SAWMILL ANALYTICS
ENTERPRISE ANALYTICS SOLUTIONS

Sawmill Reference Guide

Enterprise Analytics

Analysing multiple different log formats from many devices and application programs can be a real headache, and a big investment in time and money for any organisation. Sawmill Analytics answers all of these problems with a single solution: Sawmill. With its open architecture of format plug-ins (currently almost 1,100) Sawmill can concurrently analyse and report on multiple projects of different formats and independent sources, providing highly attractive reports that can be drilled, filtered and interrogated on the fly, and with user-definable alerts

This unrivalled flexibility has resulted in Sawmill being selected as the analytics solution of choice by literally tens of thousands of users, including major international companies, financial institutions and banks, security vendors and service providers, small businesses, self-employed persons, consultants, government departments, web design houses, law enforcement agencies and educational establishments in almost every country of the world.

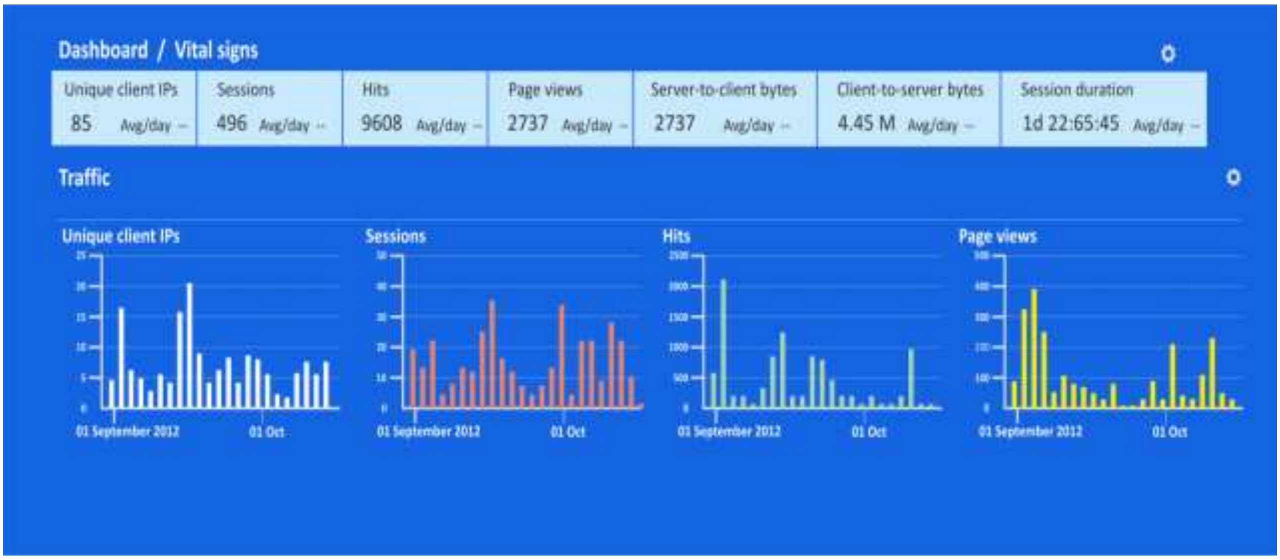

Dashboard / Vital signs

Unique client IPs	Sessions	Hits	Page views	Server-to-client bytes	Client-to-server bytes	Session duration
85 Avg/day –	496 Avg/day –	9608 Avg/day –	2737 Avg/day –	2737 Avg/day –	4.45 M Avg/day –	1d 22:65:45 Avg/day –

Traffic

No other product comes close to Sawmill for enterprise-wide analytics.

Sawmill architecture

Sawmill is multi-discipline multi-tasking software. A single copy can concurrently analyse and report on many different and separate tasks, providing business management and security intelligence data to all interested parties throughout an organisation using nothing more than a standard modern browser for access (provided permission is granted). The benefits of fast accurate data correctly shared to the right people brings major benefits for business efficiency and security.

