

MANAGING SOFTWARE LICENCES FOR A LARGE RESEARCH LABORATORY

N. Høimyr, M. Katzararova, P.Saiz, A.M. Osborne, A. Thys, P. Bähler, H. Hansbakk

CERN, Geneva, Switzerland.

Abstract

The Product Support (PS) group of the IT department at CERN distributes and supports more than one hundred different software packages, ranging from tools for computer aided design, field calculations, mathematical and structural analysis to software development. Most of these tools, which are used on a variety of Linux/Unix and Windows platforms by different user populations, are commercial packages requiring a licence. The group is also charged with licence negotiations with the software vendors.

Keeping track of large number and variety of licences is no easy task, so in order to provide a more automated and more efficient service, the PS group has developed a database system to both track detailed licence configurations and to monitor their use. The system is called PSLicmon (PS Licence Monitor) and is based on an earlier development from the former CE group[1].

PSLicmon consists of four main components: report generation, data loader, Oracle product database and a PHP-based Web-interface. The licence log parser/loader is implemented in Perl and loads reports from the different licence managers into the Oracle database. The database contains information about products, licences and suppliers and is linked to CERN's Human Resources database. The web-interface allows for on-the-fly generation of statistics plots as well as data entry and updates. The system also includes an alarm system to notify about licence expiry.

Thanks to PSLicmon, the support team is able to better match licence acquisitions with the diverse needs of its user community, to keep licence cost within limited budgets and to be in control of migration and phaseout scenarios between different products and/or product versions. The tool has proved to be a useful aid when making decisions regarding product support policy and licence acquisitions, in particular ensuring the provision of the correct number of often expensive software licences to match CERN's needs.

INTRODUCTION

Commercial software packages today are often complex products with many different features, certain of these modules for very specific applications. The management of a commercial software package includes the installation and support of different versions of a

software distribution to the user community, the provision of user training, as well as support- and licence negotiation with vendors. For each licensed product there is a need to find a balance. We need enough licences for all our users, but in order to keep spending under control, we should not pay for too many or superfluous licences for features that are not used. Another aspect of software distribution is version control, we need to ensure that users migrate to a new version of a product before removing older versions.

In the past, many of the different products were purchased and supported by local units at CERN close to the user-community or installed directly by the users themselves. Over the last years, management of commercially-licensed computer applications has been centralised on a set of central licence servers. At the same time, the distribution and support of the software products from a wide range of application domains, from computer aided engineering to software development, has been centralized in the Product Support group of IT Department. This re-organisation rationalises the support of the licence servers and software installation. But the support team is no longer close to the user community and as the IT Product Support team does not have intimate knowledge of the applications themselves, metering of use has become essential to validate the need for licence renewal and extensions.

In order to keep track of licence use, the PS group has developed a system to monitor the usage of these products. This system is linked to a database with information about the products and suppliers. The system is called PSLicmon (PS licence Monitoring) and is based on an earlier development from the former IT/CE group. The former system was based on the SAMreport¹ wrapper and report-generator from Macrovision to gather statistics linked to an Oracle database with a forms interface to enter information about products. The system had proved very useful, however certain applications could not be wrapped and were not covered by the system. Maintenance of information in the database was also rather complex and required a dedicated person. In order to avoid constraints on the distribution of applications, the new system was designed to use log data directly from different licence managers, thus avoiding application wrapping. (Application wrapping is still used to monitor products developed in-house or to gather statistics of the use of older versions of

¹SAMreport is a report generation and log-analysis tool from Macrovision for the FlexIm[3] licence toolsuite.

certain products.) To simplify updates of information in the database, a user-friendly web-interface allows for data entry and updates as well as viewing statistics. Furthermore an administrative interface has been set-up for licence entry and maintenance of the system itself.

This paper will first present some commonly-used terms and concepts of software licensing, then we describe PSLicmon and its main features.

Definitions

A *product* is a software package. Many of the products supported by our group are commercial ones, such as Ansys, Mathematica, Euclid and Mafia. A product may have several components or modules, each of them called a *feature*. Different features in many cases correspond to different binaries. For instance, Mathematica has two different features: a front-end and the kernel. In other cases a feature corresponds to a certain functionality of a program, e.g. EUCLID surface modeller or Ansys CatiaV5 interface.

A *licence* is a key to control the use of a software product and its features. A licence is usually stored in a file. There are different types of licences. First of all, there are site licences, that allow anyone in the organisation to use the product. In a site licence, there are no restrictions on the maximum number of users. Another type of licence, and the most common nowadays, is the floating licence. In this case, there is a maximum number of concurrent users for each feature of the product. If that number is reached, and a new user tries to use the product, access to the feature will be denied, and the user will have to wait until one of the current users stops using the product and frees a licence. Another type of licence is the node-locked licence, which is associated with a specific computer and the product can be executed only there. In order to allow for sharing of software products, CERN tries to avoid node-locked licences where possible.

A *licence server* is a computer that host the licence files and keeps track of the number of concurrent users. For each licence there is a process in the licence server, called *vendor daemon*, that handles the licence for a product. When a user tries to use a feature of a product, the application will first try to contact the vendor daemon of that licence and check if it is allowed to run. The licence server normally holds the *licence log files* of product and feature use. This scheme is illustrated in Fig. 1.




Figure 1: Floating licence scheme and licmon database.

If anything fails during the communication with the licence server, the users will not be able to use the product. This makes licence servers a critical element of computing infrastructure: if for any reason a licence

server has problems, all the people using products under that licence server will be unable to work. To prevent this situation, the PS group has set up a redundant licence server configuration: a set of three machines which deal the same licences. Those three machines are in constant communication with each other, and, as long as at least two of them are operating correctly, they can serve licences. Thanks to the redundant configuration, we can do maintenance on the licence servers without disturbing the end users. On our licence servers we use a locally-developed alarm system Uimon[2], to ensure that all the vendor daemons are up and running. The alarm system also monitors the operating system, disks and another aspects of the server infrastructure. To make the installation of a licence server as automatic as possible, a procedure that installs from scratch all the necessary components to run our licence servers has been created.

There are many different schemes for the implementation of floating licences. The PS group currently host licence servers for Flexlm[3], LUM[4], Mathlm[5] and Elan. The most widely-used licence manager is Flexlm, which is becoming a de-facto standard. These licence managers have some common features: they allow for sharing of a given number of licences among a large user population while ensuring the product suppliers that their licence is not being misused. They provide basic tools for analysing usage, including when limits are reached. This information is very important to provide a good service and is the basis for our product and licence management system.

PSLICMON

The core of PSLicMon is a product database, where we keep information about all the products supported by the group. The database contains essential information about each product, such as contacts at the supplier company, CERN power users, support web pages etc.

A public "Product Index" [6] Web page populated with information from the database provides a list of all supported products. The Product Index allows users to search for a software product within a given application domain and includes links to web-pages with information about how to install the product, support contacts, the Web pages of the supplier and user-portals on the Internet. The index page is shown in Fig.2.

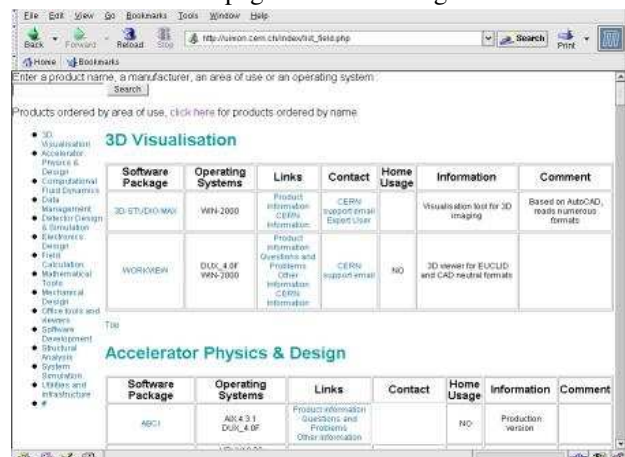


Figure 2: Product index.

The Product Index is dynamically generated from the product database and provides the user community with an up-to-date repository of applications.

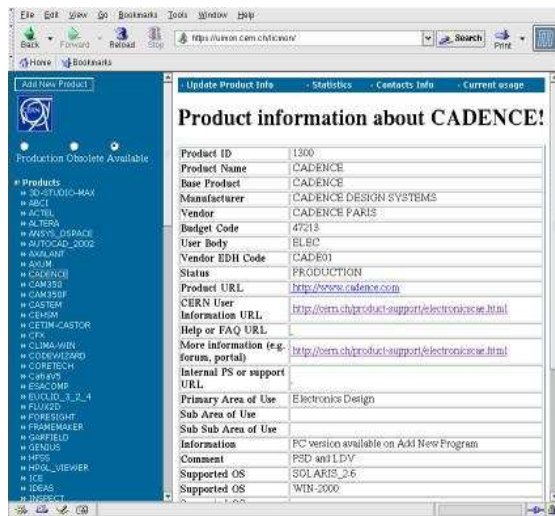


Figure 3: Product information page.

The rest of the information in the product database is related to installation, support and purchasing etc. and is reserved for the support team. (Fig.3.) The database contains information about product licences and the expiry date for all product features. Automatic reminders are sent to the person responsible for the product in time before a feature expires, in order to allow the support team to take action and renew licences in due time.

database. This part of the database is populated with information extracted from the licence server logs. Each Monday, an automatic procedure creates reports on the usage during the previous week for all the products that are monitored. A standard ASCII format has been defined for these reports, so that we can monitor products under different licence software technologies. At the moment, we can monitor four different types: Flexlm, LUM, Mathlm and SMS[7]. However, adding products based on other licence managers is straightforward, as long as some log facility to report the usage of the licensed products is offered. The PSLicmon database is linked to the CERN Human Resources (HR) database so that the account of the user from the licence server logs is linked to the user's affiliation in the Organisation. Thanks to this feature, we can detect non-authorized use of our commercially licensed applications.

Parsing of licence logs has been implemented with Perl. Interaction with the Oracle database is done via PerlDBI. The design of the database has been made as generic as possible, so that other institutes interested in using the system could do so. For more details about the structure of the database, please consult the PSLicmon Reference Manual[8].

A web interface[9] to maintain and browse the database has been implemented. From this web page, the person responsible for the product can display and update the information stored in the database. New products can be easily inserted in the system. The web interface allows

