



Sawmill v.0.2.0

for the Planetary Data System

Table of Contents

1 Sawmill Guide	
1.1 Overview	1
1.2 Release Notes	2
1.3 Installation	3
1.4 Operation	6

1.1 Overview

About Sawmill

The Report Service provides functionality for capturing and reporting metrics. Sawmill is the commercial product chosen to satisfy the PDS reporting requirements.

Please send comments, change requests and bug reports to the [PDS Operator](mailto:pds_operator@jpl.nasa.gov) at pds_operator@jpl.nasa.gov.

1.2 Release Notes

Release Notes

The purpose of this section is to provide a description of a Report Service release with respect to the Sawmill product including any impact that the new or modified capabilities will have on the Discipline Nodes or the PDS user community. If viewing the web-based version of this document, a somewhat itemized list of changes for each release can be found on the [Release Changes](#) page.

Release 0.2.0

This release of Sawmill is a component of the integrated release [1.2.0](#) of the PDS 2010 System. This release is intended as a prototype release in support of the assessment of the PDS4 standards and the system components to date. This release represents a migration of this project to a sub-project of the overall Report Service project.

Release 0.1.0

This release of the Report Service is a component of the integrated releases [1.0.0](#) and [1.1.0](#) of the PDS 2010 System. This release is intended as a prototype release in support of the assessment of the PDS4 standards and the system components to date. The new or modified capabilities for this release are as follows:

- Purchased and installed the Sawmill software package.

The liens for this release are as follows:

- Need to develop a procedure and software for retrieving metric logs from the Nodes.
- Need to configure Sawmill for the various profiles necessary to support PDS reporting.

1.3 Installation

Installation

This section describes how to install the [Sawmill](#) software package. This package serves as the majority of the Report Service for the PDS 2010 system. The following topics can be found in this section:

- [System Requirements](#)
- [Software Installation](#)

System Requirements

The software that makes up this project consists of a commercial package that is available for download and installation. The current version of the Sawmill package is 8.1.7. If installing the package in [Web Server Mode](#), there are no further system requirements. If installing the package in [CGI Mode](#), the following software is required to be installed on the host machine:

- Sun Java Standard Edition (J2SE) 1.6.X
- Apache Tomcat 6.0.20

Software Installation

This section details the download and installation of the Sawmill software package. At this time, the [CGI Mode](#) installation is preferred but if we procure a dedicated machine for this service in the future the [Web Server Mode](#) installation will be the best approach.

Download Software

This section details how to download the software package from the Sawmill web site.

1. Go to <http://www.sawmill.net/downloads.html>.
2. Select the *Proceed* button (the form fields are optional).
3. Select the desired version from the drop-down box, and select the *Proceed* button ("production (x.x.x)" preferred).
4. Select the appropriate operating system for the host machine, and select the *Proceed* button.
5. Select the appropriate hardware for the host machine, and select the *Proceed* button.
6. The download will begin automatically.

7. Verify that the *Operating System* and *Hardware Architecture* noted at the top of the page match that of the host machine.
8. Once the download completes, copy the compressed Sawmill download file (.tar.gz) from the download location to the desired host machine.

CGI Mode Installation

This section details the CGI mode installation. In CGI mode, Sawmill runs as a CGI program under another web server. The instructions that follow specify this installation for an Apache Tomcat application server.

1. Log onto the host machine with root privileges.
2. Configure Apache Tomcat to support CGI binary executables.
 - a. Add the following XML blocks to Servlet Definition and Servlet Mappings sections of `$_CATALINA_HOME/conf/web.xml`:

