

What's New In Sawmill 8

"Why Should I Upgrade To Sawmill 8?"

Sawmill 8 is a major new version of Sawmill, the result of several years of development. Nearly every aspect of Sawmill has been enhanced, and many major new features have been added. This document describes the major improvements introduced in Sawmill 8, and compares Sawmill 8 to Sawmill 7. Some features mentioned below are not available in lower licensing tiers of Sawmill, but all are available in Sawmill 8 Enterprise.

Enterprise Database Support: Microsoft SQL Server and Oracle

New Profile Wizard

Back Next Cancel

Database options

Please choose the preferred database.

Database server type: MS SQL (ODBC) Internal MySQL MS SQL (ODBC) Oracle (ODBC)

DSN:

Username: (optional)

Password: (optional)

Database name*: (leave blank to use the profile name)

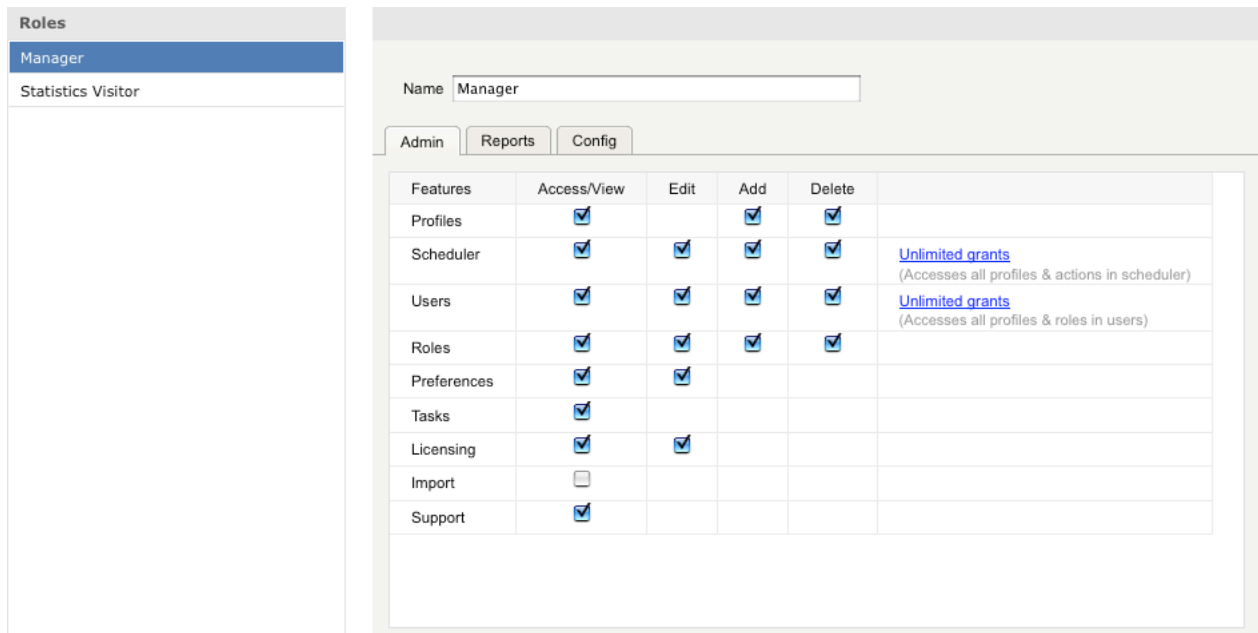
SQL table name prefix: (optional)

SQL table name suffix: (optional)

Sawmill 8 adds support for Microsoft SQL Server and Oracle, to the MySQL and internal databases supported by Sawmill 7. Like Sawmill 7 with MySQL, this “back-end database” use is transparent once it is configured--a profile which uses Microsoft SQL Server or Oracle works just the same as a profile which uses MySQL or the internal database, with all the same configuration and reporting capabilities. The only change is the location of the database storage, and the database engine used to run queries.

Use of an enterprise database engine with Sawmill can provide better performance through clustering, as well as improved reliability, redundancy, backup and recovery, and database management.

Role-Based Authentication Control (RBAC)



The screenshot shows the 'Roles' configuration page in Sawmill 8. On the left, a sidebar lists roles: 'Manager' (selected) and 'Statistics Visitor'. The main area shows the configuration for the 'Manager' role. At the top, there is a 'Name' field containing 'Manager'. Below it are tabs for 'Admin', 'Reports', and 'Config'. The 'Admin' tab is active, displaying a table of permissions for various features.