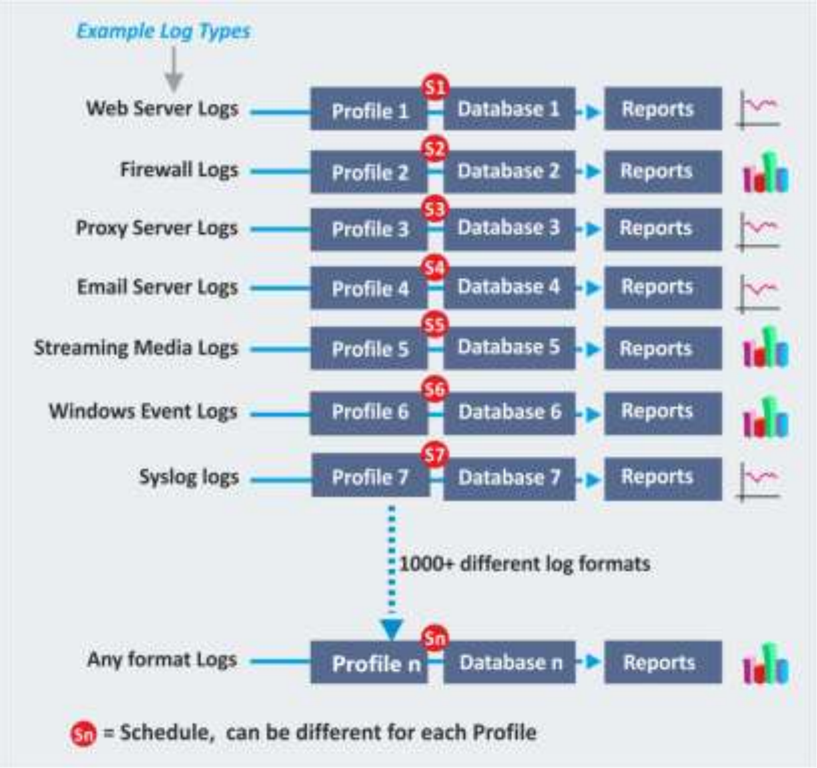

Who uses Sawmill

Anyone who needs to understand how their business works, how safe it is, how company's resources are being used, how their customers are being served, their internet bandwidth, who is consuming the bandwidth and how, what visitors are looking at on their website, what is being purchased, who has accessed which fileservers and which files, who is on their network and who is logged-in, websites being surfed by their employees, how many threats were blocked, etc. etc. etc. This type of information is needed by every manager or business owner hoping to keep better control of his business and his employees in this interconnected world by reducing risks and staying legal.

User Access Rights

User access can be controlled by the system administrator with the powerful ARBAC (Authentication & Role Based Access Control) system built in to Sawmill Enterprise. Users are authenticated before access, and then allowed to perform certain tasks as granted by the system administrator. Unlimited roles can be defined, and unlimited users grated various degrees of access.

On-Premise or Cloud Installation

Sawmill is designed as an on-premise tool running on a dedicated Sawmill server under the control of the licensee. Keeping log data and reports private are the non-negotiable aims of any responsible organisation, and that includes the wrongful use, accidental or otherwise, by a supposedly friendly service provider or cloud operator. For these reasons we continue to position Sawmill as an 'on premise' solution and not a cloud based solution. However it should be noted that Sawmill can access and report on data wherever it is stored, and that includes data stored in the cloud.

What host platforms are supported

The downloadable trial version contains binaries for Linux, Windows and Macintosh, plus encrypted source code for compiling to other platforms. Hardware specifications for the Sawmill server are dependent on the size of the log files and the 'live data' retention period, plus the pattern and type of use (live reports, static reports, report requests and request frequency etc.). Memory should be 2GB for each processor core, but more memory is always a bonus.   Good processor platforms for Sawmill are Intel or AMD, with Windows and Solaris on Intel/ AMD also very good. Sawmill is also installable in a virtual machine environment

