

2004  
NordU  
USENIX

DKUUG Nyt  
nr. 146

Invitation  
Program

# The sixth NordU/USENIX Conference

January 28-31 and February 1-3, 2004  
Copenhagen



**USENIX**  
THE ADVANCED COMPUTING SYSTEMS ASSOCIATION  
*fuug*



# SA Saturday 31th.

January 2004

# SU Sunday 1st.

February 2004

## 08:00-08:30 Registration

08:30

**KSA1** page 14  
**Keynote: Impact of Open Source Security Tools**  
 Wietse Venema

09:15

## 09:00-09:15 Registration

<b>SU1</b> page 23 <b>Development MySQL: Past, Present and Future</b> David Axmark	<b>SU2</b> page 23 <b>Clusters Harnessing distributed computing power</b> Carlos Justiniano
--	---

## 10:00-10:15 Break/Coffee

10:15

<b>SA1</b> page 14 <b>Development OpenCM: Secure, Distributed Configuration Management</b> Jonathan S. Shapiro, Ph.D.	<b>SA2</b> page 15 <b>Operating Systems Windows to Linux migration strategies</b> Marcel Gagné	<b>SA3</b> page 15 <b>Papers Creating virtual "soft" devices with User-mode Linux</b> Sapan Bhatia & Laurent Reveillere
---	--	---

## 10:00-10:15 Break/Coffee

<b>SU3</b> page 24 <b>Desktop Gnuskole</b> Jens Karso and Christian Hansen	<b>SU4</b> page 24 <b>Clusters Achieving Maximum Cluster Efficiency with HPC Middleware</b> Henry Gabb
--	--

## 11:00-11:15 Break/Coffee

11:15

<b>SA4</b> page 16 <b>Development Libxml and libxslt, a portable XML and XSLT toolkit</b> Daniel Vaillard	<b>SA5</b> page 16 <b>Operating Systems Enhancements to the fast Filesystem</b> Kirk McKusick	<b>SA6</b> page 17 <b>Papers Managing sensitive machine configuration on a large-scale</b> Vladimir Bahyl
---	---	---

## 11:00-11:15 Break/Coffee

<b>SU5</b> page 25 <b>Desktop The International open Office.org: How it Works</b> Louis Suarez-Potts	<b>SU6</b> page 25 <b>Clusters The Grid - the new way to use the Internet</b> Brian Vinter
--	--

## 12:00-13:15 Lunch

13:15

<b>SA7</b> page 17 <b>Development A physicist's take on version control</b> David Roundy	<b>SA8</b> page 18 <b>Operating Systems Handling Multiple User Databases on a single machine</b> Volker Lendecke	<b>SA9</b> page 18 <b>Mailsystems What's New in OpenLDAP 2.2</b> Howard Chu
--	--	---

## 12:00-13:15 Lunch

<b>SU7</b> page 26 <b>Desktop Sun a11y</b> Bill Hanemann	<b>SU8</b> page 26 <b>Clusters Commodity High Performance and Throughput Computing</b> Joseph Landman
--	---

## 14:00-14:15 Break/Coffee

14:15

<b>SA10</b> page 19 <b>Development DocBook</b> Norman Walsh	<b>SA11</b> page 19 <b>Operating Systems The Gnu Hurd</b> Marcus Brinckmann	<b>SA12</b> page 20 <b>Mailsystems Courier Mail Server</b> Sam Varshavchik
---	---	--

## 14:00-14:15 Break/Coffee

<b>SU9</b> page 27 <b>Operating Systems Calory Restricted OpenBSD</b> Wim Vandeputte	<b>SU10</b> page 27 <b>Operating Systems Security &amp; Ethics: Who's interests do we represent?</b> Hugh Daniel
--	--

## 15:00-15:15 Break/Coffee

15:15

<b>SA13</b> page 20 <b>Development Network Telescopes: Global Internet Security Threats</b> David Moore	<b>SA14</b> page 21 <b>Operating Systems DarwinPorts as a modern way to provide ported Software</b> Guldberg and Kronlage	<b>SA15</b> page 21 <b>Mailsystems What's the State of Spam?</b> Eric Allman
---	---	--

## 15:00-15:15 Break/Coffee

**KSU1** page 28  
**Keynote: Content Wars: Fighting the Gatekeepers**  
 J. D. "Illiad" Frazer

## 16:00-16:15 Break/Coffee

16:15-17:45

**KSA2** page 28  
**Keynote: UNIX, the Internet and Linux: Birthday and Crises**  
 Dr. Peter H. Salus

## 16:45-17:00 Closing Session

## 17:00-??? Beer ...at the corner bar



## Content

	<b>Welcome</b>	<b>3</b>
W	<i>Tutorials</i> Wednesday	4
T	<i>Tutorials</i> Thursday	8
F	<i>Tutorials</i> Friday	10
SA	<i>Conference</i> Saturday	14
SU	<i>Conference</i> Sunday	23
M	<i>Tutorials</i> Monday	5
TU	<i>Tutorials</i> Tuesday	5
	<b>Paper P3 paper/BoF</b>	<b>22</b>
	<b>General Information</b>	<b>29</b>
	<b>Conference Addresses</b>	<b>29</b>
	<b>Registration Guidelines</b>	<b>30</b>
	<b>Payment</b>	<b>30</b>
	<b>Scientific Information</b>	<b>31</b>
	<b>Hotel Information</b>	<b>31</b>
	<b>Committees</b>	<b>33</b>
	<b>Registration Form</b>	<b>34</b>

DKUUG Nyt nr. 146

ISSN: 1395-1440

Number printed: 8000

Managing Editor: Kristen Nielsen, DKUUG

Layout: www.seifertdesign.dk

Thanks to [www.wonderful-copenhagen.dk](http://www.wonderful-copenhagen.dk) for royalty free photos of Copenhagen.

Register by December 22, 2003  
and **SAVE!!!**

<http://www.nordu.org/NordU2004/>



## Welcome

We would like to welcome everyone to the sixth NordU/USENIX conference - NordU 2004 which this year is being hosted by DKUUG, the Danish UNIX-Systems User group, at Symbion Science Park in Copenhagen.

We are proud to present you with an exciting program, packed with news from the global UNIX community. NordU is known as the leading UNIX conference in Scandinavia and as an informal forum for system administrators, developers and UNIX professionals at large to meet, learn and exchange ideas on every aspect of computer and network management.

This year we have lowered the price for the conference considerably, in order to meet the recessing economy of Computer business, and thus hoping that people will have an extended possibility of attending the conference.

This year's opening keynote will be delivered by Wietse Venema. Wietse is probably best known for his longtime contribution to the open source community e.g., Postfix and TCP Wrapper. Wietse will be talking about security tools in open source projects. We will also be looking into different UNIX operating systems, the development within the cluster/grid community and desktop computing.

We still have a program containing 5 days of tutorials, 2 days of Conference with technical sessions and BoFs (small informal Birds-of-a-Feather sessions).

The technical sessions cover the latest development, in theory as well as in practice and the invited speakers will provide a broad selection of overviews on timely and entertaining topics.

The social events include a welcome reception, a conference dinner and a chess match between the grid based chess engine ChessBrain and the Danish grand master, who also is the sole danish qualifier for the next world championships, Peter Heine.

Once again - we hope to see you in Copenhagen for NordU 2004!

On behalf of The Organizing Committee

Kristen Nielsen

Welcome



Tutorial  
Program

W

Tutorial  
Program

Wednesday, January 28  
08.30-17.00

WTF1

08:30-17:00

**FreeBSD 4.8 Kernel Internals: Data Structures, Algorithms, and Networking**  
Dr. Marshall Kirk McKusick (3 days)

This course will provide a firm background in the FreeBSD 4.8 kernel. The POSIX kernel interfaces will be used as examples where they are defined. Where they are not defined, the FreeBSD interfaces will be described. The course will cover basic kernel services, process structure, scheduling, signal handling, and virtual and physical memory management. The kernel I/O structure will be described showing how I/O is multiplexed, special devices are handled, character processing is done, and the buffer pool is managed. The implementation of the filesystem and its capabilities including soft updates and snapshots will be described. The filesystem interface will then be generalized to show how to support multiple filesystem types such as Sun Microsystems's Network File System (NFS). The course will also cover the FreeBSD socket-based network architecture, layering and implementation. The socket communications primitives and internal layering will be discussed, with emphasis on the interfaces between the layers; the TCP/IP implementation will be used as an example. A discussion of routing issues will be included. The presentations will emphasize code organization, data structure navigation, and algorithms. It will not cover the machine specific parts of the system such as device drivers.

#### Who Should Take this Course

This course provides a broad overview of how the FreeBSD 4.8 kernel implements its basic services. It will be most useful to those who need to learn how these services are provided. Individuals involved in technical and sales support can learn the capabilities and limitations of the system; applications developers can learn how to effectively and efficiently interface to the system; systems programmers without direct experience with the FreeBSD kernel can learn how to maintain, tune, and interface to such systems. This course is directed to users who have had at least a year of experience using a UNIX-like system and the C programming language. They should have an understanding of fundamental algorithms (search-

ing, sorting, and hashing) and data structures (lists, queues, and arrays). Students will not need to prove relationship with a source license holder, as source code examples will be taken from the freely distributable FreeBSD system.

**Day 1 morning** Kernel Overview • Kernel terminology • Basic kernel services • Process structure

**Day 1 afternoon** Kernel Resource Management Scheduling • Signals • Virtual memory management

**Day 2 morning** Kernel I/O structure • Special files Terminal handling • Multiplexing I/O • Autoconfiguration strategy • Structure of a disk device driver

**Day 2 afternoon** Filesystems • Filesystem services • Block I/O system (buffer cache) • Filesystem implementation • Soft Updates • Support for multiple filesystems

**Day 3 morning** Interprocess Communication • Network File System (NFS) • IPC Concepts and terminology • Basic IPC services

**Day 3 afternoon** Networking Implementation • System layers and interfaces • Routing issues • Internet protocols (TCP/IP)


#### Course Text

Marshall Kirk McKusick, Keith Bostic, Michael J Karels, and John S. Quarterman, "The Design and Implementation of the 4.4BSD Operating System", Addison-Wesley Publishing Company, Reading, Massachusetts, 1996, 608 pages.

#### The instructor

##### Kirk McKusick

*Dr. Marshall Kirk McKusick writes books and articles, consults, and teaches classes on UNIX- and BSD-related subjects. While at the University of California at Berkeley, he implemented the 4.2BSD fast filesystem and was the Research Computer Scientist at the Berke-*





ley Computer Systems Research Group (CSRG) overseeing the development and release of 4.3BSD and 4.4BSD. His particular areas of interest are the virtual-memory system and the filesystem. One day, he hopes to see them merged seamlessly. He earned his undergraduate degree in Electrical Engineering from Cornell University and did his graduate work at the University of California at Berkeley, where he received Master's degrees in Computer Science and Business Administration and a doctoral degree in Computer Science. He is a past president of the Usenix Association, and is a member of ACM and IEEE.

In his spare time, he enjoys swimming, scuba diving, and wine collecting. The wine is stored in a specially constructed wine cellar (accessible from the Web at <http://www.mckusick.com/~mckusick/>) in the basement of the house that he shares with Eric Allman, his domestic partner of 24-and-some-odd years.

**MT1** **FreeBSD 5.X**  
**- An Intensive Code Walkthrough/2 days**  
 08:30-17:00 **Dr. Marshall Kirk McKusick (2 days)**

This course provides an in depth study of the source code of the FreeBSD 5.X kernel. This course is aimed at users with a good understanding of the algorithms used in the BSD kernel that want to learn the details of their implementation. Students are expected either to have taken a FreeBSD Kernel Internals class taught by the instructor or to have experience working with a BSD-based kernel (FreeBSD, NetBSD, OpenBSD, or BSD/OS). Students are also expected to have a complete background in reading and programming in the C programming language. This course will not cover the entire FreeBSD kernel. Rather it will focus on the specific areas outlined below.

**Day 1:**

- Overview of FreeBSD 5.X organization
- Process and thread organization
- New system daemons
- Code reading of fork, exec, exit, and wait

**Day 2:**

- The UFS2 filesystem
- Code reading of write through UFS2

The course will have some descriptive slides, but will primarily involve the instructor displaying and discussing FreeBSD kernel source code. While bringing a laptop to the class is not strictly necessary, it will be easier to follow along. Students will be given the lecture notes at the time that they sign up. They are very strongly urged to read the functions detailed in

the lecture notes prior to taking the class. Internet connection at the classroom will be available via Cable or Wireless.

About The Instructor see WTF1

**W2** **FreeS/WAN: Getting more network security out of less sysadmin work!**  
 08:30-17:00 **Hugh Daniel**

Few hosts or networks on the Internet currently have good site to site privacy or authentication, largely due to the high cost of design & implementation of classical VPN systems.


FreeS/WAN provides not only classical VPN technology, but also

"Opportunistic Encryption"(OE) that can make setting up network privacy very simple and quick and allow far more extensive network privacy than VPN's do. Besides VPN's and OE we will also cover WAVEsec (OE and VPN over wireless networks) and the design and long term maintenance of IPsec networks.

**The instructor**

**Hugh Daniel**

*Mr. Daniel has been a general systems consultant on and off for over 20 years during which he worked on The Well, Project Xanadu and helped start the Cypherpunks. He has put over five years into managing the Linux FreeS/WAN project to develop ubiquitous Internet privacy technology. He is now also the CTO of Xelerance.com, a commercial FreeS/WAN custom development and deployment company.*



*Every summer the Copenhagen Jazz Festival fills the city with joyful tunes.*



Tutorial Program  
**W**



**W3**

08:30-17:00

**New Features in MySQL**

David Axmark

- This tutorial will go through MySQL
- Which version are available and how to choose.
- How to choose between different storage engines like MyISAM, InnoDB and Heap/Memory.
- What parts of "Standard SQL" MySQL supports.
- New developments in the recent versions (4.0, 4.1 and 5.0) and some examples on how to use the new features.
- Examples of practical user setups.
- A short comparison to the 2003 SQL standard.

With over 33000 server downloads per day only through MySQL.com and an estimated 4 millions installations the MySQL Database system (TM) is one of the world's most used SQL databases.

**The instructor David Axmark**

*David is one of the Founders of both the MySQL Project and the company behind it (MySQL AB). David has been working with MySQL since well before it had a name. Before MySQL took over all time David worked as a consultant for over 15 years. The things he did included A State of the art Market research system (CommonLISP+CLOS + MySQL's ISAM) and a Advanced Business graphics package (in 32k RAM). He has written many lines of code in 6502 and Z80 assembler, BASIC, C, CommonLisp, (Bourne)-Shell and Perl. His involvement with MySQL started with the idea to upgrade a internal terminal based db tool (UNIREG) to a OpenSource SQL server. For MySQL David has worked with strategy, commercial and business aspects, installation, documentation and holding talks and tutorials. Hobbies include mountain hiking, discgolf and ultimate. David lives in Uppsala, Sweden with his plants and computers.*



**W4**

08:30-12:30

**Introduction to Perl 6**

Allison Randal

This talk is an introduction to Perl 6, the next major version of Perl. The primary focus is on the syntax of Perl 6: what's changed from Perl 5, what's the same, and what it will be like to develop in the language. Topics include operator changes, the extended importance of closures, smart matching, switches and loops, exceptions, argument binding, aliasing, properties, and exciting new features like junctions and grammars.

Beyond the syntax, we'll briefly look at some of the reasons behind the changes. What makes one feature more desirable than another? What makes it Perl-ish? What makes it DWIM? How on earth does Larry decide? You'll leave with a better understanding the changes in Perl 6, and the fundamental nature of Perl itself.