Note: These XML blocks are already included in the default Tomcat web.xml and only need to be uncommented. If so, be sure to add the *executable* parameter in the Servlet Definition):

```

<!-- ===== Built In Servlet Definitions ===== -->
<servlet>
  <servlet-name>cgi</servlet-name>
  <servlet-class>org.apache.catalina.servlets.CGIServlet</servlet-class>
  <init-param>
    <param-name>debug</param-name>
    <param-value>0</param-value>
  </init-param>
  <init-param>
    <param-name>cgiPathPrefix</param-name>
    <param-value>WEB-INF/cgi</param-value>
  </init-param>
  <init-param>
    <param-name>executable</param-name>
    <param-value></param-value>
  </init-param>
  <load-on-startup>5</load-on-startup>
</servlet>

<!-- ===== Built In Servlet Mappings ===== -->
<servlet-mapping>
  <servlet-name>cgi</servlet-name>
  <url-pattern>/cgi-bin/*</url-pattern>
</servlet-mapping>

```

- b. Set attribute `privileged="true"` for the Context element in `$_CATALINA_HOME/conf/context.xml`.

```

<Context privileged="true">
  .
  .

```


Web Server Mode Installation

This section details the web server mode installation. In web server mode, Sawmill runs its own web server, and serves statistics using it.

1. Log onto the host machine.
2. Navigate to the directory where the Sawmill compressed package was copied, and enter the following command from the command line:

```
% gunzip -c (sawmill.tar.gz) | tar xf -
```

A directory called *sawmill* is created containing the Sawmill executable file (*sawmill*) and a few other files.

3. To run Sawmill, execute the following commands:

```
cd sawmill  
./sawmill &
```

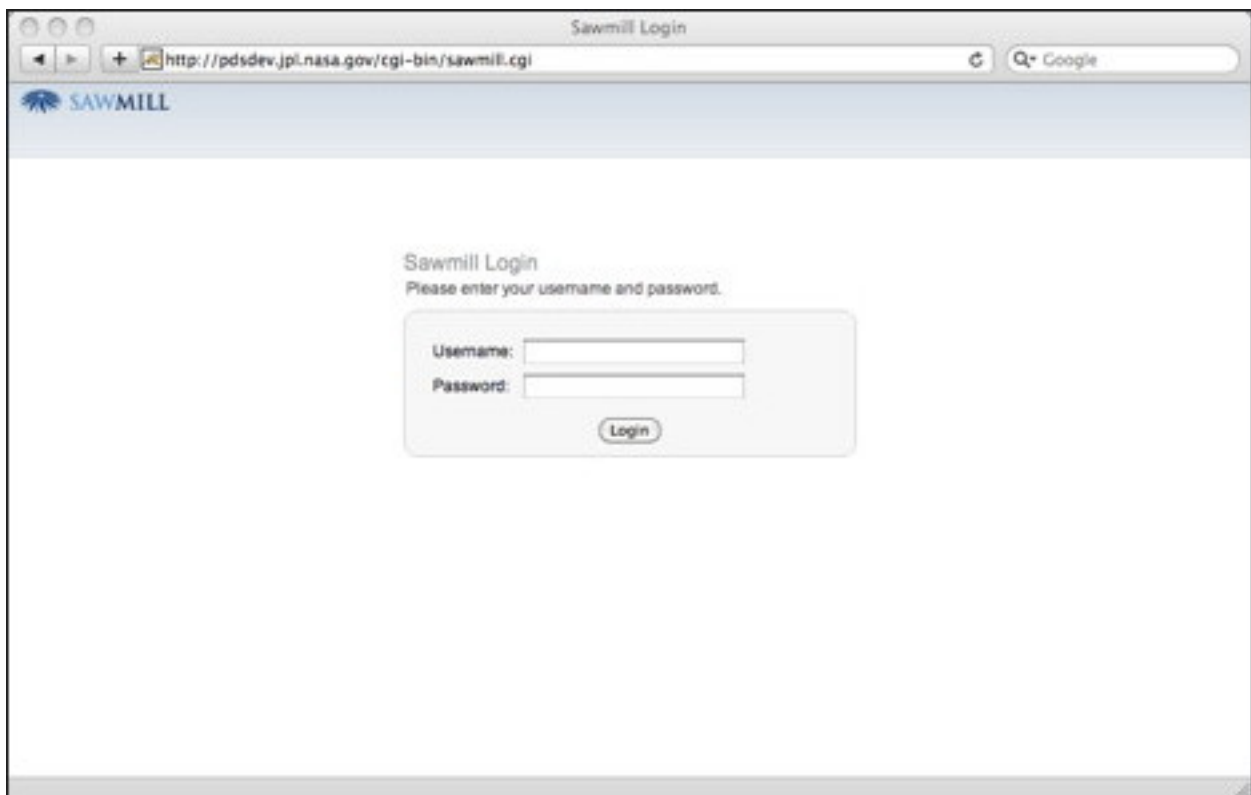
Leave off `&` if you prefer to not run Sawmill in the background.