Sawmill versions

- Sawmill Lite
- Sawmill Professional
- Sawmill Enterprise

see feature comparison table here:
http://www.sawmill.co.uk/matrix.php

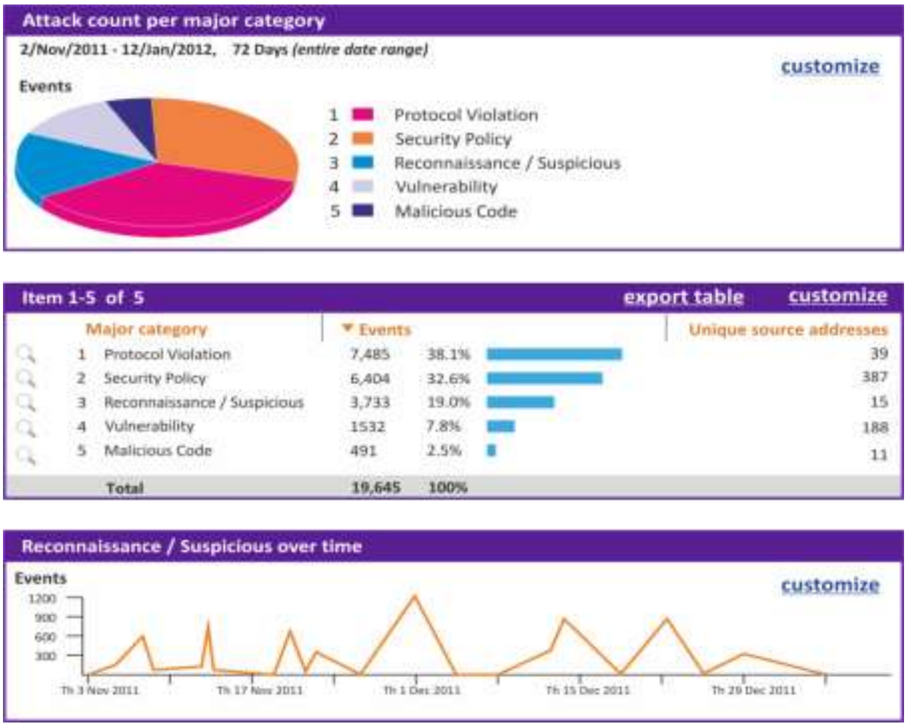

Attack count per major category

2/Nov/2011 - 12/Jan/2012,   *72 Days (entire date range)*

Item 1-5 of 5

	Major category	▼ Events		Unique source addresses
1	Protocol Violation	7,485	38.1%	39
2	Security Policy	6,404	32.6%	387
3	Reconnaissance / Suspicious	3,733	19.0%	15
4	Vulnerability	1532	7.8%	188
5	Malicious Code	491	2.5%	11
	Total	19,645	100%	

Reconnaissance / Suspicious over time

How does Sawmill access log files

Sawmill is agent-free. Logs stored on a network accessible drive are accessed pointing and clicking on the file containing the logs to give Sawmill the path. When a Profile runs Sawmill will automatically import the incremental logs not already imported. Logs stored remotely can be accessed by ftp/sftp.   When logs are imported they are pre-processed by a plug-in and parsed into the normalised Sawmill internal format and entered into the database. Each Profile creates and maintains its own database.

Using syslog to collect logs

As an alternative to direct network access to a specific file where the logs are it is also possible to use a syslog product to collect and forward to a syslog server that Sawmill can access. Syslog can bring many valuable benefits to secure data collection, plus it frees the Sawmill server to concentrate on processing of the logs. Sawmill already supports all popular syslog servers and can read their formats and we can recommend and supply excellent syslog solutions on request

What host services does Sawmill need

Sawmill is a free standing program requiring only the services of an operating system. It even includes its own web server, but can use an external web server for publication if preferred

How does Sawmill access Windows Event logs

The Sawmill SIEM extension exports windows event logs directly from all windows containers on a schedule set up by the Sawmill system administrator. As the logs are exported they are also converted on-the-fly from a binary format to a text base format for processing by Sawmill

REPORTS AND ALERTS

What is a Report

Log files consist of fields of data. The Sawmill plug-in will interpret and convert the majority of the fields into *report pages* - one for each field, including computed fields such as geo-location or session information*.* A *report page* may contain one or more elements such as a line graph, bar chart, pie diagram, and/or table of data depending on the type of data being presented by Sawmill (Sawmill makes this initial decision). A *Report* is the collective name for all the individual *report pages* and *dashboards*

What is a Dashboard

A Dashboard is normally the default report in Sawmill and consists of a group of *report pages* combined on a single screen. The user can determine how a dashboard is constructed (i.e. what *report pages* it contains) so that the most important data is collectively displayed on a single screen for immediate assimilation and action by the user. A Report can have multiple dashboards

Publishing of Sawmill Reports

Sawmill contains its own dedicated web server for publishing reports (Webserver mode), or it can make use of any other web server (CGI mode).

