# NEWSLETTER

## New name for EUUG!

## Now: EU


European UNIX® systems Us

Volume 10, No. 4
Winter 1990

## CONTENTS

# The European Forum for Open Systems

# NEWSLETTER


*Volume 10, Number 4*
*Winter 1990*



ISSN 1011-4211

**Editorial Team:**

Frances Brazier
Publications Executive

Department of Cognitive Psychology
Vrije Universiteit
de Boelelaan 1115,
1081 HV Amsterdam.
The Netherlands
frances@psy.vu.nl

Alain Williams
Editor

Parliament Hill Computers Ltd.
7 Prospect Street
Caversham
Berkshire RG4 8JB
U.K.
addw@phcomp.co.uk
Tel: +44 734 461232
Fax: +44 734 474194

**Typesetting:**

Gina Baikenycz

The Instruction Set Ltd.
City House
190 City Road
London EC1V 2QH
U.K.
gina@inset.co.uk

**Printed by:**

Rank Xerox (Copy Bureaux) Ltd.
68-74 Rochester Place
London NW1 9JX
U.K.

**Published by:**

EurOpen
Owles Hall
Buntingford
Hertfordshire SG9 9PL
U.K.
europen@EU.net
Tel: +44 763 73039
Fax: +44 763 73255

# Editorial

*Alain D. D. Williams*
*addw@phcomp.co.uk*

*Parliament Hill Computers Ltd*

## The EUUG becomes EurOpen

To a hushed audience in a Nice auditorium Michel Gien announced the change of name.

*Why?* people wanted to know.

UNIX has matured and changed, and so has our user group. UNIX has graduated from University and is earning a healthy living in Business.

UNIX no longer belongs just to the hackers, but to computer users everywhere.

UNIX is not simply UNIX any more, but is the central plank of the Open Systems movement which has *opened* the industry up to real competition.

### New Needs

It is right that we change to accommodate the needs of our new fellow members.

Gone are the days when everyone had source and every site a resident guru. It is a sign of success that todays users are not just programmers, and have disparate needs.

Specialist Working Groups are being set up to cater for these needs, read about this on page 12.

### Growing Success

The EUUG has grown well from the initial meetings between the Dutch and the British, and is now going to provide a central focus for the vigorous Open Systems community in Europe.

European issues in Open Systems are a major concern. The problems of character sets is a good example and is discussed by Keld-Jorn Simonsen.

### New Member Groups

In Nice the EurOpen Governing Board accepted membership applications from two new groups:

- Spain
  We welcome the Spanish group who held a successful national event the week before.
- The Soviet Union
  Not to be outdone the Soviet group had it's own first conference the week after the EurOpen Nice conference.

This brings membership of EurOpen to 20 countries.

### Cooperation with UniForum

There has long been cooperation with several User Groups – notably USENIX, and now a new alliance has been announced with UniForum. This is symbolic of the growing spirit of unity to be found all areas of the Open Systems community.

Three initial agreements were announced:

- A joint European show and conference will be organised in Autumn 1992 in Amsterdam.
- Cooperation in publications related to the UNIX and Open Systems communities.
- Cooperation in the preparation of contributions to the international standards bodies.

### Onwards

These changes and alliances will provide us with the strength and purpose to become a true focus for the user community and a force in the Open Systems arena.

Delegates were provided with an EurOpen towel to mop up their tears at the passing of the name EUUG into the annals of history.

### Personal Bits

Harriet was born to Nigel and Sally Martin on 19th October. It will not be long before he has to think about buying a larger boat.

Henry and Emma were born to Jill and myself on 3rd August. Both are well – although we could do with a little more sleep.

# Multi Character Set Mail

*Keld-Jorn Simonsen*
*keld@dkuug.dk*

*Center for Anvendt Datalogi*
*University of Copenhagen*

Keld Simonsen is the chairman of the Danish UNIX User Group. He is currently occupied with standards and networking activities.

The current international R&D mailing systems seldom provide support for more than the English language. The paper describes a system in the TCP/IP and UUCP world capable of supporting many character sets.

## Support for Character Sets in Current Mail Systems

Today many character sets are used in the Research & Development World, coming from the usage of many different computer systems and brands, and also from the different needs for character set support in different parts of the world. With the advent of world-wide electronic mail, which is also gatewayed between different mail systems, a need for multi-character set electronic mail has arisen.

Mail systems capable of supporting more than one character set have rarely been seen. IBM systems have been capable of supporting full 8-bit EBCDIC, but to what extent the different EBCDIC versions (Danish, Spanish...) have been supported is unknown to me. TCP/IP can only support 7-bit ASCII according to the specifications. UUCP can support full 8-bit, but rarely does that wholeheartedly, as UNIX mail often passes TCP/IP mailers. DECnet can support 8 bit, but then just one 8-bit character set. When exchanging mail between these systems, the usual solution has been using the largest common subset of characters, which is normally identical to support for ASCII or a subset thereof.

The system described here takes the opposite direction, that is (instead of providing the common subset) to provide the union of all character sets to all the actors in the world-wide network.

## Character Sets

The character sets in use today can be split up into several families:

1. The ISO 646 7-bit family

   ASCII is here the most significant member. ISO 646 leaves 10 positions undefined, and two positions (# and $) can have alternate representation. This has been used to define a lot of national variants, which has succeeded to a varying extent in getting used. In the Nordic countries (except Iceland) the national ISO 646 character sets have played a major role, while in for example Germany and France the usage has been of lesser importance. As 8 bit

equipment is taking over the market, the importance of 7-bit support is diminishing, but a lot of software and hardware will still need to be supported for some years to come. Especially ASCII support would be needed for a long time. The European Computer Manufacturers' Association (ECMA) is registering 7-bit (and 8-bit and more) character sets, more then 20 7-bit character sets are registered here.

2. The ISO 8859 family

   This family consists of about 10 8-bit character sets (there are proposals for some new ones on the way). It covers different parts of Europe and other parts of the world using Latin characters, Cyrillic, Greek, Arabic and Hebrew. Common to all the parts is the first 128 positions, which are identical to ASCII, and there is room for control characters such as defined by ISO 6429 in code 0-31 and 127-159. In 8859 all different characters have a code value of just one octet.

3. Other 8-bit character sets

   Different vendors have defined their own character 8-bit sets, notably DEC, HP and Apple, and IBM as the most inventive. IBM has defined a number of PC code pages, but without room for the upper control characters of ISO 6429. Also IBM has defined a lot of EBCDIC 8-bit character sets for use on their mainframe computers, the EBCDIC encodings are without any resemblance to ASCII, while all the rest of these vendor defined character sets has the ASCII codes as a proper subset.

4. The ISO 6937 8/16 bit family.

   This family is also 8-bit in nature, and has also the ASCII encoding as a subset, but a lot of accented characters are represented by a non-spacing accent followed by the base letter. Thus you need to look at either 8 or 16 bits to handle a character. This makes life complicated and support for this in processing on computers has been low. It is also primarily intended for communications, and has found some use there. Members of the family are the ISO 6937-2, the T.61 and the TELETEX character sets. The OSI protocols are using these character sets, although I have not seen them in full use in any OSI system.

5. The 16 bit character sets

   The East-Asian cultures need support for more characters than can be handled in 8 bits, and they have defined their own 16 bit character sets. Several exist for Japanese, Chinese and Korean languages. They often also support some use of Latin, Cyrillic and Greek characters. Most of these allow room for the ISO 6429 control character set, but at least some IBM PC 16 bit character sets do not comply with this.

   A new 16-bit code, UNICODE, defined by some vendors from the USA, is trying to accommodate most characters used in the computerised world. It does not leave space for the ISO 6429 control characters, though, and thus breaks the structure of most character sets, which adhere to the ISO 2022 standard.

6. The ISO 10646 32-bit character set

   The latest standard character set development is the ISO 10646 standard, whose aim is to include all the characters in the world and give them one unique 32 bit value each. As the most used characters can be represented in either 8 bits or 16 bits, the standard specifies ways of compacting the code, into 1, 2 or 3 octets. ISO 10646 has ASCII and ISO 8859-1 as proper encoded subsets.

## Character Naming, Tabulation and Conversion

One fundamental part of my work has been to give a lot of characters in the above character sets a short name, and most of the character sets mentioned above have then been tabulated using these short character names. Currently (Oct 1990) about 1000 characters have been given a short name, and over 60 character sets have been tabulated. No 16 bit character set has been tabulated, though, partly due to lack of knowledge of East Asian characters.

The namings have been used to build conversion tables between the character sets. If the character occurs in both character sets, the conversion is simple; but if the character to be converted does not appear in the new character set, what to do then? The solution followed here is to retain the

information by representing it with the short name of the character and a preceding escape character. The escape character is selectable. The short character names are quite mnemonic, and thus quite readable if encountered in a character set not capable of showing them directly.

All the character sets have been given names, and some also nicknames, for example ISO 8859-1 has the nickname "latin1".

This naming and tabulation of characters and character sets provides a convenient way to handle multiple character sets together. Routines written in C have been made for string conversion. Most character sets in use in the Western hemisphere has been tabulated for this purpose.

## The Multi Character Set Sendmail

The conversion routines have been built into Sendmail (5.61 and 5.64) together with handling of specifications of the character sets. Via a Sendmail mailer specification and the mailertable file it can be specified on a site basis as to which character set the site is employing. New headers (X-Charset and X-Char-Esc) specify alongside the letter how the body of the letter is encoded. The headers themselves are currently untouched.

There are plans to do a SMTP negotiation protocol, where two SMTPs can agree to exchange ISO 10646 instead of ASCII (ISO 646). This work is mainly undertaken by Dan Oscarson, LTH, Sweden. In the Sendmail I am running, there is only provision for known sites to run anything else than ASCII, no negotiation is possible. The default character set used for all unknown sites is ASCII, encoded with an escape character of decimal 29, which happens to be invisible to most standard UNIX pagers. It should only be allowed to send out anything other than ASCII if the recipient site has agreed.

## Experiences with the Multicharacter Set Sendmail

The multicharacter set sendmail has been in operation at dkuug.dk since February 1990. dkuug.dk is the main internet and UUCP backbone machine for Denmark, serving over 100 sites on DKnet - the Danish branch of the European Unix network (EUnet). About 10 of the 100 sites are running something different than the Danish default character set - which is 7-bit Danish ISO 646 with some Danish national letters. There are both ISO 8859-1 and IBM codepage machines using the conversion, and the conversion between the 8-bit sites works fine. Just recently the first international 8-bit connection was established to a site in Iceland.

The 7-bit Danish ISO 646 causes some trouble, because one binary code actually may be used for two things: sometimes it means a Danish letter, but at other times it is used for an ASCII character (for example in C programs). When converting such a letter to 8 bit, the ASCII character meaning is chosen to avoid problems with C programs. A procedure looking at the text and trying to determine if it is natural language text in some national variant of ISO 646 (say Swedish 7-bit) or program code in ASCII could remedy this problem.

It would be nice to have some of the headers converted too. My thoughts are to convert all headers, I think it is secure enough.

## Future Plans

It is planned to make support for user mailers (User Agents - UAs) for different character sets. First support for the "elm" mailer is planned. This will allow a user connecting to a mail host via a PC to run to the PC code page, while another terminal running say ISO 8859-1 on the same mail host can conveniently run his/her character set, and the third user on the same machine using an old 7-bit terminal can get the best out of it too.

Also in this way, it is possible to get the multi character support on machines not supporting anything fancy in the mailer daemon (or Mail Transfer Agent - MTA), as the user may install his/her own UA, or there could be installed a system wide UA with multi character support - without even changing the MTA. The design with encoding in ASCII means that the message can pass a lot of dumb MTAs and the information still be retained for presentation for the recipient.

Similar capabilities are planned for netnews in the *nn* program made by Kim Storm, Texas Instruments, Denmark.

The EUnet has planned to include the multi character set support on its backbones, and also RARE is considering multi character set support, where the work described here may be employed.

# User Interface Management Systems and Application Portability*

*Robert Seacord*
*rcs@sei.cmu.edu*

*Software Engineering Institute, USA*

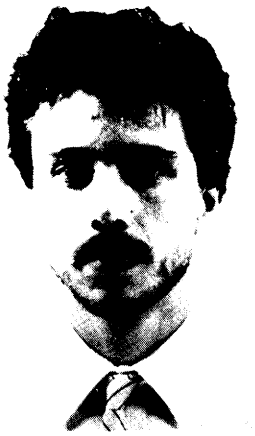

Robert Seacord is a Member of the Technical Staff in the User Interface Project at the Software Engineering Institute (SEI).

Before joining the SEI, Seacord was employed by the IBM Corporation where he worked in the areas of software engineering, processor development and communications.

Seacord earned a bachelor of science degree in computer science from Rensselaer Polytechnic Institute. Since that time he has taken courses at CMU where he is pursuing a master's degree in software engineering. Seacord is a member of the IEEE Computer Society.

Higher level programming languages and standard operating systems now provide greater portability of application software than previously possible. Software developed in C for Unix, for example, can be easily ported to a variety of different architectures and machines. Developing language and operating system standards such as ANSI C and IEEE POSIX will further enhance application portability. At a quick glance it appears as though open systems are finally becoming a reality, but are they really?

As porting software to different architectures becomes more and more a matter of simply recompiling the software for that architecture, it is apparent that a serious problem in portability is with the user interface of these systems. Now, more than ever, when a customer buys a computer platform they are also buying a "look and feel" associated with that system. When using an Apple Macintosh, for example, the user expects to be able to perform a variety of actions using a single button mouse. When working with an MS DOS application, the user expects to be able to perform the same actions using the keyboard. When running OpenLook, Motif, or NeXTStep the user expects the application to provide a defined look and feel. Porting an application to simply run on a different platform is insufficient; there is a requirement for the application interface to behave in a similar fashion to other applications developed for that environment.

Software designs are usually not extensible enough to allow the integration of different user interface toolkits, particularly if these toolkits employ significantly different models in their application interfaces. Changing toolkits or integrating new toolkits usually requires major modification to the application, which then requires extensive re-testing of the application.

Current standards activities are just beginning to address these problems. To date, standards bodies have attempted to define an abstract user interface toolkit that can be implemented in different ways. Where this approach provides some degree of device independence it does not allow for the removal of stylistic concerns from the application. For example, where one user interface style guide may call for a pull-down menu, another may call for a command line interface.

One approach for addressing the problem of application portability across multiple look and

feel platforms is the definition and implementation of a method for separating the application from the user interface. This separation makes changing the user interface of the system practical. It also makes it possible to change user interface toolkits without modifying the application software.

## User Interface Management Systems

The term user interface management system (UIMS) was first coined at a 1982 Workshop on Graphical Input Interaction Technique (GIIT) [17]. UIMSs are, among other things, intended to encourage the separation of a software system into an application portion and a user interface portion. The application portion of a system implements the core functionality, while the user interface portion implements the user interface dialogue.

UIMSs provide facilities for defining both presentation and the computer-human dialogue components of a user interface. A UIMS also may provide facilities to support prototyping, encourage a design that allows for easy modification of the user interface, support implementation and maintenance of the user interface, and allow for the evolutionary incorporation of new user interface technologies.

Most UIMSs are based on the Seeheim architecture [7] (see Figure 1). This architecture uses a layered approach similar to the one used in International Standards Organisation (ISO) Open Systems Interconnection (OSI) standard. The architecture is intended to encourage the separation of functionality between the application and the user interface portions of a software system. The three different layers of the architecture provide differing levels of control over user input and system output.

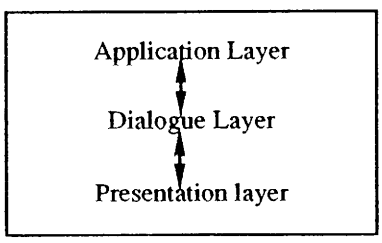


Figure 1: UIMS Architecture

The application layer consists of the core functionality of the application that can be described in a presentation independent manner. For example, in a calculator program this would include the underlying math subroutines library.

The dialogue layer specifies the presentation dependent portion of an application system including the dynamic behaviour of the user interface. The dialogue should allow the display and removal of interaction objects without application involvement and support cascading menus, direct manipulation, and other user interface styles and techniques. The dialogue provides the mapping between the presentation and application layers. The user interface dialogue may be specified using a user interface definition language (UIDL) or by an interactive technique.

The presentation layer controls the end-user interactions and generates low-level feedback. The presentation layer consists of a collection of interaction objects defined for a given user interface technology.

## Existing Approaches

Since the 1982 GIIT Workshop, there have been a number of efforts to build UIMSs that achieve the goal of separation of concerns, while remaining a practical approach to software development. These systems have been classified by the model used in dialogue specification. Some of the more successful approaches have been: event-driven, declarative, object-oriented, data-driven, and interactive layout systems.

In the event-driven approach, inputs are considered as events which are immediately handled by event handlers. Event handlers can cause output events, change the internal state of the system, and call application routines. Examples of event-driven systems include the University of Alberta UIMS [6], ALGAE [5], Sassafras [9], and the TeleUSE UIMS [16].

Another approach is the use of declarative languages in which constraints are defined in order to specify what the system should do rather than specify how it should be done. An example of a system that takes this approach is Cousin [8]. In this category, there is a class of systems which automatically generate the user interface based on a definition of the semantic commands supported by the application. Examples are the UofA* UIMS [12] and MIKE [10].

An object-oriented approach uses objects for defining user interactions, transforming data and

interacting with the application. A good example of a commercially available system that uses an object-oriented approach is Open Dialogue [1] developed by Apollo Computer.

In the data-driven approach, the application communicates with the UIMS in terms of shared data elements. The UIMS behaves like an active database in that it provides a mapping between application data and user interface toolkit objects, and notifies the application of changes to application data resulting from user interactions. This approach was implemented in the Serpent UIMS [2] developed at the Software Engineering Institute at Carnegie Mellon University and the George Washington UIMS [11].

Interactive layout systems allow the user to build user interface by direct manipulation. Examples are Menulay [3], DialogEditor [4], vu [13], and TAE+
[15].

## UIMS Study Group

The IEEE P1201 working group was formed in January of 1989 and chartered to develop standards that would further application and user portability in the X Windows Environment [14]. Since P1201 was formed, Open Software Foundation (OSF), Sun and AT&T have independently developed toolkits for X Windows. Much of the P1201 effort has been spent trying to decide if any of these toolkits can serve as a basis for a standard or if a "virtual" toolkit approach can be used.

In August of 1989, a UIMS study group was begun in P1201 to determine if UIMS technology was sufficiently advanced to solve the problem of application portability across multiple look and feel platforms and to define the scope of a UIMS standard.

The group identified two components where standardisation would be beneficial to the industry. The first of these is an application programmers interface (API) that would:

1. Provide a standard application programmers interface across changes in the underlying toolkit.
2. Support the separation of an application into presentation independent and presentation dependent layers corresponding to the application, dialogue, and presentation layers of the Seeheim architecture.
3. Allow the development of applications that are presentation independent (i.e., the underlying windowing system or user interface toolkit).

The second component is a UIMS interchange format (UIF). The purpose of a standard UIF is:

1. To enable a wide variety of UIMSs to use a single format to store and exchange their data.
2. To allow vendors to develop compilers or interpreters that could "execute" the UIF on their platforms in a manner analogous to postscript printers.

The P1201 UIMS study group has evaluated a number of user interface management systems including Serpent, TeleUSE, and TAE+. The consensus of the group is that the state of the practice is sufficiently advanced to warrant a standards effort. It is believed that a UIMS standard would enhance both application portability and the state of the practice in user interface development.

---

Dr. Randy Pausch of the Computer Science Department here at University of Virginia, USA has actually implemented a User-Interface Toolkit named SUIT (Simple User Interface Toolkit) that is in fact portable across diverse platforms including X-Windows, MS-Windows, the Apple Macintosh, and others. One nice feature is that using SUIT doesn't lock you into any Look and Feel and the Look and Feel can be changed without recompiling.

Folks interested in UIMSs might well want to get more information on SUIT. Sending a request via e-mail to: graphics@cs.virginia.edu should suffice.

## References


[1] Apollo Inc.
Open Dialogue: Designing Portable, Extensible User Interfaces. 1987

[2] Bass, L.J., et al.
Serpent: A User Interface Environment. In Proceedings, Winter 1990 USENIX Technical Conference. Washington, D.C., January, 1990.

[3] Buxton, W., Lamb, M.R., Sherman D., Smith, K.C.
Towards a Comprehensive User Interface Management System. In Computer Graphics: SIGGRAPH'83 Conference Proceedings, pages 35-42. Detroit, MI., July, 1987.

[4] Cardelli, L.
Building User Interfaces By Direct Manipulation. In Proceedings, ACM SIGGRAPH Symposium on User Interface Software, pages 152-166. ACM, New York, NY, 1988.

[5] Flecchia, M.A. and Bergeron, R.D.
Specifying Complex Dialogs in ALGAE. In Proceedings SIGCHI+GI'87: Human Factors in Computing Systems, pages 229-234. Toronto, Ont., Canada, April, 1987.

[6] Green, M.
A Survey of Three Dialogue Models. ACM Transactions on Graphics 5(3):244-275, July, 1986.

[7] Green, M.
Report on Dialogue Specification Tools. User Interface Management Systems :9-20, 1985.

[8] Hayes, P.J., Szekely, P.A., Lerner, R.A.
Design Alternatives for User Interface Management Systems Based on Experience with COUSIN. In Proceedings SIGCHI'85: Human Factors in Computing Systems, pages 169-175. San Francisco, CA, April, 1985.

[9] Hill, R.D.
Event-Response Systems – A Technique for Specifying Multi-Threaded Dialogues. In Proceedings SIGCHI+GI'87: Human Factors in Computing Systems, pages 241-248. Toronto, Ont., Canada, April, 1987.

[10] Olsen, D.R.
The Menu Interaction Kontrol Environment. ACM Transactions on Graphics 5(3):318-344, 1986.

[11] Sibert, J.L.
An Object-Oriented User Interface Management System. In Computer Graphics: SIGGRAPH'86 Conference Proceedings, pages 259-268. Dallas, Texas, August, 1986.

[12] Singh, G. and Green. M.
A High Level User Interface Management System. In Proceedings SIGCHI'89: Human Factors in Computing Systems, pages 133-138. ACM, New York, NY, 1989.

[13] Singh, G. and Green. M.
Designing the Interface Designer's Interface. In Proceedings, ACM SIGGRAPH Symposium on User Interface Software, pages 109-116. ACM, New York, NY, 1988.

[14] Mehta, S.
User Interfaces and the IEEE P1201 Committee. Unix Review 8(1):14-20, January, 1990.

[15] Szczur, L. and Miller, P.
Transportable Applications Environment (TAE)+: Experiences in Objectively Modernising a User Interface Environment. In OOPSLA'88 Conference Proceedings, pages 58-70. San Diego, CA, November, 1988.

[16] TeleLOGIC.
TeleUSE Reference Manual. 1989

[17] Thomas, J.J. and Hamlin, G.
Graphical Input Interaction Technique (GIIT) Workshop Summary. Computer Graphics 17(1):5-30, January, 1983.


Volume-Number: Volume 21, Number 179

EUUG

COMPUTING LABORATORY
**UNIVERSITY OF KENT**
AT CANTERBURY

# UNIX Courses at UKC

UNIX SYSTEM ADMINISTRATION
3 day Residential Course
25th – 27th March 1991

A classic complaint about UNIX documentation is that it assumes that you already know the answers, especially in the area of system administration. The course will be of particular use to those who are newly responsible for the day to day running of a UNIX system. It will not only provide a structured introduction to the techniques used in the daily housekeeping chores but there will also be an in depth coverage of the philosophy behind the techniques. Both lectures and practical sessions will be used to present the following topics:

- Documentation Overview
- The File System and its Implementation
- The Shell, Regular Expressions and Editors
- Operating Systems Fundamentals, Reboot and Shutdown
- User Administration, System Backups
- Security

Some exposure to UNIX will be assumed. No knowledge of the C language is required, although the ability to understand simple shell programs will be an advantage.

The course will be held in the Computing Laboratory at the University of Kent, which is housed in a modern building on a 300 acre campus a mile outside the city of Canterbury in South East England. The department has been using UNIX since 1976. In 1980, a VAX11/780 was purchased to fulfill the needs of interactive computing on the Kent campus. This machine ran UNIX, making Kent the first British university to offer UNIX as part of its mainstream computer service. The Kent campus is now served by an extensive network of UNIX based machines ranging from "mini-mainframes" to workstations

The tutors will be Peter Collinson and Alan Ibbetson. Peter worked for Kent for 15 years and was responsible for "making UNIX work" on the campus. He is now a freelance consultant, and has been made Consultant to the Computing Laboratory. Alan heads the team currently responsible for maintaining UNIX at Kent

The fee for this 3 day course is £350, including course notes and lunches. Bed and breakfast accommodation can be provided in single study bedrooms on campus as an optional extra at £16.35 per night. Evening meals may be purchased separately. The number of available places is limited, and places will be allocated in the order that bookings are received. Booking forms are available from:

Mrs J Farmer
Computing Laboratory
The University
Canterbury, Kent CT2 7NF
England
Tel: +44 227 764000 Ext 7695
Email: jcf@ukc.ac.uk

# EurOpen Executive Report

*Helen M W Gibbons*
*europen@EU.net*

*The European Forum for Open Systems*

Helen Gibbons is the business manager of EurOpen and is contactable at the EurOpen secretariat.

## A New Era Has Begun For The European Unix systems Users Group

### New Name

At the Governing Board meeting held in Nice on 21st October 1990, the decision was taken to change the name of the EUUG to **EurOpen**, "The European Forum for Open Systems".

The change of name reflects the development of the European UNIX community which has come to maturity over the past few years. It testifies to the significant contribution that UNIX has brought to the computer industry in general, and at the same time open the way to the era of Open Systems far beyond the early vision of the UNIX pioneers.

It is intended that EurOpen will expand its services to cover all aspects of Open Systems, including technical as well as strategic and market interests with the objective to satisfy the expectations of Open Systems Users.

To mark the change, a new logo has been designed. This logo will take over from the old EUUG "green sausage". A little "EurOpen" pin will be sent to every member.

The National Groups are now looking to see what can be done to bring their names and logos into line – so look forward for changes!

In line with the change of name the network address of the secretariat has changes to *europen@EU.net*, although the old mailbox will continue to work for a while.

### New Groups

Two new National Groups have recently affiliated to EurOpen, from the Soviet Union and Spain.

