



Integration of User Services using the World Wide Web

R G Evans, M K Boparai, J C Gordon, M M Curtis

Rutherford Appleton Laboratory, Oxfordshire, UK

ABSTRACT:

The Rutherford Appleton Laboratory houses one of the UK's National Centres for academic research supercomputing, providing facilities for a mixed community of users across the UK. In order to provide a uniform interface to our users and to minimise our own support cost we have integrated as many as possible of our support services to a WWW interface. As well as new documentation we use Perl scripts to present old unstructured documents in a consistent way. Resource control and access to our Helpdesk system is through WWW Forms and dynamic monitoring of machine performance on our J932 and other computing services is also presented via a browser interface. Some concerns over security lead us to restrict some Web pages to the local domain and we have implemented the text only browser Lynx on the J932 to assist in this.

1. BACKGROUND

1.1 RAL

The Rutherford Appleton Laboratory (RAL) is part of the government funded Central Laboratory of the Research Councils (CLRC) which is one of Europe's largest multidisciplinary research support organisations. The main areas of CLRC's research and support provision are: astronomy and planetary science, computing, networking and supercomputing, particle physics, earth observation, radio communications, microelectronics, high powered laser facility, a synchrotron radiation source and the world's most powerful pulsed neutron and muon source, known as 'Isis'.

1.2 Atlas Centre

The Atlas Centre at RAL is host to a number of the computing facilities provided by CLRC including a Cray J932-8192, a 6-processor Digital Alpha 8400, a Digital Alpha 7000 VMS service (4 processors), a HP Farm (30 nodes) and a Digital Alpha Farm (5 nodes), as well as extensive data storage facilities and a Virtual Reality Centre. The major areas of research carried out on the Atlas Centre's Cray include atmospheric and oceanographic modelling, quantum chemistry and computational fluid dynamics.

2. SUPPORT & THE WORLD WIDE WEB

The Atlas Centre provides scientific support as well as general helpdesk support to users from across the UK, of which there are approximately 1000 registered on the Cray J932 service. The first RAL Web pages went live in June 1994, when information on the various departments and research projects was made available. The Laboratory's WWW is now used to provide more diverse information for staff as well as its external users; this includes WWW interfaces to the staff diary system, access to library services, financial data system, online stores catalogue and local site services. For users of our central computing services the WWW was already one of the main sources of documentation two years ago; now it is the source of a variety of information to our users - in addition to various forms of documentation,

this includes:

- resource control
- news
- access to helpdesk
- performance monitoring
- automated operations system

These are briefly discussed below. We serve a very mixed community of users, so our main aim is to provide technical information as simply as we can. Thus we have been deliberately slow in introducing the latest innovations like Java, in order to keep it available to the *majority* of our users.

3. USE OF THE WORLD WIDE WEB & ISSUES ARISING

3.1 Documentation

Our documentation falls into 3 main categories:

- **Legacy:** this is information that was traditionally available in a local documentation directory mounted onto all our main service machines. We have used Perl scripts to present these old unstructured but useful documents in a consistent way.
- **New** and future documents will be available only on our Web pages.
- **CRI** documentation: We have made CrayDoc files available via a WWW interface using the Cray DynaWeb server (based on the commercial Web server, DynaWeb, from EBT that transforms SGML source files to HTML and serves them to a Web browser).

NB: Due to Copyright and licensing restrictions, CRI documents cannot be posted to an external server or made publicly available. However, we can allow users to browse CRI documents on an internal server, until the licensing issues are resolved. We have suggested to Cray that we restrict access to the proprietary documents to the UK academic domain. We would be interested to hear how similar sites have overcome this.

3.2 News

Due to difficulties of exporting local newsgroups to a wide range of other institutions, we provide a WWW interface so that remote users can use a Web browser to read selected newsgroups from our Web pages. The Web server is not directly hooked into the news server, so the news is not dynamic - it is converted to HTML every 30 minutes. The local Cray newsgroup is available via a link from the Atlas Cray Service Home Page. We also did not want to broadcast our local Cray news to the whole world, partly because it is of little use outside RAL, so we restricted this to the local domain only. However, this restricts use to only those users who have access to a local machine running a Web browser (we have the text only browser Lynx installed on the J932). Many of our Cray users are also registered on another service at the Atlas Centre where they can run a Web browser -on which it is useful to read news of, for instance, Cray unavailability if the Cray is currently down! These restrictions ensure some security and obviates the general concern - do we really want to broadcast any sensitive Cray-related information, such as bugs, world-wide?

3.3 Resource Control

The Resource Management Section at the Atlas Centre is concerned with allocation and monitoring of the resources of the central computers listed in Section 1 above. This involves administering Research Grant information, registration of users and accounting their use of various resources. Access is provided via a WWW Form-based interface to the various services, including

- querying account details such as CPU usage on main computing services
- userid information and user contact details
- search facility on userid/surname
- For staff use only, there are (password protected) forms to query the grant information database
- The WWW Form interface links directly to an INGRES database.