.

Live Reports

Sawmill reports are generated directly from the database and are therefore 'live'. Live reports can be drilled, filtered and queried and re-published on demand at any time. The significant benefit of this approach is that every report will contain the latest imported log data.

Static Reports

Sawmill can also produce static html reports for viewing only, saving processing power etc.

Custom Reports

Custom reports can be created and log streams correlated (combined) where possible to create derived reports of very high value.

Real-world names/headers/titles

External metadata tags and labels can be imported into the reports during report generation to make the resulting reports far more readable and understandable.

Alerts

Sawmill can generate alerts based on content or events identified in the log file

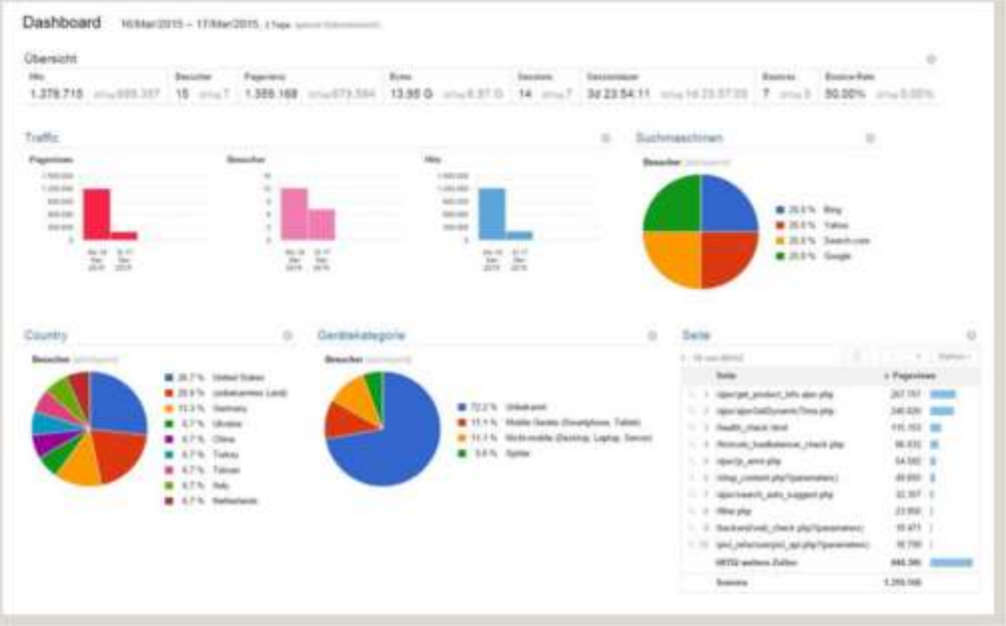

SAWMILL LICENSING

Sawmill licensing for End Users

End users pay a single fee for a perpetual license. Unless stated in writing to the contrary the licensee fee is for a single installation. The license fee is calculated according to the number of Profiles. Profile packs are available in the following sizes: 1, 5, 10, 25, 50, 100, 500, 1000. Additional Profiles can be added to an existing licensed copy at any time, without losing the original investment or data

Sawmill licensing for Service Providers

Service Providers pay a fee for a 12 month license. The license fee is calculated according to the number of Profiles. Profile packs are available in the following sizes: 1, 5, 10, 25, 50, 100, 500, 1000. Additional Profiles can be added to an existing licensed copy at any time, without losing the original investment or data

Licensing limits in the License Agreement (EULA)

- Number or users - no limit
- Size of log files - no limit
- Database size - no limit

PLUG-INS & PROFILES

What is a Plug-in

A plug-in is a log filter that parses and normalises the incoming log data into the Sawmill internal format for input to the database. There are almost 1100 plug-ins in Sawmill, each one developed for a specific log or event format to extract the usable data from the log.

New Plug-ins

New plug-ins are created by Sawmill Analytics when a new log or event format is encountered or the original vendor of a system or application program modifies his logging strategy. The creation of a new plug-in can take anything from 1 day to 1 week depending on the complexity of the format. Skilled sysadmins may also be able to develop their own custom plug-ins

What is a Profile

A log files is imported and the reports generated under the control of a Profile, and each Profile will only import logs of a single format type. Sawmill Professional and Sawmill Enterprise can run multiple different Profiles concurrently, accessing and importing logs of different formats from different locations. For Professional and Enterprise editions Profiles are available in the following pack sizes: 1, 5, 10, 25, 50, 100, 500, 1000. Sawmill Lite can only have a single Profile and import logs from a single source.