The Soviet UNIX Systems Users Group (SUUG) is based at 24 Vavilova Str, Moscow, 117812, USSR (Tel: +7 095 939 0056). It is not on the network at present. The Chairman is Vladas Leonas, and it has 28 members.

The Usarios de UNIX de España (UUES) based at Telefonica I&D, C/Emilio Vargas 6, 28043 Madrid, Spain. Tel: +34 1 337 42 22. Net: jmanas@dit.upm.es . It was formed on 11th July 1990 with 26 members. The chairman is Jose A Manas.

### The Governing Board Meeting

The total number of national groups now affiliated to EurOpen is 20, and all but Iceland were present

at the October meeting in Nice to receive the report from the Chairman, Michel Gien, and to consider the budget proposals for 1991.

Building on AFUU experience, it was decided to encourage the creation of Working Groups on various common interest subjects (technical and strategic) within the National Groups and to set up their co-ordination at the European level.

It was also decided to organise the representation of EurOpen towards standards organisations and to co-ordinate the building of a common European Users view on Open Systems standards. There will be more to report on both of these activities shortly.

It was with regret that the Governing Board learned of the resignation from the Executive Committee of Philip Peake and Neil Todd. Both had resigned due to pressure of work, and they were heartily thanked for the time and efforts that they had given to the EUUG in the past.

## EurOpen Directory

The EurOpen directory is now in the course of preparation, thanks to the Hungarian Group who have prepared the data base and will arrange for the printing of it.

The directory will be sent free of charge to all members of the National Groups by the end of the year.

## Software Distribution

Both tapes and cartridges are now available from CWI. Four new tape distributions have been prepared during the past 6 months:

- Nice conference tape
- a new GNU tape
- the Starter Kit
- PP 5.0 (X400 mailer for ISODE 6.0)

Any queries regarding tapes should be directed to Frank Kuiper, please refer to his column elsewhere in this newsletter.

## EUnet

The EUnet executive now consists of Yves Devilliers (Chairman), Axel Pawlik (Vice Chairman), Bjorn Eriksen (Treasurer) and Ted Lindgren (Secretary).

The group meets frequently to discuss EUnet related matters and EurOpen representatives are invited to attend both the EUnet executive meetings and the backbone meetings.

The InterEUnet service is now being provided in France, Germany, Holland and Italy, but the service is still in a pilot phase.

## Agreement with UniForum

EurOpen has agreed to co-operate with UniForum on three major topics:

- to co-sponsor a joint European Show and Conference in October 1992, in Amsterdam.
- to co-operate in the area of publications related to the UNIX and Open Systems communities.
- to co-operate in preparing concerted contributions to international standards bodies on matters relating to UNIX and Open Systems.

## EurOpen Conferences

Future planned EurOpen conferences include:

| | |
|---|---|
| Tromsø, Norway | 20 - 24 May 1991 |
| Budapest, Hungary | 16 - 20 September 1991 |
| Jersey | 13 - 17 April 1992 |

Some copies of the Munich proceedings are still available from the EurOpen secretariat at a cost of £20 sterling, and the Nice proceedings are available at £25 sterling.

# Working Groups

Following a proposal put by the French National Group, *AFUU*, the EurOpen Governing Board has agreed to a scheme to help National Groups who wish to set up Working Groups (WGs) and Special Interest Groups (SIGs) on topics of interest to their members.

Reflecting the fact EurOpen is an umbrella organisation, co-ordination of activities between groups will be undertaken at the European level. The AFUU experience shows that properly organised Working Groups are a valuable service to National Group members and thus encourage new members into the group.

This experience is to be put at the disposal of other National Groups through the choice of *Jean-Michel Cornu* as co-ordinator, working in conjunction with the Executive Committee though their National Grups liaison person. Working Groups set up within any National Group will be provided with information on the work of all similar groups by Jean-Michel. He will also help National Groups to set up workable structures for these groups.

It is intended to hold a preliminary meeting in Brussels in December to outline the international infrastructure that will be needed to support the national groups.

## What can I do ?

The success of Working Groups depends on the membership (ie you) using them to work together.

- Decide on a topic which interests you and which you feel could benefit from a Working Group.
- See whether other national members of your National Group share your interests.
- Contact your National Group committee and ask them to circulate your ideas to other members.
- If there is sufficient interest, ask your National Group to contact either Jean-Michel Cornu or Norman Hull to help with the arrangements for setting up such a group.
- Start working. Liaise with Jean-Michel, and benefit from the work being done by other groups as they will benefit from your work.

*Remember*, this is a service which EurOpen is helping your own National Group to provide, while still offering you the benefits of Europe-wide collaboration.

## What is a Working Group ?

A Working Group is a collection of interested individuals from many back-grounds (Manufacturer, Software Houses, End Users, ...) who have a shared interest.

This group will work towards satisfying the needs of its members. By working together the members can achieve better results than individuals or people from a single organisation.

The members of the group will be in regular contact and *actively* pursue the interests of the group.

A Special Interest Group has the same aims as a Working Group, but its members feel that they cannot devote the time needed to actively pursue the goals. They will, however, receive full information on what comes out of related Working Groups.

*Remember*, the more wide ranging the needs of the members in the group, the more successful the group is likely to be.

## What Sorts of Topics ?

Anything on which there is a high enough level of interest, such as:

| | |
|---|---|
| Benchmarks | Security |
| Workstations | Portability |
| Users Requirements | Real Time |
| Databases | User Interfaces |

This in not an exclusive list, just a few to give you the idea.

How would you like to work at the European level with others also interested in your favourite topic?

## Contacts

EurOpen contacts:

| Working Group Coordinator | EurOpen Executive Member |
|---|---|
| Jean-Michel Cornu | Norman Hull |
| 69 rue de Seine | Q2 Relational Software Ltd |
| 91130 Ris Orangis | Court Place |
| France | Carlow |
| | Ireland |
| Tel: +33 1 69 43 48 47 | Tel: +353 503 31745 |
| Fax: +33 1 60 78 00 98 | Fax: +353 503 43695 |
| Net: Jean-Michel.Cornu@enst.fr | Net: norman@q2rs.ie |

You should also contact your national group as they will know about Working Group organisation in your country.

# News from Italy

*Valeria Schiavi*
*valeria@siemil.uucp*

*i2u*

After 10-years of technical and organisational experience at the Computer Center of the University of Milan, Italy, Ms. Valeria Schiavi works now as freelance. Her task usually consists in assisting the technical staff of EDP sectors involved with Unix administration and Data Communication. Inside i2u she is responsible for : - keeping contacts with the Italian press; - informing EurOpen of the most important i2u activities; - gathering technical documentation about Unix/C/C++.

## i2u Convention

The i2u annual Convention will take place on 15-16-17 May 1991. The Convention will cover technical and applications issues, and is organised in the three traditional layers: Conference, Exhibition, Tutorial.

### Conference: 16-17 May 1991

Topics will include:

- Binary Compatibility standards
- Performance Evaluation
- Open Systems Architecture
- OLPT
- DBMS
- Security
- Unix and OSI
- Multiprocessing and Parallel Programming
- Image processing
- Standard
- Desktop
- New technical trends (Plan 9, Mach etc.)
- Market reports
- User experiences in the Italian area

### Tutorials: 15th May 1991

Advanced and basic technical themes will be covered.

### Exhibition: 15-17 May 1991

Access to the Exhibition will be independent from access to the Conference, but of course themes and demonstration covered by the Exhibition are in accordance with those of the Conference. Over 1500 visitors are expected.

### Location

The Convention will take place at the following address:

Centro Congressi Milanofiori
Assago
Milano
Italy

## i2u Catalogue

The i2u Catalogue of products and companies which deal with Unix in Italy was officially presented on 5th October 1990, at the International Fair "SMAU", Milan, Italy. The Catalogue includes a profile of about 100 companies and 400 products, and is a prototype of a more definitive "Version 1" scheduled for May 1991. "Version 1" will be distributed to attendants in the i2u Convention 1991. The Catalogue will be continuously enriched and updated, and reflects the information of the i2u on-line Data-Base.

# E-mail in the Alps - An Introduction to the CHUUG

*Patrik Eschle*
*eschle@physik.unizh.ch*

*Physik University Zürich*

Patrik Eschle is the secretary of the CHUUG. He works on his thesis at the Physics Institute of the University of Zürich where he tries to bounce laser beams off each other in sodium vapour.

The Swiss Unix Systems Users Group CHUUG has existed now for almost a year and has been a member or EurOpen since Spring 1990. Although our logo has only just made it to the back cover of the EurOpen newsletter, it's time to introduce ourselves to the EurOpen community.

Many years ago, there once was a group of Swiss Unix users. That group was also a member of the EUUG and so were the members of that Swiss group. Then, one day in Spring 1989, the members of the Swiss group received a letter from the EUUG saying that the Swiss group was no longer a national group of the EUUG and that the members of the Swiss group had to become EUUG direct members.

The reasons for the separation of the Swiss group from the EUUG are not important here. Fact was, that there was no longer a Swiss National group and that the Swiss users of EUnet services were no longer represented in the EUUG.

The Swiss backbone at that time was cernvax and was run by Dietrich Wiegandt from CERN. For financial reasons, CERN wanted to stop the service and so the need for a new Swiss backbone arose.

This announcement from CERN was followed by long e-mail discussions that clearly showed the need for a forum of the e-mail users. In the end it was decided that the backbone should be run by SWITCH, a foundation of the Swiss Universities with the primary goal of providing telecommunication services to the academic community.

So Switzerland once again had a backbone, but still no national group. It was Georges Schild who took the initiative in Autumn 1989: CHUUG was founded on October 30 and accepted as a EUUG member in Spring 1990.

So much for the historical part. What's the situation of the CHUUG today, one year after it was founded?

We have about 70 members, roughly half of them are individuals and the other half are commercial members. About a third of them are former EUUG direct members, many of the others are new to electronic mail and news. We were quite surprised by the large demand for electronic communication; there really seems to be a need for e-mail services and knowledge about it.

Most of the CHUUG members and the executive board are from the Swiss-german speaking part of Switzerland but we would really like to see French or Italian speaking members on the executive board. Our statutes are written in

English, German, French and Italian, and all letters to the members have at least an English translation.

As mentioned above the Swiss backbone is not run by the CHUUG itself, but by SWITCH. This is not only due to historical reasons but also due to the fact, that SWITCH has good connectivity, offers professional services and that with Simon Poole the backbone is run by an excellent postmaster.

On the other hand, SWITCH's services are not cheap and there are some political controversies between SWITCH and EUnet. The CHUUG therefore plans to run it's own backbone. No dates are fixed yet nor do we have the money at the moment, but the project exists.

A separate CHUUG backbone would be most interesting for all private and commercial e-mail users. The academic world is fully connected to the internet and has all wishes fulfilled by SWITCH. The private and commercial sites on the other hand use, with a few exceptions, UUCP over telephone or X.25 and about 50 of them are directly connected to the SWITCH central system chx400. A CHUUG backbone could provide cheaper and more decentralised services to the Swiss EUnet community.

What are we doing beside planning backbones? We've held our first tutorial on October 8, where Simon Poole introduced us to the secrets and traps of sendmail configuration and Joseph Gaehwiler talked about UUCP (both talks are available from the CHUUG secretary at nominal costs). We also plan to start publishing a newsletter next spring.

Our general meeting will be held in December at Olten, the exact date is not fixed yet - I hope to meet many of you there!

# News from DKUUG

*Keld-Jorn Simonsen*
*keld@dkuug.dk*

*Center for Anvendt Datalogi*
*University of Copenhagen*

Keld Simonsen is the chairman of the Danish UNIX User Group. He is currently occupied with standards and networking activities.

What has happened since April in Denmark? Read on and see.

We finally got our login service up and running. The DKUUG login machine is running mail and news, and you just need a modem and a terminal or PC to connect to it. You need to be a DKUUG member and then you pay something like 60 ECU a year to be connected - there is no extra charge for any mail sent or received. Reading news costs 1 DKK a minute, however.

DKnet has got its site number 100, and new volume-dependent charges have been introduced for news. We still have the fixed prices for news, but for small volumes it is cheaper just to pay by volume under the new price scheme.

Some people go and some come... We have had to find a new newsletter editor and a new backbone manager.

DKUUG had a stand at the recent office & computer exhibition in Copenhagen. Quite some people visited us and had a look at mail and news.

Our member meetings have been quite successful, attracting on average more than 100 people each time. Themes have been: mainframe UNIX, financial UNIX, UNIX and public services, desktop UNIX, system development, administrative systems and more.

The club-evenings in Copenhagen have had some problems with missing speakers, but hope to be on the track again. Themes have been standards and networking.

The DKUUG-nyt newsletter is out once a month (usually).

DKUUG is very active in POSIX standardisation, and we have been involved in producing the first national profile for POSIX. This profile covers Danish, Faeroese and Greenlandic, it is character set independent, and actually is known to cover more than 60 character sets. It builds on some results from the multi character set mail described elsewhere in this issue, and as the first national profile, it is included in the POSIX standard as an example in an appendix.

That completes the Danish story for now.

# Recent EUnet Development in Finland

*Pekka Keskiivari*
*pnr@ngs.fi*

*Nixu Oy, Finland*

Pekka is a freelance consultant and a member of the FUUG board. He also gives lectures at the Helsinki University of Technology. In the past couple of years he has been mostly involved with operating systems and communications. His interest, however, is shifting towards user interfaces and psychology, since he has become convinced that computers are, well, blunt and dull instruments.

Recently quite a lot has happened within the Finnish part of EUnet. The backbone has moved from Tampere University of Technology to Otaniemi Science Park, and is now administrated directly by the Finnish Unix Systems User Group. Due to this, the administration etc. has also, of course, changed. With the new backbone we have started testing ISDN as the transport media for UUCP-traffic. We have also started investigating how it would be possible to use these ISDN lines for TCP/IP traffic. Recent events in the USSR and Estonia have allowed us to get UUCP connections to Moscow and Tallin. So, quite a lot is going on.

## New Backbone

Last spring Digital Equipment Corporation Finland donated the Finnish Unix systems User Group (FUUG) a DecSystem 5400. The new machine came with 32 Megabytes of main memory, 1.2 Gigabytes of disc space and 16 serial lines. This donation finally made it possible to transfer the official backbone from Tampere University of Technology to the FUUG. This had been a goal for a long time, mainly motivated by the increased load to TUT, and their request to move the backbone.

The new backbone is located in Otaniemi Science Park. It currently has five 9600 bps modem lines and one ISDN line. The ISDN line can carry two B-channels, which means that we can have two simultaneous connections through the ISDN line. We also expect to get another ISDN line within a couple of weeks.

The backbone, which, by the way, is named "fuug.fi", is directly connected to Otanet, the Otaniemi wide local area network. This 10 megabit fiber optic Ethernet connects "fuug.fi" directly to Helsinki University of Technology, Technical Research Centre of Finland, and State Computing Center. It also provides connections to FUNET, the Finnish University and Research Network, and NordUNet, the Scandinavian wide Academic network.

Currently the possibility of carrying TCP/IP traffic to various places is limited because of political difficulties. Full TCP/IP traffic is only allowed inside Otaniemi region, which is just a campus area, maybe two or three kilometers wide. We hope that with the new agreement between EUnet and NordUNet we should become able to carry full TCP/IP traffic to the rest of EUnet. Currently we are only allowed to carry SMTP, NNTP, UUCP and Berkeley BIND protocols.

The new backbone quickly proved its importance. At the beginning of August, after we had succeeded in resolving all the initial problems, including some nasty hardware problems that were apparently caused by too high a temperature in the machine room during the hot summer days, we started to move the UUCP sites to the new

backbone. Currently nearly all UUCP sites are already using the new machine.

## New Administration

With the new machine we acquired new administration staff, naturally. Petri Ojala, a bright, young student of technology at Helsinki University of Technology, started to study the skills required for UUCP backbone administration last January. Today he is one of the most prominent UUCP and networking specialists in Finland. He is the person mainly responsible for the day to day operations and administration of the backbone.

Pekka Nikander (me), a freelance consultant, is the member of the FUUG board that is responsible for the network. However, due to my other duties I have been too busy lately and therefore Petri has taken care of nearly everything to do with the net.

These two people, Petri and Pekka, together with Marko Kaittola of Finnish University and Research Network, also a member of FUUG board, form the kernel of Finnish EUnet administration and maintenance. As mentioned, Petri takes care of the day to day operations and administration. Marko is the key person in X.400 and X.500 development. This work is mainly funded by FUNET, and only partially FUUG or EUnet business. The coordination of work and new developments are the responsibilities of Pekka.

## ISDN

The local Helsinki Phone Company, HPY (Helsingin Puhelin Yhdistys), is actively testing new ISDN technology. EUnet Finland is one of the pilot customers testing this technology. Currently we have got one Basic Rate ISDN line to the backbone. This means that we can have two 64 kb/s lines open simultaneously. Several of the EUnet sites in the Helsinki region have also got a Basic Rate connection. HPY has promised to provide the backbone with another Basic Rate connection within a couple of weeks. So, when you read this story, we should already be able to have four 64 kb/s connections open simultaneously.

Our ISDN development project is planned to proceed in a number of phases:

1. Connect individual sites to the backbone using ISDN-RS-232 converter boxes and using UUCP protocol. This is working now and we routinely run at 19200 bps asynchronous using V.110 conversion.
2. Replace the usage of UUCP with SLIP. We are actively working on this. The remaining problem is mainly due to some difficulties with Ultrix. Ultrix version 3.1 SLIP contains some bugs, and we cannot run Ultrix 4.0 due to errors in serial port drivers. However, it is well possible that this will be in place when this story gets printed.
3. Raise the speed to full 64 kb/s. We need synchronous i/o cards and drivers. Petri Alhola of Digiware is trying to find equipment for this stage.
4. Replace external adapters with internal ISDN cards. Digital has developed an ISDN card that could be used in DecStation 5400, but they have currently software only for VMS, not for Ultrix. There is also a number of test products for other Unix machines. I am aware of the following ones:
   - — software and card for 386 Unix by BinTec Computersysteme, Germany
   - — software and card for Sony workstations by Sony
   - — packet driver software for MS-DOS by Michigan University
   - — ISDN-TCP/IP gateway software for MS-DOS by Systec, Germany
   - — University College London Primary Rate ISDN card and software for Suns and VME-bus

If anyone knows of other products or solutions, please let us know. Also, if you have any experience with any of these, please let us know. As I write this, I have not yet seen any of these products in working, and I do not know very much about them.

## USSR

The recent developments in the USSR has made it possible for them to make modem connections to the outside world. Petri Ojala has done really good work in this field. He has put up UUCP connections to both Demos, Moscow, Russia, and Estonian Academy of Science, Tallin, Estonia. The Soviet Union has also got a new top level domain for themselves: *.su*. If you are interested,

you can contact Alexander Shcmundak of Estonian Academy of Science <sasha@ioc.ew.su>, or Dmitry Volodin of Demos Co-operative <dvv@hq.demos.su>

### Conclusions

There is quite a lot happening with EUnet here in Finland. I just wish that I and the others had a little more time. This all would go ahead a bit faster.

### Congratulations

Johan Helsingius got married to Jaana on the 21st September. We wish him as much vigour in this new enterprise as he has displayed with the FUUG.

---

## The Ultimate Guide to the Vi and Ex Text Editors,

Hewlett Packard, Benjamin/Cummings, 1990, ISBN 0-8053-4460-8. (UK) Price £19.75, Soft Back, 334 pp. Reviewed by Joseph Bustard of Picturefast Computing Ltd.

Every new machine has it's own set of programmer tools for software development. Moving from one system to another can often involve learning a completely new set of commands and options. This is particularly true when it comes to editors with graphical user interfaces which depend on the capabilities of the host's display hardware. Despite it's age, many programmers still prefer to use "vi", the closest thing there is to a standard visual editor. As an integral part of Unix, it is always available, works on almost any terminal and is powerful enough to handle any task. It has a rich command set that few programmers ever completely master, usually learning by picking up bits of information from colleagues. A definitive guide is obviously long overdue. The "Ultimate Guide" sets out to fill this role and does so reasonably well. It is aimed at both new users, who need to know the basic commands for text manipulation, and at experienced programmers who are interested in it's more obscure capabilities.

The basic commands are covered in great detail, many of them are more flexible than most users realise. An entire chapter is devoted to starting and stopping the editor, including how to recover partially edited files after a system crash. Later chapters discuss text manipulation, cursor positioning, manipulating cut and paste buffers, search and replace operations, shell commands and editor configuration. New users are lead gently by the hand starting with the basic operations needed to do everyday editing on program files, through to more complex techniques that make a programmers life a whole lot easier. Throughout, there are lots of little gems. Command sequences that you would never think of yourself but that are obvious when you see them. These vary from the simple "dwwP" (delete word, move to next word, paste buffer before the current character) which swaps two words around, to shell commands such as "`!} sort | pr -t4`" which sorts the file into a four column table. I found the description of tag files particularly interesting. These allow vi to start up and position itself at a particular routine without the user having to name the source file or search through it for the right line. This facility is great if you editing functions in a program that is split over a number of files but is hardly mentioned in the standard Unix documentation. Editor configuration is another area where even experienced vi users tend to be vague. Most know that auto-indentation can be turned on or off by putting "set autoindent" or "set noautoindent" in the .exrc (editor resource) file but there are lots of other options that no one seems to know about. The "Ultimate Guide" devotes an entire chapter to these.

Overall, the book is written in an easy to read style and is well organised. It is very easy to find what you want. Beginners will find it useful as a tutorial while more experienced programmers will tend to use it as a "dipping" resource, using it to solve specific problems or to learn specific techniques that they know vi should be able to do, if only the right buttons are pushed. I doubt if a book about a programmers' editor like vi will ever make the bestsellers list but I'm sure that this one should find a home in many programmer's desks and many company libraries.

# UKUUG News

*Mick Farmer*
*mick@cs.bbk.ac.uk*

*Birkbeck College*

Mick is a lecturer at Birkbeck College (University of London) and the Secretary of the UKUUG. His interest is in all aspects of Distance Learning and he is the Senior Consultant (Software) for LIVE-NET, an interactive video network connecting London's colleges. He is also a member os the University's VLSI Consortium, mainly because the design tools draw such pretty pictures.

## Start Bit

The UKUUG wishes our umbrella organisation, EurOpen, all the best with its new drive into Open Systems. The new logo was out in force during the recent conference in Nice where the change of name was announced. Apparently, the grid represents the network connectivity within our community (the large circle). Wow!

## Membership Figures

Our membership continues to increase and now stands at 565, broken down into the categories given below.

| Category | Number |
|---|---|
| Academic | 124 |
| Commercial | 356 |
| Honorary | 5 |
| Individual | 67 |
| Student | 13 |
| Total | 565 |

Our recent London Conference helped the boost in membership so it's up to us to keep their interest going!

## Forthcoming Events

- Our Winter '90 Technical Meeting will take place at Queen's College, Cambridge (England) on 17-19 December, 1990. As with all our winter meetings, this one will have a strong networking flavour.
- Our Summer '91 meeting will be held in Liverpool (England) in June or July 1991.
- Our Winter '91 meeting will take place in Edinburgh (Scotland) in December 1991.

Please note that EurOpen members and members of other national groups are always welcome to attend our events.

## Workshop Videos

We have produced two video programmes on relevant material to those working in the UNIX community. Both of these are the result of successful one-day workshops organised by the UKUUG.

- UNIX Security

  A three hour video discussing the following topics:

  — The HACKMAN project; System V/MLS; An analysis of the Internet worm.

— The Sun Yellow Pages system; Secure RPC; Some myths and facts about UNIX security.

- And more...

UNIX System Administration

A four hour video discussing the following topics:

— POSIX developments; System management; Managing X.400 mail systems.

— Project Athena; System administration in a heterogeneous environment.

- And more...

Each video costs £60.00 (plus VAT in the UK) and can be ordered from the UKUUG Secretariat, by e-mail from ukuug-videos@ukc.ac.uk, or directly from:

Birkbeck College Video Services
Department of Computer Science
Birkbeck College
Malet Street
London WC1E 7HX
UK

+44 71 631 6351

## London UNIX User Group (LUUG)

This lively group continues to meet on the last Thursday of every month (except December). Contact Andrew Findlay for details (Andrew.Findlay@brunel.ac.uk).

## Midlands UNIX Group (MUG)

A successful first meeting was held on 17 October when Julian Onions of X-Tel Services Ltd. (formally from Nottingham University) spoke about X.400/PP to a captivated audience. Contact Kevin Hopkins (mug-request@cs.nott.ac.uk) or David Osbourne (cczdao@clan.nott.ac.uk) for details of forthcoming meetings.

We wish them, and any other embryonic group, all the best. The UKUUG is always willing to assist the formation of local UNIX groups throughout the country.

## FaceSaver Project

This service is now in full swing and we intend taking it to most major conferences in the future. At the recent conference in Nice our stand took pride of place next to the (fledgling) EurOpen stand.

This project is being supported by Acorn Computers Ltd. of Cambridge (England) as part of their on-going commitment to UNIX .

## Software Distribution Service

We currently have over 600 megabytes of (compressed) source archives on-line for UKUUG members to access. The distribution service is offering the same service to EurOpen members (through their national groups) at the same price. E-mail enquiries to ukuug-software@doc.ic.ac.uk.

Please note that there was an error in the article describing this in the last newsletter, the ordering contained 3 items, depending on the type of tape that was ordered. These should have been (in order):

- 6250
- 1600
- QIC24

## Stop Bit

Open Systems is the name of the new game in town. UKOpen, UK Open Systems Group, Open/UK, EurOpen-UK, UK UNIX and Open Systems Group, Open UK Systems, ...

What do you think?

# CALL FOR PAPERS

NLUUG Spring conference 1991
8 May 1991

in "De Reehorst", Ede
The Netherlands

## Invitation

The program committee for the bi-annual NLUUG conferences invites speakers to make a contribution on the theme:

## SYSTEM MANAGEMENT
### Can Open Systems be Secure ?

The purpose of the conference is to discuss recent developments in the field of System Management from different points of view; day-to-day reality, technical difficulties, scientific approach.

System Management currently is going though some rapid changes. Systems are more and more "Open", connected through "Open" networks, and the capacity of systems increases tremendously. The trend towards more Open-ness in both systems and networks may well be in conflict with security. Also, the procedures for system management are undergoing changes. from "several system managers for one system" to "several system managers for several systems" to "one system manager for (many) more than one system". Network management becomes an integral part of system management.

The theme covers, but is certainly not restricted to, the following topics:

— User administration

— Monitoring

— Access Control

— Software Installation

— Back-up procedures

— Accounting techniques

## Contacts:

Jos Alsters <josal@sci.kun.nl>
Katholieke Universiteit Nijmegen
Afd CCZ,
Toernooiveld 1,
6525 ED Nijmegen,
The Netherlands
+31 (0) 80-653535

Emile van Dantzig <dantzig@tmu.uucp>
Transmediair Solutions bv
Postbus 297,
3720 AG Bilthoven,
The Netherlands
+31 (0) 30-281820

Jack Jansen <jack@cwi.nl>
Centrum voor Wiskunde en Informatica
Kruislaan 413,
1098 SJ Amsterdam,
The Netherlands
+31 (0) 20-5924098

Hans Doorenbos <hado@hpuamsa.uucp>
Hewlett-Packard Nederland B.V.
Startbaan 16,
1187 XR Amstelveen,
The Netherlands
+31 (0) 20-5476911

Dick Wiersma <dick@cnps.philips.nl>
Philips C@P services, Ned. Philips Bedrijven B.V.
Postbus 218,
5600 MD Eindhoven,
The Netherlands
+31 (0) 40-732354

## Important Dates:

| | |
|---|---|
| Deadline for submissions: | 1 January 1991 |
| Deadline for abstracts: | 1 March 1991 |
| Deadline for papers: | 15 March 1991 |

## Conference Details:

| | |
|---|---|
| Time-schedule of conference: | 09.30 - 17.00 hrs |
| Date of conference: | Thursday, 8 May 1991 |
| Location of conference: | Ede, "De Reehorst" |

The time slot for a presentation will be 30-45 minutes.

To have proceedings ready on the day of the conference, speakers will be asked to submit the full text of the paper in advance (deadline 15 March).

## General:

In general, questions can be put to:

Patricia Otter <nluug@nluug.nl>
buro BLUUG
Burg. Verderlaan 15X,
3454 PE De Meern,
The Netherlands
+31 (0) 3406-22055

# Book Review

## UNIX Shell Programming,

Lowell Jay Arthur, John Wiley, 1990
ISBN 0-471-51821-2.
(UK) Price £22.95, Soft Back, 272 pp,

Reviewed by Rob Henley of Siemens SDG, Reading, UK. (rob@siesoft.co.uk)

*Achilles and the Tortoise have met at the Cafe Clytemnestra just outside the gates of Troy. They are discussing a new book on one of the Tortoise's favourite subjects: shell programming.*

*Achilles:* I must say I liked the book - all that talk of "power users" and "travelling the avenues of knowledge and wisdom"!

*Tortoise:* Yes, but I wouldn't get too excited - it goes on to say that a "gung-ho power user" is someone who uses a *.exrc* file! And didn't you find some of the claims rather misleading. I can't really see the shell is very object oriented for instance. How about polymorphism?

*A:* Certainly very attractive, but I can never quite remember her face ...

*T:* Moving quickly on, how about "one line of shell can do the work of 100 lines of C" or - apparently from Brad Cox's antipodean brother Bruce - "10 lines of C++"!

*A:* This is true in some specific cases - for example there's a good example of a concordance program, but I suppose the claim is not very general

*T:* Or "the syntax of English is subject-verb-object".

*A:* What?

*T:* Exactly - and you must have seen the gaff about UNIX files being null terminated!

*A:* Really! *(thinking)* I wonder if I could use all those spare bytes to give me a little more room on my Pomegranate? After all, the're clearly redundant and I must have thousands dotted all over the disk!

*T:* Well, that's not the only problem you'll have if you follow this advice. Straight after advocating the shell as a rapid prototyping tool which should always be used before diving into C, it recommends liberal use of the sticky bit to speed things up.

*A:* Sticky bits and diving into the sea?

*T: (becoming vexed now)* Frankly the book is not well organised, the treatment of many of the commands is rather shallow and the examples and exercises are uninspiring to a Tortoise. Moreover, there's a notable lack of continuity and ...
*(just then some shifty looking characters wander past pushing a large wooden horse on wheels)*

*A:* A wolf in sheep's clothing ... *(watching the horse)*

*T:* Indeed - and an empty vessel ... *(looking at the picture of a nautilus)*

*A:* Oh - hardly empty I think. But anyway you must admit it is helpful to have the Bourne, C and Korn shells expounded side by side - especially in areas like flow control

*T:* True - although one of my heathen cousins from Mercia recently came up with a rather good alternative called WASH

*A:* I suppose they'll need it over there with all that wode to get off - but would a heathen know all the sophisticated tricks of the trade given in this book - or think of treating the shell as a 4GL with RDBMS facilities.

*T:* I agree there are some good ideas in the book, but it nevertheless seems more of a hard sell of UNIX to you captains of industry than an inspiration to us "techies" (humph!). I can't help feeling Kernighan and Plauger's "The UNIX Programming Environment" takes some beating.
*(the horse is drawn into the gates as the friends go their separate ways)*

# Publications

## EurOpen Publications

EurOpen publications may be ordered from the Secretariat at Owles Hall.

The publications available are listed on the next page, it is planned that EurOpen members will be able to also order National Groups publications from one single point: the secretariat.

The EurOpen publications available are:

EurOpen Newsletter

Your EurOpen National Group membership gives you one free copy of this newsletter, you may order extra copies of the newsletter for distribution within your organisation. Please contact your national group who will arrange this for you.

Proceedings

Proceedings from some past conferences are still available. The list of those that you can order is on the next page.

European E-Mail directory

This is a must if you are a serious e-mail user. It provides you with a way of referencing e-mail sites thoughout Europe by means of several different indexes.

The second edition to this has been prepared and will be available soon.

## USENIX Publications

We receive requests for USENIX publications and, as a result of our close cooperation with that organisation, we are pleased to announce a service by which a range of publications can be ordered through EurOpen. The range includes:

;Login:

This is the USENIX newsletter which is published every two months, and carries articles on a variety of topics of interest to UNIX users everywhere and usually runs to around 40 pages.

Computing Systems

This is the journal of the USENIX Association, a quarterly publication which is devoted to the analysis and understanding of advanced computing systems. It is perfect bound with a printed spine for ease of reference and usually runs to around 100 pages.

Proceedings

Some proceedings from past USENIX conferences and workshops are available from stock – the current list is enclosed with the attached order form.

4.3 BSD Manuals

These manuals have proved popular with EurOpen members but our stock is now exhausted. However, USENIX can continue to supply them direct and the only difference from the past ordering procedure is that you may have to wait a little longer to receive them direct from the States.

Ordering procedure. An order form is enclosed with details of how payment can be made by credit card, by direct payment to the bank or by certain types of cheque or banker's draft.

# Publications available though EurOpen

| | | | |
|---|---|---|---|
| :Login: | The USENIX newsletter | 6 issues/year | £20.00 |
| Computing Systems | The USENIX Journal | 4 issues/year | £30.00 |
| USENIX Proceedings | Anaheim Conference | June '90 | £25.00 |
| | C++ Conference | Apr '90 | £30.00 |
| | Washington DC Conference | Jan '90 | £27.50 |
| | Graphics Workshop V | Nov '89 | £19.00 |
| | Dist & Multiproc Workshop | Oct '89 | £33.50 |
| | Large Inst Sys Admin III Workshop | Sept '89 | £15.00 |
| | Baltimore Conference | June '89 | £24.00 |
| | UNIX Trans Proc Workshop | May '89 | £13.50 |
| | Software Management Workshop | Apr '89 | £24.00 |
| | San Diego Conference | Feb '89 | £33.50 |
| | C++ Conference | Oct '88 | £33.50 |
| | C++ Workshop | Nov '87 | £33.50 |
| | Graphics Workshop IV | Oct '87 | £17.00 |
| | Washington DC Conf | Jan '87 | £20.00 |
| | Graphics Workshop III | Dec '86 | £17.00 |
| 4.3 BSD Manuals (EurOpen Members Only) | User's Manual Set<br>Programmer's Set<br>System Manager's Manual | | Contact EurOpen at Owles Hall for prices |
| EurOpen Proceedings | Dublin | Autumn '83 | £2.00 |
| | Nijmegen | Spring '84 | £5.00 |
| | Cambridge | Autumn '84 | £5.00 |
| | Paris | Sprint '85 | £5.00 |
| | Copenhagen | Autumn '85 | £10.00 |
| | Finland/Sweden | Spring '87 | £20.00 |
| | Dublin | Autumn '87 | £20.00 |
| | Munich | Spring '90 | £20.00 |
| | Nice | Autumn '90 | £25.00 |
| Directories | European E-Mail directory, 1st edition | | £18.00 |
| | European E-Mail directory, 2nd edition | | £20.00 |

All prices include post and packing.

The price for ;Login: and Computing Systems is for a one year subscription.

ORDER FORM: This page may be photocopied for use. Please print!

Name: .....

Company name: .....

Address: .....

.....

E-mail address: .....

I would like to order the following: .....

.....

I enclose my payment in the sum of £ .... and understand that a receipt will be sent.

I would like to receive membership details of EurOpen: YES/NO ....

Payment may be made in one of 4 ways:

1. By UK Cheque or Bankers Draft, made payable to EurOpen and drawn on a UK bank. Eurocheques are acceptable, but each cheque must be for £100 or less.
2. By Direct Payment to EurOpen's bank. Please tell your bank that you will pay all charges so that EurOpen will receive the full ammount due.

   The Bank of Scotland
   61 Grassmarket
   Edinburgh
   Scotland EH1 2JF

   Account Number: 00613997
   Bank Sort Code: 80-31-50

3. ☐ by VISA ☐ by ACCESS/EUROCARD/MASTERCARD

   Name as it appears on the card (block capitals) .....

   Address of card holder .....

   .....

   Card Account No ..... Date of Expiry .....

   Signed ..... Date .....

4. Payment can also be made in ECUs by direct payment to EurOpen's bank, which is:

   The Bank of Scotland
   International Division
   Operations Dept
   PO Box 86
   120 St. Vincent St.
   Glasgow
   Scotland G2 5DZ

   Account Number: 41791 ECU 01
   Bank Sort Code: 80-20-13

EUUG

# EurOpen Spring '91 Conference and Exhibition

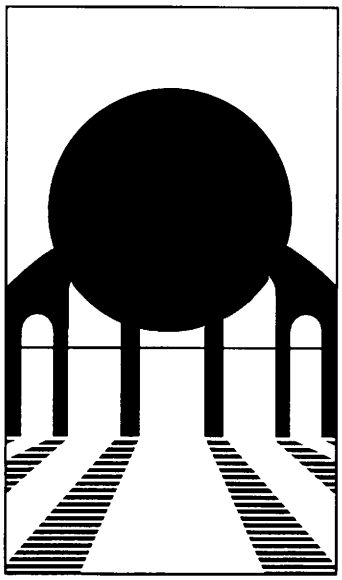

Tromsø, Norway
20-24 May 1991

## Announcement

The Spring 1991 conference of The European Forum for Open Systems will be in Tromsø in Norway from 20th to 24th May 1991.

The conference will concentrate on Distributed Open Systems in Perspective. In addition to an overview of the issues involved in the design of distributed open systems the conference will address problems encountered and solutions found when distributed open systems are employed.

The three day Conference with Commercial Exhibition on Open Systems, UNIX and related subjects will be accompanied by Technical Tutorials on the Monday and Tuesday, followed by the main conference on Wednesday to Friday.

## Warm Clothes not Needed

Although this conference will take place North of the Arctic Circle; the warmth of friendship of fellow delegates, not to mention, the heat (flames?) of the technical discussions, will be more than enough to keep you from feeling the cold.

5th November, 1990
17th December, 1990
4th February, 1991
11th February, 1991

# USENIX Association News for EurOpen Members

*Donnalyn Frey*
*donnalyn@frey.com*

*Frey Communications*

Donnalyn is the USENIX Association Press Liaison. She provides members of the press, USENIX Association members, and EurOpen members with information on the activities of the USENIX Association.

## Ah Ha, the Missing Column!

Due to a printing error, this column was missing from the last issue of the EUUGN. The column therefore combines the Autumn and Winter 1991 news.

## The 1990 Summer USENIX Association Conference

Dennis Ritchie, of AT&T Bell Laboratories and co-author of the UNIX operating system, presented the keynote address at the USENIX Association 1990 Summer Technical Conference and Exhibition on 11 - 15 June at the Anaheim Marriott Hotel and Convention Center in Anaheim, California. This is reported by Marc Donner.

### OPEC vs. the Medellin Cartel

Dennis Ritchie, looking distinguished in a rarely worn suit and tie, delivered the USENIX Conference and Exhibition keynote address entitled "What Happens When Your Kid Turns 21?". This was Dr. Ritchie's first conference address in over three years. Departing from conventional expository style, his keynote was in the form of a telephone interview.

Ritchie began by acknowledging the influence of Multics, a project on which both he and Ken Thompson had participated in the 1960s. Features of UNIX inspired by Multics include the tree-structured file system and the concept of the shell as a separate program. He mentioned that he is a bystander in UNIX and C today.

Ritchie answered questions on the current state of UNIX and C, particularly recent standardisation efforts. He noted that the ANSI C standard seems to be a good standard, having clarified and updated many items without breaking them. In addition, he observed that while the core of the POSIX effort seemed sound, he was uncomfortable with the lack of coherent network facilities in the proposed standard and he found Real-Time Extensions work to be troublesome. In general, he mentioned he was ambivalent about standards as they seem to create a tension between standardisation and innovation. The real problem, he noted, is when is the time right to standardise?

Ritchie briefly commented on the contest between the Open Software Foundation and UNIX International, observing that the squabbling was harmful but, given the development history of UNIX, inevitable. The future movie could be called "OPEC vs. the Medellin Cartel" or perhaps "Ninja Turtles vs. the Fantastic Four." He mentioned that when he told Ken Thompson about the OSF, Ken noted dryly "Imagine, IBM

and DEC in the same room, and we did it." Ritchie also said that the OSF could be criticised for not producing less than might be expected, but agreed that the bygone days of easily available source code were definitely better.

The imaginary interviewer on the telephone asked Ritchie his opinion of C++. Ritchie declined to comment, citing an "agreement with Bjarne [Stroustroup] that I don't give lectures on C++ and he doesn't talk about old C."

The final part of the keynote "interview" was more philosophical in tone. Asked about his greatest satisfaction from the UNIX work, Ritchie cited the influence it has had in creating new companies and new directions for old companies. The advent of inexpensive high-performance microprocessors was brought about, in his view, by the wide availability of portable software to run on them, eliminating the need for an expensive software development effort.

Asked about things he would do differently, Ritchie responded that Ken [Thompson] always said he would spell 'creat' with an 'e.' "More seriously," he said, "there are lots of small sins." He mentioned the experiment in declaration syntax embodied in C had probably not succeeded. He complained that many implementations of UNIX had become bloated, observing that the 10th Edition, the latest version at Bell Laboratories, required only 140K of text space. Overall, Ritchie said, "I have very few qualms about UNIX as a whole; it has turned out to be very adaptable."

When asked to discuss his regrets, Ritchie noted that the opportunity to create a viable alternative to X was lost when the Blit work was not aggressively pursued. He also regretted that the UNIX work had not satisfied AT&T's need for a system to manage very large software development, like the 5ESS system. He mentioned that either the small-is-beautiful development model of simple, elegant powerful tools isn't appropriate to large system development, or the message didn't get across.

Asked for a message to attendees, Ritchie observed that UNIX, as a child that has come of age, was independent and that all he could do is let it run its own life and wish it well.

With that message, the interviewer asked Ritchie what he had been doing lately? After mentioning having written an ANSI preprocessor for C, Ritchie proceeded to show a videotape of an infamous practical joke that Ritchie, Rob Pike, and magicians Penn and Teller played on their boss, Nobel Prize winner Arno Penzias.

### Concurrent Sessions

A second track of the conference featured invited talks on such topics as: Peter Langston of Segue Software and Mike Hawley of MIT Media Lab demonstrating how computers make music, and how are they being used in music production, arrangements, and composition; Andrew Hume repeating his popular talk on regular expressions; Craig Hunt, of the National Institute of Standards and Technology discussing TCP/IP system administration; Rob Kolstad of Sun Microsystems moderating a system administration problem solving panel; and Neil Groundwater doing an introduction to X Window concepts.

## Scholarship Winners

Because of the excellence of the applicants this year, the USENIX Association awarded two scholarships for 1990-91. The winner of the competition is John Heidemann of the University of California at Los Angeles. The runner-up award has been given to Matthew Blaze of Princeton University.

## 1991 Winter Conference in Dallas, Texas

The 1991 USENIX Winter Technical Conference will be held in Dallas, Texas on January 21 - 25, 1991. The theme of this conference is "What's next: by the year 2010, evolution or revolution? Unix derivative or Something Else?"

A complete program is in this issue.

## Summer Conference and Exhibition in Nashville, TN

The 1991 USENIX Association Summer conference and Exhibition will be held in Nashville, Tennessee on June 10 -14, 1991. The theme of the conference is "Multimedia - Interfaces of the Future". A call for papers is included in this issue.

## 1991 USENIX Workshops

Upcoming workshops include:

**Software Development Environments in UNIX** on January 16 - 18, 1991 at Grand Kempinski Hotel in Dallas, Texas cosponsored with the SIGMA Project of Japan.

**Distributed/Multiprocessor Systems Symposium**, co-sponsored with SERC of Purdue University, ACM SIGCOMM and SIGOPS and the IEEE-CS Technical Committees Distributed Processing, Operating Systems, and Software Engineering on March 21 - 22, 1991 in Atlanta, Georgia.

## Monograph Series on Advanced Computing Systems

Marc Donner, of the IBM Thomas J. Watson Research Center, has been selected as the editor of the USENIX Association's Monograph Series on Advanced Computing Systems, to be co-published by USENIX and the MIT Press. Marc has written the following request for books and monographs from the computing community:

The USENIX Association and the MIT Press intend to co-publish books and monographs on the general topic of computing systems. The intended audience for these books is the community of system designers, builders, users, and scholars. Our intent is to publish material of lasting interest and importance, with an emphasis on actual systems. Subjects may include design, implementation, history, and analysis of real systems. While we are inspired by UNIX and UNIX-like systems, we do not expect to limit our attention to such systems in any way, as we see ourselves responsible to the entire systems community.

We see several specific needs that we would like to satisfy and for which we solicit manuscripts. The needs fall in two areas - books in traditional styles and formats about topics important to the systems community and things new or unusual.

Among things new or unusual, we are interested in exploring at least these ideas:

**Significant systems** - many significant systems are documented, if at all, only in reference manuals or user guides. Journal publications often concentrate on narrow specific details, as is appropriate for focussed technical audiences. What is lost is the broad description of the design and its evolution, with consideration of the success and failure of specific features and lessons learned.

**Code** - We are interested in exploring the possibilities of publishing code to read. A truism among the programming community is that one learns to write good programs by reading good and bad programs. Sadly, there is little code available to read. The recent interest in public-domain code and open systems has increased the quantity of high-quality source code available. Many open questions in the publication of code remain to be explored. The conventional codex form, long accepted as appropriate for literary works and texts, may not be the right one for programs. Very few experiments have been made with this form, something that we hope to encourage. The audience for published code includes serious students of systems, including both the undergraduate and advanced levels, and practitioners involved with development, modification, and analysis of actual systems.

**Important technical reports** - many important technical reports, issued in small numbers by industrial organisations, research labs, or university departments, are not disseminated as widely as they merit. This is often because the originating organisation doesn't have the resources or the will to publish them more widely and because the material is deemed inappropriate by commercial publishers because of its narrow scope or limited size. Many technical reports are too large for journal publication and too small for conventional book publication. We hope to provide a means of publication and distribution of the best of these.

Authors will enter into a contractual arrangement with the USENIX Association. The Association is in the process of selecting a publisher to handle marketing and distribution. We hope that you will consider this arrangement a viable option for your next manuscript.

To submit a manuscript or proposal for consideration for the Monograph Series, send a copy to:

Marc D. Donner, Editor USENIX Association 2560 Ninth Street, Suite 215 Berkeley, CA 94710, USA

or send electronic mail to

books@usenix.org

## Computing Systems Publishes Compact Disk Supplement

The second issue *Computing Systems*, Volume 3:2, complete with a compact disk containing all the musical illustrations associated with the papers in this issue is now available. Original computer music is presented by Peter Langston and Mike Hawley, with accompanying papers which discuss the music. This issues also includes paper By Tim Thompson on on Keynote, a language and extensible graphic editor for music. Lastly, the issue includes a controversy discussion on portability by Stuart Feldman and W. Morven Gentleman. Subscriptions to Computing Systems are available at a special discount rate to EurOpen members for $25.00 USD for a one-year subscription (four issues). Single copies, including the CD, are available for $11.00. To place an order, contact the University of California Press, Journals Department, 2120 Berkeley Way, Berkeley, CA 94720, USA.

## Further Information on Conferences and Workshops

If you need further information regarding USENIX conferences or workshops, contact the USENIX Conference Office at

22672 Lambert Street
Suite 613
El Toro
CA 92630
USA

Email to judy@usenix.org or {uunet,ucbvax}!usenix!judy
Tel: +1 714 588 8649
FAX: +1 714 588 9706

## Further Information about the USENIX Association

If you would like information on membership, or would like information on ordering USENIX publications (proceedings, manuals, Computing Systems, the Monograph Series, or the Association's newsletter, ;login:, please contact the USENIX Association Executive Office at

2560 Ninth Street
Suite 215
Berkeley
CA 94710
USA

Email to office@usenix.org
Tel: +1 415 528 8649
FAX: +1 415 548 5738

# CALL FOR PAPERS: SUMMER 1991 USENIX CONFERENCE

Opryland Hotel, Nashville, TN, USA
10-14 June 1991

## Multimedia – Interfaces Now and the Future

Systems designers and developers must face the challenge of how to support and deliver the new types of interfaces --voice, video, animated graphics, touch and music -- that users are demanding now. While adding immeasurably to the power of the computer for communication and expression of ideas, the new media multiplies system complexity and magnifies the challenge of providing easy access to computer resources.

What are the technical engineering requirements to enable your operating systems to process effectively the new types of data? What must you, and your organisation, do to prepare?

Moreover, what of the users of multimedia? What are the requirements and the particulars of the approaches they take. And how can we use new multimedia interfaces to improve information handling? These are some of the questions tackled by presenters and attendees at the USENIX Summer 1991 Conference.

### Formats for Presentations

The USENIX Summer 1991 Conference will provide a variety of forums in which attendees can explore multimedia issues, as well as more general operating system and environment questions.

We invite submissions of your papers, describing new and exciting work, multimedia-related or not, for the technical track. Please target a sophisticated technical audience, particularly knowlegeable of operating systems issues but keenly interested in new, interesting projects in many areas. Suggested topics include, though are not limited to:

- Multimedia

— applications and research
— systems integrating voice, video, audio, touch or music
— data compression technology
— user interface/human factors

- Hypermedia

— authoring systems
— hypermedia/multimedia documents
— Operating systems issues
— multiprocessor systems
— distributed systems
— secure systems
— fault tolerant systems
— systems for novel architectures
— distributed file systems

- Communications and Networking

— protocols
— performance
— administration and security

- Programming Environment

— user interfaces, windowing, graphics
— compilers and language technology
— software development and other support tools
— testing and debugging

- Sophisticated Applications

— databases
— transaction processing
— instructional
— scientific, biological, medical, etc.

## Form of Submissions

Submissions to the technical track should represent new work and be in the form of an abstract and outline. Your submission should be complete enough to provide details of the approach and give confidence in the final paper. Full papers will be accepted as well. A submission should be from 3-5 pages and include:

1. Author name(s), postal addresses, telephone numbers and e-mail addresses.
2. Abstract: 100-300 words.
3. Outline: 2-5 pages giving enough details of the approach or algorithms to allow the committee to understand and judge the submission.
4. References and citations to relevant literature. Please show you are aware of previous work (and not reinventing the wheel).
5. Time needed for presentation. Slots are usually 30 minutes but adjustment can be made when in-depth background or audio-visual support is desirable.
6. Audio-visual presentation requirements. We are happy to provide assistance and equipment in making you presentation as audio and visually appealing as possible.

Authors whose submissions are accepted will receive instructions for the preparation of final papers, to be published in the conference proceedings. We are looking into possibilities for making audio and video materials as well.

## Relevant Dates

| | |
|---|---|
| Abstracts and outlines due: | **6 February 1991** |
| Notifications to authors: | **1 March 1991** |
| Final papers due: | **19 April 1991** |

## Address

Please submit one hard copy and one electronic copy to the address below:

Deborah Scherrer
Nashville USENIX Technical Program
mt xinu
2560 Ninth Street
Berkeley, CA 94710, USA

| | |
|---|---|
| Internet: | nashville@usenix.org |
| UUCP: | uunet!usenix.org!nashville |
| Telephone: | +1 415 644-0146 |
| FAX: | +1 415 644-2680 |

Be sure to include your postal and electronic mail addresses in all correspondence.

# USENIX Association 1991 Winter Conference

Grand Kempinski Hotel, Dallas, Texas, USA
21-25 January 1991

## TUTORIALS

| Monday, 21 January | |
|---|---|
| An Introduction to the TCP/IP Protocol Suite | Richard Stevens, Consultant |
| An Introduction To C++ | Robert Murray, AT&T Bell Laboratories |
| UNIX System V Release 4.0 Internals I - File, VM and Process Subsystems | Steve Buroff, AT&T; Michael Scheer, ProLogic Corporation. |
| Programming The X Window System, Version 11 | Oliver Jones, Saber Software, Inc. |
| An Introduction to 4.3/4.4BSD Internals | Thomas W. Doeppner Jr., Brown University |
| The Mach Overview | Avadis Tevanian, Jr., NeXT, Inc. |
| Introduction to the Internals of the GNU C Compiler (GCC) | Richard M. Stallman, GNU Project |
| UNIX on Modern Architectures | Curt F. Schimmel, Amdahl Key Computer Labs |
| An Introduction to Object-Oriented Programming | Dave Taenzer, US West Advanced Technologies |
| UNIX Technologies of Japan | Jun Murai , Keio University Hiromichi Kogure, UNIX System Laboratories Pacific, Ltd. |
| Network Security | Dan Geer, Digital Equipment Corporation; Jon A. Rochlis & Jeffrey I. Schiller, MIT |

| Tuesday, 22 January | |
|---|---|
| UNIX Network Programming | Richard Stevens, Consultant |
| Using C++ Effectively | Andrew Koenig, AT&T Bell Labs |
| UNIX System V Release 4.0 Internals II Session & Streams and Subsystems and Code | Steve Buroff, AT&T; Mike Scheer, ProLogic Corporation. |
| Introduction to Programming With the X Toolkit Intrinsics | Paul Kimball & Chuck Price, Digital Equipment Corporation |
| New Kernel Facilities in 4.3BSD-Reno | Marshall Kirk McKusick and Michael J. Karels, U. C. Berkeley |
| Mach Virtual Memory Internals | Nawaf Bitar, Hewlett-Packard Company |
| Advanced Topics in Systems Administration | Evi Nemeth, University of Colorado Rob Kolstad, Sun Microsystems |
| Programming in PERL | Instructor: Tom Christiansen, CONVEX Computer Corporation |
| Parallel Programming and Scalable Software | Stephen C. Johnson, nCUBE |
| Network Computing System and Architecture: Overview and Tutorial in Writing Distributed Applications | Nathaniel Mishkin & Paul J. Leach, Hewlett Packard; Richard Mackey, Open Software Foundation |

| Tuesday, 22 January — Half day | |
|---|---|
| C++ Programming Style | Tom Cargill, Consultant |
| C++ Tactics | Robert Murray, AT&T Bell Laboratories |

# USENIX Association 1991 Winter Conference

Grand Kempinski Hotel, Dallas, Texas, USA
21-25 January 1991

## PRELIMINARY TECHNICAL PROGRAM

Wednesday, Thursday and Friday

| Wednesday, 23 January | |
|---|---|
| 9:00 - 10:30 | OPENING REMARKS AND ANNOUNCEMENTS<br>*Lori S. Grob, Chorus Systemes* |
| | Keynote Address: Eben Ostby, Pixar<br><br>Eben Ostby joined the Pixar Animation Research and Development Group (then the Lucasfilm Computer Graphics Project) in 1983. With a background in computer science in design, Mr. Ostby has designed and implemented animation and modelling systems for three-dimensional computer graphics. He has also worked on a number of films. He was Technical Director on Knickknack, Tin Toy and Red's Dream, and a technical contributor to Luxo jr, Young Sherlock Holmes, Flags and Waves and The Adventures of Andre and Wally B. He produced and directed the film Beach Chair, a computer animated mini-travelogue, which is considered a classic in its genre. His current research areas include the procedural generation of plaids. |
| 11:00 - 12:30 | KERNELS 1<br>*Chair: Barry Gleeson* |
| | Processors, Priority and Policy: Mach Scheduling for New Environments: David L. Black, Carnegie Mellon University<br><br>A 2nd Generation Kernelised Unix: Marc Guillemont, Jim Lipkis, Doug Orr, Marc Rozier, Chorus Systemes<br><br>Partitioned Multiprocessors and the Coexistence of Heterogeneous Operating System Environments: Nick Vasilatos, Concurrent Computer Corporation |
| 11:00 - 12:00 | INVITED TALK: TOOLKIT GRAPHICS<br>*Doug Blewett, AT&T Bell Laboratories* |
| 12:30 - 2:00 | LUNCH |

| | |
|---|---|
| 2:00 - 3:30 | FILE SYSTEM PERFORMANCE<br>*Chair: Trent Hein* |
| | Extent-like Performance from a Unix File System: Larry McVoy and Steve Kleiman, Sun Microsystems<br><br>Smart Filesystems: C. Staelin and H. Garcia-Molina, Princeton University<br><br>Lessons Learned Tuning the 4.3BSD Reno Implementation of the NFS Protocol: Rick Macklem, University of Guelph |
| 2:00 - 3:00 | INVITED TALK: TROFF MACRO PROGRAMMING<br>*Sharon Murrel, AT&T Bell Laboratories. Jaap Akkerhuis, mt Xinu* |
| 4:00 - 5:30 | THREADS & NETWORKS<br>*Chair: Deborah Scherrer* |
| | Sun OS Multi-thread Architecture: S.R. Kleiman, M.L. Powell, S. Barton, D. Shah, D. Stein and M. Weeks, Sun Microsystems<br><br>Bringing the C Libraries With Us Into A Multithreaded Future: Michael B. Jones, Carnegie Mellon University<br><br>A Tree-Based Packet Routing Tabe for Berkeley UNIX: Keith Sklower, CSRG, University of Ca - Berkeley |
| 4:00 - 5:30 | INVITED TALK: UNIX SECURITY TODAY AND TOMORROW PANEL<br>*Pat Bahn, organiser, GTE Government Systems*<br>*Bill Cheswick, moderator, AT&T Bell Laboratories* |

| Thursday, 24 January | |
|---|---|
| 9:00 - 10:30 | INTERFACE TOOLS<br>*Chair: Tom Duff* |
| | An X11 Toolkit Based on the Tcl Language: John K. Ousterhout, University of Ca - Berkeley<br><br>User Interface Construction Based On Parallel and Sequential Execution Specification: Toshiyuki Masui, Carnegie Mellon University<br><br>$HOME MOVIE - A Home Movie System for Producing Demos on a Sun: Stephen A. Uhler, Bellcore Computer Systems Research Division |
| 9:00 - 10:30 | INVITED TALK: SYSTEMS ADMINISTRATION FORUM - Part 1<br>*Rob Kolstad, Sun Microsystems* |
| 11:00 - 12:30 | AWK PAPER AND KERNEL PANEL |
| | Awk As A Major Systems Programming Language: Henry Spencer, University of Toronto<br><br>*Panel* — Kernel Directions (1 Hour)<br><br>Rick Rashid, Carnegie Mellon Univ., Mike Powell, Sun Microsystems, Michael Karels, University of Ca- Berkeley, Michel Gien, Chorus Systemes, Moderator TBA |
| 11:00 - 12:30 | INVITED TALK: SYSTEMS ADMINISTRATION FORUM - Part 2<br>*Rob Kolstad, Sun Microsystems* |
| 12:30 - 2:00 | LUNCH |
| 2:00 - 3:30 | PROGRAMMING TOOLS<br>*Chair: Marc Donner* |
| | Program Loading in OSF/1: Harminder G. Singh, Larry W. Allen, Kevin G. Wallace and Melanie B. Weaver, Open Software Foundation<br><br>Compiling from Saved State: Fast Incremental Compilation with Traditional Unix Compilers: Alastair Fyfe, Ivan Soleimanipour and Vijay Tatkar, Sun Microsystems<br><br>A New Hash Package for UNIX: Margo Seltzer and Ozan Yigit, University of Ca - Berkeley |
| 2:00 - 3:00 | INVITED TALK: USING DISTRIBUTED OBJECTS<br>*Vinny Cahill, University of Dublin* |

| | |
|---|---|
| 4:00 - 5:30 | FILE SYSTEMS<br>*Chair: Steve Bourne* |
| | Evolutionary Path to Network Storage Management: Antony W. Foster, Robert K. Israel, Arun Taylor, Tracy M. Taylor, Neil Webber, Epoch Systems<br><br>A Highly Available Network File Server: Anupam Bhide and Stephen P. Morgan, IBM Research; Elmootazbellah N. Elnozahy, Rice University<br><br>The OSF/1 Unix Filesystem (UFS): Susan LoVerso, Noemi Paciorek and Alan Langerman, Encore Computer Corp; George Feinberg, Open Software Foundation |
| 4:00 - 5:30 | WORK IN PROGRESS SESSION<br>*Chair: Lisa Bloch* |

| **Friday, 25 January** | |
|---|---|
| 9:00 - 10:30 | OBJECTS IN ACTION<br>*Chair: Michel Gien* |
| | Advancing Files to Attributed Software Objects: Andreas Lampen, Technische Universitat Berlin<br><br>Organising Tools in a Uniform Environment Framework: Axel Mahler, Technische Universitat Berlin<br><br>The Process File System and Process Model in UNIX System V: Roger Faulkner and Ron Gomes, Sun Microsystems |
| 11:00 - 12:30 | INSECURITY<br>*Chair: Michael Karels* |
| | Limitations of the Kerberos Authentication System: Michael Merritt and Steven Bellovin, AT&T Bell Laboratories<br><br>UNIX Password Encryption Considered Insecure: Philip Leong and Chris Tham, University of Sydney<br><br>An Authentication Mechanism for USENET: Matt Bishop, Dartmouth College |
| 11:00 - 12:00 | DISTRIBUTED FILE SYSTEMS PANEL |
| | Mike Kazar, Transarc; John Osterhaut, University of CA -Berkeley; Rafael Alonso, Princeton University; Brian Palowski, Sun Microsystems Moderator: Peter Honeyman IFS/ University of Michigan |
| 12:30 - 2:00 | LUNCH |

| | |
|---|---|
| 2:00 - 3:30 | KERNEL II<br>*Chair: Jan Edler* |
| | An Experimental Implementation of Draft POSIX Asynchronous I/O: A. Lester Buck and Robert A. Coyne, Jr., IBM Federal Sector Div.<br><br>The Parallelisation of UNIX System V Release 4.0: Mark Campbell, Richard Barton, Jim Browning, et.al., NCR Corporation<br><br>An Overview of the Integrity S2 NonStop-Ux Operating System: Peter Norwood, Tandem Computers |
| 2:00 - 3:00 | INVITED TALK: DEBUGGING X AND X TOOLKIT APPLICATIONS<br>*Paul E. Kimball, Digital Equipment Corporation* |
| 4:00 - 5:30 | DISTRIBUTED PROCESSING<br>*Chair: Max Meredith Vasilatos* |
| | DRUMS: A Distributed Statistical Server for STARS: Andy Bond and John H. Hine, Victoria University of Wellington<br><br>Experience Building a Process Migration Subsystem for UNIX: Dan Freedman, University of Calgary<br><br>A Modular Architecture for Distributed Transaction Processing: Michael Wayne Young, Dean Thompson and Elliot Jaffe, Transarc Corporation |

For additional USENIX Conference Information, or pre-registration materials contact:

USENIX CONFERENCE OFFICE
22672 Lambert St., Suite 613
El Toro, CA 92630, USA
Tel: +1 714 588-8649
FAX: +1 714 588-9706

# Calendar of UNIX Events

This is a combined calendar of planned conferences, workshops, or standards meetings related to the UNIX operating system. The information here is collected by those listed below after an idea by John S. Quarterman of Texas Internet Consulting. The information comes from the various conference organizers, ;login:, Communications of the ACM, CommUNIXations, and many others. We encourage others to reuse this information, but we ask for proper acknowledgment, for example by including this statement.

If you have a UNIX related event that you wish to publicise then contact either John S Quarterman at *jsq@tic.com*, Alain Williams at *addw@phcomp.co.uk*, Susanne W Smith at *sws@calvin.wa.com*, or Carolyn Carr at *carolyn@usenix.org*. giving brief details in the style that you see below.

Abbreviations:

| | |
|---|---|
| APP | Application Portability Profile |
| C | Conference or Center |
| CC | Computer Communication |
| G, MD | Gaithersburg, Maryland |
| GM | General Meeting |
| LISA | Large Installation System Administration |
| MHS | Message Handling Systems & Application Layer Communication Protocols |
| OSE | Open Systems Environment |
| S | Symposium |
| SEDMS | Symposium on Experiences with Distributed and Multiprocessor Systems |
| T | Tradeshow |
| U | UNIX |
| UG | User Group |
| W | Workshop |

| mon days | conference | location |
|---|---|---|
| | 1990 | |
| Dec | CHUUG GM | Olten, Switzerland |
| Dec 2-5 | SunUG CT | San Jose, USA |
| Dec 4-5 | JUS UNIX Fair '90 | Tokyo, Japan |
| Dec 4-7 | IETF | IAB, U. Colorado, Boulder, CO, USA |
| Dec 10-12 | Sinix C, Unix Asia '90 | Singapore |
| Dec 10-14 | JTAP | Washington, D.C., USA |
| Dec 10-14 | DECUS S | Las Vegas, NV, USA |
| Dec 13-16 | Sinix T, Unix Pavillion '90 | Singapore |
| Dec 17-19 | UKUUG C | Cambridge, UK |
| | 1991 | |
| Jan 7-11 | IEEE 1003 | New Orleans, LA, USA |
| Jan 16-17 | Multi-User C Show for Gov't UniForum | Ottawa, ON, Canada |
| Jan 16-18 | USENIX, Software Devel. Environments | Grand Kempinski, Dallas, TX, USA |
| Jan 21-24 | UniForum | Infomart, Dallas, TX, USA |
| Jan 21-25 | USENIX | Grand Kempinski, Dallas, TX, USA |

| mon days | conference | location |
|---|---|---|
| Feb 13-14 | DKUUG W Network -high level course | Denmark |
| Feb 18-22 | DECUS S | Ottawa, Canada |
| Feb 21 | DKUUG C CAD/CAM | Denmark |
| Feb 27 | IEEE CS SCC/SAB | San Francisco, CA, USA |
| Mar | IETF | IAB, Wash. U, St. Louis, MO, USA (tentative) |
| Mar 13-20 | CeBIT 91 | Hannover, Germany |
| Mar 21-22 | Distributed Systems - S | Atlanta, Georgia, USA |
| Mar 21 | DKUUG C Management Systems | Denmark |
| Mar 26-29 | AFUU C | CNET Paris La Defense, France |
| Apr | NCR Unix User Group C | San Antonio, Texas, USA |
| Apr 1-5 | Integrated Net. Man. S | IFIP, IEEE, Arlington, VA |
| Apr 10-12 | IEEE 1003, USENIX, Uniforum, EurOpen | Miami, FL, USA (Tentative) |
| Apr 15-19 | IEEE 1003 | Swissotel, Chicago, IL, USA |
| Apr 18 | DKUUG C GUI and multimedia | Denmark |
| Apr 22-26 | ISO/IEC JTC1 SC22 WG15 | Netherlands |
| Apr 22-26 | DECUS Muenchen Symposium | Hannover, West-Germany |
| May 6-10 | DECUS S | Atlanta, GA, USA |
| May 9 | APP/OSE Users Forum | NIST, G, MD, USA |
| May 15-17 | i2u convention | Milano, Italy |
| May 15-17 | Multi-User C Show | UniForum Canada, Toronto, ON, USA |
| May 20-24 | EurOpen | Tromso, Norway |
| 29 May | DKUUG C Comms and Network | Denmark |
| 30 May | DKUUG C Systems development and CASE | Denmark |
| May 29-Jun 2 | ENA C | Seattle, WA, USA |
| Jun/Jul | UKUUG C | Liverpool, UK |
| Jun 10-14 | USENIX | Opryland, Nashville, TN, USA |
| Jun 17-19 | Sun User Group | Atlanta, GA, USA |
| Jun 17-20 | INET '91 | Copenhagen, Denmark |
| Jun 20 | DKUUG C Standardisation | Denmark |
| Jul 8-12 | IEEE 1003 | Santa Clara, CA, USA |
| Aug 5-8 | Interex C | SanDiego, CA |
| Sept 10-12 | European Sun User Group CT | NEC, Birmingham, UK |
| Sept 16-20 | EurOpen | Budapest, Hungary |
| Sept 24-27 | AUUG CT | Darling Harbour, Sydney, Australia |
| Oct | ISO/IEC JTC1 SC22 WG15 | Sweden |
| Oct 10-11 | Multi-User C Show | UniForum Canada, Montreal, Quebec |
| Oct 21-25 | IEEE 1003 | Southern Europe (location tentative) |
| Oct 30 | IEEE CS SCC/SAB | Nashville, TN, USA |
| Nov 14 | APP/OSE Users Forum | NIST, G, MD, USA |
| Dec | UKUUG C | Edinburgh, UK |
| Dec 9-11 | Sun User Group | San Jose, CA, USA |
| Dec 9-13 | DECUS S | Anaheim, CA, USA |
| | 1992 | |
| Jan 13-17 | IEEE 1003 | Orlando, FL, USA (location tentative) |
| Jan 20-24 | USENIX | Hilton Square, San Francisco, CA, USA |
| Jan 20-23 | UniForum | Moscone Center, San Francisco, CA, USA |
| Spring | EurOpen | Jersey, UK |

| mon days | conference | location |
|---|---|---|
| Mar 11-18 | CeBIT 92 | Hannover, Germany |
| Apr 20-24 | IEEE 1003 | Montreal, PQ, Canada (location tentative) |
| May 4-8 | DECUS S | Atlanta, GA, USA |
| Jun 8-12 | USENIX | Marriott, San Antonio, TX, USA |
| Jun 22-24 | Sun Users Group | Washington (DC), USA |
| Jul 13-17 | IEEE 1003 | Alaska, USA (location tentative) |
| Sept 8-11 | AUUG C T | World Congress Centre, Melbourne, Australia |
| Autumn | EurOpen/UniForum | Amsterdam, Netherlands |
| Oct 19-23 | IEEE 1003 | Scottsdale, AZ, USA (location tentative) |
| | 1993 | |
| Jan 25-29 | USENIX | Town & Country, San Diego, CA, USA |
| Mar 15-18 | UniForum | Moscone Center, San Francisco, CA, USA |
| Mar 24-31 | CeBIT 93 | Hannover, Germany |
| Jun 21-25 | USENIX | Cincinnati, OH, USA |
| | 1994 | |
| Jan 17-21 | USENIX | Hilton, San Francisco, CA, USA |
| Feb 7-10 | UniForum | Dallas Convention Center, Dallas, TX, USA |
| Mar 16-23 | CeBIT 94 | Hannover, Germany |
| Jun 6-10 | USENIX | Boston, MA, USA |
| | 1995 | |
| Mar 6-9 | UniForum | Dallas Convention Center, Dallas, TX, USA |
| | 1996 | |
| Mar 11-14 | UniForum | Moscone Center, San Francisco, CA, USA |
| | 1997 | |
| Mar 10-13 | UniForum | Moscone Center, San Francisco, CA, USA |

## Organising Bodies

NIST/NBS/POSIX
Roger Martin
National Institute of Standards
and Technology
Technology Building, Room B266
Gaithersburg, MD 20899, USA
+1-301-975-3295
+1-301-975-3295
rmartin@swe.icst.nbs.gov

IEEE Computer Society
P.O. Box 80452
Worldway Postal Center
Los Angeles, CA 90080, USA
+1-202-371-0101

UniForum (was /usr/group)
2901 Tasman Drive
Suite 201
Santa Clara CA 95054, USA
+1 408 986 8840
+1 408 986 1645 fax

/usr/group/cdn
241 Gamma St.
Etobicoke, Ontario M8W 4G7
Canada
+1-416-259-8122

Tracy MacIntyre
Exhibition Manager
EMAP International Exhibitions Ltd.
Abbot's Court
34 Farringdon Lane
London EC1R 3AU
United Kingdom
+44-1-404-4844

AUUG Inc
P.O. Box 366
Kensington
N.S.W. 2033
Australia
uunet!munnari!auug
auug@munnari.oz.au
tel: +61 2 361 5994
fax: +61 2 332 4066

AMIX, c/o IPA
P.O. Box 919
Ramat-Gan
Israel, 52109
+972-3-715770
+972-3-715772
amix@bimacs.bitnet
amix@bimacs.biu.ac.il

Japan UNIX Society (JUS)
#505 Towa-Hanzomon Corp. Bldg.
2-12 Hayabusa-cho
Chiyoda-ku, Tokyo 102
Japan
bod%jus.junet@uunet.uu.net
+81-3-234-5058

UNIX Fair '88 Association
1-1-1 Hirakawa-chu,
Chiyoda-ku, Tokyo 102
Japan

Singapore Unix Association - Sinix
20 Bideford Road #11-05
Wellington Building
Singapore 0922
+65 734 3256

DECUS U.S. Chapter
219 Boston Post Road, BP02
Marlboro, Massachusetts 01752-1850
USA
+1-617-480-3418

DECUS Europe
1-3, chemin Anneville, Box 176
CH-1213 Petit-Lancy 1
Switzerland
tel: +41 - 22 - 709 42 64
fax: +41 - 22 - 792 25 03

DECUS Munich (for Germany, Austria, Hungary):
DECUS Muenchen e.V.
Freischuetzstr. 91
D-8000 Muenchen 81
Germany
tel: +49 - 89 - 95 91 - 44 30

USENIX Association Office
2560 Ninth St., Suite 215
Berkeley, CA 94710
USA
+1 415 528 8649
office@usenix.uucp

National Expositions Co., Inc. (UNIX EXPO)
15 West 39th Street
New York, NY 10018
USA
+1-212-391-9111
fax: +1-212-819-0755

Sun UK User Group
Sue Crozier
Sun Microsystems, UK
+44 276 20980

Sun User Group the Netherlands (SUG-NL)
Clementine Voest
Sun Microsystems Nederland
+31 33 501234

Sun User Group, Inc.
Peter H. Salus
PO Box 167
Cambridge, MA 02142
USA
+1 617 739-0202
peter@uunet.uu.net

USING
P.O. Box 1077
Lisle,
Illinois 60532, USA

UniForum NZ Secretariat
PO Box 585
Hamilton
New Zealand

TRUUG
Esref ADALI,
Professor of Control and Computer Engineering,
TRUUG UNIX '90 Chairman,
Istanbul Technical University
Ayazaga
Istanbul
Turkey
+90 1 176 3586

EurOpen National group addresses can be found on the back cover of this newsletter.

Here is a list of acronyms that you might find useful:

| | |
|---|---|
| ACE | Advanced Computing Environments |
| ACM | Association for Computing Machinery |
| AFUU | The Association Française des Utilisateurs d'UNIX |
| AUUG | The Australian UNIX systems Users Group |
| DECUS | The Digital Equipment Computer Users Society |
| EurOpen | The European Forum for Open Systems, (was EUUG) |
| FNUG | Federation of NCR User Groups |
| GUUG | The German UNIX Systems User Group |
| IEEE | Institute of Electrical and Electronics Engineers |
| IETF | Internet Engineering Task Force |
| Interex | The International Association of Hewlett-Packard Computer Users |
| JUS | Japan UNIX Society |
| MCNTI | Moscow International Center of Science and Technical Information |
| NCR UUG | NCR UNIX User Group, Inc. |
| NIST | The National Institute of Standards and Technology |
| NLUUG | The Netherlands UNIX Users Group |
| NSF | National Science Foundation |
| SAB | Standards Activities Board |
| SERC | NSF/Purdue/Forida Software Engineering Research Center |
| SUUG | Soviet UNIX Users' Group |
| Sinix | The Singapore UNIX Association |
| UKUUG | The United Kingdom Unix systems Users' Group |
| USENIX | The Professional and Technical UNIX Association |
| UniForum | The International Association of UNIX Systems Users |

# Windows Column

*Sarah Keefe*
*sarah@x.co.uk*

*IXI Limited*

Sarah Keefe is marketing manager at IXI Limited, a software house specialising in X software, training and consultancy. With over 5 years experience in the computer industry, she has worked at Torch Computers and was involved in Torch's UNIX-based graphical user interface, OpenTop.

## The Lazy Man's Route to X

I have taken over the role of Windows columnist and would very much welcome ideas and articles from readers with views on broad ranging aspects of windowing technology.

## Introduction

The emergence of the X Window System from MIT and its acceptance as the *de facto* standard for controlling windows and graphics on workstations, PCs and intelligent terminals has resulted in it receiving significant attention from UNIX hardware vendors. But how easy is it for software authors to pick up and run with it?

## X Features

X is the outcome of a project initiated by MIT to investigate the educational application of a large network of workstations in a multi-vendor environment . The portability of applications was crucial so to meet this aim, X was developed with the following features:-

- It operates on any bitmapped screen architecture (not just character based)
- Applications are device independent
- The system is network transparent
- Multiple applications can operate concurrently on one screen
- Supports overlapping and obscured windows
- Applications can use many windows

Since X was the result of a collaborative project between MIT and leading UNIX vendors, it is deemed that no vendor has a lock on the technology. In fact, X is publicly accessible software and is attractive to hardware vendors because it not only meets industry needs for an open windowing system but hardware vendors don't have to pay royalties to MIT when bundling the software. Consequently, all major UNIX vendors now bundle X with their offerings.

Software developers always try to widen the potential market for their software by selecting a portable and popular operating system such as UNIX but until the arrival of X, there were various proprietary windowing interfaces which made the transfer of software from one system to another difficult and expensive. X is now sweeping away that fragmentation, leading to a huge opportunity for software developers.

The down side of all this is that 90-95% of applications in existence today are written to run for character terminals. Software houses have developed expertise in designing and implementing character based software; many of

these companies would like to jump on the X bandwagon, but have discovered the following difficulties:

- X expertise is scarce and therefore expensive to recruit.
- Writing X applications from scratch is difficult and time consuming. Programmers have to learn a new paradigm of programming, consequently it takes approximately 6-9 months before an programmer is fluent in X.
- Redesigning an existing application to use X is very difficult because it has to change to being event driven.
- Many programs cannot be modified to use X because they are written in languages which do not have an X library (such as COBOL or FORTRAN) or which run on operating systems such as PICK which do not support X directly.

Users are still locked in to character-based computing, and enormous numbers of applications continue to be developed for alphanumeric terminals. By the end of 1993, there will still be an installed base of around eight times as many alphanumeric terminals as X displays. Software houses will need to support character terminals as well as X displays so applications should be enhanced in ways which will not only improve their functionality on character terminals but also for users of X displays.

## Terminal Emulation

One of the ways in which these problems will be eased is to make it easier for character based applications to run well on X displays. Xterm is a standard X client which provides VT100 and Tektronix emulation to allow character applications (including the Unix shell) to run in a window. This gives users the ability to have multiple applications running and to cut and paste information between them. This in itself is a big step forward, but the applications themselves will only run in exactly the same way as they do on a VT100 screen. The applications are no easier to use. What is really needed is a way for these applications to use the facilities of a graphical user interface (scrollbars, pull-down menus, buttons, mouse clicks and drags and the like) without having to become X applications.

Another approach is to use an intelligent terminal emulator that gives character-based applications many Graphical User Interface (GUI related features but without the need to learn X programming. The terminal emulator should provide multiple fonts (handling proportional fonts correctly), colour, line drawing and shading characters. All of the usual cursor motion, line and character insertion and deletion and scrolling area functions should also be included.

The terminal emulator interprets character escape sequences coming from the application and converts them into X protocol messages the X display can recognise. User input from button presses and menu selections is sent back to the application as escape sequences. Some of these sequences may be replaced by programmable text strings where that is more convenient for an application.

Applications can also set up fields of the display which are sensitive to the mouse, use menus, buttons and scrollbars and display bitmaps. Cut and paste may be handled automatically, or the application can control the X selection mechanism directly. The use of the mouse is configured in terms of clicks and drags which are to be reported back to the application. The terminal emulator filters the actual X events and throws away sequences which are not interesting to the application. User interaction may be enhanced by using the pop-up dialogs which are built into the terminal emulator.