4. Open a web browser, and navigate to the URL displayed when Sawmill was executed (will not be running on port 80).
5. Enter *License Key* when prompted.
6. Proceed through the steps to create an admin username and password.

1.4 Operation


Operation

The Sawmill software offers a web-based interface for configuring and generating reports with the software. The login screen is as follows:

A screenshot of a web browser window titled "Sawmill Login". The address bar shows the URL "http://pdsdev.jpl.nasa.gov/cgi-bin/sawmill.cgi". The page features the "SAWMILL" logo in the top left corner. The main content area contains the text "Sawmill Login" and "Please enter your username and password." Below this is a login form with two input fields: "Username:" and "Password:". A "Login" button is positioned below the password field.

If viewing this document in online form, click the image for a larger version.

Once the user is successfully authenticated, they are presented with the main screen as follows:



If viewing this document in online form, click the image for a larger version.
 More details to come regarding how to configure and generate reports.

Generate Logs

This section details the log format we will use for the Report Service, and how to setup your Apache or Tomcat web servers to produce the logs.

Log Format

The log format we will be using is the Apache NCSA Combined Format, with the format string as follows:

```
%h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-agent}i\"
```

Where each value is as follows (a "hyphen" in the output indicates that the requested piece of information is not available):

%l	The RFC 1413 identity of the client.
%u	The userid of the person requesting the document as determined by HTTP authentication.
%t	The time that the request was received [day/month/year:hour:minute:second zone]
%r	The request line from the client is given in double quotes, including, in order, the method used by the client, the requested resource, and the protocol used.
%>s	The status code that the server sends back to the client.
%b	The size of the object returned to the client, not including the response headers.
%{Referer}i	The "Referer" HTTP request header that gives the site that the client reports having been referred from.
%{User-agent}i	The User-Agent HTTP request header that is the identifying information that the client browser reports about itself.

We are using this log format because it is a commonly used format that provides more useful metrics that allow Sawmill to ignore extraneous log entries (e.g., bots, web crawlers, worms).

Setup Tomcat

Tomcat produces access logs that record all requests processed by the web container. The following are the steps needed to configure the server to create the Combined Format logs:

1. Log into the machine hosting the web server you would like to gather metrics from.
2. Open `$(CATALINA_HOME)/conf/server.xml` for editing.
3. Find the "Host" XML node. It should resemble the following:

```
<Host name="localhost"  appBase="webapps"
      unpackWARs="true" autoDeploy="true"
      xmlValidation="false" xmlNamespaceAware="false">
  .
  .
  .
</Host>
```

4. Within this block, add the following XML to enable the access logs:

```
<Valve className="org.apache.catalina.valves.AccessLogValve" directory="logs"
       prefix="localhost_access_log." suffix=".txt" pattern="combined"
       resolveHosts="false"/>
```

- Be sure to check if this Valve is enabled elsewhere in the Host block. If so, it will now produce a second set of access logs.
- The value for *directory* is the absolute or relative pathname of a directory in which log files created by this valve will be placed.
- Verify the prefix values are different so they do not overwrite each other.
- See the [Tomcat Configuration](#) to set other configuration attributes including path, file name suffix, etc.

5. Restart Tomcat.

Setup Apache

Apache produces access logs that record all requests processed by the server. The following are the steps needed to configure the server to create the Combined Format logs: