes Event Logs

Since no agent currently exists for Lotus Notes, this objective attempts to connect to a target Domino server, and download the log.nsf (MiscEvents, MailRoutingEvents, ReplicationEvents and NNTPEvents), catalog.nsf, and names.nsf databases, and insert the resulting data into appropriate data stores on the Snare Server.

User and Group information, plus notes access controls are also downloaded. Depending on your log volume, and data retention settings within Lotus Notes, you may need to modify some settings within Domino, in order for Domino to return appropriate results back to the Snare Server. Within the Domino web server configuration page is a section named "Conversion/Display". From the Domino Administrator, click the Configuration tab, expand the Web section and click Internet Sites.

1. Choose the Web Site document you want to edit and click Edit Document.
2. Click the Domino Web Engine tab. Under "Conversion/Display", the default settings are: Default lines per view page: 30 Maximum lines per view page: 1000. These values should be configured as follows:
Default lines per view page: 250 Maximum lines per view page: 0

 The objective will attempt to download event log data tagged with a date/time of yesterday's date, by Lotus Notes. It is recommended that this objective be configured to run once per day.

User and group information will be used by user/group snapshot objectives.

8.2.2.8. Solaris Users and Groups

Retrieve users and groups by connecting to all, or specific, Snare for Solaris Agents that have sent data to the Snare Server, and requesting a dump of the user and group data.

User and group information will be used by Solaris objectives to convert numeric user and group ID information into user/group names, and to implement user/group snapshot objectives.

In order to run this objective successfully, you should have at least one 'Snare for Solaris' agent installed on a server that has full NIS visibility, with 'remote control' activated.

8.2.2.9. Windows Users and Groups

Retrieve users and groups by connecting to all, or specific, Snare for Windows Agents that have sent data to the Snare Server, and requesting a dump of the user and group data.

User and group information will be used by Windows objectives to convert SID information into user names, and to implement user/group snapshot objectives.

In order to run this objective successfully, you should have at least one 'Snare for Windows' agent installed on a Domain Controller or Member Server, with 'remote control' activated.

8.2.3. Remote Management

A complete guide to using the Agent Management Console can be found in the "Snare Server v7 Agent Management Console" user guide.

The Remote Management section under Snare Agents provides the ability to audit and manage the configuration of the Snare Agents within your environment.

By default it contains a single 'Manage Agents' objective, but this objective can be cloned, renamed, and deleted, to support as many different combinations of agent configurations as required. Simply *Right-Click* on the objective and choose from the options in the menu.



No other objectives within the Status or System menu can be manipulated in the same way.

8.2.3.1. Configuring Agent Management

To set up an Agent Management objective, open one of the Remote Management objectives (the default is *Manage Agents*) and go to the Configuration section (click 'Configure' in the top icons menu).

8.2.3.1.1. Snare Agent Type

Specify the type of the Snare Agents to be managed through this objective. Different agent versions have been grouped based on Operating System and major version changes to prevent incompatible configurations from being compared and saved. *Remember, you can easily Clone the objective for each type of Snare Agent you use.*

The objective will attempt to contact any Agents that match the required Operating System based off the log tables they have reported events for in the last 3 months. It restricts these results based off the reported version number from the Agent itself.

For example, the "Snare Agent for Linux (1.x.x - 2.1.x)" type will not match a "Snare Agent for Windows" or even a "Snare Agent for Linux v2.2.0" Agent. However, it will match a "Snare Agent for Linux v2.1.0" or a "Snare Agent for Linux v1.8.2" Agent.

8.2.3.1.2. Hostname Filter

To help filter the Agents being managed through the objective, a *hostname filter* can be specified that will restrict the managed Agents to those matching the filter. The filter supports by default * as a wildcard, but this can be changed to support Regular Expressions if the option is enabled.

The objective will not attempt to contact any Agent which is excluded due to the hostname filter.

For example, a filter of '*.intersectalliance.com' will manage 'agent001.intersectalliance.com', but not 'agent002.dni.gov.au'.

8.2.3.1.3. Version String filter

To further restrict the managed agents, a version string can be specified using the same matching rules as the Hostname Filter (i.e. * wildcard, or regular expressions). This filter is based off the reported version from the agent, which it can only obtain by attempting to contact the agent during each regeneration.

For example, a regular expression filter of "4\.\0\.[01]\.d" will match any version between 4.0.0.0 and 4.0.1.9.

 This will not override the version restrictions set in the **Snare Agent Type** selection. If the specified version string is incompatible with the Agent Type, no agents will be managed.

8.2.3.1.4. Non-reporting Agents

Snare Agents that do not report to the Snare Server can be specified within the Non-reporting Agents box. This will add them into the list of hosts to query and be managed if online.

Non-reporting Agents bypass the *Hostname Filter* selection, however they will still be checked by the *Snare Agent Type* and *Version String Filters*.

Non-reporting Agents can be listed with each agent on a new line, with the IP address and hostname separated with a comma.

For example, a block of non-reporting agents would look like this:

```
10.0.0.100,CUSTOM-0-100.SNARE.IA
10.0.0.101,CUSTOM-0-101.SNARE.IA
10.0.0.102,CUSTOM-0-102.SNARE.IA
10.0.0.103,CUSTOM-0-103.SNARE.IA
```

Alternatively, an IP address range can be specified by clicking the 'Add IP address range' button. Simply specify the IP range in the format: 10.1.1.1-10.1.1.5, and the domain to append to the IP to form the hostname.



Adding an IP address range inserts all of the specified IP addresses into the Non-reporting Agents field. So it is easy to remove specific IP addresses from the middle of the range as required.

8.2.3.1.5. Alternate Passwords

The objective, by default, uses the Agent password specified in the Configuration Wizard when it attempts to communicate with each Agent. If that password fails, it will attempt each of the Alternate Passwords until it finds one that works. This allows you to support legacy configurations, periodically change the Snare Agent password, or use different passwords for different groups of Agents, without stopping the objective from communicating with the agents.

8.2.3.1.6. Alternate listening port

Likewise, the objective uses the Agent listening port specified in the Configuration Wizard when it attempts to communicate with each agent. If that port fails, it will attempt to use the alternate port specified here. The purpose of this field is the same as the alternate password fields - legacy and change support.

8.2.3.1.7. Management Mode

Only highlight differences between Master config and Agent config

The objective will only highlight the differences between the Master config and each Agent configuration. These differences will need to be manually resolved.

Push Master config to all managed Agents on schedule

The objective will attempt to update each non-matching Agent with the configuration template from the Master. This method requires no manual intervention to sync up managed.



This option is only supported by some Snare Agent versions.

8.2.3.1.8. Comparison Options

Ignore Agent version mismatch in configuration differences report.

By default the Agent Version is not saved in the configuration lists so it doesn't report as a mismatch during normal operation of the objective. Disabling this option will compare version numbers, which can be useful when upgrading the fleet to track down any Agents that have been missed in the upgrade.

Ignore offline/uncontactable agents.

Normally any Agents that are uncontactable are highlighted on the report, and trigger an email notification (if enabled). While this is useful when scanning a known number of Agents to ensure availability, it can cause needless notifications when scanning a whole network for a small number of Agents within the network. Enabling this option will ignore offline Agents by removing the highlight and disabling the email notifications.

8.2.3.2. Understanding the Objective

The objective needs to be configured and regenerated at least once before any of its functions are available.

8.2.3.2.1. Snare Agents

The screenshot shows the 'Snare Agents' interface with the following sections:

- Agents matching the master configuration:** 10-1-2-12.DNI.GOV.AU* (10.1.2.12, Windows, v4.1.0)
- Agents with configuration different to the master configuration:** 10-1-2-19.DNI.GOV.AU* (10.1.2.19, Windows, v4.1.0), 10-1-2-3.DNI.GOV.AU* (10.1.2.3, Windows, v4.1.0), 10-1-2-4.DNI.GOV.AU* (10.1.2.4, Windows, v4.1.0)
- Agents that cannot be contacted:** CLIPPER.INTERSECT.LOCAL (10.0.3.11), FLASH.INTERSECT.LOCAL (10.0.3.10)
- Agents ignored by hostname filter: *.intersect*:** ALBATROSS, ALSDFKJSD, CASSOWARY, CATBIRD, COOT, CURLEW, CURRAWONG, FINCH, GANDALFV0, GIMLIV0

The first content tab provides a summary overview of the status of each Agent (both managed and ignored), grouped by status. The hostname of each agent is listed, along with IP address, Agent Type and version.

Non-reporting agents will be postfixed by a *.

- **Agents matching the Master configuration**
 - These Agents are online and completely match the Master Configuration, with no differences.
- **Agents with configuration different to the Master Configuration**
 - These Agents are online but their configuration is different to the Master Configuration.
- **Agents that cannot be contacted**
 - Agents that match the Operating System (from Agent Type filter), and Hostname filter, but cannot be contacted. Some of these Agents may have the wrong version, but since this information is identified once the Agent is contacted it is impossible to determine this information while they are offline.
- **Agents ignored by version string filter**
 - Online Agents excluded from management due to the version string filter.
- **Agents ignored by hostname filter**
 - Agents excluded from management due to the hostname filter.
- **Agents ignored by type filter**
 - Agents excluded from management due to the Agent type filter.

8.2.3.2.2. Master Config

The screenshot shows the 'Master Config' interface with the following sections:

- Refresh Master Config:**
 - Non-managed Agent Configuration from GIMLI.FRITZ.BOX (IP/Hostname)
 - Existing Agent Configuration from - select -
 - Refresh Master Configuration**
- Master Config Status:**
 - Master Configuration last retrieved from GIMLI.FRITZ.BOX at 2012-10-16 09:55:51
- Master Configuration from GIMLI.FRITZ.BOX:**

Parameter	Value
Version	1.8.0-src1
OS	Linux
Allow	1
WebPort	6161
Restrict	
RestrictIP	
Password	
PasswordSnare #1	
PasswordSnare #2	
PasswordSnare #3	
Network #0	10.0.4.60:6161

The Master Config tab provides a way to view the existing Master Configuration, refresh it with imported config from an Agent, or clear it completely.

The **Refresh Master Config** box allows the Master Configuration to be imported from either a custom Agent,

specified by an IP address or hostname, or from an existing agent specified in the dropdown.

Refreshing the Master Configuration will not re-analyse the current Agent configurations, and as such the differences list may not be accurate until the objective is next regenerated.

The current Master Configuration will be listed for reference. Depending on the Agent type, it may be possible to manually edit some of the configuration fields without needing to make the changes on the Master Agent and re-import. If any fields have been manually edited, the hostname will be listed as 'manual update'.

Finally, the Master configuration can be cleared from the objective if it is no longer required/accurate.

8.2.3.2.3. Config Differences

Agent Configuration from LINUX001.GANDALF.FRITZ.BOX		
Parameter	Value	Master
Objective #2		
event	execve	(login_auth,login_start,logout)
uid	*,(root)	Not Configured
Objective #3		
criticality	2	3
event	(login_auth,login_start,logout)	(mount,umount,umount2,settetimeofday,clock_settime,swapon,swapoff,reboot,setdomainname,create_module,delete_module,quotactl)
Objective #4		
criticality	3	Not Configured
event	(mount,umount,umount2,settetimeofday,clock_settime,swapon,swapoff,reboot,setdomainname,create_module,delete_module,quotactl)	Not Configured

Agent Configuration from LINUX010-TEST.GANDALF.FRITZ.BOX		
Parameter	Value	Master
Objective #2		
event	execve	(login_auth,login_start,logout)
uid	*,(root)	Not Configured

The Agent Config Differences tab lists each online Agent that doesn't match the Master configuration exactly, with a list of the parameters that do not match for review. Only the Agent parameters which do not match the Master parameters are listed, and all others will be an exact match.

The first column is the configuration parameter name, the second is the value set on the Agent, and the third is the Master configuration value. Fields which are missing from either side will be marked with 'Not Configured'.

8.2.3.2.4. Refresh Specific Agents

Select Agent(s) to refresh

CUSTOM-0-1.FRITZ.BOX* CUSTOM-0-2.FRITZ.BOX* CUSTOM-0-3.FRITZ.BOX* GIMLI.FRITZ.BOX LINUX001.GANDALF.FRITZ.BOX
 LINUX010-TEST.GANDALF.FRITZ.BOX LINUX010.GANDALF.FRITZ.BOX SAURON.FRITZ.BOX

Refresh Selected Agents

This function will trigger a Regeneration of the Objective. Only the selected agents will be retrieved, and updated, with cached data being used for the other agents to work out the differences list.

[Refresh Selected Agents](#)

When there are a large number of Agents managed in a single objective, it takes time to Regenerate everything. This makes debugging a small subset of Agents a problem if you have to wait for 30+ minutes after each change. To get around this, the Refresh Specific Agents tab provides the ability to specify which Agents you wish to regenerate on-demand.

Simply select the Agents to be regenerated by ticking the Checkboxes and clicking 'Refresh Selected Agents'. The objective will regenerate, and only retrieve data from the selected Agents. It will use the previously retrieved config values to work out the differences reports for each Agent that wasn't refreshed.

8.2.3.2.5. Agent Processing Errors

Agent	Error
CUSTOM-0-1.FRITZ.BOX*	Unable to find a listening port to connect to 10.0.0.1 on, agent could be offline. (Tried: 6161, 6163)
CUSTOM-0-2.FRITZ.BOX*	Unable to find a listening port to connect to 10.0.0.2 on, agent could be offline. (Tried: 6161, 6163)
CUSTOM-0-3.FRITZ.BOX*	Unable to find a listening port to connect to 10.0.0.3 on, agent could be offline. (Tried: 6161, 6163)

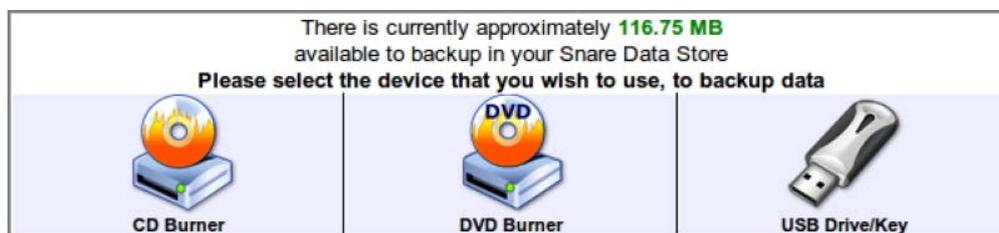
Any errors encountered while regenerating the objective will be listed on this tab. All of the Agents listed under "Agents that cannot be contacted" will be listed on this page with the reason why they could not be contacted. Likewise, when Config Push is enabled, Agents with a configuration that does not match, will have a reason listed here too.