Features	Access/View	Edit	Add	Delete	
Profiles	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
Scheduler	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Unlimited grants (Accesses all profiles & actions in scheduler)
Users	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Unlimited grants (Accesses all profiles & roles in users)
Roles	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
Preferences	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			
Tasks	<input checked="" type="checkbox"/>				
Licensing	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>			
Import	<input type="checkbox"/>				
Support	<input checked="" type="checkbox"/>				

Sawmill 8 adds Role-Based Authentication Control, for highly granular user permission management. In Sawmill 7, there were only Administrators, who could do anything, and non-Administrators, who could only view the reports of specified profiles (and could not configure profiles). Sawmill 8 extends this by allowing the creation of any number of roles, each with specific permissions, and any number of users in one or more role. For instance, it is possible to create a role that can manage a particular profile, but not other profiles; and that can edit the log filters of that profiles, but not delete them; etc. The permissions of each user can be controlled at a very detailed level.

Real-Time Importing/Reporting

Sawmill 7 always operated in a “batch” mode: log files were imported periodically, often nightly; reporting was not available during import; and reports were then generated from that snapshot of the data, until the time of the next import. Sawmill 8 supports true real-time importing and reporting. A profile can be configured to read a continuous stream of log data, adding each line to its database as it appears on the stream. Reports can be generated at any time, and show the latest data in the database as of the moment of report generation. This allows a profile to show up-to-the-second reports at any time.

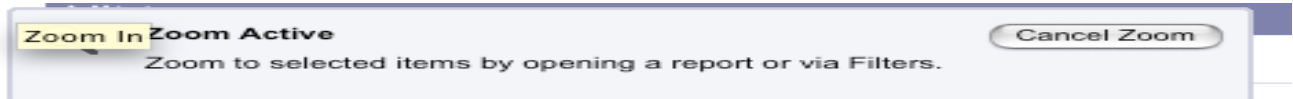
Improved Memory Management

Sawmill 7 could run out of memory on 32-bit systems, due to a number of memory management approaches which could cause high memory (or address space) usage for large datasets. Sawmill 8 improves memory management to keep memory usage low for even very large datasets, allowing enormous datasets to be processed on systems with limited memory, and on 32-bit systems.

Built-In SQL Database

Sawmill 8's internal database supports a subset of SQL, allowing information to be queried directly from the internal database using SQL statements (Sawmill 7's internal database does not support SQL queries). Sawmill 8 also uses a unified set of SQL queries internally to perform database operations on all supported databases (the internal SQL database, MySQL, Oracle, and Microsoft SQL Server), including the internal one, for better reliability and performance.


Enhanced Reporting User Interface



Item 1 - 5 of 5

Hit type	Hits	Page views
1 hit	3,699 74.0 %	0
2 page view	1,264 25.3 %	1,264
3 broken link	32 0.6 %	0
4 spider	3 0.1 %	0
5 error	2 0.0 %	0
Total	5,000 100.0 %	1,264

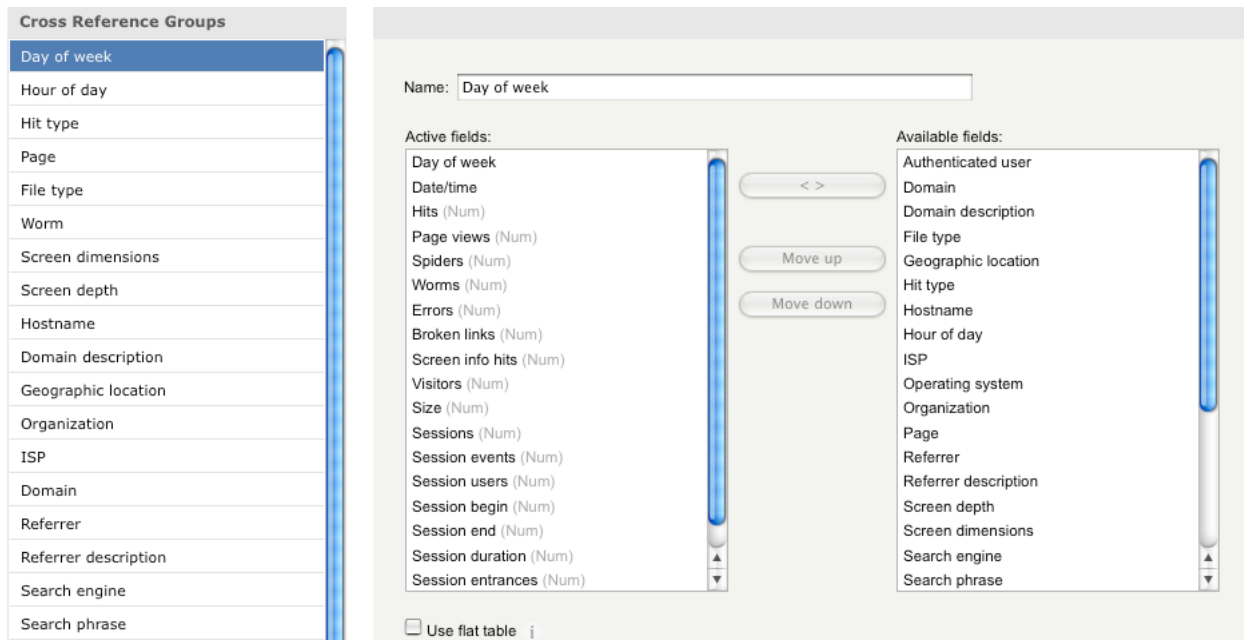
Page views (descending)



Sawmill 8 has a completely redesigned reporting web interface. The interface will be familiar and an easy transition for Sawmill 7 users, but has many improvements to simplify and enhance report viewing. Improvements over Sawmill 7 include:

- Frames have been replaced with multi-column pages, for better integration with larger environments.
- Zooming is easier, due to a new approach that involves clicking an item, and immediately clicking a report name in the left menu to zoom to it (instead of using a “zoom to report” menu like Sawmill 7).
- Zoom supports clicking on multiple items to zoom simultaneously on several rows.
- Filters can be grouped into Filter Sets, and named and managed easily.
- A new Macros feature saves the current report and/or filter set, so it can be easily called up again later.
- Report customization from the Reports page is greatly enhanced, providing most of the functionality of Sawmill 7’s Report Editor without leaving the Reports page.
- Graphing improvements include 3D pie charts and smoother (anti-aliased) graphs.
- The current report can be emailed immediately.
- It is now possible to filter on numerical fields with Report Filters, for instance to select all events with bytes > 1024.
- Session fields are integrated with other fields, so any report can include a “sessions” column, a “session duration” column, or other session-related columns; and session reports like Session Pages can contain non-session columns like “referrer.”
- New “report table filters” can filter the table after it is generated, allowing dynamic suppression of table rows; e.g., show only rows with more than 15 page views.
- Pivot tables can now be created dynamically from the Reports interface; for instance, an indented “audio codec” column can be added to a “video codec” report without leaving Reports.
- An improved date filter picker allows relative date selections like “last 3 months” or “yesterday.”

Enhanced Administrative User Interface

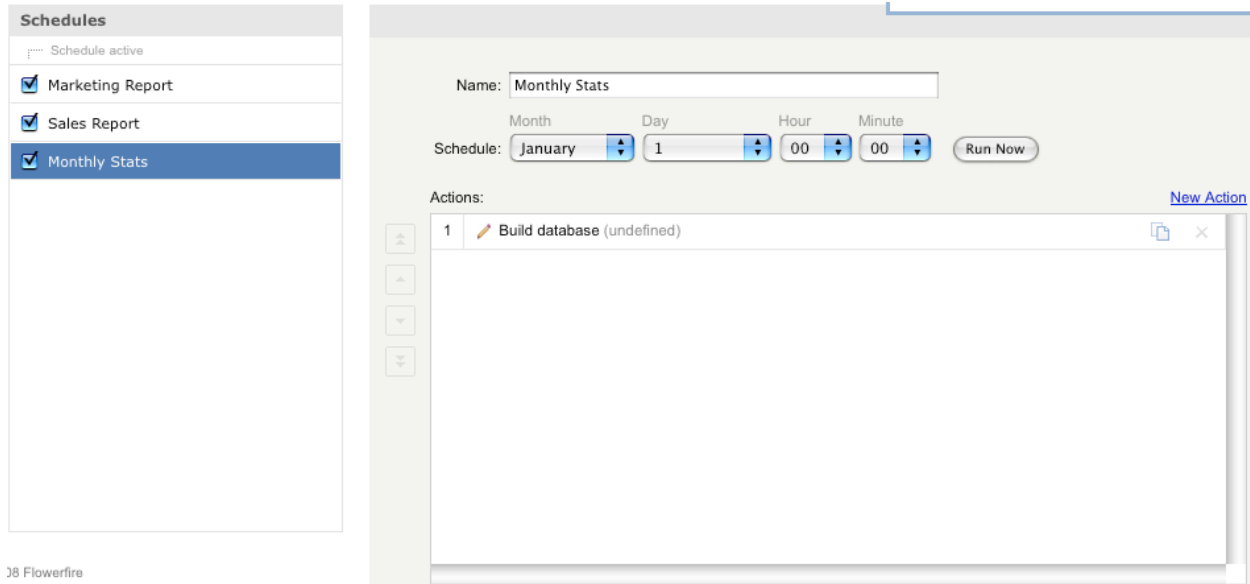


Sawmill 8 has a completely redesigned Admin and Config user interface. The interface will be familiar and an easy transition for Sawmill 7 users, but has many improvements to simplify administration and configuration. In addition, there are several major new components that were not present in Sawmill 7:

- The Log Fields Editor
- The Database Fields Editor
- The Session Fields Editor
- The Report Fields Editor
- The Cross-Reference Groups Editor
- The New Field Wizard. This new wizard makes it much easier to create custom fields; it creates a new log field, an associated database field to track it, a report field for use in reports, a log filter to set its value, a cross-reference table to cache it, and a new report to display the results, all in a single step.

Sawmill 8 allows you to edit *all* aspects of the profile from the web interface. This is a marked improvement over Sawmill 7, where many of the advanced customizations required direct editing of the CFG files. For instance, adding a custom report in Sawmill 7 required at least five separate edits of the profile CFG, using a text editor; in Sawmill 8, it can be done in a single step with the New Field Wizard, or the five steps can be done separately with the new web editors, without any text editing.

Improved Scheduler



The Sawmill 8 Scheduler has a number of improvements over Sawmill 7. Most notably, there is a Run Now button to immediately run any scheduled task, and scheduled tasks can list multiple actions that run in sequence. For instance, a scheduled task can update the database, then remove data older than 30 days, and then email a report. In Sawmill 7, that would have required three separate scheduled tasks, and they would have had to be carefully spaced so they didn't collide.

Simplified Date Filter Syntax

Sawmill 8's data filter syntax is vastly expanded over Sawmill 7. In Sawmill 7, date filters could be 15/Jan/2008-15/Feb/2008, or similarly for months and years; in Sawmill 8, a wide range of intelligently interpreted options are also available. Some examples:

- last 10 months
- 3 months ago-2 months ago
- q1/2008
- q3/2008-Feb 2009
- yesterday

and much more. These can be used in the Scheduler, or on the command line; similar options are available in the Reports through the Date Picker.

Database Import/Export

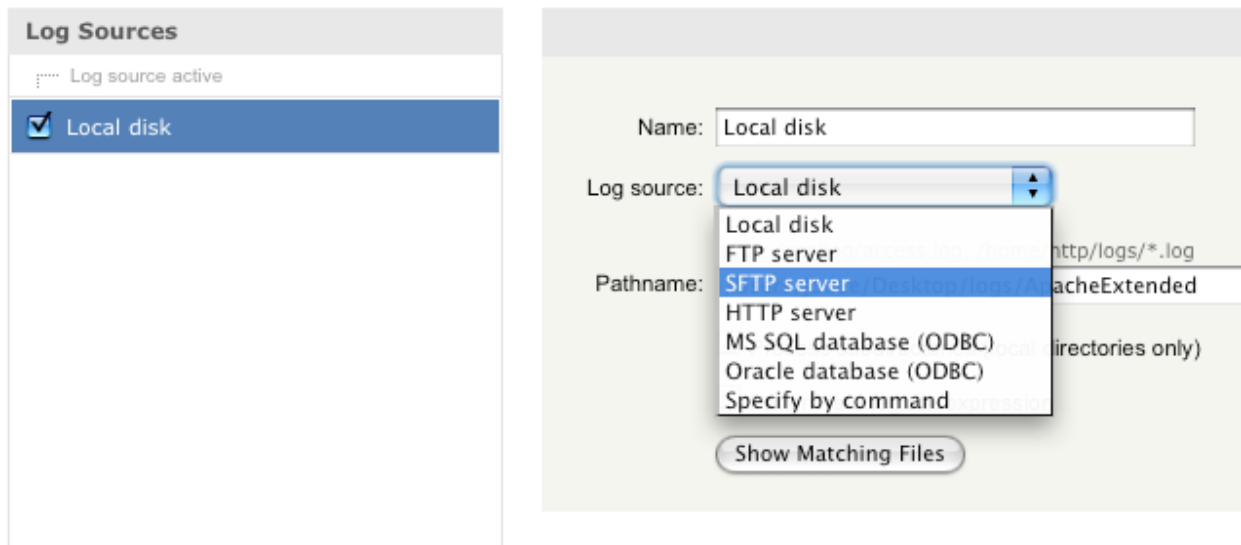
Sawmill 8 can export an entire database to a text format, and import a database from that text format (Sawmill 7 provided no database import or export capabilities). This makes it easy to move a database between platforms, or from one database server to another, without having to rebuild it from the log data.

Improved Report Caching And Performance

Sawmill 8 adds several additional levels of report caching, beyond the HTML report cache used by Sawmill 7. In Sawmill 7, any change to a report would require it to be regenerated; in Sawmill 8, many changes can be regenerated from the cache, including paging through the report, changing the sort, or changing the visible columns. This makes some report operations much faster in Sawmill 8.

Sawmill 8 also takes advantage of multiple processors or cores, when available, to improve the performance of reports. In Sawmill 7, report tasks always used only one processor. This can result in much faster report generation for very large datasets.

Additional Log Source Options



Sawmill 8 adds support for SFTP and SQL log sources, to the list supported by Sawmill 7 (local file, FTP, command, HTTP). SFTP provides a more secure, and more reliable, method for downloading log data from a remote system. SQL log sources can read log data directly from a SQL database. Sawmill 8 can also recursively process a hierarchy of folders on an SFTP or FTP server using a single log source; Sawmill 7 could only download the contents of a single FTP folder, not its subfolders.

Improved Log Import Performance

Sawmill 8 automatically splits log processing across multiple processors, to improve import performance. Sawmill 7 also supports splitting database builds across processors, but Sawmill 8 detects the number of processors automatically to do the split, and Sawmill 8 also uses an efficient channel of communication between processes to reduce disk contention between threads, for better performance and scalability of the initial data import step. Sawmill 8 can also be configured to split log processing across multiple servers in a cluster, for even higher performance (Sawmill 7 supports only multiprocessor on a single server).

Sawmill 8 can be configured to build cross-reference tables, indices, hierarchy tables, and session tables, on-demand, deferring them until they are needed. When it is

configured in this way, the initial database import step is the *only* step of a database build or import, so reports are available as soon as the import is done (or sooner, if it's using real-time; see above). That makes database updates much faster, because they do not need to build all the support tables; those tables will be built later, when a report requests them, if at all. This provides an opportunity for configuring a report-speed vs. import-speed tradeoff that does not exist in Sawmill 7, where all support tables are always built automatically during database build and update.

Sawmill 8 uses improved SQL queries for building cross-reference tables. Cross-reference table generation, which is the largest part of database build time, are much faster in Sawmill 8 than they were in Sawmill 7, especially during database updates. Sawmill 7 rebuilt all cross-reference tables from scratch after each update; Sawmill 8 performs an incremental update of the cross-reference tables, which is much faster.

Other Enhancements

Sawmill 8 includes many other enhancements over Sawmill 7. Some of these are:

- “Report fields” for more flexibility and fine-tuning of report elements and table data.
- Direct URL access to reports; URLs can include profile name, report name, date filter, filter expression, filter comment.
- “One-level slice” of a hierarchical field in a report column; for instance, a months report (across all years).
- New Calendar report.
- Independent sort field and sort direction of the drill down field.
- Option to include row numbers and aggregation rows when exporting a CSV table within the reports interface.
- PDF report generation from the Scheduler.
- Minimum and maximum aggregation rows to tables.
- Default date filter per profile.
- Table row selection (to mark a row in yellow).
- User-created actions (-a options), with fully customizable parameters and behavior.