Creating a Profile

To create a Profile the user can use the 'profile wizard'. This guides the user through the process of creating a Profile by requiring answers to five simple questions: 1. are the logs are local or remote, 2. define the path to the logs, 3. internal or external database, 4. define the fields required in the report, 5. give the Profile a name. On completion the user can run the Profile manually, or set up a schedule to run the Profile on a regular basis.

How Many Profiles are needed

Count how many reports you need to produce on a regular (scheduled) basis and that is how many Profile you will need as a minimum. You may need more Profiles for administration or development purposes, or to produce multiple different reports from the same original log file data. As an example, if your project involves processing 10 different log files, and producing 2 reports from each, then as a minimum you will need 20 Profiles, so you would purchase the 25 pack. This will give you 5 spare for development/administration.

OPERATIONAL CONSIDERATIONS

Best Operational Practice

Sawmill can run multiple profiles, the number limited only by the license (the activation key) and the hardware. Invariably the database will grow and good housekeeping will be needed, but when new areas of analytical analysis are added together with more profiles then it may be better and more resilient to use multiple copies of Sawmill on multiple servers, instead of increasing the size of a single server . By dividing the processing into naturally occurring segments the exposure to failure or risk is reduced.

Which database: Sawmill internal or external SQL

All Sawmill products support the internal Sawmill proprietary database which scales well and is fast, being designed especially for Sawmill and the storage/manipulation of log data - handling large log files in excess of 1GB per day, and potentially without limit. If the customer's preferred database be used the Sawmill Enterprise edition supports ORACLE, MS SQL-SERVER, MySQL. Sawmill Professional and Sawmill Lite are limited to the internal database only.

How big is the Sawmill Database

Standard Sawmill includes xrefs and indices into the database along with the parsed log files. Including these allows for faster report generation and filtering/drill-down, but they can be switched off in the Profile set-up. If left as default (i.e. xrefs and indicies included) then the database can grow in size between 2x - 10x the size of the original log file, and take longer to build. Our general advice is therefore to use a fast hard disk drive with enough space for both the database and the raw log files when aggregated over the retention period.

Cluster Processing and High Availability

Sawmill can be operated in a cluster where multiple copies process and produce identical reports based on the same dataset/database. Full details are contained on our website. Adding a Load Balancer in front of a cluster will create a high availability cluster, diverting traffic away from failed nodes to live nodes.

Sawmill Customisation Rights

Customization rights are granted by the Sawmill EULA (End User License Agreement) and as summarised below. The EULA should be consulted at all times when modification of the user interface is being considered. A copy of the EULA is always available on request.

- **Sawmill Lite** - does not allow customization of the user interface
- **Sawmill Professional** - allows the text attributes (colour, fonts etc.) to be modified. Two areas of the user interface are also made available for the placement of 'white label' logos or other graphic items.
- **Sawmill Enterprise** - allows the total customization of almost everything on the user interface, attributes, placement, and content. Everything except the Sawmill logo and Copyright notice

Important note: the Sawmill logo and the copyright notice must never be changed, moved, modified or obscured without the prior written permission of Sawmill Analytics. Any attempt to do this will invalidate the customer's license to use the software as defined in the Sawmill EULA.

Sawmill Analytics Support programs

- **Installation Support** – free priority support for the first 30 days
- **Free Support** - first year only. No SLA, lowest priority
- **Premium Care Support** - expedited high priority support by email, telephone, Skype and TeamViewer etc. and includes free access to all available updates. Annual fee

Sawmill Analytics Consultancy services

- **Special Plug-in Development** - a professional consultancy service for the creation of new or modified plug-ins
- **Special Profile Development** - a professional consultancy service for the creation of new Profiles
- **Installation and Consultancy** - a professional consultancy service providing on-site or remote installation/configuration
- **Training Course for system administrators** - a professional consultancy service providing remote training for system administrators

Any Log • Any Format • Any Platform • Anywhere

@ Sawmill