Example errors:

- Unable to find a listening port to connect to 10.0.0.199 on, agent could be offline. (Tried: 6161, 6163)
- Unable to find a password to authenticate to 10.0.4.20:6163 on.
- Reported Agent Version at 10.0.4.30:6161 doesn't match expected type 'Windows'.

8.3. Data Backup

8.3.1. Data Backup



Snare can backup data to optical, or removable USB media.

Select a device type to continue to the data archival process.

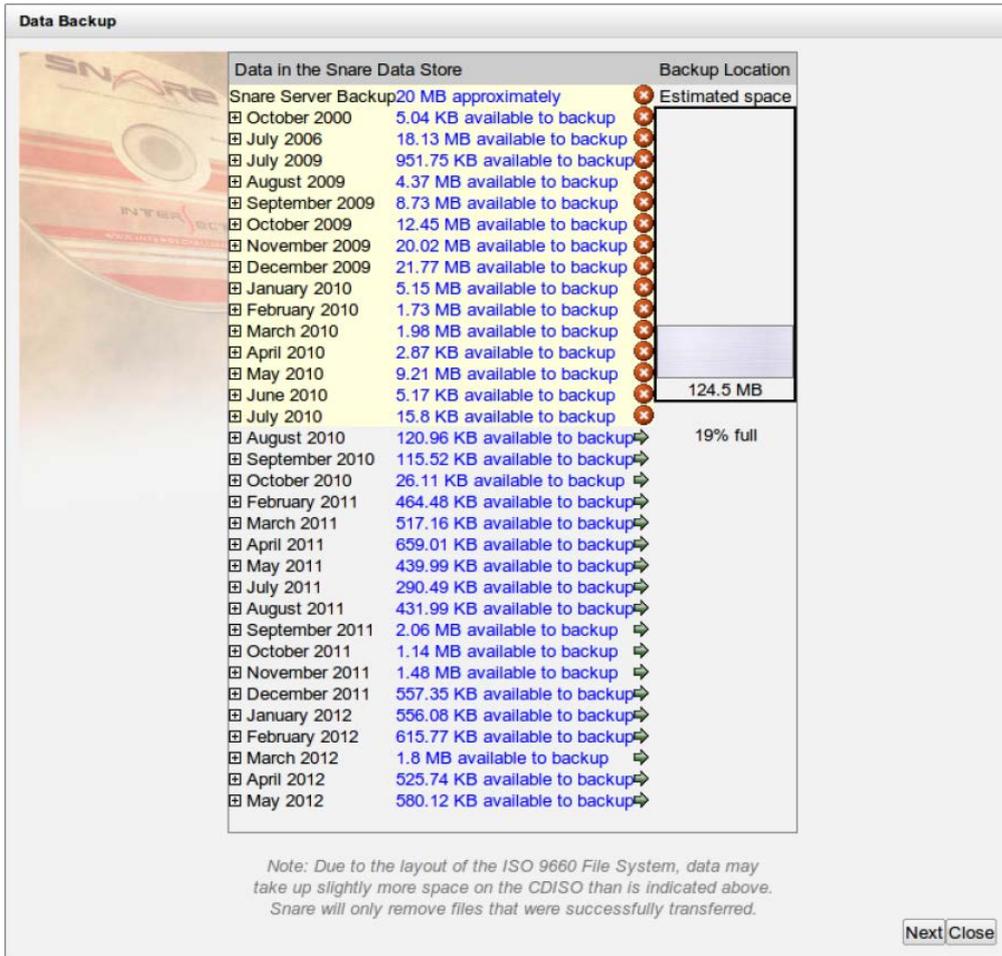
8.3.1.1. Optical Media - Interactive

Selecting either the CD or DVD options will present an option to generate either:

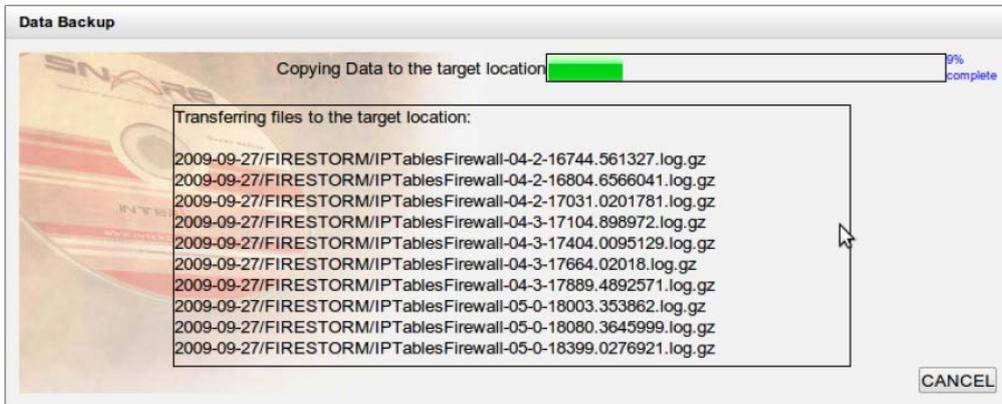
- A CD or DVD, mounted in the local Snare CD/DVD burner, or
- An ISO image of an appropriate size that is filled with eventlog data from your Snare data archive, which can be downloaded to your local workstation to burn to optical media.

A dialog will appear that asks you which months/days to transfer. Months can be expanded in order to include or exclude individual day's worth of data. Months or days can be added to the archive by clicking on the associated arrow button to the right of the date. Months or days can be removed by clicking on the red cross that is associated with the date.

A 'fuel gauge' indicator is available to the right of the dialog, and provides an indication of the current fill state of the CD or DVD.

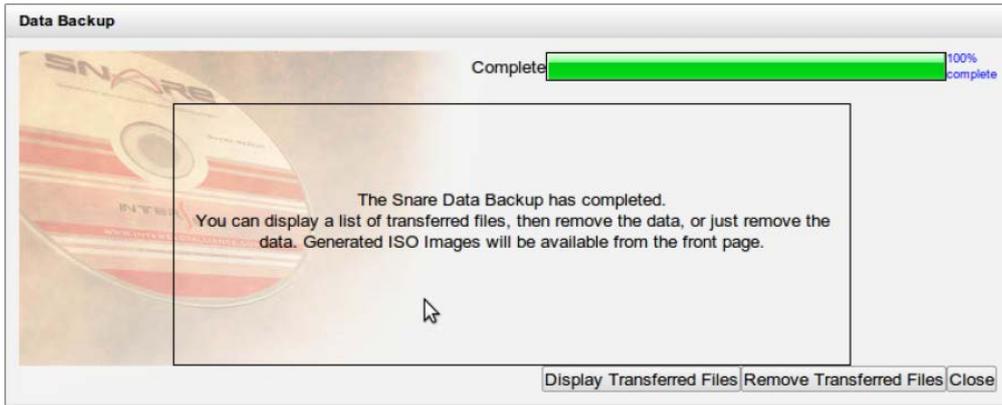


After clicking the 'Next' button, a new dialog will appear, showing the status of the transfer.



Once the process has completed, the dialog will offer you the opportunity to display, or remove the files that have been transferred to CD/DVD.

- ✓ Snare validates the CD or DVD after generation, to make sure that files of the correct name and size have been copied to the optical media. However, for peace of mind, it is highly recommended that the physical media, and contents, be checked on another server before the files that have been migrated off the server, are removed from the Snare data archive.



If you have chosen to generate an ISO image, the image file will be available for download from the front objective output page. You can also choose to remove the CD or DVD from the dialog that pops up when you select the download link, or request an MD5 checksum of the image, to provide a level of assurance that your download matches the image generated by the Snare Server.

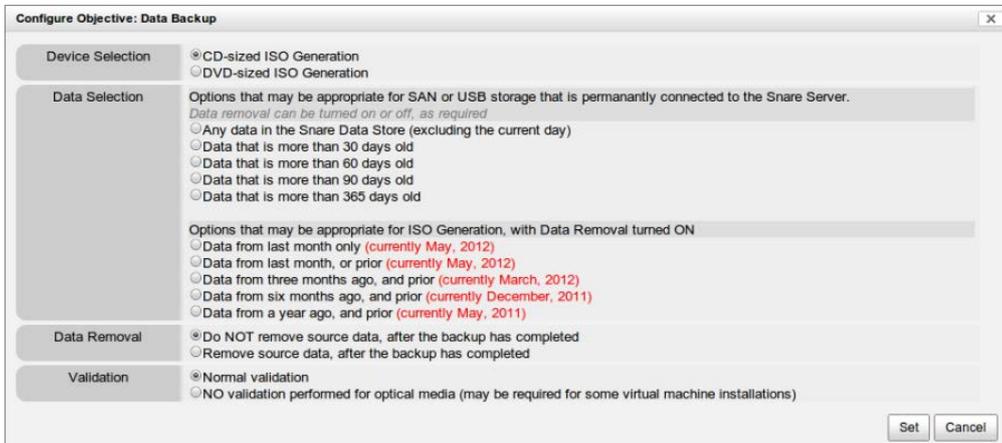


8.3.1.2. Optical Media - Scheduled

When run as a scheduled task, the objective will check the configuration settings for your preferred optical media type (CD or DVD). On regeneration, the objective will create a CD or DVD sized ISO image, which will be available to you to download and burn to a local CD/DVD drive.

Automated ISO generation is only of practical benefit when combined with automated data removal; otherwise, the CD/DVD image generated during each scheduled run, will contain practically the same data as the previous scheduled run. The configuration settings dialog allows you to choose to archive:

- Data from 'last month' only.
- Data that is more than 30, 60, 90 or 365 days old.



8.3.1.2.1. USB Media

Choosing the 'USB Drive/Key' button will allow you to synchronise all, or a portion of your current event log data, with a USB device.

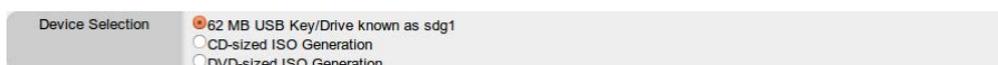


Existing data already present on the device will be compared against the current contents of your data archive, and only new, or changed, data will be copied across to the target device.

Data that already exists on the target device, but has been removed from the Snare Server data store, will not be touched.

- ✓ 1 terabyte external USB drives are common, and reasonably cheap. A 1 terabyte external USB drive can hold somewhere near 40-50 terabytes of compressed snare log data - which is roughly equivalent to a year's worth of data at 5,000 events every second, for the entire year.

When a USB key or hard drive is plugged into the Snare Server, the configuration settings dialog will list the device as a selectable option.



Choosing a USB device as a target device, and setting the objective to regenerate nightly with all data other than the current day, will provide an automated external backup solution for eventlog data. Once you have either filled the external drive, or wish to swap to other media, any data that has been copied over to external storage can be removed manually, and the USB media synchronisation reestablished for the new device.

8.3.2. Remove Data

The Remove Data objective provides the ability to remove data by date, log type or agent.

Date Selection	Log Type Selection	Agent Selection
ANY DATE	ANY LOG TYPE	IRI_100004
2000-10-02	ACF2Log	IRI_100005
2006-07-25	AIXAudit	KINGFISHER
2009-07-01	CISCORouterLog	LAPWING
2009-07-02	CKEPosLog	LOCALHOST
2009-07-03	CyberGuardFirewallLog	MARTIN
2009-07-04	Firewall1Log	NIGHTJAR
2009-07-05	GauntletFirewallLog	PELICAN
2009-07-06	GenericLog	QUAIL
2009-07-07	GenericSyslog	ROSELLA
2009-07-08	IPTablesFirewall	SHEARWATER
2009-07-09	ISAFWSLog	SHELDUCK
2009-07-10	IrixSAT	SILVEREYE
2009-07-11	LinuxAudit	SKYLARK
2009-07-12	MSWinEventLog	SPINEBILL
2009-07-13	MailLog	SPOONBILL
2009-07-14	NetScreenFirewall	SWAN
2009-07-15	NetgearFirewallLog	SWIFT
2009-07-16	NetgearRouterLog	TEAL
2009-07-17	NortelVPNRouter	UNKNOWN SYSTEM

Remove the selected data

Selecting a date (or range of dates), will update the "Log Type Selection" column, to display a list of log types that are available for the chosen date(s). Choosing a log type will update the list of agents that are available for the chosen log types and dates.

Once you are satisfied with your selection, clicking the 'Remove the selected data' button will start the process of removing the actual underlying files, and regenerating the metadata associated with those particular dates, log types and agents.

It may take up to 15 minutes for the changes made by the file removal process to reflect in the list of dates/log types and agents displayed by this objective, or other objectives that rely on the Snare Server metadata subsystem.

8.4. Data Restore

8.4.1. Arbitrary Data Import

The Snare Server can attempt to import arbitrary log data that is text-based, and uses newline (or newline/carriage-return) characters to mark the boundary between different lines. Logs of this format will be imported to either the 'GenericLog' or 'GenericSyslog' data sources, with dates either derived from the uploaded data (if available), or specified within the import form.

Choose up to 9 files to import at once.

Upload files							
File	Log Type	Time	Day	Month	Year	System	
<input type="text"/> Browse...	GenericSyslog	From Log	From Log	From Log	2012	Unknown	
<input type="text"/> Browse...	GenericSyslog	From Log	From Log	From Log	2012	Unknown	
<input type="text"/> Browse...	GenericSyslog	From Log	From Log	From Log	2012	Unknown	
<input type="text"/> Browse...	GenericSyslog	From Log	From Log	From Log	2012	Unknown	
<input type="text"/> Browse...	GenericSyslog	From Log	From Log	From Log	2012	Unknown	
<input type="text"/> Browse...	GenericSyslog	From Log	From Log	From Log	2012	Unknown	
<input type="text"/> Browse...	GenericSyslog	From Log	From Log	From Log	2012	Unknown	
<input type="text"/> Browse...	GenericSyslog	From Log	From Log	From Log	2012	Unknown	
<input type="text"/> Browse...	GenericSyslog	From Log	From Log	From Log	2012	Unknown	
<input type="text"/> Browse...	GenericSyslog	From Log	From Log	From Log	2012	Unknown	

[Go](#)

8.4.2. Snare Data Import

Data that has been exported to optical, or USB media, can be called back into the Snare Server for forensics analysis by this objective.

Alternatively, in situations where a Snare Agent has been configured to log to a local file, rather than, or in addition to, sending log data directly back to a Snare Server for analysis, such files can be uploaded to the Snare Server from this interface, by selecting the 'Upload Snare Agent exports' button.

Please select the device that you wish to use, to restore data



DVD Reader



62 MB USB Key/Drive on /dev/sdg1



Upload Snare Agent exports to the Snare Server

Examples of situations where this option is of benefit are:

- Field laptops that are not generally connected to the local organisation network, or are connected to demilitarized 'safe zones'.
- Systems that have been taken offline due to virus contamination, where forensic analysis of log data may help reveal the infection source.

Log data uploaded via the 'Upload Snare Agent exports' capability can be added file-by-file using an upload form, or alternatively, logs can be zipped together, and uploaded as a single file.

9. Output Modification Modules

The Snare Server provides several layers of increasing flexibility. Although a majority of Snare Server customers will be fully comfortable with creating objectives in the user interface, there are also opportunities for advanced users to change the way that Snare reports data.

Snare uses a range of computer languages to accomplish tasks. High-speed languages such as C and C++ are used in places where speed is critical, such as the front-end collection system, or the code that interrogates the data store. Higher level 'scripting' languages such as Perl or PHP are used for presentation tasks.

 PHP-based modules can be created or modified by customers who are familiar and comfortable with PHP programming in particular, and the UNIX operating system in general. It is HIGHLY recommended that only experienced users attempt this, as there is a risk that an error in your code may disrupt the normal presentation functionality of the Snare Server.

Output modification modules can accomplish several tasks:

- Change the colour of an entire row, based on the content of the row.
- Change the colour of a particular field, based on the content of the field, or the content of a row element.
- Change the content of a field, based on the previous content of a field, or the content of a row element.

Example

- For Windows Security logs, use the USERNAME field to scan for the user in question in the corporate personnel directory. If the user exists, append an image link, so that the user's personnel photo is displayed alongside the user name.
- When an account is created on a Windows server, highlight the user name in green. When an account is removed, highlight the user name in red.
- For PIX firewall events that indicate a packet has been blocked, highlight the entire row in red.

Output modification modules should share the same name as the field, or Token, for which they are designed to modify, but in uppercase characters, and appended with ".php". So, for example, if you wished to create an output modification module for the 'USERNAME' field, you would create a file called 'USERNAME.php'.

If you wished the output modification module to change ANY field called 'USERNAME', regardless of the data source (eg: regardless of whether it was a Windows Security data source, an IP Tables Firewall Log data source, or a Squid Proxy data source), then you would create the file in

```
/data/SnareUI/Global/Modules/USERNAME.php
```

If you wished the output modification module to be specific to a data source, you would create the file in `/data/SnareUI/Global/Modules/DATASOURCENAME/USERNAME.php` - where DATASOURCENAME is the name of the data source (eg: CISCORouterLog, WinSecurity, Tru64Audit)

An example follows that:

- Is stored in the WinSecurity directory (`/data/SnareUI/Global/Modules/WinSecurity/DESTUSER.php`).
- Modifies the 'DESTUSER' field; a Token defined in several Windows Security related objectives.
- Changes the colour of the entire row to RED, if the Windows EventID is '999'.
- Changes the colour of the DESTUSER field to:
 - Green, if the event is related to user creation.
 - Red, if the event is related to user removal.
 - Blue, if the user has been modified or enabled.
 - Orange, if the user account has been disabled.
- Modifies the contents of the DESTUSER field so that the text is surrounded by the HTML "strikeout" elements, if the event is related to user removal.

```

<?php
class DESTUSER
{
    function Colour($text,$row)
    {
        if(in_array($row["EVENTID"],array(624,4720))) {
            return("green"); # User Created
        }
        if(in_array($row["EVENTID"],array(630,4726))) {
            return("red"); # User Removed
        }
        if(in_array($row["EVENTID"],array(625,626,642,4720,4730))) {
            return("blue"); # User modified/enabled
        }
        if(in_array($row["EVENTID"],array(629,4725))) {
            return("orange"); # User account disabled
        }

        # Fallback.
        if(strpos($row["STRINGS"],"Created")) {
            return("green");
        } else if(strpos($row["STRINGS"],"Deleted")) {
            return("red");
        } else if(strpos($row["STRINGS"],"Changed")) {
            return("blue");
        }
    }

    function PrintData($text,$row)
    {
        if(in_array($row["EVENTID"],array(630,4726))) {
            $text="<strike>$text</strike>";
        }
        return($text);
    }

    function RowColour($row)
    {
        if($row["EVENTID"]==999) {
            return("red");
        }
    }
}
?>

```

10. Pre-Processed Tokens - FTokens

Similar in general to output modification modules, pre-processed tokens (or FTokens) change the contents of a particular field. The key difference however, is:

- Normal output modification modules change the content at output-time.
- FTokens change the actual data as it is being processed, before the modular objective matches are applied.

The distinction may not be easy to understand at first, but the differences can be significant.

If, for example, you wished to create an objective that did the following: "[Display out-of-hours logins for any Windows users, who work on the second floor of the building](#)", then you would find it difficult to create an appropriate objective in Snare, without resorting to FTokens.

By default, Snare will be able to help you with the out-of-hours part of the query - each and every event sent to the Snare Server will incorporate a date/time. However, Windows does not know where users are physically located within the building, so there is nothing in the event for Snare to use, to detect whether a user is on the first, second or any other floor of the building.

This is where FTokens come into play. In order to determine on which floor a particular user works you will need to utilise the user name to interrogate an external database (eg: a web-accessible personnel application).

The first step in the process is to create a normal Snare Token called, for example, USERFLOOR. The source field would be the appropriate Windows security log username field (USERNAME, in most cases, but possibly DESTUSER depending on the events that you are interrogating). The regular expression would be simple, and all-inclusive: "(.*)".

The next step is to create a file, with the same name as the Token you wish to use as an FToken, appended with ".php" - eg: USERFLOOR.php. The file can be stored in `/data/SnareUI/Global/FTokens`, if this FToken should apply to ANY field called 'USERFLOOR', regardless of the log data source - or, it can be stored in `/data/SnareUI/Global/FTokens/DATASOURCENAME` - where DATASOURCENAME is the name of the data source (eg: CISCORouterLog, WinSecurity, Tru64Audit)

The screenshot shows a dialog box titled "Add/Modify Field". It has three main sections: "Field Name", "Configure the Field", and "Field Information".

- Field Name:** A text box contains "USERFLOOR". To its right is a button labeled "Remove this Field".
- Configure the Field:** A dropdown menu shows "USERNAME" selected. To the right of the dropdown is the text "What field contains the information you are interested in extracting?".
- Field Information:** A text area contains the following text: "The name you have entered for the Token, is a reserved word in Snare. It will cause a small routine called a FUNCTIONAL TOKEN (created by your Snare Server support team) to be invoked. The functional token will change the content and output of this field during processing." Below this text is a mouse cursor.

At the bottom right of the dialog box are two buttons: "Modify Field" and "Cancel".

This file will utilise a function called 'GetData' to attempt to retrieve the floor that the user is a member of. The GetData function will replace the username, with the floor/level number, BEFORE the modular objective match query runs.

As such, despite the fact that the USERFLOOR token is, initially, merely a copy of the USERNAME/DESTUSER field, you can actually treat it as though it includes the user's floor information when you are constructing your search criteria.

Here is an approximation of the USERFLOOR FToken, using an imaginary web server that supplies username to floor information.

```

<?php
class USERFLOOR
{
    var $Table;
    var $UserFloorData;

    function USERFLOOR($table)
    {
        $this->UserFloorData=array();
        $this->Table=$table;

        # Read in our username to floor correlation information.
        # Assume DumpFloorData returns username,floor number
        $FloorInformation=file("http://personnelserver.dni.gov/DumpFloorData");

        if(!$FloorInformation) { return; }

        foreach($FloorInformation as $data) {
            list($username,$floor)=explode(",",$data);
            $this->UserFloorData[strtoupper($username)]=$floor;
        }
    }

    function GetData($username)
    {
        $username=strtoupper($username);
        return($this->UserFloorData[$username]);
    }
}
?>

```

Using this information, we can now construct a query in our modular objective something along the following lines:



11. Snare Operational Checklists

The purpose of this section is to detail the checklist items required to ensure the correct operation of the Snare Server. This section details those items which are common to all Snare Server installations, regardless of the objectives which have been set.

The next few paragraphs discuss the major components that need to be checked on a regular basis. Clearly, different sites will have their own specific requirements, based on their site, architecture, risk profile, system(s) in use, and so on. For these reasons the following discussion on checklists should be taken as a list of recommendations, which may or may not be adopted. Ultimately, it is up to the user(s) of the Snare Server to be guided on the best strategy for the following Snare Server system checklist items.

11.1. Service (Event Collection) Status

The Health Checker objective checks whether the main services are running. These services are crucial in that they ensure the proper collection of events and without them being active, the Snare Server is unable to collect events. It is strongly recommended that the services be checked on a daily basis, preferably via email notification. Failure of any of the key services will disable the collection of event logs.

11.2. Agent Status

It is not uncommon that agents stop reporting. This may be due to the host being disconnected from the network, the host being re-imaged, the Snare Service being stopped, some third party product 'killing off' the Snare Service, and so on. It is therefore important that these instances be investigated, with a view to correcting the situation. One way to check the Agent status is to use the Manage Agent objective found within System -> Snare Agents -> Remote Management. Simply clone the default objective for each type of Snare Agent running in your network, set the Configuration, and regenerate. It will quickly tell you which Agents cannot be contacted and what versions they are all running.

11.3. Configuration Retrieval

The accurate reporting abilities of a number of objectives relies on the successful retrieval of configuration information, namely the users and group related information for Solaris and Windows (and Cognos to a much lesser extent). It is therefore important to check that this task has been scheduled adequately, and will collect all the necessary configurations. The objectives related to User and Group retrieval, inside the Retrieve Data container within 'Snare Agents' under the System category, should be configured to collect the required user and group information by specifying the server to collect the information (should be a domain controller, but may be a member server in the case of Windows). The password is the password set for the Snare Agent remote control. In order to check that this objective is running correctly, the objective may be 'regenerated' on the fly, and the results checked to verify correct operation.

11.4. Date-Time Completion of Objectives

The date-time completion of objectives provides, at a glance, quick information on whether the objective has been scheduled to regenerate at regular intervals. For those objectives that have been scheduled to regenerate at midnight as part of a scheduled run, the time noted in the objective may indicate if a run is taking too long. For instance, objectives scheduled to run at midnight (which includes all objectives scheduled to run on a 'daily' basis) will usually have a time of 00:20:00, meaning it has completed at 20 minutes past midnight. If a time states something like 03:30:00, then the load on the Snare Server may be becoming onerous, which could indicate the data store is becoming too big or the query is too general.

Alternatively, it may indicate other problems which may require troubleshooting.

11.5. Receipt of Emails

In some instances, scheduled emails may not have been received. This may be due to a number of reasons, such as:

- The objective was not scheduled.
- The objective was scheduled, but no one was listed to receive the email.
- The email was sent out, but the email server rejected it.
- The email was treated as SPAM by an anti-SPAM device.
- The email was sent to the incorrect email address.
- An error with the Snare Server scheduling system.

It is important to check that emails have been received on a regular basis, as determined and required by the agency. The objective 'Modify Objective Schedules' within the Administrative Tools area of the System section, provides a single place from which to view those staff who are receiving scheduled reports, and allows changes to be made directly from this objective.

This objective may also be used to scan the list every so often to remove users that no longer have access, and should not be receiving Snare Server reports.

11.6. Statistics

The objectives within the Status section provide a powerful data management tool. Some objectives in this category are refreshed each time they are accessed, but it is recommended that checks be made to ensure that the following objective(s), as a minimum, are set to 'daily' automatic regeneration:

- General Statistics
- System Status

It is not required that the client receive a daily email report on the above objectives, so long as the objective is scheduled to run and/or viewed on a regular basis. In this way, the report will be available to those who require it, and will not require regeneration. This is important if the data store is of considerable size, and requires a significant amount of time to regenerate.

11.7. Windows Agents

There have been notable problems with the Microsoft Active Directory product changing the audit settings through the use of Group Policy Objects. This means that whilst the Snare Agents are expecting to set and collect audit events based on established and agreed policy, the Active Directory will reset the audit values to those determined by the Domain Administrator. It is therefore important to check that the Snare Agent objective configuration has not changed and if Group Policy is used to configure the audit settings, check that these settings are in line with the information required by the Snare Agent.

11.8. Data Archival

It is strongly recommended that the data archival process be manually checked on an irregular basis to ensure that any media used has not been corrupted. There are two aspects to the archival process, namely the archiving of Snare settings, and the archiving of actual audit event data. Both of these specific objectives are available in the Data Backup section of the System category. Since the CD or DVD write process relies on a number of hardware and Linux related applications, there is always the possibility that the process may not work correctly, and/or may corrupt the write process. Since this process is out of control of the Snare Server, it is strongly recommended that a manual check be undertaken every so often to independently and manually verify the correct archival of information and data.