Due to security concerns we have restricted some of these pages to the local domain; this allows access to external users registered on any of our central computers. For users who are only registered on the Cray, Lynx is installed and can handle Forms (and is 'lightweight' enough to justify running on the Cray).

3.4 Helpdesk

A WWW Form is available to our users to contact support staff by sending queries to our Helpdesk - a proprietary Windows-based system from Remedy Corporation.

Forthcoming developments for the Helpdesk will allow users access to a WWW interface to our Helpdesk to view FAQs - this serves as an extra source of information for miscellaneous, recurring information asked for and required by users, which is not necessarily covered by the formal documentation. This should reduce time spent by user support staff who otherwise duplicate effort in responding to some recurring queries. This information can gradually filter into the main documentation at periodic reviews.

3.5 Performance Monitoring

It is often desirable and useful for support staff and users alike to be able to monitor machine performance in terms of CPU usage, memory usage, I/O rates and even NQS queues (perhaps with a view to ascertain why a machine is responding poorly!). We currently provide a WWW interface to the following statistics:

Cray J932 Statistics

CPU Usage: Daily & Last 7 days

Interactive Response

Number of Login Sessions

IO Rates on the J932

Memory and Swap statistics

NQS Queue Summary

J932 Weekly Queues

DEC 8400 Statistics

CPU Usage

Interactive Response

LSF Queues

Atlas Datastore Statistics

Basic Statistics

J90 Details

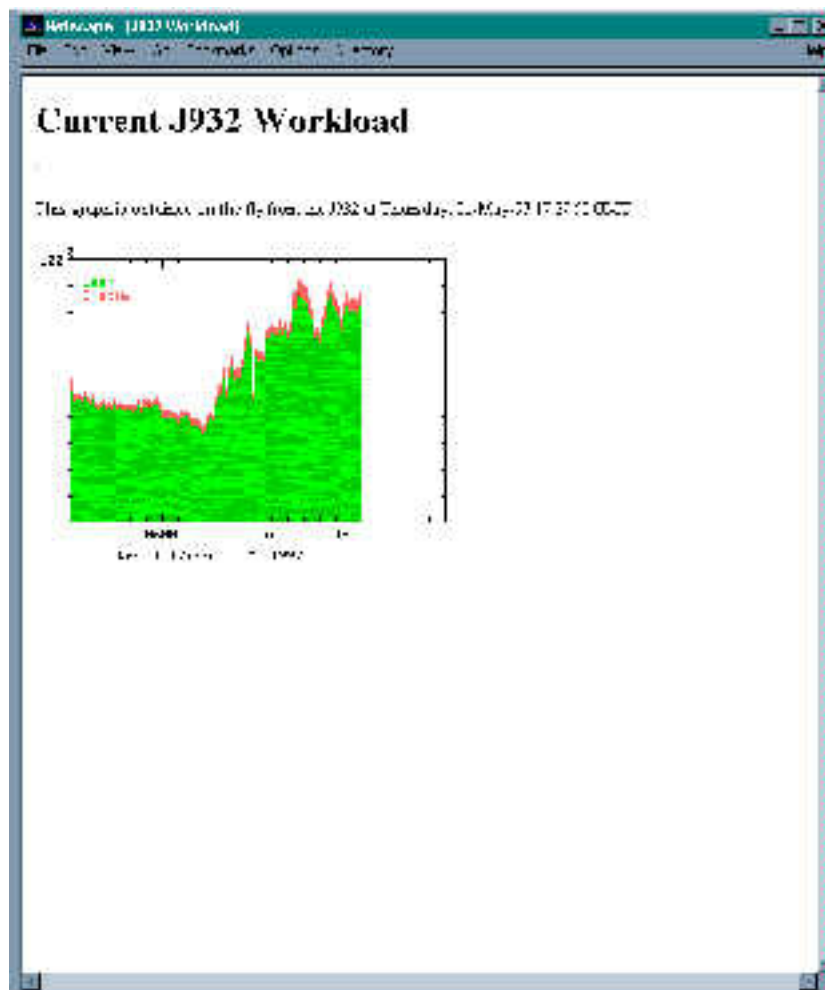
Import Export Query

The technique currently used to produce these dynamic performance monitoring

figures and graphs is somewhat 'home-made': we use a combination of Perl scripts and crons to produce data, that is used by the 'gd' GIF-manipulating library to generate graphs and charts 'on the fly', which are delivered to the Web browser. We would therefore be interested to hear of any alternative methods employed by other sites.

As this involves programs that execute in real-time to produce dynamic output, this clearly raises some concerns over security. These pages are consequently still on 'test' until the security issues can be resolved.