The advantages of the terminal emulator approach include:

The application can still work with character terminals. This reduces installation and support problems in mixed installations. Since the (GUI is contained entirely within the emulator, changing from an OSF/Motif look and feel to Microsoft Windows 3.0 or Open Look can be made by switching in another emulator, with no changes to the application.

Obviously, different character-based applications will be sending unique escape sequence messages to the terminal emulator so a way of overcoming this had to be found. The answer is either to change the application's libraries so that correct escape sequence calls are made to the intelligent terminal emulator or build a higher level language that sits between the application and the terminal emulator and filters out any inappropriate sequences.

The first approach has been taken by IXI's X.deskterm technology and has already been used by Uniplex to produce an X version of their market leading office automation software. Since the Uniplex suite of software comprises a total of 30 applications, it was far quicker for Uniplex to change application libraries common to all 30 applications than change each application in turn. Uniplex estimated that it would take 40 man years to move over to X, instead it took approximately 2 man years, enabling them to bring their product to market ahead of their competition. X.deskterm allows a phased migration to X. The product can change slowly, meaning that existing users are not confused by a seemingly new product.

The high level language approach has been used by two products:

- Deskterm Soft Option
  This is from IXI. This contains technology licensed from Cambridge Connectivity.
- ALEX
  This is from System Strategies.

Both products can be described as interface builders and are most suited to cases where the developer does not have access to source code. Deskterm Soft Option uses script files, written in an English like language, to build the user interface – no knowledge of Motif, X or C is required. ALEX on the otherhand is a C based language that is in many ways comparable to UIL and incorporates X concepts like widgets.

There are many variables to consider before deciding which lazy route to take to X. These include the nature and size of application, existing skill level, time constraints and access to source code to name a few. These products could play an important role in the growth of X software over the next few years, enabling users to reap the productivity benefits that a (GUI offers.

For those who are interested, addresses of the companies mentioned are:

| | | |
|---|---|---|
| IXI Limited | System Strategies | Cambridge Connectivity |
| 62-74 Burleigh Street | Waterman's Yard | 197 Hight Street |
| Cambridge | 32a The Mall | Cottenham |
| CB1 1OJ | London W5 3JJ | Cambridge CB4 4RX |
| England | England | England |
| Tel: +44 223 462131 | Tel: +44 81 556 2307 | Tel: +44 954 51920 |

# USL Column

*Peter Newland*

*USL*

For furtherr information on this column please contact Gill Mogg on *gill@uel.uucp*. Gill is Market Communications Manager at USLE. The guest writer this issue is Peter Newland.

Peter Newland is a Project Manager for UNIX System Labs Europe based in London. He is responsible for the USL development work carried out in Europe much of which concentrates in internationalisation and localisation of UNIX System V. In addition, Mr Newland represents UNIX International on the X/Open Internationalisation Working Group and participates in the UniForum technical sub-committee for Internationalisation.

## UNIX System V Release 4 Supporting an International Market

Peter Newland discusses the internationalisation features of UNIX System V Release 4 and looks particularly at the UNIX System V Release 4 European Language Supplement and the UNIX System V Release 4 Multi-National Language Supplement USL handles internationalisation development primarily through USLE and USL Pacific (USLP). The localisation work for Europe has been undertaken by USLE.

With the advent of UNIX System V Release 4 (SVR4) significant steps have been made towards providing a single global Operating System - that is, an operating system capable of simultaneously supporting multiple users each interfacing with the system in their own language. Here we explain the Native Language capabilities offered by SVR4 and by a number of additional SVR4 products that are now available.

As we shall see these products provide extensive support for European languages as well as for the Far Eastern markets such as Japan.

### Introduction

Much has been said over the years about the portability of the UNIX Operating System. This originally concentrated on source code portability of the Operating System itself and of UNIX applications. Today we are looking for binary application portability - not only between different machines but between different users each working in their own language.

To be cost effective in today's global markets, computer manufacturers and application developers need to be able to develop a single product with worldwide portability. With SVR4 and the associated localisation products discussed in this article this objective is now possible.

### How does the computer know how to speak my language?

The UNIX Operating System and the applications that it supports need to interface with the users in their own language , using the local conventions with which they are familiar, and working with an appropriate character set. In addition, this must be offered on a per user basis so that a single binary application or Operating System utility will concurrently support multiple users working in different locales. The 'locale' is the combination of the language, code set and local conventions used.

For an application to support multiple 'locales' it must be internationalised - that is, any assumption about the language, code set or local conventions

that will be used at run time must be removed, The language specific information that will be required at run time must be stored separately from the application. In this way, each time the application is invoked, the user's environment is examined and the appropriate 'locale' information is used.

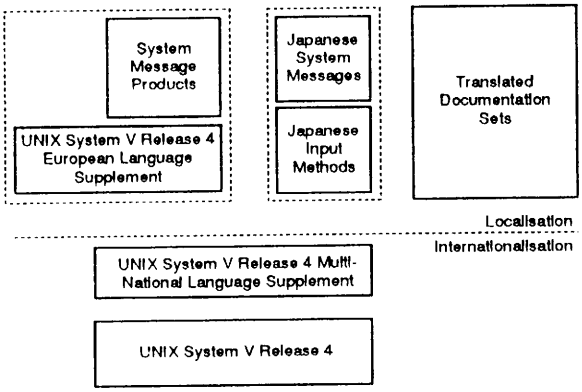


Figure 1. UNIX System V Release 4 International Product Set.

Much of this locale dependent information will be supplied by the Operating System and accessed through standard interfaces (see the SVR4 European Language Supplement). Other information such as the program messages are specific to the application. These messages can be stored in message catalogues and retrieved through the message retrieval interfaces as defined by X/Open( in the X/Open Portability Guide Issue 3.

Since the application is independent of the language, code set and local conventions used at run time, the application developer requires no specialist knowledge of the target languages that the application will eventually use. Only when the software development is complete and the application is 'localised' - that is, when the message catalogues are translated, are the language skills required. Thus, the application development and the application localisation can occur in two separate stages.

Clearly, not all applications can be converted to have international applicability. There is for example no standard way of calculating taxation that is internationally portable. A few applications will remain niche national applications but the vast majority of applications will in future be developed once and once only for a global market.

## UNIX System V Release 4 International Products

Before we look in detail at the features that the SVR4 platform offers for the international marketplace, let's take a quick look at how these individual products fit together.

The products separate into two areas either supporting internationalisation or localisation. The internationalisation products - SVR4 and the SVR4 Multi-National Language Supplement (SVR4 MNLS) supply the features needed to develop applications for the international market. These products do not themselves offer any 'locale' information. It is the localisation products which contain the run time support for the application, providing the data needed for the application to function correctly in a particular locale.

This product separation allows considerable flexibility. The customer can select the localisation package that meets their particular requirements and new packages can be developed in a timely manner as and when there is an identified demand. The internationalisation products form a base for application development. Any internationalised applications can make use of any of the run time environments provided by the localisation products.

As figure 1 shows, the base for all these products is SVR4 which provides a number of important features for internationalisation. It supplies the development tools to build applications operating with single-byte code sets, which is sufficient for most European applications. The SVR4 European Language Supplement, ( SVR4 ELS) the primary localisation product for Europe, requires no additional internationalisation support.

The SVR4 MNLS provides further internationalisation support. Specifically this product provides support for translated system messages for a number of UNIX utilities and gives additional support for languages such as Japanese, which use multi-byte code sets. When the SVR4 Enhanced Security product becomes available early in 1991 it will include all the features currently offered through the SVR4 MNLS. All the localisation products such as the SVR4 Japanese System Messages, Japanese Input Methods, European System Messages Products or applications that exploit the MNLS multi-byte features will be portable across these two platforms.

## The UNIX System V Release 4 European Language Supplement

The UNIX System V Release 4 ELS is the fundamental localisation product for the European market. This product provides support for European code sets, locales, fonts, keyboards and printers - offering full support for European applications and users.

### Code Sets

The SVR4 ELS enables applications to work with a number of different code sets. This gives the user the flexibility to use many existing applications which assume a particular code set is in use. The user is therefore able to customise the environment to his own needs.

### Internal Code Sets

The internal code sets are those that an application can use when processing data.

Ten different internal code sets are available when using the SVR4 ELS, these are shown in table 1.

| The 10 internal code sets supported are: | |
|---|---|
| ISO 8859-1 | Latin alphabet No.1 |
| ISO 8859-2 | Latin alphabet No.2 |
| ISO 8859-3 | Latin alphabet No.3 |
| ISO 8859-4 | Latin alphabet No.4 |
| ISO 8859-5 | Latin/Cyrillic |
| ISO 8859-7 | Latin/Greek |
| PC 437 | United States |
| PC 850 | Multilingual |
| PC 860 | Portugal |
| PC 863 | Canada - French |
| PC 865 | Norway |

Table 1: SVR4 European Language Supplement Internal Code Sets

### External Code Sets

The external code sets are the code sets that can be used by terminals, printers etc. connected to the system.

The code set used 'externally' by the terminal or printer does not need to be the same as that used by the application. By using the SVR4 MNLS kbd STREAMS module together with the SVR4 ELS mapping tables, any of the code sets shown in table 2 can be used as the 'external' code set. This mapping can be done transparently in the I/O STREAM allowing the application to use one code set while the terminal or printer operates in another.

The code sets that can be used by terminals with the SVR4 ELS are:

- All the internal code sets
- ISO 646 National Variants for:
  - Denmark
  - France
  - Great Britain
  - Germany
  - Italy
  - Norway
  - Portugal
  - Spain
  - Sweden
  - Yugoslavia
- DEC Multilingual
- Roman 8
- EBCDIC

Table 2: SVR4 European Language Supplement, External Code Sets.

### Console Display

All the internal code sets mentioned in table 1 can be displayed on the integrated console for the Intel i386/i486.

### Code Set Mappings

Code set mapping is the ability to convert data from one code set representation to another.

The SVR4 ELS provides extensive support for mapping data between the various internal and external code sets.

| LANGUAGE | COUNTRY | CODE SET |
|---|---|---|
| Danish | Denmark | ISO 8859/1, PC 850, PC 865 |
| Dutch | Belgium | ISO 8859/1, PC 437, PC 850 |
| Dutch | Netherlands | ISO 8859/1, PC 437, PC 850 |
| English | Australia | ISO 8859/1, PC 437 |
| | | ISO 646, PC 850 |
| English | Canada | ISO 8859/1, PC 437 |
| | | ISO 646, PC 850 |
| English | UK | ISO 8859/1, PC 437 |
| | | ISO 646, PC 850 |
| English | US | ISO 8859/1, PC 437 |
| | | ISO 646, PC 850 |
| Finnish | Finland | ISO 8859/1, PC 437, PC 850 |
| French | Belgium | ISO 8859/1, PC 437, PC 850 |
| French | Canada | ISO 8859/1, PC 850, PC 863 |
| French | France | ISO 8859/1, PC 437, PC 850 |
| French | Switzerland | ISO 8859/1, PC 437, PC 850 |
| German | Austria | ISO 8859/1, PC 437, PC 850 |
| German | Germany | ISO 8859/1, PC 437, PC 850 |
| German | Switzerland | ISO 8859/1, PC 437, PC 850 |
| Greek | Greece | ISO 8859/7 |
| Icelandic | Iceland | ISO 8859/1, PC 850, |
| Italian | Italy | ISO 8859/1, PC 437, PC 850 |
| Italian | Switzerland | ISO 8859/1, PC 437, PC 850 |
| Norwegian | Norway | ISO 8859/1, PC 850, PC 865 |
| Polish | Poland | ISO 8859/2, |
| Portuguese | Portugal | ISO 8859/1, PC 850, PC 860 |
| Russian | USSR | ISO 8859/5 |
| Serbo-Croatian | Yugoslavia | ISO 8859/2 |
| Spanish | Spain | ISO 8859/1, PC 437, PC 850 |
| Swedish | Sweden | ISO 8859/1, PC 437, PC 850 |
| Turkish | Turkey | ISO 8859/3 |

Table 3: SVR4 European Language Supplement Locales Supported

## Locales

The SVR4 ELS provides 75 locales each representing different combinations of language, territory and code set. The locale includes information in the following areas:

— Character classification and conversion

— Numeric format

— Date and time format

— Collation order

— Monetary information

This allows applications that are written using the ANSI and X/Open interfaces provided in SVR4 to function correctly in the supported combinations of language, code set and country. The locales supplied by the SRV4 ELS shown in table 3, support 18 languages across 25 countries and 10 code sets.

## Postscript Printers

The SVR4 ELS provides extensive PostScript support, so that when working with any of the internal code sets (table 1) the data can output directly to a PostScript printer through an enhanced version of the SVR4 postprint(1) utility.

This PostScript Support is also designed to be compatible with most recent versions of troff.

## Keyboard Layouts

For the Intel i386/i486 implementation, a number of different keyboard layouts are supported. These are shown in table 4. In addition a utility that

generates a diagram of the current keyboard layout in PostScript is provided.

| |
|---|
| Belgian |
| Danish |
| Dutch |
| English |
| French |
| French-Canadian |
| French-Swiss |
| German |
| German-Swiss |
| Greek |
| Italian |
| Latin-American Spanish |
| Norwegian |
| Portuguese |
| Spanish |
| Swedish |

Table 4: Keyboard layouts supplied with the SVR4 ELS

### String Extraction Compiler

An automated string extraction facility is provided. This tool will allow you to recompile an existing source code application (pre ANSI C), without modification and extract strings into an X/Open conformant message catalogue, ready for translation.

### Internationalisation Features of SVR4

The SVR4 ELS described above is a localisation product, that is, it builds on the internationalisation features already provided in SVR4. These include the following:

- C compilation system conformant to the ANSI C Standard. ANSI C gives the definitions for a number of locale dependent functions, and specifies the wide character data type, wchar_t, (capable of storing any single or multi-byte character) and some related functions.
- SVR4 is conformant to the X/Open Portability Guide Issue 3 (XPG3) and has obtained a source code brand. XPG3 specifies additional functions to support internationally portable applications including the message catalogue interfaces.
- SVR4 saw the completion of the 8-bit clean up and the first steps towards multi-byte support in the utilities.
- SVR4 provides STREAMS based TTY sub-system supporting Extended UNIX Codes (EUC). This feature allows additional STREAMS modules to be placed on the I/O STREAMS to support particular localisation requirements. For example, the code set mappings needed to support some European terminals, or the processing required to support one of the Japanese Input Methods can be carried out on the STREAM in a way that is transparent to the application.
- Integrated console support for the INTEL i386/i486 implementation providing utilities and system support for downloadable fonts, and to change the keyboard layout.

### UNIX System V Release 4 Multi-National Language Supplement

While SVR4 already meets the requirements for a typical European application, additional features are needed for languages that use multi-byte code sets. SVR4 MNLS gives the additional features needed to develop applications that have worldwide portability supporting both single and multi-byte code sets. In addition SVR4 MNLS supplies enhancements to one hundred utilities both for multi-byte 'locales' and to support localised USL System Message products.

All the features that are currently offered in the SVR4 MNLS product will be provided in the base in the Enhanced Security Release.

### Multi-byte Character Support

The multi-byte character support provided by SVR4 MNLS is based on extended UNIX Codes (EUC) (See Figure 2.). It allows application developers to write new applications in a language independent manner and to 'internationalise' existing applications with minimal charges to program logic. Such applications will have worldwide portability supporting languages which require either single byte or multi-byte code sets.

The EUC multi-byte functions are contained in two libraries:

- The wide character library (libw), provides multi-byte analogues and extensions to the C library functions. This library includes functions which are wide character equivalents of the standard I/O functions, and wide character classification and conversion functions.

- The curses library (libcurses) has been enhanced to support multi-byte character handling. SVR4 MNLS provides enhancements and additions to existing functions in the library.

The functions in these two libraries are used in over one hundred of the UNIX utilities to provide support for multi-byte locales.

- The Extended UNIX Code (EUC) provides a template for supporting up to four code sets at the same time; the US ASCII code set and three additional code sets, each of which may be single or multi-byte. The particular code sets are determined by the locale of the application.
- EUC multi-byte characters are converted to a wide character data type, wchar_t, as defined by the ANSI C Standard, for character processing within applications.

Figure 2: Multi-byte Encoding Scheme - EUC

### STREAMS based Code Set Conversion and Character Mapping

The SVR4 MNLS provides a STREAMS code set and character mapping module. This can be used for code set mapping transparently in the input STREAM, or to provide compose and dead key functionality.

### Translated System Message Products

The SVR4 MNLS and the Enhanced Security Release will support translated system messages for over one hundred of the UNIX utilities. The first of these products will become available at the beginning of 1991.

### Translated Manuals

An important part of SVR4 is its associated documentation. Translation of this documentation into a number of European languages is now in progress. The German version of SVR4 documentation will be available by the end of 1990, and the corresponding French translations are due early in 1991 (see Figure 3). Later in 1991 we anticipate translations in additional languages to become available.

1. UNIX System V/386 Release 4 ANSI C Translation Guide
2. UNIX System V/386 Release 4 Migration Guide
3. UNIX System V/386 Release 4 Network User's and Administrator's Guide
4. UNIX System V Release 4 Programmer's Guide: ANSI C and Programming Support Tools
5. UNIX System V Release 4 Programmer's Guide: Networking Interfaces
6. UNIX System V Release 4 Programmer's Guide: POSIX Conformance
7. UNIX System V Release 4 Programmer's Reference Manual
8. UNIX System V Release 4 System Administrator's Reference Manual
9. UNIX System V Release 4 User's Guide
10. UNIX System v/386 Release 4 User's Reference Manual

For further information on translated documentation contact Prentice Hall on: 0442 231900

Figure 3: Translated documentation shortly to be available in French and German

# Report on ISO/IEC JTC1/SC22/WG15 (POSIX)

*Dominic Dunlop*
*domo@tsa.co.uk*

*The Standard Answer Ltd.*

Dominic Dunlop has been hanging around the European UNIX scene for years now, industriously researching into why things won't work and assiduously spreading the bad news among those who will listen. Involved with POSIX during its growth from a single, short document to a multiplicity of legalistic tomes, he is currently paid to report to EurOpen (formerly EUUG) and USENIX members on the progress of the ISO POSIX working group. He rushed this screed out so as to meet the EurOpen Newsletter deadline, and apologises for the resulting lack of polish...

## Meeting of 23rd – 26th October, 1990, Orcas Island, Washington, U.S.A.

### Introduction

As far as ISO is concerned, POSIX, like Gaul, is divided into three parts. Forget all those proliferating IEEE 1003 POSIX working groups (eighteen at of them at the last count), and just think of three standards: IS 9945–1 for a definition of the services offered by operating system; IS 9945–2 describing the shell and tools; and IS 9945–3, system administration.

The good news is that you can now buy the first edition of the first of these.[1] The bad news is that all three ISO standards projects are running into scheduling difficulties. And there is even more bad news if you are an Ada™ fan: in order to ease its difficulties, the ISO POSIX working group has put a serious road block between your favourite language and an international standard defining how you may use it to access POSIX services. Why did we do this, and why don't we feel bad about it? Read on...

### 9945–3 — System Administration

As you are probably aware, the IEEE P1003.7 working group on system administration has decided that current UNIX administrative tools and practices are sufficiently obsolete, inadequate and diverse that they are not worth standardising. Instead, the group has elected to define a new, object-based administration scheme which views a system as a collection of objects to be administered, and a network of systems simply as a larger collection of such objects.

Although this approach grafts neatly onto the network administration work which has been going on in the internet and Open Systems Interconnection (OSI) communities, it will be a while before it produces any results. As we shall see in connection with 9945–2, when ISO delegates responsibility for the development of a standard to another body, as it has done with the POSIX standards, it likes the documents to be in a relatively stable state before they are submitted for use as ISO Working Document (WDs). 1003.7 thinks that it will have something suitable for ISO

1. From the IEEE, which has agreed to print the world's first combined IEEE/ANSI/ISO standard — on ISO standard A4 paper. Ask for ANSI/IEEE Std. 1003.1:1990. It will cost you $52.50 if you are an IEEE member, $75.00 otherwise. Add $5.00 for surface mail to Europe. In the U.S.A., call (800) 678-IEEE; elsewhere, +1 908 981 1393. IEEE accepts major credit cards.

to start work on by 1992.

Unfortunately, ISO rules state that, unless a project that it has authorised has resulted in a WD within three years of authorisation, the authorisation stands in danger of automatic withdrawal. The only way out is for a national standards organisation participating in the development of the standard to call for a vote on project continuation before the time limit expires. The time limit for the production of a draft for 9945–3 has almost been reached, with no prospect of the deadline being met.

It seems inevitable that the twenty-four countries participating in the ISO POSIX project will be formally balloted as to whether they think that the authorisation to develop a system administration standard should stand, despite the missed deadline. This is not a particularly big deal: an examination of ISO's information technology standardisation work reveals that around twenty percent of projects miss one deadline or another. (OSI standards have a particularly poor track record.)

Nevertheless, it is embarrassing when the managerial finger is pointed at one's own project. Already, the special pleading has started; the SC22 Advisory Group, which makes recommendations on policy issues to the ISO Programming Languages subcommittee, has suggested that in general standards developed within SC22 are larger and more complex than most others, and the time limits given in JTC1 directives... will therefore often be too short'[1]. This may be true — although work elsewhere in ISO suggests that SC22 has no monopoly on large projects. However, it seems to me that the time limits given by the directives cannot reasonably be relaxed: if no visible progress has been made on a project after three years, those involved had better be given an opportunity to ask themselves why, and to consider whether they wish to continue giving their support. I am sure that, if it comes to a vote, the result will favour the continuation of the system administration project. Just don't hold your breath waiting for the final standard.

## 9945–2 — Shell and Tools

The shell and tools standard work is not crowding a deadline as closely as is system administration, but is not clear of trouble either. At least we have a CD (one step beyond a WD), corresponding to draft 9 of 1003.2, but have failed to move it forward to the next stage, a Draft International Standard (DIS). According to the directives, we have four years in which to do this before serious questions get asked, and the ISO directorate makes a decision about project termination. Although our progress to date has not been rapid, we have some time in hand.

Our first attempt to register the 1003.2 draft as a DIS failed. (See my report on WG15's Paris meeting[2]). The problem was that, while the technical content of a DIS is supposed to be essentially the same as that which will appear in the ultimate International Standard (IS), we all knew that the content of 1003.2 was still undergoing rapid and sometimes radical change. There was no way that draft 9 could been accepted as a DIS. (The U.S. National Institute for Standards and Technology (NIST) ultimately decided not to base a Federal Information Processing Standard (FIPS) on draft 9 for similar reasons.)

Draft 11 (or later) of 1003.2 will be passed to ISO in January, 1991 (or later), in the hope that it can be registered as a revised CD, and will stand more chance of clearing the the remaining hurdles which separate it from ISO status. Until this happens, we have a situation described by one normally restrained working group member as a pure disaster,. We are about to suggest that draft 6 of 1003.2A, the User Portability Extension, due early in 1991, is registered as a Proposed Draft Amendment (PDAM) to 9945–2, without having a stable document to amend![2]

In this situation, somebody may ask us why we don't just roll the amendment into the next, hopefully more stable, version of the CD. The practical answer to the question is that the IEEE is treating 1003.2 and 1003.2A as two separate documents, and we would prefer to do the same. How much weight such an argument might carry with the ISO secretariat is another question...

## 9945–1 — Operating System Interface

Now that 9945–1:1990 operating system interface definition is an international standard,

2. The UPE to 1003.2 describes interactive utilities for program development, supplementing 1003.2's description the non-interactive tools used in shell scripts.

international standards work on POSIX has reached the end of its beginning. What do we do next? The problem is that we are spoiled for choice. An embarrassing number of the 1003 projects represent extensions to, or restatements of, the services described in 9945–1:

1003.1A: A 1003.1 extension draft, covering covering tweaks such as symbolic links, will be ready for us early in 1991. We will probably vote to register it at our next meeting.

1003.1LI: (Provisional name.) This is the language-independent specification of the services defined by the current 1003.1 standard in terms of their C language interface. It may be ready in late 1991, provided that enough people can be found to work on it.

1003.1C (Provisional name.) Building on the definition provided by 1003.1LI, these C bindings will correspond exactly to the C interface defined by the current 1003.1. Again, a draft may be ready late in 1991.

1003.2: The shell and tools standard defines C language interfaces to regular expression handling, filename expansion, argument string parsing and more. Arguably, these belong in 9945–1. They are also candidates for language-independent specification.

1003.4: We have requested that the draft of 1003.4, real-time extensions to the portable operating system interface, is registered as a PDAM to 9945–1. The first international POSIX standard has only just hit the streets, and already we are trying to amend it!

1003.4A: The 1003.4 working group considers that draft 5 of its threads (lightweight process) standard will be ready for submission to ISO at the same time as 1003.4. As yet, we have made no decision to accept it.

1003.4B: This is simply a language-independent specification for the services described by 1003.4 in terms of their binding to the C language. The IEEE working group does not know when it will be ready, and we don't yet know when we will be ready to accept it. The two issues are connected: if we say we want work on it to be accelerated, it is likely to be ready more quickly.

1003.5: The Ada description of the portable operating systems interface is well on its way to becoming an ANSI/IEEE standard. Expect to see it in 1992. Sadly, for reasons explored below, 1003.5 is unsuitable as a basis for an ISO standard.

1003.6: The security extension to the operating system interface will be ready for us to have a look at in January of 1991, although it will be a while before it is mature enough for PDAM registration.

1003.8: Transparent file access, that is, transparent access by a process hosted on one system to files held by another, is making rapid progress after narrowing down its goals until it identified an achievable target. The IEEE working group expects to have a document suitable for ISO review by mid 1991.

1003.9: The FORTRAN[3] bindings to the operating system interface definition are written in a manner which is more to ISO's taste than the Ada description of the same services, and will be ready for ISO review in late 1990. However, we have elected not bring it forward to international standards status in the near future. Again, our reasons are explored below.

1003.16: This recently-authorised IEEE project aims to produce C language bindings to some future language-independent specification of the POSIX operating system interface. Like Ada and FORTRAN, it is tied up with the whole issue of language independence.