Sample System Checklist

Priority	System Checklist Item	Description
HIGH	Snare Services	It is strongly recommended that the services crucial to the running of the Snare Server be checked on a daily basis.
HIGH	Configuration Retrieval	It is strongly recommended that the Snare Server be configured to collect system configurations (through the 'Data Retrieval' objectives) on a daily basis, and that this be checked to ensure its accuracy.
HIGH	Email Transmission	It is strongly recommended that user(s) of the Snare Server check the health checker objective is received on a daily basis, and that the other objective-based emails are received on a regular basis.
MEDIUM	Agent Status	It is recommended that the servers running Snare Services (such as 'Snare for Windows') be monitored to determine whether they are still reporting events. It is recommended that this be done on a regular basis, depending on the risk profile of the agency using the Snare Server.
MEDIUM	Removal of Old Data	It is recommended that data be removed on a regular basis, as required.
MEDIUM	Objective Completion	It is recommended that the completion times of some objectives Times be checked to ensure that the Snare Server is operating within reasonable limits.
MEDIUM	Statistics	It is recommended that the 'General Statistics' objective be set to regenerate on a weekly basis.
MEDIUM	Windows Agents	A Windows network, especially one using Active Directory, may propagate group policies down to hosts which contain the Snare Service. This may disrupt the audit event collection process, and compromise some of the objectives. It is recommended that periodic correction of group policy audit settings and Snare Agent settings be undertaken.

11.9. User Checklist

The previous discussions on the System Checklist can also be applied to the User Checklist. The purpose of the Snare Server is to report on system events, and provide summaries of specific security objectives. These objectives are realised through the use of the various parameters available to the Snare Server, including the list of objectives described in this document. In any given agency, the list of objectives, their frequency of reporting and the requirements of reporting will vary greatly.

A generic 'User Checklist' therefore cannot be developed that covers the requirements of all agencies. The purpose of this checklist is to remind users of those items that should be checked every so often, based on the organisational risk profile and infrastructure requirements. It is strongly recommended that a checklist be developed for each agency that has a requirement for the Snare Server.

i Sample User Checklist

Security Objective	Task Description	Checklist
Health Checker	If any 'Problems' or 'Warnings' are shown on the Health Checker objective, they should be investigated immediately.	Daily
Check Domain Administrators Group	This configuration item shows the authorized and unauthorized members of the Domain Administrators group. Notify <position/person> to report those members who are not authorized.	Weekly
Accounts recently created or deleted	This objective reports on all accounts that have been created or deleted. Check that unauthorized users are not creating or deleting accounts, or that newly created, suspicious accounts are reported immediately. Report any incidents to <position/person>.	Weekly
Groups recently created or deleted	Same as above, except for Group accounts. Report any incidents to <position/person>.	Weekly
Modifications to the account of sensitive users	This objective reports on selected, sensitive accounts that have been changed. If any changes to the specified users have been changed, then this must be reported. Use Email Template 2 to report an incident.	Weekly
Firewall Failed Connections by Destination Addr	This objective looks at the failed connections on the 'Company X' firewall. Any failed connections that appear via a yellow or other colored 'square' should be immediately reported. Use Email Template 3 to report an incident.	Weekly
Archival	It is recommended that data be archived on a weekly or monthly basis as required.	Monthly
General Snare Server Check	Check that all agents are reporting back to the Snare Server. Check that retrieval of Windows, Cognos and/or Solaris information is taking place. Note the growth in the size of the data store and, if necessary, trim the data store by removing data.	Monthly

12. Third Party Data Sources

12.1. Snort Sensor

Organisations that use the Snort network intrusion detection system can send data to the Snare Server via the syslog protocol. Snare will be able to collect, interpret, and report on the events. The following information provides an overview of the steps required to configure the Snort sensor to send eventlog data back to the Snare Server. Note that there is no configuration required on the Snare Server.



What you need

The IP address or DNS name of the Snare Server



How to..

On the host that is acting as a Snort collection sensor:

- In the file `/etc/syslog.conf`, add the following two lines:

```
# Send all SYSLOG events to the Snare Server
*. *@12.23.34.45
```

- Please substitute the IP address, or the DNS name, of the Snare Server for the string "12.23.34.45"
- Modify the file `/etc/snort/snort.conf` to include the following line:

```
output alert_syslog: LOG_AUTH LOG_ALERT
```

- An existing (or possibly, multiple) 'output' line may already exist in the file - that is acceptable. Snort will be able to send output to both targets.
- Restart your snort network intrusion detection system and syslog daemon. Depending on your distribution this may be one of:
 - `/etc/init.d/snortd; /etc/init.d/syslog restart`
 - `service snortd restart; service syslog restart`

12.1.1. Troubleshooting Snort

Checking for Snort Sensor errors:

- Look in `/var/log/messages` for errors.
- Run manually:
 - `/usr/sbin/snort -D -i "ppp0" -c /etc/snort/snort.conf`
- ..then look in `/var/log/messages` for errors

12.2. Collecting ACF2 Data

The Snare Server is able to collect ACF2 processed reports, via FTP transfer. The processed reports need to be transferred to a particular directory on the Snare Server, which will then be uploaded by Snare Server processes, on a daily basis.

The ACF2 processed reports are based on specific utilities, provided with the ACF2. The utilities produce formatted reports on the following activity on a mainframe, which can then be collected by the Snare Server, and used for reporting:

- ACFRPTLL Logonid Modification Log
- ACFRPTRL Dataset Rule Modification Log
- ACFRPTEL Infostorate Modification Log
- ACFRPTDS Dataset Violation/Logging
- ACFRPTRV Resource Violation/Logging
- ACFRPTPW Invalid Password Authority Log

The end of this chapter contains a listing of an example JCL which could be used to run, extract and send the

ACF2 processed reports to the Snare Server. This sample job has been set up for the Logonid Modification Log report, but could easily be configured for all the reports listed above. Each step in the sample job below performs the following steps. Note that a fixed transfer library name is used because a reference to this library is stored in an FTP parm library which cannot be changed with each run. Some of the programs used in this job are defined below.

- Deletes previous day's FTP transfer library.
- Runs ACF2 report, placing output in a GDG (7 generations kept).
- Allocate new FTP transfer library and copy report from GDG created in previous step.
- FTP the transfer library to Snare Server. The 'snarexfer' FTP user must be used. This user defaults to the "/data/SnareCollect" directory on the Snare Server. The ACF2 processed reports must be placed in the "ACF2Log" sub-directory. So the full path becomes: "/data/SnareCollect/ACF2Log". Member level security is used to protect the FTP lid password.

The IEBGENER program used in the sample job is an IBM-supplied utility program designed to generate copies of data sets when disk storage or tape is involved. The IKJEFT01 program is the TSO/E program, and is used to perform a TSO function within a batch job.

CSCSNR01

```
***** Top of Data *****
//CSCSNR01 JOB (P,SCF81),ACT.SECURITY,CLASS=C,MSGCLASS=J
/*JOBPARM SYSAFF=PROD
//-----
//*
/* JOB TO PRODUCE ACF2 LIDMOD REPORT FOR XFER TO SNARE SERVER
/*
/*----- DELETE TEMP XFER LIB -----
/*
//STEP1 EXEC PGM=IKJEFT01,REGION=8192K
//SYSPRINT DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSTEM DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSTSIN DD *
DELETE 'CSC.SNARE01.LIDMODS.XFER'
/*
/*----- ACF2 LID DB MODIFICATION LOG REPORT -----
/*
//STEP2 EXEC PGM=ACFRPTLL
//SYSPRINT DD DSN=CSC.SNARE01.LIDMODS.REPORT(+1),
//
DISP=(,CATLG),
//
VOL=SER=BTCH52,
//
UNIT=SYSDA,
//
SPACE=(TRK,(60,5),RLSE),
//
DCB=(GDGMODEL,RECFM=FB,LRECL=142,BLKSIZE=27974)
//SYSUDUMP DD SYSOUT=*
//REC01 DD DSN=CTF.SMFJR,DISP=SHR
//SYSIN DD *
MASK(*****)
DETAIL
NOUPDATE
SYSID(****)
/*
```

```

/*----- COPY REPORT FROM GDG TO XFER LIB -----
/*
//COPY
EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSUT1
DD DSN=CSC.SNARE01.LIDMODS.REPORT(+1),
DISP=SHR
//SYSUT2
DD DSN=CSC.SNARE01.LIDMODS.XFER,
//
DISP=(NEW,CATLG,DELETE),
//
VOL=SER=BTCH52,
//
UNIT=SYSDA,
//
SPACE=(TRK,(60,5),RLSE),
//
DCB=*.SYSUT1
/*
DCB=(RECFM=FB,LRECL=142,BLKSIZE=27974)
//SYSIN
DD DUMMY
/*
/*----- FTP XFER FILE TO SNARE SERVER -----
/*
//STEP4 EXEC FTP,
//
SERVER='CSCSNARE',
//
FTPUSER='SNAREXFER',
//
FTPCMDS='CSCSNR01',
//
ENV='PROD',
//
SOUT='*'
/*
/*----- Notify Security Monitoring Team if job fails -----
/*
/*JOBFAIL IF ((RC > 4) | (ABEND)) THEN
/*
//SENDMEMO EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSUT1 DD *
HELO NCC
MAIL FROM:<PSC0SCHD@AGENCY.COM>
RCPT TO:<ITSECMON@AGENCY.COM>
DATA
TO:ITSECMON<ITSECMON@AGENCY.COM>
SUBJECT:SNARE REPORT FTP JOB FAILURE: JOB CSCSNR01
PLEASE CHECK SDSF OUTPUT FOR THIS JOB ASAP AND DETERMINE WHY.
>> THIS E-MAIL IS GENERATED BY A BATCH JOB RUNNING ON THE
>> AGENCY'S MAINFRAME ENVIRONMENT.
.
QUIT
/*
//SYSUT2 DD SYSOUT=(B,SMTP)
//SYSIN DD DUMMY
/*

```

```
//JOBFAIL ENDIF  
//*=====
```

12.2.1. Lotus Notes / Domino

The Snare Server is able to connect to a Domino server to retrieve eventlog data from log.nsf. It can also retrieve user and group information, plus access controls. However, some of the default settings in Lotus Domino can cause problems with the Snare Agent; please modify the server as follows: From the Domino Administrator page, click the Configuration tab, expand the Web section and click Internet Sites.

1. Choose the log.nsf and click Edit Document.
2. Click the Domino Web Engine tab. Under "Conversion/Display" complete these fields:
 - a. Default lines per view page: 250 (default 30)
 - b. Maximum lines per view page: 0 (default 1000).

13. Regulatory Reporting

13.1. Additional Information

InterSect Alliance provides detailed implementation guides for several major national, international, and industry-specific security frameworks.

More information is available from the web site:

<https://www.intersectalliance.com/strategic-issues/security-regulations/>

13.2. Sarbanes-Oxley (SOX)

Sarbanes-Oxley (SOX) is an Act of the US Congress, designed to mandate specific controls over financial information. The implications for IT systems are more implied than specified, and are governed by the sections detailed below. The ideas presented below are provided as a guide, and should therefore be treated as guidance to be used in conjunction with an agency's risk assessment and security policy.

13.2.1. Section 404: Management Assessment of Internal Controls

SEC. 404. MANAGEMENT ASSESSMENT OF INTERNAL CONTROLS.

(a) RULES REQUIRED.—The Commission shall prescribe rules requiring each annual report required by section 13(a) or 15(d) of the Securities Exchange Act of 1934 (15 U.S.C. 78m or 78o(d)) to contain an internal control report, which shall—

(1) state the responsibility of management for establishing and maintaining an adequate internal control structure and procedures for financial reporting; and

(2) contain an assessment, as of the end of the most recent fiscal year of the issuer, of the effectiveness of the internal control structure and procedures of the issuer for financial reporting.

(b) INTERNAL CONTROL EVALUATION AND REPORTING.—With respect to the internal control assessment required by subsection (a), each registered public accounting firm that prepares or issues the audit report for the issuer shall attest to, and report on, the assessment made by the management of the issuer. An attestation made under this subsection shall be made in accordance with standards for attestation engagements issued or adopted by the Board. Any such attestation shall not be the subject of a separate engagement.

The intent of the above section of the Act is to ensure that controls are in place to safeguard the financial reporting components of an agency, and an assessment is made in relation to the effectiveness of the controls. The Snare Server is used to report on certain events created by the native operating systems, applications or appliances. It is recommended that a risk assessment be conducted when developing an audit and event monitoring strategy. However, the following ideas are presented as a guide as to how the Snare Server can be configured to meet the basic guidelines of the Sarbanes-Oxley Paragraph 404 requirements. It is **strongly recommended** that these ideas be reviewed against the specific requirements of an agency.

Those IT systems within an agency that do not directly support financial and/or audit functions should be monitored for general administrative activity. These systems may include devices and hosts such as network appliances (e.g. routers, firewalls), servers that provide secondary support functions such as DNS, print services, general file services, and workstations. Those servers and other hosts that are used to support key financial systems and applications should be closely monitored. In these instances, there will be additional event collection such as file auditing (to monitor key files and directories), and perhaps event collection from the financial database and/or custom applications that generate the audit events. These will be entirely dependent on the function of the server which may be hosting the financial application or data.

The following is therefore a recommendation on the events to be collected only from these nodes. Again, we do strongly recommend that this collection profile be guided by the outcomes of an agency's risk assessment and security policy.

13.2.2. Network Devices

All management and security events, and failed connections. The management events should include events such as general reconfiguration, reboots and password changes. Usually, events produced by these devices are sent out via SYSLOG, and not controlled by a Snare Agent, in which case, the device should be configured to send administrative/general events and failed connections.

13.2.2.1. General Workstations and Servers

All management and security events, logins and logouts both failed and successful, accounts created and deleted, should be logged from workstations and servers that do not directly support financial activity. The Snare Agents used for collection of such events should thus be configured to collect only those events to support this requirement. In other words, there is no need for process monitoring or file access auditing on these servers and workstations.