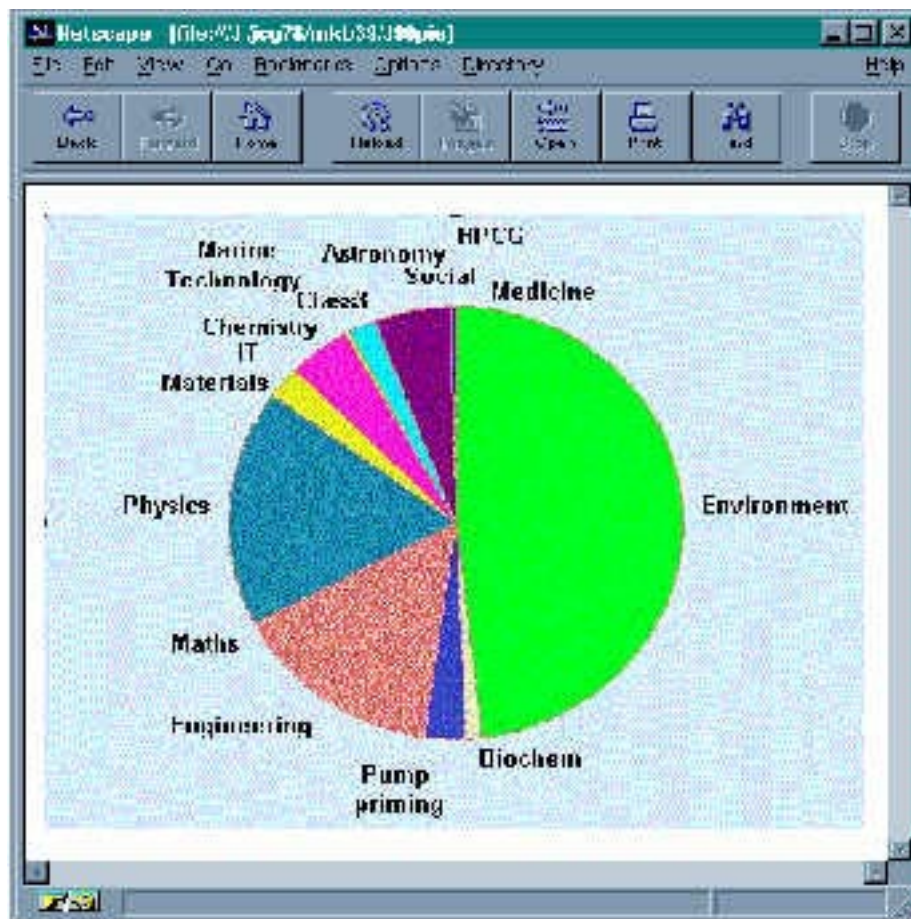
3.6 The Atlas Centre Automated Operations System



The Atlas Centre Automated Operations System monitors the status of the computers providing our production services. Its aim is to alert operations staff and platform managers to various categories of problems. The central monitoring system is built around two closely coupled systems: a workstation (RS6000) running the SURE package developed at CERN, and a PC running the AUTOMATION POINT XC package from Computer Associates (Legent). The SURE system 'pings' the remote computers, runs specific commands to collect information from them and displays status. It passes error messages and codes to the Automation package in the PC, which contains a rulebase to determine what action to take. This system is particularly useful in identifying problems with inter-dependent services, e.g. the Cray Data Migration facility is set up on our system such that small files are written to disks mounted on another machine, and if that machine is down it has a direct effect on the Cray. It is planned to make this available to all users and a WWW interface is available. We believe it is extremely useful to have Web access to this especially at times when there are operational problems with any of our computer services: the last thing a user needs at times like this is to be trying to call our telephone service line along with hundreds of other users!

3.7 Atlas Services & Machine Usage Statistics

A series of reports on the last six months usage of our services are available on the WWW, which are updated regularly. The page below shows a pie chart illustrating how our Cray service is utilised by the various scientific communities. Also available are graphs showing monthly Cray CPU utilisation.



4. SUMMARY AND CONCLUSION

We have integrated many of our user services to a WWW interface: three categories of Cray documentation (legacy, new and CRI proprietary), local Cray news, resource control, helpdesk access, dynamic performance monitoring and our Automated Operations System. By doing this, we have provided a uniform interface to our users as well as reducing our future support costs. Because we serve a very mixed community of users, we have deliberately kept our Web pages as simple as possible (e.g. no Java) in order to target the majority of our users, at least for the time being.

5. CONTACT INFORMATION

For further information, please see 'Computing Facilities' under

<http://www.dci.clrc.ac.uk/>

or contact us at

r.g.evans@rl.ac.uk, m.k.boparai@rl.ac.uk, j.c.gordon@rl.ac.uk, m.m.curtis@rl.ac.uk

[Table of Contents](#) | [Author Index](#) | [CUG Home Page](#) | [Home](#)