I wrote last time about the background to the language independence debate.[2] ISO strongly

3. Obscure style note: one is supposed to refer to the proposed 1990 version of the language as Fortran; to older versions as FORTRAN. 1003.9 is a binding to FORTRAN 77.

favours language independent service specifications, but very few people in the U.S. are interested in writing them. ISO has delegated development responsibility for POSIX to ANSI, which in turn has passed the buck to the IEEE — an organisation which ISO cannot officially talk to or aid. As a result, IEEE is saddled with a problem which it is ill-equipped to solve.

At our Paris meeting, we signaled our disappointment with the IEEE's progress towards a language independent specification for POSIX by refusing to register drafts of 1003.4, .5 and .9. The list above shows that we have relented on 1003.4, but have left .5 and .9 out in the cold.

The difference between this meeting and the last is that they are now definitively out in the cold, and will not be let into the ISO process until we are very close to having a language-independent version of IS 9945–1 for them to bind to. Here, with a few interpolations in square brackets, is the resolution that says why:

> Language Independent Specifications:
>
> *Whereas*, SC22 AG [the advisory group mentioned above in connection with 9945–2] has recommended that the production of language independent specifications and language bindings for POSIX be carried out in such a way that it does not delay the standardisation of the C language bindings [1] [Thank you. That's just what most of us wanted to hear]; and
>
> The production of a language independent specification corresponding to IS 9945-1:1990 and subsequent C language-based amendments, together with a C-language binding to that language independent specification, is required by the Division of Work Item JTC 1.22.21 [A Division of Work Item is an ISO mechanism for splitting an authorised project into several sub-projects]; and
>
> The production of further language bindings to the language independent specification corresponding to 9945-1:1990 as subsequently amended is ultimately desirable; and
>
> WG15 considers that thin, language bindings (which must be read in conjunction with a service definition) are suitable candidates for standardisation, but thick, bindings (those which incorporate a service definition duplicating and possibly conflicting with the service definition provided by another standard) are not [The terms thin, and thick, derive from the number of pages in the document in question. 1003.5 is a thick, binding, so we do not like it much; 1003.9 is a thin, binding, but to the C language specifications of the current 1003.1];
>
> *Therefore*, JTC1/SC22/WG15 requests the U.S. member body [ANSI, which in turn gets the IEEE to do the work] to provide a schedule for the delivery to WG15 and SC22 of 9945–1-related documents which is subject to the following constraints (listed in order of precedence, highest first):
>
> 1. The incorporation or development of language independence features shall not be on the critical path(s) for the production of C language-based documents;
> 2. The ultimate goal is the production of an extended [extended, that is, by 1003.4, 1003.6, 1003,8...], language independent 9945-1 and accompanying thin, binding to the C language at the earliest possible date;
> 3. Every attempt shall be made to observe JTC1/ISO rules on the bringing forward of amendments etc., with the need to seek waivers being highlighted if this appears necessary in order to satisfy the constraints above;
> 4. Language bindings, other than those for the C language, shall not be brought forward to WG15 or SC22 for any purpose other than review and comment before the language independent 9945-1 has been registered as a DIS; and
> 5. Where possible in the light of other constraints, C language-based documents shall include a informative annex setting out a language independent definition of the services defined by the normative body of the document;

> The schedule shall identify time frames for at least the following document circulation and registration milestones:
>
> 1. Thick, C bindings for amendments to 9945–1:1990;
> 2. Language independent specifications corresponding to 9945-1:1990 and subsequent amendments;
> 3. Thin, C bindings to language independent specifications corresponding to 9945–1:1990 and subsequent amendments;
> 4. A combined language independent 9945–1 and accompanying thin, C binding to the services that it defines; and
> 5. Thin, bindings for further languages to the whole of the combined language independent 9945–1, or to supersets or subsets of the services which it defines.

I hope that your eyes have not glazed over: public statements of policy get get convoluted and legalistic at this level, but all of this verbiage actually represents a concise description of the problem and what we see as a route to its solution.[4] For the first time, this tells the IEEE exactly what type of document ISO working group wants to see, and in which order:

a. C-based standards first.

b. Language independent standards with a corresponding thin, C binding second.

c. Thin, (and only thin) bindings to other languages no sooner than b.

d. Examples of language independent specifications (as opposed to definitive standards for them) any time that IEEE can manage to produce them.

e. All in accordance with ISO rules on the publication of amendments and revisions to standards (we hope).

There was some understandable objection from the U.S. to micro-management, — if we were defining so many goals, constraints and checkpoints, why didn't we just write the schedule ourselves? The answer is that there is still quite a lot of flexibility allowed: the IEEE has a dozen or more documents to bring forward, and it can decide on the ordering. And the dates. We just want to know when those dates are, and to disallow certain orderings.

The amount of resource that the IEEE can muster to work on language independent specifications determines when step b can occur. Does anybody want to volunteer to make it sooner than 1995?

The real victim of our newly-defined policy is Ada. It is clear that that there will be an ANSI/IEEE standard for an Ada definition of the POSIX operating system interface, probably within two years. It is now equally clear that, because it will be a thick, binding, this standard cannot move forward to gain international status. There may ultimately be a thin, Ada binding to a future language-independent 9945-1. It may or, more likely, may not define an interface identical to that defined by 1003.9. We could face the unpalatable prospect of an ISO standard which differs from the corresponding ANSI standard.

Why don't we feel too bad about this? Well, it seems that the main requirement for an Ada POSIX standard comes from within the U.S. 1003.5 will fill this need, and we should not seek to delay it. The need for an international standard in this area is less clear, but we have now given clear guidelines on the form that it should take, just as soon as anybody wants to develop it.

That said, it is clear that we still have much to learn about...

## Coordination.

One aim of the IEEE and ISO POSIX projects is that the international standards that result should be identical to the corresponding U.S. standards. Another is that ISO publication should not lag IEEE publication by too long. IS 9945–1 is a benchmark in both cases: by dint of the IEEE agreeing to print for both organisations, content is identical, and publication is simultaneous. This will be a hard act to follow, not least because there are thousands of pages of IEEE drafts in the pipeline, all of which must undergo international review before they can even start going through the three-stage ISO mill which grinds documents into international standards.

4. Although I could be biased: I drafted the resolution.

It has been the policy of the IEEE not to submit documents to ISO until they reach their first IEEE ballots — that is, until they are reasonably mature and complete. In view of our rejection of 1003.2 draft 9 because we did not consider it mature enough, this seems like a prudent approach. The trouble is that by the time the IEEE considers a document mature, it is likely also to be difficult to change in any significant manner. If we had earlier visibility into the subject matter and approach of the IEEE's work, we could comment on its future acceptability to ISO. For example, we could have suggested that 1003.5 pursued a thin„ rather than a thick„ binding.

To try and get out of the hole that we have dug for ourselves, we have requested early visibility„ of IEEE draft standards. Seeing standards when they are young and small will also aid international understanding of the larger, more mature versions when they appear.

## OSCRL

The POSIX project bears a growing similarity to an ancient yet still inhabited castle: some parts are old and crumbling; others require constant repair if they are to remain habitable; and, all the time, new walls, doors and towers are being added. 1003.7 even seems to be demolishing a few unsightly outbuildings. Coordination should ensure that nobody builds a wall where someone else wants a door. Or a whole new tower when all that was needed was a new entrance to an existing one, as happened in the case of 1003.5.

No castle is complete without a ghost, and POSIX has one: OSCRL — Operating System Command and Report Language. Started in the early eighties, it was (to simplify to an almost indictable extent) an attempt to define a common Job Control Language for the large computers of the day. It found a home in SC21, which looks after OSI, just before it became apparent that UNIX was going to become the open„ operating system of choice. Work on OSCRL, initially enthusiastic, tailed off. It was all but non-existent by the time the ISO POSIX project was set up. However, it is ISO policy when starting new projects to examine any related work which it may have undertaken, and the search turned up OSCRL as covering topics to be addressed by 9945–2 and 9945–3.

SC21 welcomed the chance to pass the project to another group, and we reluctantly agreed to take it over. Then the ISO central secretariat lost all the paperwork. (It is a fact of life that all bureaucracies lose paperwork.) We had an excuse to prolong OSCRL's spell among the undead, provided that we could put up with the periodic howls from its (few) proponents.

These howls recently resulted in a polite suggestion from the SC22 AG that we should do something. That something might be the massaging of the existing material (if it can be found) in to a Technical Report — a type of document which few people ever read, and the production of which is discouraged by ISO. But a TR may just be the sort of headstone that OSCRL lacks. We will be trying to nail the coffin lid down at our next meeting, which takes place in the Netherlands from 14th – 17th May, 1991.

# EurOpen Software Distribution

*Frank Kuiper*
*euug-tapes@EU.net*

*Centre for Mathematics and Computer Science*
*Amsterdam*

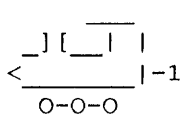

In our continuous effort of trying to offer as much as possible to the EurOpen members, yet again the announcement for new tapes. Yes, plural. There are no less than three new tapes, and one new version of an existing tape: EUUGD3 STARTER KIT, EUUGD21 "Nice conference", EUUGD22 PP5.0, and a revision of EUUGD16 Brussels - GNU.

As I write this, the STARTER KIT. tape is not yet ready, but it will be by the time you read this. Yep, here it is! Quite a number of moons ago, it was decided that a tape like this should be made available for the EurOpen community. I won't go into details of why this tape is issued much later than hoped for. Just say that the EurOpen is comprised of a number of people who do their EurOpen-work as a side job, which results in some ''amateuristic''-like side effects.

This said, let's now concentrate on what will be on the STARTER KIT distribution. The distribution is a combination of mail, news and network utilities, combined with some tools that might come in handy and (hopefully) a lot of documentation, like:

mail:
- mh
- elm2
- sendmail (5.61 currently)
- smail
- ida
- mailway

network:
- ka9q
- snmp
- uupc
- uucp over x25

news:
- bnews
- cnews
- nn
- nntp
- rn
- vms
- tmnn

gnu-tools:
- gawk
- grep
- make
- mh
- tar

ducmentation:
- iso3166
- ethernet-info
- several RFC's

misc:
- kermit
- tn3270
- vacation

With regards to the documentation a note. It is our intention to have National Groups supply information for the tape, in the local language, about how their network is set up, how to connect, what is and what is not (yet) possible, and some further guidelines.

We still have to build this database of local information. This STARTER KIT will be (hopefully frequently) updated and new software and documentation incorporated over time. This of course largely depends on the amount of time that I can make available for this and how much

input the local groups give. I will do my best.

The second new tape, EUUGD21, is the one specially made for the Nice conference. This time the conference tape is a collection of some odd looking items. First there is the dtree software, generously donated to the community at large, by no less then David Tilbrook <nixtdc!dt@uunet.uu.net>. David wrote a lot of this and urged others to do the same. "The D-Tree is a source distribution consisting of over 250 tools, many of which are designed to maintain software hygiene in the development of medium-size software distributions."
Other ''odd'' ends are: ABC, an easy to learn, but powerful language, developed at CWI (contact abc-list-request@cwi.nl, to join a mailing list).
PP-5.0Postman Pat (see later). As PP requires ISODE6.0, I have included that as well.
Furthermore, I have included new versions of: ka9q, sendmail, smail, vmh, elm, mailway, pop, popper, and slip, ''just'' to fill the empty space on the tape, and to provide these to those who are already connected to EUnet, but have no real use for the STARTERKIT.

The third new tape, EUUGD22, contains the already mentioned PP-5.0. PP is a Message Transfer Agent, intended for high volume message switching, protocol conversion, and format conversion. It is targeted for use in an operational environment, but may also be useful for investigating Message related applications. Good management features are a major aspect of this system. PP supports the 1984 and 1988 versions of the CCITT X.400 / ISO 10021 services and protocols. Many existing RFC 822 based protocols are supported, along with RFC 1148 conversion to X.400. PP is an appropriate replacement for MMDF or Sendmail.

And, last but not least, I have made a new version of the Brussels EUUG 1989 spring conference, aka EUUGD16. With the latest versions of all the GNU programs I could find.

Ok. Below you'll find the usual list of available distributions. As always, anyone is invited to make their own tools, games, etc available for publication on an EurOpen tape. Please contact me for more details. Don't hesitate, just put the results of many nights of serious programming and hacking in the public domain, and you might even become famous!

This is a list of all the current (October 1990) EurOpen Software Distributions. It is a short description of the available tapes. Any changes to the contents of the tapes, as well as announcements of new tapes will be placed in the EurOpen Newsletter. I am working on a method so you can automagically, by e-mail, easily find out which program is on which distribution. For the moment you will have to e-mail, call, or write me to find out.

Prices of the tapes are in Dutch guilders (DFl), and do not include VAT-taxes. Prices do include postage cost for surface mail within Europe. Any special shipment costs, like with DHL, will be billed through.

The first price listed with each distribution, is for 1/2", 9-track, reel tapes in tar 1600 bpi format, the second one is for distributions on 1/4" cartridge tapes in (Sun) QIC-24 format. Prices for 800 bpi reel tapes and QIC-11 cartridges may differ from the ones listed.

Note that you have to be an EurOpen member (or a member of a local UUG) to obtain tapes at list prices. Non-members will have to pay an extra DFl 300,- per tape.

EUUGD1 R6: UNIX V7 system, specially made for small DEC PDPs (11/23, 11/34, etc). The Kernel supports the UK terminal driver. V7 source licence minimum.
Price: DFl 130.-/190.-

EUUGD2: Early Pascal compiler of the Free University of Amsterdam. V7 source licence minimum.
Price: DFl 130.-/190.-

EUUGD3: STARTER KIT It finally is here! The tape contains many public domain mail, news, networking utilities, tools, attributes and information. It is THE tape for all who want to share in the joy of really feeling connected to the world in general, and the Unix TM world in particular!
Price: DFl 130.-/190.-

EUUGD4: Software tools, sampled by the Software Tools Users Group. Most of the software is written in Ratfor, for which a Fortran support tool is included. This tape is available in different formats: DEC RSX, DEC VMS, UNIVAC, IBM MVS, UNIX tar, MIT line feed format, and MIT card format (80 columns).
Price: DFl 160,-/190,-

EUUGD5: Currently not available. See tape EUUGD20 for new benchmark software.

EUUGD6: (USENIX 83.1)
USENIX tape, containing contributions from various UNIX System Group Members. This is a licence dependent distribution: V7, V32, SIII, V6 or no licence disclosure available.
Price: DFl 250,-/310,-

EUUGD7: UNIX|STAT Version 5.2. A collection of about 25 data manipulation and analysis programs written in C by Gery Perlman.
Price: DFl 70,-/190,-

EUUGD8: A collection of useful software, based on the so called Copenhagen tape (EUUG UNIX conference Autumn 1985).
Price: DFl 130.-/190.-

EUUGD9: A collection of useful software, based on the so called Florence tape (EUUG UNIX conference Spring 1986).
Price: DFl 160,-/190,-

EUUGD10: MMDFIIb. Multichannel Memo Distribution Facility (version IIb). This is a powerful, domain oriented mail system with access control and the ability to communicate over a variety of network systems including TCP/IP, JANET, UUCP, PHONENET, etc. It has been ported to a variety of UNIX's including but not limited to 4.[123] BSD, 2.9 BSD, System III/V on a variety of different hardware. You should first obtain a licence agreement by sending a message to euug-tapes@EU.net. Return the signed licence with your order.
Price: DFl 100,-/190,-

EUUGD11: This is the 'Boat' tape; the Helsinki EUUG 1987 spring conference. It contains about 25 Megabytes of programs, games, etc. Including: jove, less, nag, news, rn , uEmacs, uuencode and larn.
Price: DFl 130.-/190.-

EUUGD12: This is the Dublin EUUG 1987 autumn conference tape. It contains about 26 Megabytes of programs, games, etc. Including: copytape, crc_plot, fastgrep, jove, kermit, notes, uupc, nethack, cron, sendmail, mh, Recipes, brl-gw, isode, pcip, pctelnet.
Price : DFl 130.-/190.-

EUUGD13: The latest conference tape for the London EUUG 1988 spring conference tape. It contains things like: cake, chat, config, copytape, graphedit, kermit, little-st, mcc, mstools, news, pd-diff, pdtar, perl, postscript, psfig, pshalf, shar, rpc, moria4.85, omega, arc, backup, smail, sush, watcher, and much, much more.

Price : DFl 130.-/190.-

EUUGD14:

This is version 6.0 of this non-proprietary implementation of some of the OSI parallel protocols suites as defined by the International Organisation for Standardisation (ISO), the International Telegraph and Telephone Consultative Committee (CCITT), and the European Computer Manufacturer's Association (ECMA).

This release is coded entirely in C, and is known to run under the following operating system without kernel modifications:

BSD 4.2 and 4.3
Ultrix
AT&T UNIX SVR2 and SVR3
AIX
HP-UX
ROS
Pyramid OsX

Since a Berkeley UNIX system is the primary development platform for ISODE, the documentation and source are somewhat slanted towards that environment. The tape contains some 12Mb of both tools and documentation in machine readable form. The EurOpen will send you a tape only.
Price: DFl 130,-/190,-

EUUGD15: Here it is! The complete X11 Windowing system, as distributed by MIT, release 4: X11R4. Do to the vast growth in user contributed software, this distributions now totals 60 Mb in compressed form. This results in two 1/2", 9-track tapes, and one 600 ft, Qic-24, 1/4" cartridge. This includes the core system, as well as much, very much user contributed software.
Price : DFl 260.-/200.-

EUUGD16: This is the Brussels EUUG 1989 spring conference tape, and consist entirely of software from the GNU project from the Free Software Foundation (not to be confused with OSF:-).
On this tape you will find: ispell, g++, awk, gcc, gdb, Cscheme, emacs, lisp-manual, libg++, binutils, bison, ghostscript, gas-dist, gawk, gnews, gnuchess, make, oops, pace, ps-emacs, scheme, sed1, tar and torture.
Price : DFl 130.-/190.-

EUUGD17: This tape contains the software for ET++. From the abstract of the "Autumn 1988 EUUG Conference Proceedings":
"ET++ is an object-oriented application framework implemented in C++ for a UNIX environment and conventional window system. The architecture of ET++ is based on MacAPP and integrates a rich collection of user interface building blocks as well as basic data structures to form a homogeneous and extensible system."
It totals some 18Mb of software that the people of the Institut fuer Informatik of the University of Zurich were so kind to let us, mere mortal souls, play with. Have fun.
Price: DFl 130.-/190.-

EUUGD18: This is the "Vienna EUUG 1989 autumn conference tape", and consists entirely of games! There is a SUN specific set, a set for the X Windowing System environment, and a general useable set. All the games supplied are working, and have been tested at CWI by our "Games Keeper <play@cwi.nl>". For many games he added additional features, not found in the originals.
Some of the games included are: for SUN: Asteroids, Mahjongg, Othello, Qix, Sdi, Tetris. For the X environment: Xtrek, Xgo, Xwanderer, Xrobots. General games: Nethack, Adventure, Atc, Empire, Reversi, Yahtzee, Trek73, Backgammon, Corewars, MazewarsV, Vtrek, and lots, lots more.
If this doesn't bring some fun back into using computers, I don't know what else can :-)
Price : DFl 130.-/190.-

EUUGD19: This is the "Munich EUUG 1990 spring conference tape", and consist entirely of graphics material. Conversion programs, display tools, toolkits to build you own display program, and off course images, lots of pictures to play around with.
Price: DFl 130.-/190.-

EUUGD20: This tape contains benchmarking software and is named "AFUU/SSBA 1.2, benchmarks". The French group have done a good job creating a tape with all the necessary tools, so you can finally bring your machine down to it's knees, and see what it is really worth.
Price: DFl 70.-/190.-

EUUGD21: This is the "Nice EUUG 1990 autumn conference tape", and consist of a number of different kinds of software, like: dtree, abc, new versions of various mail and news utilities, and PP5.0.
Price: DFl 130.-/190.-

EUUGD22: This is the "Postman Pat PP5.0" distribution. PP is a Message Transfer Agent, intended for high volume message switching, protocol conversion, and format conversion. It is targeted for use in an operational environment, but may also be useful for investigating Message related applications. Good management features are a major aspect of this system. PP supports the 1984 and 1988 versions of the CCITT X.400 / ISO 10021 services and protocols. Many existing RFC 822 based protocols are supported, along with RFC 1148 conversion to X.400. PP is an appropriate replacement for MMDF or Sendmail.
Price: DFl 130.-/190.-

# EurOpen Software Distributions Order Form

If you want to order any tape, please write to:

EurOpen Software Distributions
c/o Frank Kuiper
Centrum voor Wiskunde en Informatica
Kruislaan 413
1098 SJ Amsterdam, The Netherlands

For information only:

Tel: +31 20 5924121 (or: +31 20 5929333)
Fax: +31 20 5924199
Telex: 12571 mactr nl
Internet: euug-tapes@EU.net

Please note that for distributions D1, D2 and D4 a copy of your source licence agreement with AT&T for at least UNIX version 7 should be enclosed. Note also that you have to be an EurOpen member (or a member of a national UUG) to obtain tapes at list prices. Non-members will have to pay DFl 300,- per tape extra as handling fee. Please enclose a copy of your membership or contribution payment form when ordering. Do not send any money or cheques, you will be invoiced.

All 1/2", 9-track, reel tapes come in tar format, 1600 bpi. 800 bpi is possible on request. Cartridge tapes come in tar format, written with dd, with a blocking of 126b. This is a so-called QIC-24 format, written on a Sun. QIC-11 is available on request.

This page may be photocopied for use. Please print!

Name: .....

Company name: .....

Address: .....

.....

.....

E-mail address: .....

I would like to order the following:

.....

.....

.....

Tape format, either 1/2" 9-track, or 1/4" cartridge: .....

Copy of EurOpen (or national UUG) membership (or payment)
form enclosed? Yes / No
Copy of AT&T source licence enclosed? Yes / No

"I declare to indemnify the European UNIX systems User Group for any liability concerning the rights to this software, and I accept that EurOpen takes no responsibilities concerning the contents and proper function of the software."

Date: ..... Signature: .....

# Call Doc Strange

*Colston Sanger*
*doc.strange@gid.co.uk*

*GID Ltd*

Time passes, and Colston Sanger is now a senior consulting infusion engineer with GID... He is also a visiting lecturer in the Faculty of Engineering, Science and Mathematics at Middlesex Polytechnic. Apart from that, he has been incredibly busy this last month organising a (successful) seminar on multimedia.

## Setting Up a Trailblazer

Our company, GID, bought a second-hand Dowty Trailblazer modem a while back, so I've recently been having a lot of fun setting it up. Yes, Trailblazers can be *a lot* of fun, and by now I can do a passable imitation of one: from the *piko-shshsh* at login to the full-steam-ahead of serious *chugga-chugga-chugga* UUCP file transfer.

So how do you set up a Trailblazer? In the hope of making it easier for anyone else who has to do it, let me describe my setup.[1] You'll remember from previous columns that I live in the System V world — specifically, AT&T System V Release 3.2 with 'Basic Networking Utilities' (that's HoneyDanBer UUCP to anyone else). I should add that I'm using standard straight-through cables without any tweaks.

### /etc/inittab and /etc/gettydefs

First, `/etc/inittab`, perfectly ordinary — no magic here:

```
11:234:respawn:/usr/lib/uucp/uugetty -r -t 60 tty11 9600CS8 # TB modem
```

*uugetty* (standard in BNU) is identical to *getty* except that it is bi-directional, *i.e.*, users can login, but if the line is free *uucico*, *cu* or *ct* can use it for dialling out. The **-r** option is always used with an intelligent modem or direct line that has a *uugetty* at the remote end (this causes *uugetty* to wait to read a character before putting out the `login:` prompt). **-t** is a timeout — in this case, 60 seconds. **9600CS8** (stands for '9600 baud, 8 bit')

1. Many thanks to jpp@specialix.co.uk, lenny@icus.ICUS.COM, olapw@olgb1.oliv.co.uk and tehm@uel.co.uk for sharing their thoughts with me.

is one of a set of homemade labels into **/etc/gettydefs**:

```
9600CS8# B9600 CS8 HUPCL CLOCAL IXOFF OPOST ONLCR #B9600 CS8 CREAD HUPCL
BRKINT IGNPAR ICRNL IXON IXOFF OPOST ONLCR TAB3 ISIG ICANON ECHO ECHOE
ECHOK #sixnine login: #2400CS8

2400CS8# B2400 CS8 HUPCL CLOCAL IXOFF OPOST ONLCR #B2400 CS8 CREAD HUPCL
BRKINT IGNPAR ICRNL IXON IXOFF OPOST ONLCR TAB3 ISIG ICANON ECHO ECHOE
ECHOK #sixnine login: #1200CS8

1200CS8# B1200 CS8 HUPCL CLOCAL IXOFF OPOST ONLCR #B1200 CS8 CREAD HUPCL
BRKINT IGNPAR ICRNL IXON IXOFF OPOST ONLCR TAB3 ISIG ICANON ECHO ECHOE
ECHOK #sixnine login: #9600CS8
```

Note that these are meant to be single-line entries — they are split here to fit the page. Also that each entry in **/etc/gettydefs** must be followed by a blank line.[2]

There were several conflicting opinions among my net advisors about what exactly to put in **gettydefs**: some said you only need one constant-speed entry between the computer and the modem; others disagreed on which initial flags are required (the baud rate and HUPCL should suffice).[3] Myself, I'm still experimenting — which is why I haven't tried a 19200 baud connection yet.

### Required Services

If you've got a Trailblazer, you probably want to make the most of it. Mine is set up to offer:

| | |
|---|---|
| Incoming: | *uucico* at 9600 PEP, 2400 and 2400 MNP,<br>plus a bit of *kermit*,<br>plus terminal emulation at 2400 MNP |
| Outgoing: | *uucico* at 9600 PEP,<br>plus *cu* at 9600 PEP |

### The Actual BNU Setup

So how is it done? Essentially by having a default, generic, dial-in setting stored in the Trailblazer, and then reconfiguring it temporarily for outgoing traffic. All this, of course, is made possible by the fact that you can have different BNU configuration files for different sorts of

2. If you are fiddling with **/etc/gettydefs**, it's probably a good idea to use *getty -c* to check your handiwork.
3. You may want to read the manual page for *gettydefs*(4) so that you can make your own interpretation.

service, as defined in **/usr/lib/uucp/Sysfiles**:

```
#ident      "@(#)/usr/lib/uucp/Sysfiles.sl 1.2 3.1 05/01/86 37353 AT&T-SF"
#
#     Sysfiles provides a means of specifying different Systems,
#     Devices and Dialers files for uucico and cu.  This is useful
#     for certain networks where, for example, requests for login
#     service might be accepted on a different address than requests
#     for uucico service (hence different Systems files), or where
#     there are different chat scripts to request each service
#     (hence different Dialers files).
#
#     Another use of Sysfiles is to split large Systems files into
#     smaller, more manageable files (e.g., local and global Systems
#     files).
#
#     FORMAT:
#     service=<service name>  systems=<systems file list> \
#                       devices=<devices file list> \
#                       dialers=<dialers file list>
#
#     Where service name is "uucico" or "cu".  Each file list is a list
#     of colon-separated file names.  File names are relative to
#     /usr/lib/uucp unless a full path name is given.  Files are searched
#     in the order that they appear in the file list.
#
#     The defaults are the usual uucp files: /usr/lib/uucp/Systems,
#     /usr/lib/uucp/Devices and /usr/lib/uucp/Dialers.
#     ...
service=uucico    systems=Systems \
            devices=Devices \
            dialers=Dialers
service=cu  systems=Systems \
            devices=Devices.cu \
            dialers=Dialers.cu
```

### The Default Dial-in Setup

Here are the bits that define the default dial-in setup. I've done it by defining a fake system 'setup' in **/usr/lib/uucp/Systems** so that I can *cu* to the modem to download the settings:

```
#ident      "@(#)/usr/lib/uucp/Systems.sl 1.2 3.0 12/06/85 37395 "
#
# Entries have this format:
#     Machine-Name Time Type Class Phone Login
# ...
# You didn't really think I was going to broadcast my Systems file, did you?
#
# setup - a fake - used to initialise the modem
setup Any ACU S9600 -
```

There is no phone number so '-' is used as a place-holder. **S9600** ('setup, 9600 baud') is a speed class — it maps back to the relevant entry in **/usr/lib/uucp/Devices.cu**:

```
#ident      "@(#)/usr/lib/uucp/Devices.sl 1.4 3.1 05/01/86 35982 AT&T-SF"
#
# NOTE - all lines must have at least 5 fields
#    use '-' for unused fields
# The Devices file is used in conjunction with the Dialers file.
# Types that appear in the 5th field must be either built-in
#  functions (801, Sytek, TCP, Unetserver, DK)
#  or standard functions whose name appears in the first
#  field in the Dialers file.
# Two escape characters may appear in this file:
# - \D which means don't translate the phone #/token
# - \T translate the phone #/token using the Dialcodes file
# Both refer to the phone number field in the Systems file (field 5)
# \D should always be used with entries in the Devices file, since the
# Dialers file can contain a \T to expand the number if necessary.
# ...
# Trailblazer stuff - C.S., Thu Sep 20 16:37:52 BST 1990

# Ordinary (!) 9600 baud PEP cu
ACU tty11,M - 9600 tb9600 \D

# Setup - to initialise the modem
ACU tty11,M - S9600 setup_TB \D
```

The ,M appended to the second field is an optional modem control flag. It indicates that the line should be opened O_NDELAY without waiting for a carrier (CD).

**Dialers.cu** is the important one, where the Trailblazer default configuration is defined:

```
# Dialers.cu
#
# Trailblazer stuff - C.S., Thu Sep 20 16:43:22 BST 1990
# &F   - Factory defaults
# Q0E1  - At first, so we can see what we're doing
# E0   - Later, don't echo commands,
# Q2   - local 'noise' (responses), remote quiet
# M1   - Speaker on when dialing (factory default)
# X3   - Show extended result codes, including MNP and PEP
# S Registers:
# s0=1       - Answer after one ring (factory default)
# s48=1 - All 8-bits are significant
# s50=0 - Auto connect speed determination (factory default)
# s51=255 - Automatic *interface* speed selection
# or
# s51=4 - Revert to 9600 baud interface speed when connection terminates?
# s52=2 - Go on-hook when DTR drops and reset to NVRAM settings
# s53=4 - DCD tracks remote carrier, DSR is on when modem is ready
# s54=3 - Pass BREAK to remote in sequence in data stream
# or
```

```
# s54=1 - discard buffered data, pass BREAK to remote?
# s55=3 - (Respond sanely to command escape sequence - jpp's words)
# or
# s55=1 - discard buffered data, send escape sequence to remote? (same as s54)
# s58=3 - Modem uses XON/XOFF flow control (factory default)
# or
# s58=4 - Modem uses CTS/RTS *and* XON/XOFF flow control?
# s65=1 - Failsafe XON/XOFF control
# s66=0 - Interface speed can change, no flow control
# or
# s66=1 - Lock interface speed - go with the flow?
# s68=255 - Modem uses whatever flow control the computer uses
# s92=0 - PEP tones at beginning of answer sequence (factory default)
# s95=2 - MNP auto-reliable mode
# s110=0 - Disable data compression
# or
# s110=1 - Use data compression if remote requests it (but not for
#          compressed news!)
# s111=255 - Accept any protocol
# &W  - write to NVRAM
#
# (This should, of course, all be on one line.)
setup_TB =W-, "" \M\pA\pA\pA\pA\pA\pT&F ""
AT&FQ0E1M1X3S0=1S48=1S50=0S51=255S52=2S53=4S54=3S55=3S58=3\r
OK ATS65=1S66=0S68=255S92=0S95=2S110=0S111=255E0Q2&W\r\m

# to find out what ROM revision...
# setup_TB =W-, "" \M\pA\pA\pA\pA\pA\pT&F "" AT&FQ0E1I3\m
# setup_TB =W-, "" \M\pA\pA\pA\pA\pA\pT&F "" AT&FQ0E1N?\m

# For ordinary cu
# 9600 baud-ish (in fact, auto speed determination), data compression
# enabled if remote is enabled, use protocol specified by remote.
tb9600  =W-, "" \MATZQ0E1S50=0S110=1S111=255 OK ATDT\T CONNECT \d\r\m
```

A couple of details here. =W-, are simple substitutions, meaning map '=' to 'W' (for the Trailblazer, that's wait for dialtone for the period specified by S6, default 4 seconds) and map '-' to ',' (pause, default 4 seconds). The \M and \m at the beginning and end of the chat-script turn on and off the CLOCAL flag;[4] and all those \pA\pA\pA ('pause A pause A pause A') are there just to wake up the modem.[5]

## To Lock or Not to Lock?

In the default dial-in configuration defined above there are several alternative settings that I've not quite decided on yet, the most important of which is S66. As I mentioned earlier, there were conflicting opinions among my net advisors on whether the interface speed between the computer and modem should be locked. (This has also surfaced elsewhere as the great 'autobauding' controversy, and has spawned at least one *uugetty* replacement that sets the interface speed directly from the CONNECT speed.) All I'm saying (for now!) is that I don't have the interface speed locked and I've not seen any problems.

4. Hence I decided to keep CLOCAL in the initial flags of **/etc/gettydefs** for the time being.

5. Different ROM revisions seem to need more or less of these.

### uucico Setup

I've included the setup I use for outgoing *cu* above (note that there is no '&W' at the end of the *cu* chat-script). What I haven't shown yet is the outgoing *uucico* setup. Here it is. First, from **/usr/lib/uucp/Devices**:

```
# Trailblazer
ACU tty11,M - 9600 tbpep9600 \D
ACU tty11,M - 9600 tb9600 \D
```

and from **/usr/lib/uucp/Dialers**:

```
#ident      "@(#)/usr/lib/uucp/Dialers.sl 1.11 3.2 01/12/88 22765 AT&T-SF"
#
# Each caller type that appears in the Devices file (5th field)
# should appear in this file except for the built in callers.
# Each line consists of three parts:
# - the name of the caller
# - the translation table for the phone number to translate from
#   the 801 codes (=-) to the code for the particular device
# - a chat script (same format and meaning as the login scripts
#   that appear in the Systems file.
#
# Meaning of some of the escape characters:
# \p - pause (approximately 1/4-1/2 second delay)
# \d - delay (2 seconds)
# \D - phone number/token
# \T - phone number with Dialcodes and character translation
# \N - null byte
# \K - insert a BREAK
# \E - turn on echo checking (for slow devices)
# \e - turn off echo checking
# \r - carriage return
# \c - no new-line
# \n - send new-line
# \nnn - send octal number
#
# Trailblazer stuff - C.S., Thu Sep 20 18:10:09 BST 1990
# Adapted from Tris Mabbs's (uel) scripts
#
# 9600 baud, PEP mode, use UUPC protocol, data compression enabled
# if remote is enabled
# (Should be on one line.)
tbpep9600 =W-, "" \MATZ\p\p "" ATQ0E1S50=255S51=255S110=1S111=30
OK ATDT\T CONNECT \d\r\m
#
# 9600 baud-ish (in fact, auto speed determination), data compression,
# use protocol specified by remote
tb9600 =W-, "" \MATZ\p\p "" ATQ0E1S50=0S110=1S111=255 OK ATDT\T CONNECT \d\r\m
```

### Finally, an Interesting Question

Finally, here's an interesting question: since Trailblazers are such super-reliable, error-checking *wonderful* things, why not use an error-free transport protocol such as 'e'(used by AT&T STARLAN) or 'f' (BSDish for use with X.25) instead of the standard 'g' protocol? If you used 'e', you wouldn't have to do all that UUCP error-checking and you would get even greater throughput!

Well, I've just tried it — and it seems to work, though not conclusively. To set it up, just add a **,e** (or maybe a **,eg** to play safe) to the Type field in the **/usr/lib/uucp/Systems** file. For example:

```
qwerty Any ACU,eg Any bluskye26879 "" \r\d\r\d in:--in: nuufred word: bldg!gbt2
```

Of course, you'll need an SVR3 UUCP neighbour that supports the 'e' protocol.

## Yet More Books to Read

The book to start with is, naturally, your Trailblazer manual, together with the section on setting up Basic Networking Utilities in the *UNIX System V Release 3.n System Administrator's Guide*. Also worth reading is the chapter on UUCP by Brian Redman (the 'Ber' in 'HoneyDanBer') in Stephen Kochan and Patrick Wood (eds), *UNIX Networking*, Hayden Books, 1989 (ISBN: 0-672-48440-4). More generally, the Nutshell handbook *Managing uucp and Usenet* (ISBN: 0-937175-48-X) is useful, and most of the books on system administration have a chapter on setting up one or another flavour of UUCP.

---

## A Book on C,

Al Kelley and Ira Pohl, Benjamin Cummings, 1990, ISBN 0-8053-0060-0.
(UK) Price £19.95, Soft Back, 519 pp.

Reviewed by Rob Henley of Siemens SDG, Reading, UK. (rob@siesoft.co.uk)

A very readable and comprehensive introduction to the C programming language. The approach is very thorough, with a wide variety of annotated examples, useful chapter summaries and exercises. The book covers all the new ANSII features in the main text and includes an appendix summarising the new features for readers already familiar with non-ANSII compilers. The authors cater for both UNIX and DOS users where reference is made to the programming environment itself.

After a short overview to get you writing, compiling and running programs, the book leads you chapter by chapter through every feature of the language, taking in many of the algorithms and data structures common to computer science curricula on the way. Starting with the lexical elements of the language, chapters proceed to describe fundamental data types, flow control and functions before moving on to the more sophisticated data structures. There are also good descriptions of the C preprocessor and topics such as profiling and timing code which are usually omitted from books of this type. Appendices cover the C standard library and language syntax.

If I found any criticism to make it was that the book didn't give much feel for the dynamic nature of C programming, the importance of pointers, memory management and the UNIX bias toward small cooperating processes. In particular the description of *fork*( ) and *exec*( ) didn't impart the feeling of elegant cooperation which they embody.

The book assumes little previous programming experience, although some knowledge of structured programming in a high level language would be beneficial. The treatment may be too slow for experienced programmers, although there are many useful code examples to dip into. It would make an excellent introduction to C for an undergraduate or anyone wanting a gentle introduction to programming in a high level language. Whilst the book is well indexed, hardened C programmers will still prefer the Kernighan and Ritchie classics as reference texts.

# News from the Netherlands

1st November 1990, the NLUUG had her autumn conference with the theme Open Systems. There were two parallel sessions of presentations and a tutorial session by Erik Hardy (CMU) on User Interface Development Tools. During the day there were UNIX/Open products on show from several vendors, and there was a publisher with a selection of publications.

*Copies can be ordered from the NLUUG (address see back cover) Price: Netherlands F 40, other countries F 50 (both incl. postage). The proceedings are predominantly in Dutch.*

Thanks are due to Patricia H. Otter <patricia@xirion.nl> of Xirion Unix Software & Consultancy bv who translated the abstracts of the presentations:

**Title:** XAA - The UNIX Application Architecture
**Speaker:** Willy Witsel
**Institute:** NCR Nederland N.V.
Buitenveldertselaan 3, 1082 VA Amsterdam

Why do vendors create Open Systems? Do they do it of their own free will, or are they forced to do so? Why do some vendors still sell proprietary systems, and will they succeed in the future? Do Open Systems provide real benefits for customers?

Questions, questions, questions !!

In this presentation we will tried to give the answers to the above questions by examining the effects of Open Systems in different industries. First we took a close look at the automobile industry. We followed technological advances in this industry since the mid 50s. We compared the evolution of Open Systems in this industry with the Open Systems Architecture of the computer industry in the 80s. Next we presented the definition of an Open Systems Architecture, and we discussed the benefits for both vendors and customers. We also demonstrated how vendors can create standards and how these standards become part of an Open Systems Architecture. Lastly we broke the Open Systems Architecture down into five basic levels.

**Title:** Distributed Application Architecture (DAA)
**Speaker:** Diana Jackson
**Institute:** Data General Nederland BV
Burgemeester Stramanweg 101,
1101 AA Amsterdam Z.O.

DAA is a long-range model of distributed computing. DAA supports integration of PSUs, workstations, servers and compute engines into client/server environments, based on industry standards, providing platforms for applications development throughout the industry. The goal is freedom of choice for IT users by providing them utility-like access to computing, and superior integration in an Open Systems (multivendor) environment.

The key elements of a distributed application architecture are:

- Common User Environment, such as OSF/Motif, but including also an Object Management Facility such as found in NewWave.
- Open Networking.
- Distributed Object Management, for management of all network-based resources. Object orientation provides the transparency needed to achieve a truly open and user-centered computing model.

The presentation included practical examples illustrating the level of integration possible with products available today.

**Title:** OS/2 a competitor of UNIX
**Speaker:** Allard Blankesteijn
**Institute:** Xirion bv
Burg. Verderlaan 15X, 3454 PE De Meern

IBM and Microsoft introduced OS/2 as the long-expected successor of MS-DOS. OS/2 promised to solve the restrictions of MS-DOS (memory restrictions and single tasking). As OS/2 is multi-tasking like UNIX, it is a competitor for UNIX on the desktop in the business world. In situations in which a choice has to be made between UNIX and OS/2 it is therefore important to recognise the consequences of such a choice. This presentation discussed a number of these consequences, with regard to the experience gathered by integrating OS/2 on one side and UNIX on the other.

**Title:** An Open model-approach of standardisation
**Speaker:** Th. Bruins or W Thieme
**Institute:** RNL (PTT strategy and research)
Postbus 421, 2260 AK Leidschendam

The word "*open*", as it is used in for instance OSI, UNIX/OPEN or ONP, is part of a trend closely connected with selective standardisation. The goal behind this all is achieving the modular assembly with products coming from different suppliers.

The word is only used in areas of telecommunications and computer science, but in fact it is a process in which a maximum of production freedom and market volume is combined with a minimum of standardisation. Usually, structured models are used.

This method is generally used in (international) normalisation bodies like ISO and CCITT.

**Title:** The SERC Network, how to survive as a System manager.
**Speaker:** Eugene Bogaart
**Institute:** Software Engineering Research Centre
Postbus 424, 3500 AK Utrecht

The computer infrastructure of the Software Engineering Research Centre (SERC) in Utrecht is an open environment based on an integrated network of workstations (from different manufacturers) and personal computers. The platform offers the researchers of the institute a wide variety of services, such as an enormous collection of software-tools. On each hardware platform a homogeneous interface hides the distribution details from users. The presentation covers the evolution and design of the infrastructure seen from a network integration point of view. The selection of standards and the technological implementation are very important. Also we looked at system management in an open network. The selection of a network model consisting of a model for the configuration management with adequate automated support, can improve daily life enormously.

**Title:** Views: An Open User-interface System
**Speaker:** Steven Pemberton
**Institute:** Centrum voor Wiskunde en Informatica
Kruislaan 413, 1098 SJ Amsterdam
steven@cwi.nl

From the user's point of view, Views is a computing environment where all actions are achieved by editing documents, so that once you have learned how to use the editor, in principle, you can work out how to do everything else.

- For instance, for document (i.e. file) management, instead of the Unix commands ls, mv, cp, rm, chmod, cat/more, edit etc., you just have a document that describes the contents of the directory (whether textually or graphically).
  - — to rename a file, you just edit its name in the document
  - — to copy it, you just use the editor's copy-paste facility
  - — to remove it you just delete the entry with the editor
  - — to chmod it, you just edit the modes field
  - — to look at the file (and edit it) you just "visit" the entry and so on.
- Similarly, instead of ps, kill, nice, etc., you just edit a (dynamic) document that describes the running processes.
- Instead of lpr, lpq, lprm, you just edit a document that describes the printer queue.
- To read your mail, you have a document - your inbox - that shows the headers of your mail. You can visit entries, copy/paste them, delete them, and so on.
- Similarly to read news, and a whole host of other applications that fit in the metaphor.

From the application builder's point of view, Views is an open-architecture computing environment, where applications are easily added, where the user-interface is guaranteed to be identical across applications, and where within applications user-interface issues are largely absent.

Just as C-programs don't have to know anything about file-name expansion in the argument list, nor anything about different sorts of input and output devices, Views applications don't have to know anything about the user-interface. The application builder has only to define the objects used in the application (or use predefined objects), giving a default definition of the external representation of those objects, and can specify invariants that have to hold between different objects, using self-defined or standard functions.

From that point on, the Views system takes over, displaying objects, and letting the user edit them. If an object that is linked by an invariant to other objects gets changed, then the other objects get updated too.

As an example, when an object is visited, a window is opened and the object being visited is connected to the window by the function "display". Apart from the object itself, display also takes as input parameter the object describing the external representation of the object, and uses this information to display the representation of the object in the window. This means that whenever the object gets changed (edited), display is called to redisplay it (maintaining the invariant between the object and its representation). Similarly, since the representation of an object is defined by another object, that too can be edited by the user, and if that happens, display also gets called to redisplay the object. This means that you can easily change the external representation of documents on-the-fly.

Views is platform-independent, by using STDWIN, a platform-independent windowing package. The current prototype version of Views runs equally well on X, the Macintosh, and the Atari ST.

An interesting aspect of the system is that (as alluded to above) the presentation of an object is independent of the object itself. So a graphic object, for instance, can be simultaneously viewed and edited in MacDraw-like way, and as pic-style textual description. Editing either view causes the other to be updated to match.

**Title:** UNIX, Standardisation, Profiles and EWOS
**Speaker:** Willem Wakker
**Institute:** Associated Computer Experts bv
Van Eeghenstraat 100, 1071 GL Amsterdam

A number of efforts are currently underway to define a standard for UNIX. Both formal standards bodies such as ISO and the IEEE, and less formal groups such as X/Open are involved. Even after the proposals from the various groups are aligned, leaving just a single standard per topic, the final result will still be massive: the sheer number of areas to be standardised guarantees this.

Within CEN/CENELEC, the European Workshop on Open Systems (EWOS) is using the concept of Profiles (or Functional Standards) to make the similarly complex OSI standards manageable, and this work is now being extended to cover UNIX standards.

This talk will give a brief overview of current UNIX standardisation efforts, and how they are progressing. The work of EWOS on profiles will then be explained, and how such profiles allow for the definition of manageable standards for particular types of applications and systems.

**Title:** Digital's View on Open Systems
**Speaker:** Peter Oosterling
**Institute:** Digital Equipment B.V.
Europalaan 44, 3526 KC Utrecht

What are Open Systems. What do X/Open and OSF mean for Open Systems. How does the market behave. What are the consequences of Open Systems for:

— clients and/or users

— software suppliers

— hardware suppliers

— system integrators

What is Digital's strategy regarding Open Systems:

— foundation

— NAS: integration and portability of applications, data and knowledge of people

— UNIX, one of the important elements supplying a complete range of services in one package of in a multi- (hardware and software) suppliers information management environment

**Title:** IBM Risc in SAA
**Speaker:** Hans de Groot
**Institute:** IBM Nederland N.V.
Gildenkwartier 199, Utrecht

With the introduction of the RISC system/6000 and the advanced interactive executive (AIX) version 3 operating system, IBM made a statement that it is willing and able to provide solutions in the world of Open Systems.

Acceptance of all relevant standards is a integral part of this willingness.

Roughly the architectures of the RISC system/6000 and the AIX version 3 operating system was covered. The relationship with the typical IBM environment, as described by the systems applications architecture (SAA), is characterised by peaceful coexistence and very good mutual contact.

**Title:** Open Systems: "Solutions for Business"?
**Speaker:** J.B. van Rongen
**Institute:** Coopers & Lybrand Management Consultants
Weena 151-161, Postbus 2755, 3000 CT Rotterdam

The view of a management consultant on Open Systems.

We approach the architecture of information-supply and automation from a managerial angle. What are the company's processes, how should these be implemented to get the best results, what is the best implementation of the information-supply, etc.

Often we advise companies against development, and we suggest they buy a package. In a good package (for instance MRPII) hundreds of man-years of knowledge are incorporated, which can never be equaled by new development.

However, functionally good packages are static, rather old and therefore expensive to use. More and more, companies are confronted with rapidly changing markets, and the need for flexible action. The classical, static information-supply becomes a tremendous burden.

Thus the demand arises for more flexible packages and flexible information-architectures. This talk covers these trends and the position that Open Systems may have in companies. We will give examples of a number of companies that have gone that direction. Moreover we point out how companies can handle "the burden of the past".

**Title:** X/Open
**Speaker:** a technical officer
**Institute:** X/Open Co Ltd.
Apex Plaza, Foxbury Road, Reading, UK.

No details known.

**Title:** The SPARC Application Binary Interface (ABI)
**Speaker:** H. Willems (SUN) and J. Sprey (ICL)
**Institute:** Sun Microsystems
Computerweg 1, 3821 AA Amersfoort
and
ICL Nederland B.V.
Het Kwadrant 1, 3606 AZ Maarssen

No details known.

**Title:** Team Computing: a concept for computers working together
**Speaker:** Bouwe van der Eems
**Institute:** Hewlett-Packard Nederland B.V.
Meidoornkade 10, Postbus 170, 3990 DD Houten

In recent years we have had a rapid development in computer technology. Many companies have big networks of computer systems, connected through LAN's. This brings an important feature called Distributed Computing. However a network infrastructure is not enough to have this system efficiently operational. The Team Computing concept aims to realise high level cooperation between computers servicing both end-user and system manager. This is achieved by tools like Task Brooker, Passwd/Etc, Omniback, etc which form the basis for the future Distributed Computing Environment of OSF. In this talk the concept of Team Computing and the tools which come with it are covered.

# EurOpen Autumn Conference Abstracts

Here are the abstracts of the papers the were delivered at the EurOpen conference held in Nice in October. Copies of the proceedings may be obtained from EurOpen at Owles Hall at a cost of £25.

Thanks are due to Stuart McRobert <sm@ic.doc.ac.uk> who typeset the proceedings and provided the abstracts.

### Project Hygiene

*Vic Stenning*

*Anshar Limited,*
*Camberley,*
*England.*

It is suggested that many of the difficulties encountered by systems and software projects are not the result of deep technical problems, but rather arise from a lack of basic *project hygiene* – from failure to enforce various elementary principles and disciplines that are self-evidently prerequisite to a successful outcome. Such disciplines are primarily concerned with the control and co-ordination of both project *activities* and project *products*.

Examination of some of the elementary principles of project hygiene suggests that the principles are often violated. However, while the failings may be obvious, the means for overcoming such failings are somewhat less so. Suggestions are made both on methods and procedures that might improve the standard of hygiene within a typical project, and on topics that demand particular attention.

Of course, simply focusing on hygiene will not of itself ensure project success. The successful development of computer systems and software demands a co-ordinated process incorporating effective methods and supported by effective tools. However, basic hygiene must be at the very heart of the process, and is essential to creating the conditions under which the various methods and tools can usefully be deployed.

Ultimately, project hygiene is rather like any other form of hygiene: nothing spectacular comes from its presence, but the effects of its absence can be dramatic.

### Washing Behind Your Ears: Principles of Software Hygiene

*Nixdorf Computer Engineering Corporation*

*David M. Tilbrook*
*John McMullen*

*Nixdorf Computer Canada Ltd.*
*Suite 1800*
*2235 Sheppard Ave. East*
*Willowdale,*
*Ontario,*
*Canada*

This paper presents a discussion of the objectives of and impediments to software hygiene, and a list of suggestions to be followed. The suggestions are aimed at mid-sized projects, although we believe the basic ideas are applicable to all sizes of efforts.

This paper is prepared to accompany Vic Stenning's "Project Hygiene". We borrow heavily from that paper with respect to the basic theme of achieving better project quality through the application of some simple principles. We apply Stenning's Principles and Suggestions to source management, as well as developing our own disciplines of software hygiene.

### Driving the Software Release Process with Shape

*Ulrich Pralle*

*Technische Universität Berlin*
*Sekretariat FR 5-6*
*Franklinstr. 28/29*
*D-1000 Berlin 10*
*uli@coma.cs.tu-berlin.de*

While conventional UNIX tools provide acceptable support for elementary software management functions, such as version control and basic configuration management, there is a substantial lack of support for higher level system management functions, such as release preparation and change management for complex software systems. One major reason for the development of the *Shape* toolkit was to improve this situation by providing a close integration of basic software object management and configuration control functions.

The toolkit is a set of version control commands and *Shape*, a configuration management tool, which are integrated on top of a dedicated version object base. Shape uses a rule based approach to identify and select particular version objects in the object base by means of arbitrary attributes. It records identified configurations in unique *configuration identification documents* making it easy to maintain histories of entire software systems. Tool-, variant-, and version selection definitions are specified together with object dependencies as versionless *abstract system models* in *Shapefiles*.

While it is generally easy to write Shapefiles for small (one programmer) projects, this task can be tricky for medium sized (or large) projects consisting of many subsystems developed and maintained by several programmers. In this paper we describe an easy to use yet powerful, user customizable product management system on top of the Shape toolkit. It enables individual programmers to maintain and develop subsystems of a larger project without loss of overall project integration. Administration, maintenance, and installation are considerably simplified by providing standardized functionality. Although the management system requires programmers to follow certain conventions, it provides good support for software system management.

### A Multi-Site Software Development Environment

*M. O. Pierron*
*P. Zerlauth*

*Alcatel Business Systems*
*Software Technical Department*
*1, route du Dr. Albert Schweitzer*
*67408 Illkirch Cedex*
*France*

Atelix has been the operational distributed software development environment used by Alcatel Business Systems, Strasbourg for software development since 1986. In 1989, a multi-site version of the environment (Atelix/multi-site) was introduced. Atelix/multi-site is an extension of Atelix. The aim of this article is to set down the features of Atelix: configuration management, project management and network transparency. It will describe the characteristics and the setting up of the Atelix/multi-site software development environment.

### Lynx: The Object-Oriented Inode Eater

*Andy Bartlett*
*Imperial Software Technology Ltd*
*95 London Street*
*Reading RG5 1QA, UK*
*ab@ist.co.uk*

*Greg Reeve*
*White Horse Computing Ltd*
*11a Charlieu Avenue,*
*Calne, Wiltshire, SN11 0ZD, UK*
*greg@ist.co.uk*

When development trees proliferate, source management has little to do with the logistics of the enormous amounts of disk space required to support everyone's personal development environment. Integration becomes a nightmarish mix of version skew and "first night" release nerves. When all the parts of the software are returned to the master sources, will they even clean build, let alone run as they had before the developer handed them over?

Lynx applies a chain-saw to this forest of Amazonion proportions.

### The Answer to All Man's Problems

*Tom Christiansen*
*CONVEX Computer Corporation*
*POB 833851*
*3000 Waterview Parkway*
*Richardson, TX 75083-3851*
*tchrist@convex.com*

The UNIX on-line manual system was designed many years ago to suit the needs of systems at the time, but despite the growth in complexity of typical systems and the need for more sophisticated software, few modifications have been made to it since then. This paper presents the results of a complete rewrite of the man system. The three principal goals were to effect substantial gains in functionality, extensibility, and speed. The secondary goal was to rewrite a basic UNIX utility in the perl programming language to observe how perl affected development time, execution time, and design decisions.

Extensions to the original man system include storing the whatis database in DBM format for quicker access, intelligent handling of entries with multiple names (via .so inclusion, links, or the NAME section), embedded tbl and eqn directives, multiple man trees, extensible section naming possibilities, user-definable section and sub-section search ordering, an indexing mechanism for long man pages, typesetting of man pages, text-previewer support for bit mapped displays, automatic validity checks on the SEE ALSO sections, support for compressed man pages to conserve disk usage, per-tree man macro definitions, and support for man pages for multiple architectures or software versions from the same host.

### The 2.6 MSD Software Development Environment

*Ed Gould*
*Bradley White*
*MT XINU*
*2560 Ninth Street*
*Berkeley, California 94710*
*USA*
*ed@mtxinu.com*

2.6 MSD is MT XINU's recently-released distribution of the Mach Operating System and other software. The distribution contains a complete BSD user environment, the X Window System, and some user-contributed software, in addition to Mach. The distribution supports three hardware families, the Vax, the Sun 3, and the IBM RT/PC.

The software development environment used to build and maintain this distribution – originally developed at Carnegie Mellon University – has a number of interesting properties. Among them are:

- Support for staged release levels
- RCS-based tools for source tree maintenance
- Release management and notification tools
- A compilation environment that includes viewpathing

This paper will detail the important features of the tools, and how we used them in producing the existing distribution. It will also describe our plans to use these same tools to concurrently maintain and support this release while developing the next one.

### ATK + 8859 = Multi-Lingual Text and Mail A Study in Expanding the Andrew Toolkit

*Thomas Neuendorffer*
*Information Technology Center*
*Carnegie Mellon University*
*Pittsburgh, Pa. 15213*
*tpn+@andrew.cmu.edu*

The Andrew Toolkit and Andrew Message system have been available on the X distribution tape since X11 release 2, providing multi-font text with embedded objects (drawings, rasters, animations, etc.) that can be edited, printed or mailed between users. The recent addition of the ISO 8859 fonts to the X-distribution has prompted the ITC to add some additional support to allow multi-lingual text and mail that takes advantage of the expanded character set.

This paper provides an overview of what facilities are available, and how they can be customized to fit the needs of a particular site. The goal of this paper is to illustrate some of the features of ATK that allow it to be adapted to specific needs or purposes. There is also a technical overview of what was involved in adding the ISO support to ATK while maintaining backward compatibility.

### NFS Servers for Filesystems on Optical Disks

*Hans Peter Klunder*
*iXOS Software GmbH*

*jan@ixos.uucp*

Optical disks are becoming more wide spread in diverse applications. Pressed CDROM-, writeable WORM- and rewriteable MO-disks are being used to replace paper, microfiches and magnetic tapes. This has given rise to the problem of simple and efficient access to the optical disks.

Our solution is based on the NFS protocol. An NFS Server maintains filesystems on the disks and presents them to clients in a network, using the NFS protocol and thus hiding the hardware and its peculiarities. The client kernel hides the communication with the Server, and applications enjoy the comfort of filesystems without perceiving that file accesses are satisfied by a user-level Server.

This paper presents the NFS servers we developed for each of

### Potential Pitfalls of a Distributed Audit Mechanism

*Chii-Ren Tsai*

*VDG Inc.*
*4901 Derussey Parkway*
*Chevy Chase*
*Maryland 20815*

*Virgil D. Gligor*

*Electrical Engineering Department*
*University of Maryland*
*College Park, MD 20742*

*Matthew S. Hecht*

*IBM Corporation*
*18213A24, 800 N. Frederick Ave.*
*Gaithersburg, MD 20789.*

*{crtsai,gligor}@eng.umd.edu*
*uunet!pyrdc!ibmsid!{crtsai,gligor,hecht}*

We identify several generic problems of distributed audit, such as cyclic generation of audit records, user ID mapping, and a consumer-producer problem for audit trails, that are not addressed by previous work. To study these problems, we built an experimental prototype of a distributed audit mechanism with centralized control based on the AIX 2.2.1 audit subsystem [IBM88a, Hec88a], TCP/IP [Tsa89a], Network File System (NFS) [San85a], and Distributed Services (DS) [Sau87a]. The prototype includes a central *auditor* role that can invoke and revoke auditing for remote hosts, can locate the *audit trail server* where audit trails are dumped, can perform audit system management, such as dynamically adding or deleting audit events on a per-user, per-group, or per-system basis, and can query the audit status of remote hosts. We highlight the AIX audit subsystem, present details of the implemented prototype, propose solutions to the identified problems, and discuss several open issues.

### The Xprint printing system

*P. Seghin*
*D. Buyse*

*Siemens Software S.A.*
*Rue de Neverlee, 11*
*B5810 Namur (Rhisnes)*
*seg@swn.siemens.be*
*dbu@swn.siemens.be*

The penetration of UNIX systems in large commercial environments is now a fact, and industry observers foresee an even higher growth for the future. At the same time, computer sites interconnect more and more heterogeneous machines from different manufacturers but with one common denominator: they all run a particular version of UNIX.

It is important to be able to access resources from any of the machines, quickly and reliably. The differences between machines and methods of accessing the resources on them should be invisible to the user. This does not mean, however, that such differences are invisible to the administrators of the organization or individual machines.

In such a situation, a printing system is no exception. Most of the existing UNIX spoolers although they are distributed, lack the administration and configuration functions required by a professional printing environment.

Therefore, Siemens has developed a networked printing system, Xprint, that is distributed, portable on the widest range of architectures, easy to install, use and administer, supports the most commonly used printer types and offers modern user interfaces (toolkit, command line and graphical interfaces). Furthermore, Xprint respects the ECMA [Ass88a] standards and uses XPG3 interfaces by preference.

### Computer Emergency Response – An International Problem†

*Richard D. Pethia*
*Kenneth R. van Wyk*

*Computer Emergency Response Team / Coordination Center*
*Software Engineering Institute*
*Carnegie Mellon University*
*4500 Fifth Avenue*
*Pittsburgh, PA 15213-3890*
*U.S.A.*

Computer security incidents during the past few years have illustrated that unauthorized computer activity does not obey traditional boundaries (e.g., national, network, computer architecture). Instead, such activity frequently crosses these boundaries not just once, but several times per incident [Sto89a].

International cooperation among computer security response groups can be an effective means of dealing with computer security issues faced today by the computer user community. This paper addresses the need for such cooperation and suggests methods by which individual computer security response groups can work together internationally to cope with computer security incidents.

### Distributed Computing for the Technical Workplace

*Daniel E. Geer, Jr.*
*George A. Champine*

*External Research Program*
*Digital Equipment Corporation*
*{geer,champine}@crl.dec.com*

A quantum change in the way in which technical computing is *best* done is well under way. Hardware progress (such as cheap cpu cycles), software progress (such as network transparent window systems), and political progress (especially the appreciation of open systems) represent substantive differences between the world of even two years ago and the world now possible. Project Athena has spent the last six years constructing one of the best examples of distributed computing environments. At M.I.T., ten thousand users are sharing a pool of thirteen hundred workstations and one hundred various servers; all figures are growing in real time. This talk will take

our experience in creating a network services model of computation subject to the real life constraints of the academic setting. Particular emphasis will be placed on wide-area systems management and the maintenance of a coherent computing environment on a large-scale, heterogeneous computing infrastructure.

We will further claim that the model Athena has settled on is the winning model – in fact that it has already won. The organizations that can grasp the idea that a particular computational activity is really a boat in a common sea of network services will survive; those that rely on single-source, stand-alone hardware will not.

**An X11-based Multilevel Window System Architecture**

*Mark E. Carson*
*Janet A. Cugini*

*IBM Secure Workstations Department*
*Building 182/3F43*
*800 N. Frederick Avenue*
*Gaithersburg, MD 20879*
*uunet!pyrdc!ibmsid!{carson,cugini}*

We have designed and are implementing a prototype multiple-security-level window system based on AIX and the X Window System [Sch86a]. This system supports discretionary and mandatory access control and information labeling on windows, properties, events and other objects. Our goal is a system which will meet a variety of user and government security requirements, in particular the U.S. Compartmented Mode Workstation Requirements [Woo87a], while maintaining upward compatibility with the current X protocol [Sch90a] and the Inter-Client Communications Conventions Manual [Ros89a]. Our approach is to implement only the basic security policy in the X server, distributing complicated or controversial aspects to special privileged clients. One result is that the server needs only the most minimal operating system security support, making feasible its use on machines such as X terminals. In this paper, we describe the server security policy and discuss the role played by the display, window, selection, and audit managers.

**Mounting filesystems in a networked environment**
**or**
**Why do I have to wait for *that* server to come up?**

*Christer Bernerus*

*Department of Computer Science*
*Chalmers University of Technology*
*S-412 96 Gothenburg*
*Sweden*
*bernerus@cs.chalmers.se*

The typical file system hierarchy dates back to the really old UNIX ages, long before the invention of networked file systems. This hierarchy has only been slightly revised through the years. This article describes a scheme for mounting NFS and local file systems in a consistent way avoiding some of the pitfalls that the networked context introduces if the tradition is followed.

**Developing Documentation to Meet the Changing Needs of UNIX Users**

*Linda Branagan*
*Marilyn J. Lutz*

*Convex Computer Corporation*
*3000 Waterview Parkway*
*Richardson, TX 75080 USA*
*branagan@convex.com*

In recent years, the UNIX operating system has found its way into fields ranging from medicine to mechanical engineering. Some UNIX users are no longer necessarily computer scientists – often they are educated and employed in other fields. To these users, UNIX is simply a tool through which an important set of tasks is accomplished. As a result, they apply a different set of strategies when they use software and the documentation that accompanies it.

Traditional forms of UNIX documentation do not serve the needs of these users particularly well. This paper discusses methods for developing documentation that does meet their needs.

**DIAMOND –**
**An Object Oriented Graphics Library**
**for Software Development Environments**

*Ralph Zainlinger*
*Gustav Pospischil*

*Technical University Vienna*
*Treitlstr. 3/1182/1*
*A-1040 Vienna*
*Austria*
*ralph@vmars.uucp*
*gusti@vmars.uucp*

The paper presents an object oriented graphics library that is particularly suitable for the construction of user-interfaces for Software Development Environments (SDEs). The library is based upon an object oriented interaction model that has been specifically tailored to meet the requirements of software engineering and human-computer interaction.

The library along with tools that support interactive construction of user interfaces forms a toolkit that can be used as part of the infra-structure of any UNIX based SDE.

**Why Open Systems are important to the Research Community**

*Jaana Porra*
*Johan Helsingius*

*Oy Penetron Ab*
*Espoo*
*Finland*

**Why Factionalism?**

*Pamela Gray*

*Marosi*
*UK*

Faction:- "A group or clique within a larger group, party, government, organisation or the like: party strife and intrigue: dissension."

The UNIX market is growing fast on all fronts, both technical and geographical. Nevertheless, the press still seems to delight in reporting the so-called "UNIX Wars", and to concentrate on the differences in the strategies of product suppliers and the flavours of their UNIX implementations, rather than the similarities that exist amongst them.

Why is there factionalism in the UNIX market today? We will address this by breaking the question into three parts:

1. Is factionalism today stronger than it was in the past?

2. Are there real reasons for concern?

3. Are there any positives in the present situation?

We will address each of these questions in turn.

### UNIX International's System V Roadmap UNIX System V Strategy for the '90s

*Andrew Schuelke*

*UNIX International – Europe*
*Ave de Beaulieu 25*
*1160 Brussels, Belgium*
*a.schuelke@uieu.ui.org*

One of UNIX International's primary tasks is to define requirements for future releases of UNIX System V. The Roadmap is UNIX International's top-level strategic plan and contains UNIX International's requirements for approximately the next five years of System V releases.

The first half of this paper explains the significance of the Roadmap's existence, the process by which the Roadmap was created, and the selection criteria used to select System V release features. The second half presents the Roadmap System V requirements, defining four major groups of functionality centered around the following themes: Unification, Enhanced Security, Multiprocessing, and Network Computing.

### X.400 Server Toolkit on UNIX

*Hans Georg Baumgärtel*

*Softlab GmbH*
*Zamdorferstr. 120*
*D-8000 Munich 80*
*West Germany*
*hat@softlab.uucp*

The establishment of the X.400 as a new and far-reaching technology demands the adaptability of X.400 products to the needs of existing environments. There is a need for an electronic post-office in mayor environments using X.400. The X.400 Server Toolkit described in this paper is the toolkit to construct an X.400 server for this purpose. An important element of the construction of this toolkit is the interface definition in terms of abstract services.

### Setting up an X.500 Directory Service

*Dr Andrew J Findlay*

*Brunel University*
*Uxbridge, UK*
*Andrew.Findlay@brunel.ac.uk*

One of the more difficult aspects of using electronic mail is finding the address of the person you want to talk to. At present there is no complete solution to this problem: a global directory is needed.

The X.500 standard describes a distributed directory service that has the potential to solve this problem. There are now several implementations of X.500 and the Directory is developing into a useful international service.

This paper introduces the X.500 system, and describes Brunel's experience with setting up and running a directory service based on ISODE/Quipu. Software configuration is briefly discussed and the data-gathering process is described. Suggestions are given for merging data from several sources to form coherent directory entries.

### Integration of Message Handling and Directory: A System Manager's View

*Steve Kille*

*Department of Computer Science*
*University College London*
*S.Kille@cs.ucl.ac.uk*

The OSI Directory is a tool which has large potential to improve the manageability of Message Handling Systems. This sort of statement has been made frequently, but surprisingly little work has been done on the exact manner in which these two services will interact. This paper considers how the OSI Directory could be used to provide this management support. It describes this usage in terms of system management, rather than by using detailed OSI Terminology.

### Extending 4.3BSD UNIX to Support Greek Characters

*Achilles Voliotis*
*Alexios Zavras*

*National Technical University of Athens*
*Greece*

The UNIX system, traditionally suited for English-speaking environments only, has been modified to support Greek characters. This modification actually involved much work that should not have been necessary if many programs had made less assumptions. The current state of the project and the changes performed are presented, without too much technical detail. Finally, we present a discussion about the implementation of a truly language-independent UNIX system.

### The Serpent UIMS'

*Erik J. Hardy*
*Daniel V. Klein*

*Software Engineering Institute*
*Carnegie Mellon University*
*Pittsburgh, PA 15213 USA*
*erik@sei.cmu.edu*
*dvk@sei.cmu.edu*

Serpent represents a new generation of User Interface Management Systems which manage the total dynamic behavior of an interface and which allow applications to remain uninvolved with the details of the user interface. Serpent is designed to manage the specification and dynamic behavior of (relatively) arbitrary toolkits. It provides for a fixed application programmer interface across changes in toolkits. This allows an application to evolve from one toolkit to another, or even to use multiple toolkits simultaneously.

Serpent is intended to be used either with an application (in a production environment) or without an application (in a prototyping environment). Prototypes can be built interactively, tested, and rebuilt very quickly. The dialogue description language is sufficiently rich to support the prototyping of reasonably complex user interfaces. This does not preclude, however, the construction of significantly more complex interfaces when used with an application. Applications and toolkits written in either C or Ada can be used with Serpent, although the interface description mechanism is designed to be extensible to other languages.

Two portions of Serpent of particular note are presented: Glue, which provides a mechanism for importing any X widget based

on the Intrinsics, and the Dialogue Editor, which allows for the interactive construction of Serpent dialogues.

### User Interface Builders: A Compromise Platform in the User Interface War

*Zoltan Hidvegi*

*Videoton Development Institute*
*Budapest*
*Voroshadsereg u. 54*
*H-1021 Hungary*
*sztakilvidilhidvegi*

Although X has become the dominant window system in the UNIX community the user interface competition is not over at all. In fact, there is an interface war among different look-and-feel specifications. This paper considers the possibility to support multiple toolkits and multiple graphical user interfaces from user interface builders (UIBs). To accomplish that it is proposed to elaborate a user interface description language and to use it as an intermediate level between the internal structures of a UIB and the C source code generated for the application.

### Putting UNIX on Very Fast Computers or What the Speed of Light Means to Your Favorite System Call

*Michael D. O'Dell*

*Consultant*

A computer with a 250 MHz clock and built from leading-edge technology works in fundamentally different ways compared with a one-chip CMOS VLSI processor clocking at less than 50 MHz. The interactions between a UNIX implementation and its supporting hardware have always been quite subtle and remain a considerable headache for those charged with porting the system. But in addition to the imprecisions of fuzzy functional definitions for some key system facilities, the laws of physics conspire to make the marriage of very fast computers and modern UNIX systems an even more interesting challenge than it would normally be. This paper discusses some of the matchmaking necessary to achieve a matrimonious accommodation.

### tmpfs: A Virtual Memory File System

*Peter Snyder*

*Sun Microsystems Inc.*
*2550 Garcia Avenue*
*Mountain View, CA 94043*
*peter@eng.sun.com*

This paper describes *tmpfs*, a memory-based file system that uses resources and structures of the SunOS virtual memory subsystem. Rather than using dedicated physical memory such as a "RAM disk", tmpfs uses the operating system page cache for file data. It provides increased performance for file reads and writes, allows dynamic sizing of the file system while requiring no disk space, and has no adverse effects on overall system performance. The paper begins with a discussion of the motivations and goals behind the development of tmpfs, followed by a brief description of the virtual memory resources required in its implementation. It then discusses how some common file system operations are accomplished. Finally, system performance with and without tmpfs is compared and analyzed.

### UNIX pipe extension on network for creating distributed functions

*Istvan Szeker*
*Zoltan Vinceller*

*Computer Centre*
*Eotvos Lorand University*
*Budapest*
*Bogdanfy u. 10/B*
*H-1117 Hungary*
*h250sze@ella.uucp*
*h1093vin@ella.uucp*

Present work highlights a method of organisation UNIX network wide distributed functions and network system tools. The Microcomputer System Software Research Group and the Information Systems Research Department of the Eotvos Lorand University Computer Centre have been taking part in a research project dealing with distributed information systems. The task of the group has been to investigate and develop methods and support tools, good for organisation of distributed information systems. One result was introduced in EUUG Spring 1990 conference (BUD Backtrackable UNIX Data Control).

### FlexFAX – A Network-based Facsimile Service

*Samuel J. Leffler*

*Silicon Graphics, Inc.*
*2011 N. Shoreline Boulevard*
*P.O. Box 7311*
*Mountain View, CA 94039-7311*
*sam@okeeffe.berkeley.edu*

The FAX machine is now a standard piece of office equipment for businesses and is becoming more commonplace in the home. FAX machines with a computer interface are relatively expensive, but FAX modems are not. This paper describes FlexFAX, an effort to provide low-cost, easy to use, FAX communication services to computer-based users through a network-based FAX server. FlexFAX uses commercially available FAX modems and can run on any UNIX-based platform to which a FAX modem is interfaced.

### The more I find out, the less I know? NFS Fileserver Benchmarks (including one that works)

*William Roberts*

*Department of Computer Science*
*QMW, University of London*
*Mile End Road*
*LONDON E1 4NS*
*England*
*liam@cs.qmw.ac.uk*

The problems of benchmarking are many, and the problems of benchmarking a distributed system such as the NFS filesystem are more numerous still. This paper examines three existing benchmark programs, one trivial and two serious. These programs will be seen to fall into various pitfalls, but one of them manages to produce a single meaningful measure of NFS fileserver performance. The two serious benchmarks are examined in some detail and results of tests are given. The paper closes with some warnings about the successful benchmark, some thoughts on vendor support for benchmarks, and the relationship between benchmarks and the Uncertainty Principle.

UUGA
unix user group austria

AUSTRIA - UUGA
Friedrich Kofler
Schottenring 33/Hof
A-1010 Wien
AUSTRIA
Tel: +43 222 34 61 84
Fax: +43 222 310 44 62
uuga@tuvie.at

DKUUG

DENMARK - DKUUG
Mogens Buhelt
Kabbelejevej 27 B
DK-2700 Brønshøj
DENMARK
Tel: +45 31 60 66 80
Fax: +45 31 60 66 80
mogens@dkuug.dk

GUUG

GERMANY - GUUG
Dr Anton Gerold
GUUG-Vorstand
Elsenheimerstr 43
D-8000 MUŃCHEN 21
WEST GERMANY
Tel: +49 89 570 7697
Fax: +49 89 570 7607
gerold@lan.informatik.tu-muenchem.dbp.de

IRELAND - IUUG
Norman Hull
Irish UNIX-systems User Group
Court Place
Carlow
IRELAND
Tel: +353 503 31745
Fax: +353 503 43695
iuug-exec@cs.tcd.ie

NORWAY - NUUG
Jan Brandt Jensen
Unisoft A.S.
Enebakkvn 154
N-0680 Oslo 6
NORWAY
Tel: +47 2 688970

EUUG-S
N
I
X

SWEDEN - EUUG-S
Hans E. Johansson
NCR Svenska AB
Box 1206
S-164 28 KISTA
SWEDEN
Tel: +46 8 750 26 03
hans@ncr.se

USSR - SUUG
Vladas Leonas
Soviet UNIX Systems Users Group (SUUG)
24 Vavilova Str
Moscow, 117812
USSR
Tel: +7 095 939 0056

BUUG
Belgian UNIX systems Users Group

BELGIUM - BUUG
Marc Nyssen
Department of Medical Informatics
VUB, Laarbeeklaan 103
B-1090 Brussels
BELGIUM
Tel: +32 2 477 4111
Fax: +32 2 477 4000
buug@vub.uucp

FUUG

FINLAND - FUUG
Outi Nyman
Finnish UNIX Users' Group
Perustie 23 A 11
00330 Helsinki
FINLAND
Tel: +358 0 4002043

HUNGARY - HUUG
Dr Elöd Knuth
Computer and Automation Institute
Hungarian Academy of Sciences
H-1502 Budapest 112, P.O. Box 63
HUNGARY
Tel: +361 1 665 435
Fax: +361 1 354 317

i2u

ITALY - i2u
Ing Carlo Mortarino
i2u
Viale Monza 347
20126 Milano
ITALY
Tel: +39 2 2520 2530

PUUG

PORTUGAL - PUUG
Legatheaux Martens
Avenue 24 de Julho, nº 134, 7º
Lisboa
PORTUGAL
Tel: +351 1 673194/609822
Fax: +351 1 7597716
puug@inesc.pt

CHUUG
Swiss UNIX Systems Users Group

SWITZERLAND - CHUUG
Patrik Eschle
Physik University Zürich
Winterthurer str. 190
CH 8051 ZURICH
SWITZERLAND
Tel: +41 1 257 45 88
eschle@forty2.uucp

YUGOSLAVIA - YUUG
Milan Palian
Parex Institute
Kardeljeva No 8
YU-61000 Ljubljana
YUGOSLAVIA
Tel: +38 61 214223
Fax: +38 61 214223
milan@parex.yu

CZECHOSLOVAKIA - CSUUG
Sekretariat CSUUG
Výpočetní Centrum VSE
Nám. A. Zápotockého 4
130 67 PRAHA 3
CZECHOSLOVAKIA
csuug@Czechoslovakia.EU.net

AFUU

FRANCE - AFUU
Miss Ann Garnery
AFUU
11 rue Carnot
94270 Le Kremlin Bicetre
Paris
FRANCE
Tel: +33 1 46 70 95 90
Fax: +33 1 46 58 94 20
anne@afuu.fr

ICELAND - ICEUUG
Marius Olafsson
University Computer Center
Dunhaga 5
Reykjavik
ICELAND
Tel: +354 1 694747
marius@rhi.hi.is

NLUUG

NETHERLANDS - NLUUG
Patricia Otter
p/a Xirion bv
Burg. Verderlaan 15 X
3454 PE De Meern
THE NETHERLANDS
Tel: +31 3406 61 990
nluug@nluug.nl

SPAIN - UUES
Jose A Manas
Usarios de UNIX de España (UUES)
Telefonica I&D
C/Emilio Vargas 6
28043 Madrid
SPAIN
Tel: +34 1 337 42 22
jmanas@dit.upm.es

UNITED KINGDOM - UKUUG
Bill Barrett
Owles Hall
Buntingford
Hertfordshire SG9 9PL
UNITED KINGDOM
Tel: +44 763 73039
Fax: +44 763 73255
ukuug@ukc.ac.uk

EurOpen

EUROPE - EurOpen
The European Forum for Open Systems
Owles Hall
Buntingford
Hertfordshire SG9 9PL
UNITED KINGDOM
Tel: +44 763 73039
Fax: +44 763 73255
europen@EU.net