13.2.2.2. Servers Used for Key Financial Tasks

All management and security events, logins and logouts both failed and successful, accounts created and deleted. Also, file auditing and process event monitoring should be considered on those directories that store financial reports. Care should be taken in employing file auditing, since this type of auditing can generate a large number of events. File auditing should thus be set on those specific directories that store financial reports. The Snare Operating System Agents can be set to file auditing using the micro web server provided with the agents. Financial database and/or custom applications will usually create log files of those users that have accessed data. These logs types should be considered for collection. If there is no SYSLOG or other facility enabled to send the events to a Snare Server, then contact your Snare Server support team to determine the most appropriate tool-set to collect such logs.

Note that the Snare Server will automatically collect any events sent to it either via UDP/TCP Port 6161 or via SYSLOG Port UDP 514. Events collected in this manner will be stored in the "Generic Log" table. The Snare Server will collect all events sent to it from the Snare Agent and the SYSLOG nodes.

13.2.2.3. Key Snare Server Reports

The following section details those reports in the Snare Server that should be monitored on a regular basis. The settings are the initial recommended settings, and should be fine-tuned once the Snare Server has been in operation for some time.

- Reports -> Operating Systems -> Administrative Activity
 - Monitor account creation related objectives for Windows and Mainframe systems. These objectives should be checked at least once per week. They should be checked to ensure that only authorized staff have been creating or deleting accounts.
 - Monitor objectives related to group modifications for those groups that are used to control access to financial information.
- Reports -> Operating Systems -> Login Activity
 - These objectives should be monitored, depending on the operating system in use at the particular agency. As a guide, the following objectives should be monitored on a weekly or daily basis, depending on the usefulness of reporting:
 - Failed Logins
 - Over a Threshold Value.
 - Per user for a set period.
 - To Locked Accounts (Windows).
- Reports -> Operating Systems -> File and Resource Access
 - These objectives are operating system specific, and can be created to suit specific reporting requirements. If file auditing is required on those servers that process financial information, then the agent must be configured to send the file events to the Snare Server.
 - Note: File auditing can generate an enormous amount of events, and so care must be taken to ensure that only those specific files and directories are audited.
- Reports -> Network
 - Specific router and/or firewall events can be monitored via the objectives in this category. Almost all of these objectives can be set to specifically monitor failed connections and management events.

- New modular objective
 - If events are also being collected from a database or an application, and the Snare Server does not have a specific template data source collection format available in which to store the incoming data, the events will be stored in the 'Generic Log' data source. Remember that the Snare Server will collect any event that is sent to it via UDP or TCP ports 6161, or UDP Port 514 (SYSLOG).
 - If the Snare Server does not have an existing pre-defined template objective to report on the particular events of interest, a new modular reporting objective, set to search the Generic Log data source, can be used to report on these events.
- Status -> Snare Health Checker
 - The Health Checker objective should be monitored daily to ensure the Snare Server is working optimally. This objective monitors all the key elements of the Snare Server, including the collection services, archival functions and agent reporting.

13.2.3. Section 501: Informational Partitions

i SEC. 501. TREATMENT OF SECURITIES ANALYSTS BY REGISTERED SECURITIES ASSOCIATIONS AND NATIONAL SECURITIES EXCHANGES.

501(a)(3) "(3) to establish structural and institutional safeguards within registered brokers or dealers to assure that securities analysts are separated by appropriate informational partitions within the firm from the review, pressure, or oversight of those whose involvement in investment banking activities might potentially bias their judgment or supervision;"

This section of the SOX Act specifies that securities analysts be separated from investment banking activities. This requirement will clearly be agency specific, and therefore the requirement for event monitoring will vary greatly on the functions of the agency. It is therefore recommended that the data owner be consulted as to which resources mandate that the separation be enacted. If, for example, certain directories need to be excluded to a certain group or individual, then the appropriate access controls should be set and event monitoring be set to ensure those users have not accessed the resources. In the case of file and/or directory access, the following Snare Server objectives could be used:

- Reports -> Operating Systems -> File and Resource Access
 - These objectives are operating system specific, and can be created to suit specific reporting requirements. In this case, the monitoring should be set so that access to specific directories is monitored only for those users that are not allowed to view files within those directories. The agent collection must be set so that file and directory access events are being collected for the server(s) in question.
- Reports -> Operating Systems -> Process Monitoring
 - Process event monitoring allows a Snare Server administrator to monitor processes on the host(s) from which those types of events are being collected.
- New modular objective
 - If events are also being collected from a database or an application, and the Snare Server does not have a specific template data source collection format available in which to store the incoming data, the events will be stored in the 'Generic Log' data source. Remember that the Snare Server will collect any event that is sent to it via UDP or TCP ports 6161, or UDP Port 514 (SYSLOG).
 - If the Snare Server does not have an existing pre-defined template objective to report on the particular events of interest, a new modular reporting objective, set to search the Generic Log data source, can be used to report on these events, and specifically, those users that are allowed or not allowed to have access to the application or database.

13.2.4. Section 802: Record Retention

Section 802

§ 1520. Destruction of corporate audit records

(a)

(1) Any accountant who conducts an audit of an issuer of securities to which section 10A(a) of the Securities Exchange Act of 1934 (15 U.S.C. 78j-1(a)) applies, shall maintain all audit or review workpapers for a period of 5 years from the end of the fiscal period in which the audit or review was concluded.

(2) The Securities and Exchange Commission shall promulgate, within 180 days, after adequate notice and an opportunity for comment, such rules and regulations, as are reasonably necessary, relating to the retention of relevant records such as workpapers, documents that form the basis of an audit or review, memoranda, correspondence, communications, other documents, and records (including electronic records) which are created, sent, or received in connection with an audit or review and contain conclusions, opinions, analyses, or financial data relating to such an audit or review, which is conducted by any accountant who conducts an audit of an issuer of securities to which section 10A(a) of the Securities Exchange Act of 1934 (15 U.S.C. 78j-1(a)) applies. The Commission may, from time to time, amend or supplement the rules and regulations that it is required to promulgate under this section, after adequate notice and an opportunity for comment, in order to ensure that such rules and regulations adequately comport with the purposes of this section.

This section is quite clear on the retention of data. The Snare Server will retain events in the database for a period defined by the Migration objective, if removal has been configured: System -> Data Backup -> Data Backup (See the "Configure" tab for this objective).

Once events are older than the period set in this objective, they will be migrated out of the database and onto either optical media, or external storage, in a compressed flat file form. Should the events be required to be reviewed, they can then be loaded onto a forensic Snare Server using the objective System -> Data Restore -> Snare Data Import.

13.3. National Industrial Security Program – Operating Manual (NISPOM)

The National Industrial Security Program – Operating Manual (NISPOM), is a US DoD manual which details the security requirements for those companies and agencies which participate in the DoD Industrial Security Program. The following IT event collection and analysis requirements have been derived from the NISPOM and are reproduced by the sections detailed below. The ideas presented below are provided as a guide, and should therefore be treated as guidance to be used in conjunction with an agency's risk assessment and security policy. The NISPOM handbook is available from the website: <http://www.dss.mil/isec/nispom.htm>

This section contains some ideas on establishing appropriate event logging controls that would be consistent with NISPOM compliance. Of note: it is very important that an organizational risk assessment and IT security policy be used to ensure they are the key drivers for NISPOM compliance. The US DoD Accreditation Authority has the final say on specific compliance objectives.

13.3.1. Section 8-602: Audit Capability

i **8-602. Audit Capability.** Security auditing involves recognizing, recording, storing, and analyzing information related to security-relevant activities. The audit records can be used to determine which activities occurred and which user or process was responsible for them.

a. Audit 1 Requirements

(1) *Automated Audit Trail Creation.* The system shall automatically create and maintain an audit trail or log (On a PL-1 system only: In the event that the Operating System cannot provide an automated audit capability, an alternative method of accountability for user activities on the system shall be developed and documented.) Audit records shall be created to record the following:

(a) Enough information to determine the date and time of action (e.g., common network time), the system locale of the action, the system entity that initiated or completed the action, the resources involved, and the action involved.

(b) Successful and unsuccessful logons and logoffs.

(c) Successful and unsuccessful accesses to security-relevant objects and directories, including creation, open, close, modification, and deletion.

(d) Changes in user authenticators.

(e) The blocking or blacklisting of a user ID, terminal, or access port and the reason for the action.

(f) Denial of access resulting from an excessive number of unsuccessful logon attempts

(2) *Audit Trail Protection.* The contents of audit trails shall be protected against unauthorized access, modification, or deletion.

(3) *Audit Trail Analysis.* Audit analysis and reporting shall be scheduled, and performed. Security relevant events shall be documented and reported. The frequency of the review shall be at least weekly and shall be documented in the SSP.

(4) *Audit Record Retention.* Audit records shall be retained for at least one review cycle or as required by the CSA.

b. Audit 2 Requirements. In addition to Audit 1:

(1) *Individual accountability (i.e., unique identification of each user and association of that identity with all auditable actions taken by that individual).* Periodic testing by the ISSO or ISSM of the security posture of the IS

c. Audit 3 Requirements. In addition to Audit 2:

(1) *Automated Audit Analysis.* Audit analysis and reporting using automated tools shall be scheduled and performed.

d. Audit 4 Requirements. In addition to Audit 3:

(1) *An audit trail, created and maintained by the IS, that is capable of recording changes to mechanism's list of user formal access permissions.*

The intent of the above section of the NISPOM is to ensure that controls are in place to safeguard sensitive DoD information. The Snare Server is used to report on certain events created by the native operating systems, applications or appliances. It is recommended that a risk assessment be conducted when developing an audit and event monitoring strategy. However, the following ideas are presented as a guide as to how the Snare Server can be configured to meet the basic guidelines of the NISPOM Paragraph 8-602 requirements. It is strongly recommended that these ideas be reviewed against the specific requirements of an agency.

Those IT systems within an accredited network that do not directly support the security of information (such as DNS servers, perhaps), may not be required to be audited to the extent detailed in Paragraph 8-602. Please

check with your accreditation authority. The following is therefore a recommendation on the events to be collected only from specific servers and/or workstations. Again, we do strongly recommend that this collection profile be guided by the outcomes of an agency's risk assessment and security policy.

13.3.1.1. Network Devices

All management and security events, and failed connections. The management events should include events such as general reconfiguration, reboots and password changes. Usually, events produced by these devices are sent out via SYSLOG, and not controlled by a Snare Agent, in which case, the device should be configured to send administrative/general events and failed connections.

13.3.1.2. General Workstations and Servers

All management and security events, logins and logouts both failed and successful, accounts created and deleted, should be logged from workstations and servers that do not directly support sensitive information. The Snare Agents used for collection of such events should thus be configured to collect only those events to support this requirement. In other words, there is no need for process monitoring or file access auditing on these servers and workstations.

13.3.1.3. Servers and Workstations Used for Storing and Processing Sensitive Information

All management and security events, logins and logouts both failed and successful, accounts created and deleted. Also, file auditing of the "/etc" directory on *nix systems should be considered. On Windows systems, full process event monitoring should be considered. Care should be taken in employing file auditing on Windows and *nix systems, since this type of auditing can generate a large number of events. File auditing should thus be set on those specific directories that store sensitive information. The Snare Operating System Agents can be set to file auditing using the micro web server provided with the agents.

13.3.1.4. Key Snare Server Reports

The Snare Server will collect all events sent to it from the Snare Agent and the SYSLOG nodes. The following section details those reports in the Snare Server that should be monitored on a regular basis. The settings are the initial recommended settings, and should be fine-tuned once the Snare Server has been in operation for some time.

- Reports -> Operating Systems -> Administrative Activity
 - Monitor account creation related objectives for Windows and Mainframe systems. These objectives should be checked at least once per week. They should be checked to ensure that only authorized staff have been creating or deleting accounts.
 - Monitor objectives related to group modifications for those groups that are used to control access to sensitive information.
- Reports -> Operating Systems -> Login Activity
 - These objectives should be monitored, depending on the operating system in use at the particular agency. As a guide, the following objectives should be monitored on a weekly or daily basis, depending on the usefulness of reporting:
 - Failed Logins
 - Over a Threshold Value.
 - Per user for a set period.
 - To Locked Accounts (Windows).
 - Out of Hours Login Activity.
 - Failed and Successful SU to ROOT (Unix only).
- Reports -> Operating Systems -> File and Resource Access / Process Execution
 - These objectives are operating system specific, and can be created to suit specific reporting requirements. In this case, the monitoring should be set so that access to specific directories is monitored, for those systems that contain sensitive information. The agent collection must be set so that file and directory access events are being collected for the server(s) in question.
 - **Note:** File auditing can generate an enormous amount of events, and so care must be taken to ensure that only those **specific** files and directories are audited. Check for access to the "/etc" directory by root or any other user on *nix systems. On Windows systems, check for process execution events by Administrator or equivalent users.
- Reports -> Network

- Specific router and/or firewall events can be monitored via the objectives in this category. Almost all of these objectives can be set to specifically monitor failed connections and management events.
- New modular objective
 - If events are also being collected from a database or an application, and the Snare Server does not have a specific template data source collection format available in which to store the incoming data, the events will be stored in the 'Generic Log' data source. Remember that the Snare Server will collect any event that is sent to it via UDP or TCP ports 6161, or UDP Port 514 (SYSLOG).
 - If the Snare Server does not have an existing pre-defined template objective to report on the particular events of interest, a new modular reporting objective, set to search the Generic Log data source, can be used to report on these events.
- Status -> Snare Health Checker
 - The Health Checker objective should be monitored daily to ensure the Snare Server is working optimally. This objective monitors all the key elements of the Snare Server, including the collection services, archival functions and agent reporting.