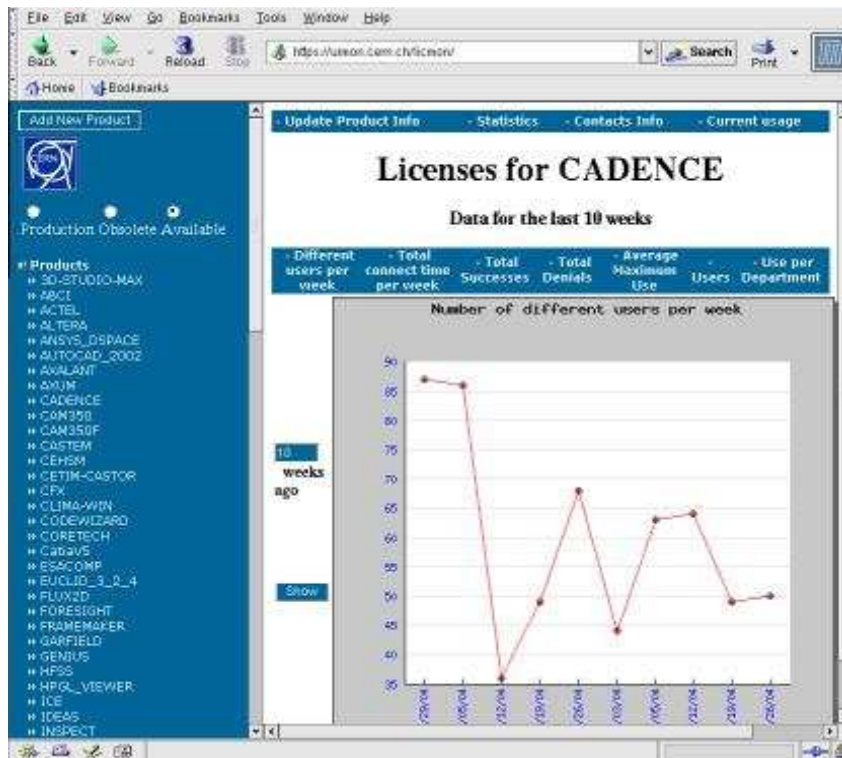


Figure 4: Example of statistics usage plot with PSLicmon, here for the Cadence EDA toolsuite.

Application and licence monitoring is the core of PSLicmon and statistics of application use is stored in the

members of the support team to browse through the usage statistics of monitored products. Different statistics views

are available, such as the number of successful accesses and denials to the product, number of concurrent users through the day, or the email addresses of the users of a particular product or feature (Fig.4).

When making a query, the user can select the number of weeks for which to query the database. Furthermore a facility to obtain specific report logs directly from the licence manager by email has been implemented for the cases where product- or feature-specific information such as client operating system is needed.

Another Web-page "ProductStat"[10] provides on-line information on the current usage of products and features. Thanks to the link to the HR database, contact details of the current users are displayed on the Web. This allows the support team to contact the users in case of blocked licences or when there is a need to warn the users about system interventions. The ProductStat web-page is made available to the user-community at CERN and allows the users to contact colleagues to free a licence when necessary, thus applying a certain peer pressure to avoid abuse of expensive licence features. (In some cases a licence may be blocked inadvertently after a system crash, and the user is not aware of the problem.)

Finally, configuration and administration of PSlicmon itself is carried out via a dedicated Web-interface for system administrators. The PSlicmon web interfaces have been written in PHP and run on an Apache Web server. Connection to the database is via a PHP-Oracle binding. Access control is ensured via Kerberos Authentication via SSL connection to the Web server.

Conclusion

The licence server infrastructure and monitoring tools constitute a platform for the support of more than 100 products in a distributed computing environment. PSlicmon is a very useful aid for members of the support team to manage applications on a day-to day basis as the database contains both statistics and reference information about suppliers, contacts and so on. Within reach of a couple of mouse-clicks, we are able to identify users of a given product and feature over the last weeks or months and contact them if necessary. This is useful when the user community should be notified about a service interruption, upcoming software upgrade or similar. The statistics and trends provided by PSLicmon have also proven to be a useful aid for presentations in coordination meetings together with the user community when making decisions regarding product support policy and evolution. In addition, thanks to PSlicmon, IT/PS has been able to detect licence abuse, obsolete products, forecast the need for extensions of licences etc. In this

way we are able to keep CERN's spending on software licences under control while matching the real needs of the user community.

REFERENCES

[1] Monitoring of Engineering Applications:

<http://documents.cern.ch/cgi-bin/setlink?base=preprint&categ=cern&id=it-2001-002>

[2] UIMON:

<http://uimon.cern.ch/uimon/>

[3] FlexIm:

<http://www.macrovision.com/products/flexim/>

[4]LUM:

<http://www-306.ibm.com/software/awdtools/lum/>

[5] MathIm:

<http://www.wolfram.com/products/network/mathim.html>

[6] Product Index:

<http://uimon.cern.ch/index/>

[7] SMS:

<http://www.microsoft.com/smsserver/>

[8] PSLicmon Reference Manual:

<https://uimon.cern.ch/wiki/bin/view/PSgroup/LicenseMonitoring>

[9] PSLicmon:

<https://uimon.cern.ch/licmon/>

[10] ProductStat Online licence usage:

<http://uimon.cern.ch/ProductStat/cestat.pl>