13.4. PCI/DSS Compliance

The Payment Card Industry Data Security Standard (PCI DSS), is a security standard developed by selected companies covering detailed security requirements for protection of credit card information and client details. The following IT event collection and analysis requirements have been derived from the PCI DSS and are reproduced by the sections detailed below. The ideas presented below are provided as a guide, and should therefore be treated as guidance to be used in conjunction with an agency's risk assessment and security policy.

The Payment Card Industry Data Security Standard (PCI_DSS) Version 2.0 Release: October, 2010 PCI DSS is available from the website: https://www.pcisecuritystandards.org/security_standards/documents.php

From the document:

i "The PCI DSS security requirements apply to all system components. In the context of PCI DSS, system components are defined as any network component, server, or application that is included in or connected to the cardholder data environment. System components also include any virtualization components such as virtual machines, virtual switches/routers, virtual appliances, virtual applications/desktops, and hypervisors. The cardholder data environment is comprised of people, processes and technology that store, process or transmit cardholder data or sensitive authentication data. Network components include but are not limited to firewalls, switches, routers, wireless access points, network appliances, and other security appliances. Server types include, but are not limited to the following: web, application, database, authentication, mail, proxy, network time protocol (NTP), and domain name server (DNS). Applications include all purchased and custom applications, including internal and external (for example, Internet) applications."

PCI DSS requirements are applicable if a Primary Account Number (PAN) is stored, processed, or transmitted. If a PAN is not stored, processed, or transmitted, PCI DSS requirements do not apply.

Audit logging capabilities underpin a range of security measures within PCI/DSS, however section 10 of the document specifically addresses logging and auditing. Requirement 10 is reproduced below for reference:

i

Requirement 10: Track and monitor all access to network resources and cardholder data.

Logging mechanisms and the ability to track user activities are critical in preventing, detecting, or minimizing the impact of a data compromise. The presence of logs in all environments allows thorough tracking, alerting, and analysis when something does go wrong. Determining the cause of a compromise is very difficult, if not impossible, without system activity logs.

10.1 Establish a process for linking all access to system components (especially access done with administrative privileges such as root) to each individual user.

10.2 Implement automated audit trails for all system components to reconstruct the following events:

10.2.1 All individual user accesses to cardholder data

10.2.2 All actions taken by any individual with root or administrative privileges

10.2.3 Access to all audit trails

10.2.4 Invalid logical access attempts

10.2.5 Use of identification and authentication mechanisms

10.2.6 Initialization of the audit logs

10.2.7 Creation and deletion of system-level objects.

10.3 Record at least the following audit trail entries for all system components for each event:

10.3.1 User identification

10.3.2 Type of event

10.3.3 Date and time

10.3.4 Success or failure indication

10.3.5 Origination of event

10.3.6 Identity or name of affected data, system component, or resource.

10.4 Using time synchronization technology, synchronize all critical system clocks and times and ensure that the following is implemented for acquiring, distributing, and storing time.

10.4.1 Critical systems have the correct and consistent time.

10.4.2 Time data is protected

10.4.3 Time settings are received from industry-accepted time sources.

10.5 Secure audit trails so they cannot be altered.

10.5.1 Limit viewing of audit trails to those with a job-related need

10.5.2 Protect audit trail files from unauthorized modifications

10.5.3 Promptly back-up audit trail files to a centralized log server or media that is difficult to alter

10.5.4 Write logs for external-facing technologies onto a log server on the internal LAN.

10.5.5 Use file integrity monitoring or change detection software to ensure that existing log data cannot be changed without generating alerts (although new data being added should not cause an alert).

10.6 Review logs for all system components at least daily. Log reviews must include those servers that perform security functions like intrusion detection system (IDS) and authentication, authorisation, and accounting protocol (AAA) servers (for example, RADIUS). 10.7 Retain audit trail history for at least one year, with a minimum of three months immediately available for analysis (for example, online, archived, or restorable from back-up).

13.4.1. Snare Server and Agent Settings for PCI DSS

The intent of the above section of the PCI DSS is to ensure that controls are in place to safeguard cardholder information. The Snare Server is used to report on certain events created by the native operating systems, applications or appliances. It is recommended that a risk assessment be conducted when developing an audit and event monitoring strategy. See our regulatory compliance document on the InterSect Alliance web page, as highlighted in "Additional Information" above, for ideas and as a guide as to how the Snare Server can be configured to meet the basic guidelines of the PCI DSS section 10 requirements. It is strongly recommended that these ideas be reviewed against the specific requirements of an organisation.

14. Snare Agents

The steps below recommend initial settings for Snare Agents. As dedicated tables exist for all of the agents mentioned below, the standard syslog systems must not be used (this will ensure the correct handling of data by the Snare Server). The use of the Snare remote control feature is recommended (the user friendly interface will maintain the appropriate syntax and formatting) and instructions to enable this service will be detailed where applicable.

14.1. Windows-Based Agents

During installation, if the machine is not part of a Windows' domain, allow Snare to take control of the audit subsystem. If the machine is part of a Windows' domain, Local and Group Policies will need to be manually edited to provide the required level of auditing.

Through the remote control interface, click Network Configuration.

1. Leave the hostname override blank. The hostname will be automatically picked up by the Snare Agent.
2. Enter the Snare Server IP address or hostname.
3. Ensure that the remote port is 6161.
4. UDP is recommended for faster and more efficient use of host and network resources.
5. Untick "Enable syslog headers".
6. Click Change Configuration. Next select Remote Control Configuration:
7. Untick "Restrict remote control of Snare Agent to certain hosts".
8. Tick "Require a password for remote access" and enter the enterprise password.
9. Ensure "Change web server default port" is unticked.

The default objectives for the Windows Agent will provide a strong auditing capability. For high traffic systems, remove the process tracking objective (Event ID Match: Process_Events) and add a new objective with the following properties:

- Logon or Logoff.
- Exclude Search Term Users: *\$,SYSTEM.

Further tuning may be required depending on the size and type of environment.

The Snare Server tables used by the Windows agent are:

1. Win Application (Application events)
2. WinSecurity (Security events)
3. WinSystem (System events)
4. MSWinEventLog (DNS, File Replication and Directory Service events).

14.2. UNIX-Based agents

14.2.1. Remote Control

For all UNIX-based agents, the following section should be included in the configuration file to enable remote control capabilities. The user friendly interface will maintain the appropriate syntax and formatting of the Snare configuration files, while also allowing the Snare Server to contact its agents to check their individual configuration settings.

```
[Remote]
  allow=1
  listen_port=6161
```

Once access is granted, please configure the enterprise remote access password through the "Remote Control Configuration" page on the micro-server.

14.2.2. Linux

The latest Linux Agent no longer requires kernel modifications to activate and gather audit information, making installation and upgrade management much easier. The Linux Agent is packaged with a comprehensive set of objectives for NISPOM, SOX and PCI compliance. Once installed, make the following changes via the Snare Remote Control Interface or the configuration file:

1. Change the remote access password to the enterprise password.
2. Remove the log file (this will avoid the risk of local storage partitions filling up).
3. Set the destination server to the Snare Server IP address or hostname.
4. Restart the 'auditd' service.
5. UDP is recommended for faster and more efficient use of host and network resources.
6. Events will be stored in the Snare Server LinuxAudit table.
7. If resource usage becomes a problem, remove all references to the "open" syscall from /etc/snare.conf and restart the agent.

14.2.3. Solaris

Ensure that BSM is properly installed and configured (/etc/security/bsmconv) before installing Snare. The Solaris agent should be installed using the "Advanced" objectives and then tuned accordingly. Once installed, make the following changes via the Snare Remote Control Interface or the configuration file:

1. Change the remote access password to the enterprise password.
2. Remove the log file (this will avoid the risk of local storage partitions filling up).
3. Set the destination server to the Snare Server IP address or hostname.
4. Restart the Snare Agent.
5. UDP is recommended for faster and more efficient use of host and network resources.
6. Events will be stored in the Snare Server SolarisBSM table.

14.2.4. Irix

The Irix Agent should be installed using option three of the install script ("Administrative, Login, process execution and file events related to system configuration files") and then tuned accordingly. Once installed (install.sh), make the following changes via the Snare Remote Control Interface or the configuration file:

1. Change the remote access password to the enterprise password (no password by default).
2. Remove the log file (this will avoid the risk of local storage partitions filling up).
3. Set the destination server to the Snare Server IP address or hostname.
4. Restart the Snare Agent using the restart script /usr/sbin/restartsnare.
5. UDP is recommended for faster and more efficient use of host and network resources.
6. Events will be stored in the Snare Server IrixSAT table.

If for any reason the Snare Agent stops, this will cause the Audit subsystem to panic and send the computer into single user mode. To recover from this mode, enter the command "init 2" from the console. This should return the system to a network accessible state. Emergency logs are available in /sat/ and can be examined using the sat_interpret command. For more details on the IRIX audit system, please visit the following URL:

http://techpubs.sgi.com/library/tpl/cgi-bin/getdoc.cgi/0630/bks/SGI_Admin/books/IA_BakSecAcc/sgi_html/ch06.ht

14.2.5. AIX

The AIX Agent should be installed using the "Advanced" objectives and then tuned accordingly. Once installed (install.sh), make the following changes via the Snare micro-web server or the configuration file:

1. Change the remote access password to the enterprise password.
2. Remove the log file (this will avoid the risk of local storage partitions filling up).
3. Set the destination server to the Snare Server IP address or hostname.
4. Restart the Snare Agent using the restart script /usr/sbin/restartsnare.
5. UDP is recommended for faster and more efficient use of host and network resources.
6. Events will be stored in the Snare Server AIXAudit table.

14.2.6. Epilog for UNIX

After installing Epilog for the first time, there are two main changes that are required:

1. Specify the log files that Epilog should monitor.
2. Set the destination server to the Snare Server IP address or hostname.
3. Syslog option must not be used when sending logs to a Snare Server so that all events are processed correctly by the Snare Server.
4. Send event to TCP or UDP port 6161.
5. UDP is recommended for faster and more efficient use of host and network resources.
6. Generally, events will be stored in the Snare Server GenericLog table.

14.2.7. Tru64

Not released outside of InterSect Alliance at this time.

15. Modular Objective Templates

The following modular objective templates are available in Snare Server version 6+. Each of these templates can be used as a basis for your own objectives.

	<p>Log Type: ACF2Log Title: Changes to Accounts Class: Account Administration</p> <p>This objective displays information relating to ACF2 account modifications. ACF2 logs can be collected from an IBM MVS mainframe, and analysed using the Snare Server. The logs from the ACF2 mainframe are collected via FTP or SCP file transfer into the /data/SnareCollect/ACF2Log/ directory on the Snare Server.</p>
	<p>Log Type: ACF2Log Title: Changes to Flags Class: Account Administration</p> <p>For each CHANGE, DELETE or INSERT, this objective displays the details of the ACF2 changes on an MVS host.</p>
	<p>Log Type: ACF2Log Title: Access to the INFOSTORE database Class: Object Access</p> <p>The Infostore database (ACF60STO) is a sensitive repository of ACF2 information. This objective displays events relating to user access to the INFOSTORE database.</p>
	<p>Log Type: ACF2Log Title: ACF2 Rule Changes Class: Rule Changes</p> <p>This objective displays events relating to changes to Rules. The ability to change ACF2 rules on MVS systems indicates privileged access. This objective is able to monitor anyone that has been modifying these rules. The changing of ACF2 (on those MVS systems that use ACF2) should be carefully monitored to ensure only authorized users are undertaking authorized activity. This objective is able to maintain a view of actions undertaken in this security management activity.</p>
	<p>Log Type: ACF2Log Title: Access to the ACF2 Resources Class: Object Access</p> <p>This objective monitors access to MVS resources that are considered to be sensitive. Use the RESOURCE field to narrow your search criteria. A ReturnCode of VIOLATION, or *VIO, indicates a failed attempt to access the resource</p>
	<p>Log Type: ACF2Log Title: Accounts Created or Deleted Class: Account Administration</p> <p>This objective displays those ACF2 users on an MVS host which have been created or deleted.</p>

	<p>Log Type: ACF2Log Title: User Login Failures Class: User Login Failures</p> <p>This objective displays information relating to ACF2 user login failures.</p>
	<p>Log Type: AIXAudit Title: Executing a Process Class: Process Objectives</p> <p>This objective is used to monitor access to processes or applications that are considered to be sensitive. Note that to use this objective, the Snare Agents must be configured to collect process events.</p>
	<p>Log Type: AIXAudit Title: Access to Sensitive Files Class: File Access</p> <p>Monitor access to file and directories that are considered to be sensitive. Note that to use this objective, the Snare Agents must be configured to report on file accesses.</p>
	<p>Log Type: AIXAudit Title: User Login Class: User Login</p> <p>This objective is used to monitor user login actions on AIX servers. Note that FTP access is also counted as a 'login', but protocols such as SSH or VNC may not generate a login event. It is important that the 'Configure' section of the objective be used to define from which system(s) the login events are required, so that the user(s) of this objective are not flooded with too many login events. This will especially be the case in agencies that are of a significant size, and are collecting events from numerous AIX hosts.</p>
	<p>Log Type: AIXAudit Title: User access to the root account Class: User SU</p> <p>This objective is used to monitor access to the root account through the /bin/su utility.</p>
	<p>Log Type: Browser Title: Messages from installed Snare Browser Agents Class: Browser Objectives</p> <p>Display configuration change and agent restart messages from Snare Browser agents.</p>
	<p>Log Type: Browser Title: Cookie Modifications Class: Browser Objectives</p> <p>Display cookie related events from Snare Browser agents.</p>

	<p>Log Type: Browser Title: Inappropriate material Class: Browser Objectives</p> <p>Display inappropriate material, accessed through a browser.</p> <div style="border: 1px solid yellow; padding: 10px; margin: 10px 0;"> <p>⚠ WARNING INAPPROPRIATE CONTENT MAY BE DISPLAYED WITH THIS RANDOM SAMPLE.</p> <p>The images are linked directly to the target site. This means that your UserID will download the images through your proxy server (if enabled), which means you may appear in your own logs.</p> </div> <p>Please also note that:</p> <ol style="list-style-type: none"> 1) The originating user may not have deliberately accessed the content in question - it may have been a popup caused by a rogue web site, and 2) The image may no longer exist on the target site, in which case, you will receive a 'no image' placeholder within your web browser.
	<p>Log Type: Browser Title: Access to Social Media and Related Sites Class: Browser Objectives</p> <p>Scan for access to social media and related sites. Please unlock this objective, and modify according to your requirements.</p>
	<p>Log Type: ConfigurationCheck Title: CISCO Pix/Router Configuration Checker Class: CISCO Configuration Checker</p> <p>Compare the current CISCO Pix or Router configuration to an authorised version. The objective will attempt to connect to the device using the 'telnet' protocol and display the current configuration. The current configuration will also be compared against an authorized 'Master' and changes will be highlighted. Two passwords in the "Configure" section (connect, and enable) are used to retrieve the configuration.</p>
	<p>Log Type: GauntletFirewallLog Title: Electronic Mail events Class: Electronic Mail</p> <p>The Gauntlet firewall generates events based on the email addresses that have been sent through the firewall. This objective allows the user to report on email information derived from the Gauntlet Firewall logs.</p>
	<p>Log Type: GenericSyslog Title: User Privilege Escalation through SU and Sudo Class: User Privilege Escalation</p> <p>This objective looks for SU or Sudo log entries in the Generic Syslog log source.</p>
	<p>Log Type: IPTablesFirewall Title: Accepted Local Network Connections Class: Local Network Connections</p> <p>Display non-dropped packets that have a source address of a non-routable IP block</p>

	<p>Log Type: IPTablesFirewall Title: Dropped Local Network Connections Class: Local Network Connections</p> <p>Display dropped packets that have a source address of a non-routable IP block</p>
	<p>Log Type: IPTablesFirewall Title: Accepted Non-Local Network Connections Class: Non-Local Network Connections</p> <p>Display non-dropped packets that have a source address of a routable IP block</p>
	<p>Log Type: IPTablesFirewall Title: Dropped Non-Local Network Connections Class: Non-Local Network Connections</p> <p>Display dropped packets that do not have a source address of a routable IP block</p>
	<p>Log Type: IrixSAT Title: General Administrative Tasks Class: Administrative Activity</p> <p>This objective reports on selected Irix audit events which indicate general administrative activity, such as <i>sat_chroot</i>, <i>sat_mount</i>, <i>sat_clock_set</i>, <i>sat_hostname_set</i>, <i>sat_domainname_set</i>, <i>sat_hostid_set</i>, <i>sat_control</i>, <i>sat_bsdipc_snoop_ok</i>, <i>sat_bsdipc_snoop_fail</i>, and <i>sat_ae_audit</i>.</p>
	<p>Log Type: IrixSAT Title: Successful Mount or Unmount Activity Class: Administrative Activity</p> <p>This objective monitors the mount or unmounting of disk volumes on Irix. This may be useful in those instances where it is required that access to specific volumes (such as floppy disks) be closely monitored.</p>
	<p>Log Type: IrixSAT Title: Executing a Process Class: Process Objectives</p> <p>This objective is used to monitor access to processes or applications that are considered to be sensitive.</p>
	<p>Log Type: IrixSAT Title: Access to Sensitive Files Class: File Access</p> <p>Monitor access to file and directories that are considered to be sensitive. Note that to use this objective, the Snare Agents must be configured to report on file accesses.</p>
	<p>Log Type: IrixSAT Title: User Login Class: User Login</p> <p>This objective is used to monitor user login actions on Irix servers. Note that FTP access is also counted as a 'login', but protocols such as SSH or VNC may not generate a login event.</p>

	<p>Log Type: IrixSAT Title: User access to the root account Class: User SU</p> <p>This objective is used to monitor access to the root account through the /bin/su utility.</p>
	<p>Log Type: LinuxAudit Title: Account Management Class: Account Management Objectives</p> <p>This objective is used to monitor account management actions on Linux Servers. Note that the Linux audit subsystem will only generate events when an account or group, is modified using account management binaries. Situations where a root user manually modifies the /etc/passwd or /etc/group files, will not be detected by this objective.</p>
	<p>Log Type: LinuxAudit Title: Accessing a File Class: File Objectives</p> <p>This objective is used to monitor access to files that are considered to be sensitive. Note that to use this objective, the Snare Agents must be configured to collect file events.</p>
	<p>Log Type: LinuxAudit Title: Group Account Management Class: Account Management Objectives</p> <p>This objective is used to monitor group account management actions on Linux Servers. Note that the Linux audit subsystem will only generate events when an account or group, is modified using account management binaries. Situations where a root user manually modifies the /etc/passwd or /etc/group files, will not be detected by this objective.</p>
	<p>Log Type: LinuxAudit Title: Executing a Process Class: Process Objectives</p> <p>This objective is used to monitor access to processes or applications that are considered to be sensitive. Note that to use this objective, the Snare Agents must be configured to collect process events.</p>
	<p>Log Type: LinuxAudit Title: User Login Class: User Login</p> <p>This objective is used to monitor user login actions on Linux servers.</p>
	<p>Log Type: LinuxAudit Title: User Account Management Class: Account Management Objectives</p> <p>This objective is used to monitor user account management actions on Linux Servers. Note that the Linux audit subsystem will only generate events when an account or group, is modified using account management binaries. Situations where a root user manually modifies the /etc/passwd or /etc/group files will not be detected by this objective.</p>

	<p>Log Type: NetScreenFirewall Title: IP Spoofing Notifications Class: Special Alerts</p> <p>Display IP spoofing notifications from Netscreen Firewalls.</p>
	<p>Log Type: NetScreenFirewall Title: Large ICMP Packet Notifications Class: Special Alerts</p> <p>Display Large ICMP Packet notifications from Netscreen Firewalls.</p>
	<p>Log Type: NetScreenFirewall Title: Port Scan Notifications Class: Special Alerts</p> <p>Display port scan notifications from Netscreen Firewalls.</p>
	<p>Log Type: NetworkMapper Title: Network Mapper Class: Network Mapping and Vulnerability Scan</p> <p>This objective allows you to scan your network for open services. New systems, or systems with unauthorized ports, will be highlighted for your attention. In addition, an optional network security scan can be conducted against any hosts that are found. The report may be displayed in tabular format, which is useful for the analysis of many hosts on any given network. In both "iconized" and "tabular" formats, an authorized "port list" can be configured on a host-by-host basis. Future scans against the host in question will highlight any changes in port activation or deactivation. The network scanner can be configured to scan both TCP and/or UDP port ranges.</p> <div data-bbox="376 1317 1449 1487" style="border: 1px solid #FFD700; padding: 5px;"> <p> Warning UDP scanning is very slow and should be used with care. This is because UDP will always have to wait for a timeout to determine if a port is closed. If this timeout is too short, then it will miss valid ports and not correctly report.</p> </div> <div data-bbox="376 1509 1449 1765" style="border: 1px solid #ADD8E6; padding: 5px;"> <p> This objective uses the free Open Source tools, NMAP and OpenVAS. NMap is used to determine the open ports on one or more hosts. OpenVAS provides the security scanning skeleton. Further details are available from: http://www.insecure.org/ and http://www.openvas.org/</p> </div>
	<p>Log Type: NortelVPNRouter Title: Configuration Changes Class: Configuration Changes</p> <p>This objective will watch for configuration changes to Nortel VPN Routers - such as the creation, destruction, or modification of particular configuration items.</p>

	<p>Log Type: NortelVPNRouter Title: Failed Logins Class: Authentication Events</p> <p>This objective will scan for failed attempts to access the VPN device by searching for events that include "Failed Login Attempt" or "failed to log in".</p>
	<p>Log Type: NortelVPNRouter Title: Successful Logins Class: Authentication Events</p> <p>This objective will scan for successful logins to the VPN device by searching for events that include "logged in from", "logged into group" or "login by using".</p>
	<p>Log Type: ObjectAccess Title: Access to the ACF2 Resources Class: Object Access</p> <p>This objective monitors access to ACF2 Objects that are considered to be sensitive. Use the OBJECT field to narrow your search criteria.</p>
	<p>Log Type: ObjectAccess Title: Access to Lotus Notes Resources Class: Object Access</p> <p>This objective monitors access to Lotus Notes Database Resources. Use the OBJECT field to narrow your search criteria.</p>
	<p>Log Type: PIXLog Title: User Authentication events Class: Authentication</p> <p>Display user authentication events</p>
	<p>Log Type: RACFLog Title: Access to RACF Resources Class: Object Access</p> <p>This objective monitors access to RACF resources that are considered to be sensitive. A ReturnCode of 0 implies a failed attempt to access the resource. Use the RESOURCE field to narrow your search criteria.</p>
	<p>Log Type: RACFLog Title: User Logins Class: User Login</p> <p>This objective displays information relating to RACF user logins.</p>
	<p>Log Type: SOCKSLog Title: Failed Authentication Class: User Authentication</p> <p>This objective looks for failed authentication events from a SOCKS server.</p>

	<p>Log Type: SolarisBSM Title: Failed access to a user account Class: User Privilege Escalation</p> <p>This objective is used to monitor failed access to a target account through the /bin/su utility.</p>
	<p>Log Type: SolarisBSM Title: Executing a Process Class: Process Objectives</p> <p>This objective is used to monitor access to processes or applications that are considered to be sensitive. Note that to use this objective, the Snare Agents must be configured to collect process events.</p>
	<p>Log Type: SolarisBSM Title: Access to Sensitive Files Class: File Access</p> <p>Monitor access to file and directories that are considered to be sensitive. Note that to use this objective, the Snare Agents must be configured to report on file accesses.</p>
	<p>Log Type: SolarisBSM Title: User Login Class: User Login</p> <p>This objective is used to monitor user login actions on Solaris servers.</p>
	<p>Log Type: SolarisBSM Title: Access to a target account Class: User Privilege Escalation</p> <p>This objective is used to monitor access to a target account through the /bin/su utility.</p>
	<p>Log Type: Tru64Audit Title: Access to Sensitive Files Class: File Access</p> <p>Monitor access to file and directories that are considered to be sensitive. Note that to use this objective, the Snare Agents must be configured to report on file accesses.</p>
	<p>Log Type: Tru64Audit Title: User Login Class: User Login</p> <p>This objective is used to monitor user login actions on Tru64 servers.</p>
	<p>Log Type: Tru64Audit Title: User access to a target account Class: User Privilege Escalation</p> <p>This objective is used track access to a target account through the /bin/su utility.</p>

	<p>Log Type: UniversalLog Title: Reading Reports Class: Report Access</p> <p>This objective allows you to search for reports that have been read.</p>
	<p>Log Type: UniversalLog Title: Print Reports Class: Report Prints</p> <p>This objective allows you to search for reports that have been printed.</p>
	<p>Log Type: UniversalLog Title: Query Term Analysis Class: Search Analysis</p> <p>This objective allows you to monitor the search terms used in Universal Log data, based on a 'Query' event in the 'Message' field.</p>
	<p>Log Type: UniversalLog Title: User Login Class: User Login</p> <p>This objective allows you to monitor user logins reported in the Universal Log data.</p>
	<p>Log Type: UserGroupSnapshot Title: Account Expiry Class: Account Expiry</p> <p>This objective displays account expiry settings (in days) by system and/or domain. Please note that this objective requires Snare for Windows version 2.6.2 or later.</p>
	<p>Log Type: UserGroupSnapshot Title: Sensitive Groups Class: Sensitive Groups</p> <p>This objective takes snapshots of the applicable group memberships and compares them to a specified list to report on authorized and unauthorized group members. The Snare Server will regularly query the specified server(s) to determine the members of all groups. This is then used by these objectives to determine which users have been authorized to be members of this group, and which are not.</p>
	<p>Log Type: UserGroupSnapshot Title: Account Flags Class: Account Flags</p> <p>This objective displays those users who have settings configured on their account that are considered sensitive or important from a security viewpoint. These attributes are queried on a regular basis by connecting to specified Snare Agents. The updated information will be displayed in these reports, on a scheduled basis, as required by the users of these objectives.</p>

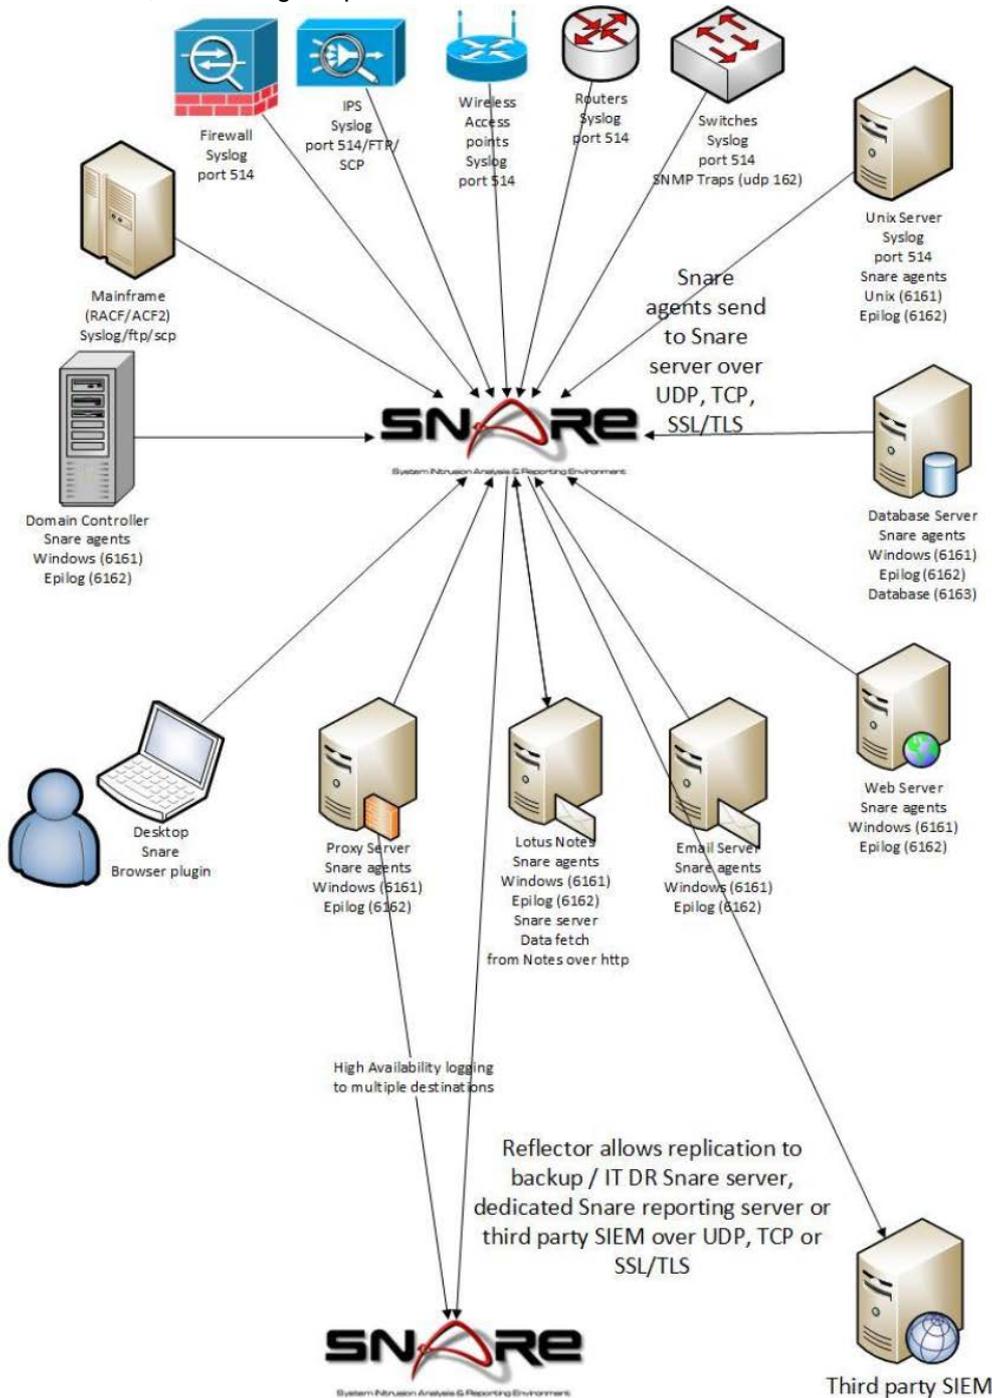
	<p>Log Type: WebLog Title: Inappropriate material accessed through a proxy server Class: Proxy Server Objectives</p> <p>Display inappropriate material, accessed through an organisational proxy server, by searching for a range of defined words, in the URLs that are logged by a proxy server.</p> <div style="border: 1px solid yellow; padding: 10px; margin: 10px 0;"> <p>Warning INAPPROPRIATE CONTENT MAY BE DISPLAYED WITH THIS RANDOM SAMPLE.</p> <p>The images are linked directly to the target site. This means that your UserID will download the images through your proxy server (if enabled), which means you may appear in your own logs.</p> </div> <p>Please also note that:</p> <ol style="list-style-type: none"> 1. The originating user may not have deliberately accessed the content in question - it may have been a popup caused by a rogue web site, and 2. The image may no longer exist on the target site, in which case, you will receive a 'no image' placeholder within your web browser.
	<p>Log Type: WebLog Title: Proxy Server Logs Class: Proxy Server Objectives</p> <p>Query Proxy Server Logs.</p>
	<p>Log Type: WebLog Title: Suspicious URL Access on your web servers Class: Web Server Objectives</p> <p>Display URLs that are generally associated with cross site scripting attacks.</p>
	<p>Log Type: WinApplication Title: NetIQ Administrative Activity Class: Administrative Activity</p> <p>This objective is used to monitor administrative activity using the NetIQ product.</p>
	<p>Log Type: WinApplication Title: NetIQ Group Administrative Activity Class: Administrative Activity</p> <p>This objective is used to monitor group administrative activity using the NetIQ product.</p>
	<p>Log Type: WinApplication Title: NetIQ User Administrative Activity Class: Administrative Activity</p> <p>This objective is used to monitor user administrative activity using the NetIQ product.</p>

	<p>Log Type: WinSecurity Title: Account Creation and Deletion Class: Administrative Actions</p> <p>This objective displays Windows accounts (in specified domains) that have been recently created or deleted.</p>
	<p>Log Type: WinSecurity Title: Group Creation and Deletion Class: Administrative Actions</p> <p>This objective displays Windows groups that have been recently created or deleted.</p>
	<p>Log Type: WinSecurity Title: Failed User Logins Class: User Login</p> <p>This objective is used to monitor failed user login actions on Windows servers.</p>
	<p>Log Type: WinSecurity Title: Windows File Access Class: Windows File Objectives</p> <p>Monitor access to file and directories that are considered to be sensitive. Note that to use this objective, the Snare Agents must be configured to report on file accesses.</p>
	<p>Log Type: WinSecurity Title: Group Modifications Class: Administrative Actions</p> <p>This objective shows modifications to specified sensitive Windows groups.</p>
	<p>Log Type: WinSecurity Title: Audit Log Cleared Class: Administrative Actions</p> <p>This objective checks to see if the Windows event logs were cleared. Note that the Snare Agent must be configured to collect these events. The clearing of an event log may indicate that a user is attempting to cover their tracks.</p>
	<p>Log Type: WinSecurity Title: Changes to the Audit Policy Class: Administrative Actions</p> <p>Although the Snare for Windows agent is able to configure the hosts audit sub-system, this objective keeps an eye on events which indicate an attempt to change the underlying audit configuration. Changes to the underlying audit subsystem may indicate a user that is attempting to hide their "tracks", or attempting to obscure their (potentially) unauthorized activity.</p>

	<p>Log Type: WinSecurity Title: Windows Process Access Class: Process Objectives</p> <p>Monitor access to applications that are considered to be sensitive. Note that to use this objective, the Snare Agents must be configured to report on process execution.</p>
	<p>Log Type: WinSecurity Title: User Modifications Class: Administrative Actions</p> <p>This objective shows modifications to specified sensitive Windows users.</p>
	<p>Log Type: WinSecurity Title: User Login Class: User Login</p> <p>This objective is used to monitor user login actions on Windows servers.</p>
	<p>Log Type: WinSecurity Title: Account Creation and Deletion on Windows and ACF2 Class: Administrative Actions</p> <p>This objective displays any differences between Windows and ACF2 account creation details. In particular, the objective will display:</p> <ol style="list-style-type: none"> 1) Account Creation and Removal for ACF2, that does not have a corresponding create/remove for Windows, and 2) Display Account Creation and Removal for Windows, that does not have a corresponding create/remove for ACF2.
	<p>Log Type: WinSystem Title: Audit Log Corrupt Class: Administrative Actions</p> <p>Display Windows machines that have reported a corrupt event log, during the reporting period. Corrupt event log reporting is only available in Snare for Windows version 3.0.0 and above.</p>

16. Collection subsystem

The Snare Server collection subsystem is a robust group of services that are capable of retrieving data from a variety of different sources, and a range of protocols.



The following services listen on the Snare Server network interfaces, and receive audit and event log data:

16.1. Snare Agent Logs: Port 6161, TCP & UDP

This is the default reception port for all Snare Agent log data.

Log data that is sent to this port, is generally assumed to be formatted in a certain way, in particular:

- Tab delimited data
- The following fields in order:
 - System name

- Log Type (eg: SolarisBSM)

By requiring this format, the Snare Server collection system on port 6161 can significantly reduce the amount of scanning that needs to be performed on each and every input line; leading to a commensurate increase in event collection per second (EPS) rates.

✔ Although the Snare Server operates a syslog collection capability, log data from Snare and Epilog agents should always be sent to this port. Sending the data to the Syslog port, will result in significant EPS rate drops, and may also result in logs that are not filed correctly in Snare's data store; leading to Windows event log data that are filed in the "Generic Log" group, and therefore cannot be used with Windows-related report templates.

16.2. Syslog: Port 514, TCP & UDP

The Snare Server operates a syslog receiver on both TCP and UDP protocols. Many devices use syslog to distribute log data to a central collection point. Network devices such as routers and firewalls often choose this log distribution method.

It is rare that data arriving via syslog provides consistent information that identifies the log source. Usually, Snare will have to scan each incoming event, line-by-line, and pattern match against a series of potential log format templates in order to match a particular event to a log format such as PIX or perhaps Unix SUDO log data. As such, the speed of collection through syslog, is approximately 20% lower than that of the primary collection port.

16.3. SNMPTrap: Port 162, UDP

SNMP traps are used for logging in a number of older network-related appliances. The Snare Server can receive snmptrap data, and make it available for analysis from within the Snare Server interface.

✔ *MIB (management information base) files are not supported by the Snare Server. SNMPTrap messages that rely on MIB files for decoding content, will still be processed - but may be presented with the original numeric content included, rather than the translated/enhanced text-based content. FTOKENs, as described in the documentation above, can be implemented to provide translations for often-monitored events.*

16.4. Browser Collection: Port 6162, TCP

Different browsers, implement slightly different default security profiles for extensions/add-ons. Whilst the Snare for Firefox browser add-on can make a TCP connection to the default Snare Server collection port on 6161, Google Chrome will only allow extensions to connect to a HTTP-compatible server, using the HTTP protocol.

As such, the Snare Server operates a HTTP simulator on TCP port 6162, and accepts encoded data sent from the Snare for Chrome browser extension.

16.5. TLS Server: Port 6163

Several newer Snare agents are capable of sending data over a TLS encrypted connection. The Snare Server TLS Server port can receive such data, and integrate the data into the normal Snare Server collection framework.

16.6. Performance

The EPS collection rates of the Snare Server is significantly dependent on the underlying hardware. In particular, single-core CPU speed is a reliable indicator of the system's ability to collect data.

On a reasonably modern, workstation-quality system (i7 or equivalent), the TCP server can sustainably collect

approximately 16,000 events per second. The UDP collection server can collect events in burst mode, of upwards of 80,000 events per second, without losing data due to CPU limitations, as long as the average EPS rates over the course of an hour, remains at or below the 16,000 EPS threshold.

Higher performance server-quality CPUs, will receive a boost in EPS figures commensurate with their single-core processing capabilities. The Snare Server uses multiple cores where available, to segregate the collection of different log sources (eg: TCP collection is on one core, and UDP is on another), so actual collection rates may be significantly higher for organisations that funnel data to the Snare Server using several different protocols (eg: TCP for servers that have high reliability requirements, UDP for workstations where guaranteed delivery of audit data is less of a concern, syslog for network devices).

17. Expert configuration

17.1. Replacing the Snare Server encryption certificate

The Snare Server generates its own, self-signed certificate for SSL/TLS. It uses this certificate for both encrypted web pages (https), and also to receive encrypted log data.

Although it is not recommended, custom certificates *can* be used to replace the self-signed certificate in the Snare Server. The following instructions assume a reasonably high degree of experience with Unix-like operating systems.



Warning

When the default certificate is replaced by a custom certificate, care should be taken to NOT overwrite the certificate. A new certificate is generally created, whenever the host name is modified within the Snare Server configuration wizard.

If you do NOT wish to overwrite the certificate, please ensure that you choose the option: "DO NOT regenerate the SSL browser certificate even if the server name has changed." in the "General" configuration section of the Snare Server wizard.

The process required to generate a self-signed certificate varies depending on your chosen certificate provider. In general, the instructions available on your providers web site for the combination of 'Linux' and 'the Apache web server', will apply to the Snare Server.

If no such documentation is available, the following general instructions may assist you with the process of requesting and installing a custom certificate:



Requesting a certificate

- Log into the Snare Server as the user 'snare'
- Run the following command:
 - `openssl req -new -newkey rsa:2048 -nodes -keyout snareserver.key -out snareserver.csr`
 - **NOTE: Your certificate provider may support key lengths greater than 2048 bytes.**
 - This command begins the process of generating two files: the Private-Key file for the decryption of your SSL Certificate, and a certificate signing request (CSR) file (used to apply for your SSL Certificate) with apache openssl.
 - When you are prompted for the Common Name (domain name), enter the fully qualified domain name for the site you are securing. If you are generating an Apache CSR for a Wildcard SSL Certificate your common name should start with an asterisk (such as *.example.com).
 - You will then be prompted for your organizational information (including geographic location).
 - The file 'server.csr' will be created. Copy and paste the contents of this file, into the certificate order form of the organisation you will be purchasing your certificate from. You are likely to need to include the BEGIN and END sections of the content.
 - The snareserver.key file, will be required later, for certificate installation.
 - As root, place the file into the directory /etc/apache2/ssl/

Installing the certificate

- You will receive a certificate from your certificate provider, and also, a 'Certificate Chain File'.
 - Upload both of these files to the Snare Server as
 - /etc/apache2/ssl/snareserver.crt <for the certificate>
 - /etc/apache2/ssl/cakeychain.crt <for the key chain certificate>
- Log into the Snare Server as the user 'snare'
- sudo to root
- Make a backup, then edit the file /etc/apache2/sites-available/000-default-ssl.conf
 - Find the following line:
 - SSLCertificateFile /etc/apache2/ssl/apache.pem
 - Modify the line to:
 - SSLCertificateFile /etc/apache2/ssl/snareserver.pem
 - After that line, add the following lines:
 - SSLCertificateKeyFile /etc/apache2/ssl/snareserver.key
 - SSLCertificateChainFile /etc/apache2/ssl/cakeychain.crt
- As root, run the following command:
 - `service apache2 restart`

Warning

Although an upgrade to your Snare Server will not overwrite your certificate, if the Apache configuration file needs to be updated for security or functionality reasons, there is a reasonable chance that your customisations to the Apache configuration file will be overwritten, returning your server back to a 'self signed' state.

After an update, please check your installation, and if required, reapply the 'Installation' information above.