**The instructor Allison Randal**

*A Perl programmer by trade, Allison is the president of "The Perl Foundation", a contributing member of the "Perl" design team, assistant project manager for Perl 6, co-author of "Perl 6 Essentials" with Dan Sugalski and Leopold Toetsch, and founder and president of "Onyx Neon". She resides in Portland, OR (USA) with 2 programmers, a writer, and a misanthropic cat.*





**W5**

13:30-17:00

**Introduction to Parrot  
(PERL 6 Virtual Machine)**

Dan Sugalski

**The instructor  
Dan Sugalski**

*Dan is a programmer, compiler writer, and past sysadmin by day, lead designer of Parrot and member of the Perl 6 design team by night, and co-author of Perl 6 Essentials with Allison Randal and Leopold Töetsch in what he laughingly called his spare time.*



Tutorial  
Program

**W**

This talk is an introduction to Parrot, a new VM designed to run dynamic languages. While its genesis lay in the Perl 6 project, it encompasses much more, aiming to provide a common platform where Perl, Python, and Ruby, amongst other dynamic languages, can interoperate. It's designed to be easy to program, embed, extend, and write compilers for.

We'll cover the basic architecture of the Parrot engine and the core features the engine provides (including threads, continuations, coroutines, proper multithread dispatch, dynamic compilation, access to the perl 6 regular expression engine, and eventual z-machine compatibility) We'll also cover writing code directly for parrot, both in parrot assembly and PIR (Parrot Intermediate Representation) code. This is surprisingly easy, as PIR is not much simpler a language than many 80's 4GLs, though whether that says more about Parrot or old 4GLs is a topic we won't cover.



*Rust*

*A popular nightclub in two floors - even on wednesdays you can have a good time...*

Photographer: Mones Bjarnhof



*The pedestrian street "Strøget" is a favourite place of the Copenhageners to shop or just to browse.*

Photographer: Reneuz Cymek



T

Tutorial Program

# Thursday, January 29 08.30-17.00

Tutorial Program

T

WTF1

08:30-17:00

**FreeBSD 4.8 Kernel Internals: Data Structures, Algorithms, and Networking**  
Dr. Marshall Kirk McKusick

See page 4

T2

08:30-17:00

**PostgreSQL**  
Bruce Momjian

### The tutorial consists of three parts:

This talk outlines the history of PostgreSQL, why people choose PostgreSQL, its current usage, and future directions. It is a good talk to introduce people to PostgreSQL and some of its uniqueness. It is not a technical talk, but give a good overview of how PostgreSQL came to exist and why it is so special.

Replication is not a single technology but rather a mix of needs and solutions that has to be tailored to each usage for optimal performance and functionality. This talk outlines the major needs for replication and major solutions, including those currently implemented in PostgreSQL and those in development.

PostgreSQL runs primarily on Unix operating systems, but there is a large base of potential users on Microsoft operating systems. While PostgreSQL already runs under Cygwin, and native port was deemed important. This talk discusses how the decision was made to do the port, and the technical challenges of porting software of this complexity from Unix to MS Windows.

### The instructor Bruce Momjian

*Bruce Momjian is a co-founder of the PostgreSQL Global Development Group, and has worked on PostgreSQL since 1996. He is the author of PostgreSQL: Introduction and Concepts, published by Addison-Wesley. Bruce is employed*



*by Software Research Associates (Tokyo, Japan) in their PostgreSQL support division. Previously, he was vice-president of Database Development at Great Bridge LLC, another PostgreSQL support company. He has spoken at many international open-source conferences. Prior to his involvement with PostgreSQL, Bruce worked as a consultant, developing custom database applications for some of the world's largest law firms. Prior to this, he was a high school computer science teacher and holds a Masters in Education.*

T3

08:30-17:00

**Python from the Ground Up**  
Fredrik Lundh

This tutorial provides an introduction to the Python programming language and its extensive standard library. The emphasis is on "how things really work"; we'll take an in-depth look at Python's object and memory model; show how to use built-in types and functions; how to work with functions, modules, and classes; and build a few small applications along the way.

### The instructor Fredrik Lundh

*Fredrik Lundh is a principal of Secret Labs AB, the creators of the PythonWorks IDE and many other Python tools. He is an expert on the use of Python with images and graphics and is also the creator of the Python Imaging Library (PIL), Python's regular expression engine, and other parts of Python. Fredrik has been an active member of the Python community for many years, and is a frequent contributor to Python newsgroups and other fora.*





**T4****Migrating Samba 2.2 server to Samba 3.0**

08:30-17:00

Volker Lendecke

Samba 3 has finally been released. This is reason for many Samba 2 users do seriously consider an upgrade from existing installations. Especially when used as a Domain Controller Samba 3 is a lot closer to being a full replacement for Windows NT than Samba 2 used to be. Samba 3 offers among other things the ability to

- map Unix Groups to NT Groups. This is necessary to control NT Member servers correctly.
- establish trust relationships to NT domains both ways
- fully migrate existing NT domains to Samba controlled domains without any changes on the domain members

As Samba 3 offers a lot of new features, some existing configurations need to be changes slightly.

This tutorial will show the new Samba 3 features with special focus on issues people might come across during the migration of existing Samba servers and NT domains to Samba version 3.

**The instructor  
Volker Lendecke**

*Volker Lendecke is a Samba-Team member for many years now. One of his early contributions to Samba was writing the Linux client smbfs for it. The course notes for the many trainings he has given have evolved into a German book on Samba, published by the dpunkt Verlag. Late contributions of Samba have been in the area of Kerberos Authentication and AFS integration.*



*Volker is co-founder of the Service Network GmbH, a Network and Security consulting company in Göttingen, Germany. SerNet hosts the yearly Samba eXPerience, the only international conference focused completely on Samba.*

*Christianborg Castle houses the Royal Reception Rooms, the Queen's Library, the audience chambers, the Sovereign in Council rooms, Parliament, the Supreme Court and the Prime Minister's Office.*



Photographer: W&amp;Co



**F****Tutorial  
Program****Friday, January 30  
08.30-17.00****WTF1**

08:30-17:00

**FreeBSD 4.8 Kernel Internals: Data Structures, Algorithms, and Networking**  
Dr. Marshall Kirk McKusick

See Page 4

**F2**

08:30-17:00

**Logging & Security: Building an Enterprise Logging Infrastructure**  
Tina Bird

The purpose of this tutorial is to illustrate the importance of a network-wide centralized logging infrastructure, to introduce several approaches to monitoring audit logs, and to explain the types of information and forensics that can be obtained with well-managed logging systems.

Every device on your network – routers, servers, firewalls, application software – spits out millions of lines of audit information a day. Hidden within the data that indicate normal day-to-day operation (and known problems) are the first clues that systems are breaking down, attackers are breaking in, and end users are breaking up. If you manage that data flow, you can run your networks more effectively.

**Topics include:**

- The extent of the audit problem: How much data are you generating every day, and how useful is it?
- Logfile content: Improving the quality of the data in your logs
- Logfile generation: syslog and its relatives, including building a central loghost, and integrating MS Windows systems into your UNIX log system
- Log management: Centralization, parsing, and storing all that data
- Legal issues: What you can do to be sure you can use your logfiles for human resources issues and for legal prosecutions

This class won't teach you how to write Perl scripts to simplify your logfiles. It will teach you how to build a log management infrastructure, how to figure out what your log data means, and what in the world you do with it once you've acquired it.

**Who should attend**

System administrators and network managers responsible for monitoring and maintaining the health and well-being of computers and network devices in an enterprise environment. Although some review is provided, participants should be familiar with the UNIX and Windows operating systems and basic network security.

**The instructor Tina Bird**

*Tina Bird is a Computer Security Officer for Stanford University, where she works on the design and implementation of security infrastructure for University systems; writing Security Alerts for desktop and server machines on the 40000-host network; healthcare information security & HIPAA compliance; and extending the university's logging infrastructure.*

*Tina Bird moderates the Log Analysis and VPN mailing lists. With Marcus Ranum, she runs <http://www.loganalysis.org>, a portal for building enterprise logging infrastructures and interpreting log data.*



*Tina Bird graduated from the University of Notre Dame with a B.S. in physics, and has a master's degree and Ph.D in astrophysics from the University of Minnesota.*

Tutorial  
Program**F**



**F3****New Technology In Sendmail**08:30-17:00  
Eric Allman

In the last few years the face of e-mail has changed dramatically. Spam, regulation, high loads, and increased concerns about privacy and authentication have caused major changes in sendmail. This tutorial assumes that you are already familiar with the basics of sendmail, including basics of how to install, configure, and operate the system. After a very brief review of sendmail functionality and terminology, we will explore some of the newer important features, including SMTP Authentication, TLS encryption, the Milter (mail filter) interface, and many of the newer policy control interfaces.

This will be an intense, fast-paced tutorial. It is highly recommended that you have read the Sendmail book published by O'Reilly and Associates, preferably the 3rd edition (but at least the 2nd edition).


**The instructor  
Eric Allman**

*Eric Allman is the Chief Technology Officer and co-founder of Sendmail, Inc. He authored sendmail while at the University of California at Berkeley, and continues to spearhead sendmail.org, the global team of volunteers that*

*maintain and support the sendmail Open Source platform. Eric was the Chief Programmer on the INGRES database management project at U.C. Berkeley and an early contributor to the Unix effort, authoring syslog, tset, the -me troff macros, and trek. He designed database user and application interfaces at Britton Lee (later Sharebase), contributed to the Ring Array Processor project for neural-network-based speech recognition at the International Computer Science Institute, and was lead developer and provided a large-scale research software infrastructure at the Mammoth project at U.C. Berkeley. He was also Chief Technical Officer at InReference, Inc. He co-authored the "C Advisor" column for Unix Review magazine for several years, is a former member of the Board of Directors of USENIX Association, and is a member of the Editorial Review Board for ACM Queue magazine. Eric received his M.S. degree in Computer Science from the University of California at Berkeley in 1980.*

**F4****Linux Migration Challenges for the Systems Administrator**08:30-17:00  
Marcel Gagné

Congratulations! You've managed to convince management that it's time to kiss the blue screen of death goodbye and migrate your organization to desktop Linux. While you may have those highly technical issues down pat, there are still the human issues to deal with. This tutorial will concentrate on several key issues to a successful Linux desktop migration starting with the technical side of the migration itself. Some highlights include:

- Your first steps after "Yes, you may go ahead"?
- Thirty minutes to a Linux terminal server deployment.
- Full steam ahead or slow boat?
- New applications to replace your old applications
- Training methodologies to flatten the learning curve
- Human administration; making your users part of the process
- Being there electronically: remote administration

**The instructor Marcel Gagné**

*Marcel Gagné is probably best known as the award-winning author of the Linux Journal "Cooking with Linux Series", for which he received the Readers' Choice award for favorite column two years in a row. His latest book is called "Moving to Linux: Kiss the Blue Screen of Death Goodbye!". His first book was the highly acclaimed "Linux System Administration: A User's Guide". One of the best known voices of the Linux world, he has written over 200 articles on Linux and open source projects for various publications including Linux Journal, InformIT, Unix Review, SysAdmin magazine, and others. A long-time systems and network administrator, he now runs a systems and network consulting firm, is a published science fiction author and editor, a pilot, an avid science and astronomy buff, and a former top 40 disc jockey. He also folds a mean Origami T-Rex.*

Tutorial  
Program**F**



**F5**

**Building Linux Clusters**

08:30-12:30

Chris Dwan

Linux clusters are being installed by groups worldwide, when significant computational cycles are needed, and budgets are constrained. Making efficient use of clusters for a particular effort requires an understanding of what they are, how they work, and how to interact with them. As clusters are aggregations of resources, and there is freedom of choice at multiple levels within the aggregations, it is important to understand the issues, choices, and upsides/downsides to the choices. Some of the design decisions are counter-intuitive, and specific choices will impact the functionality and utility. In this tutorial, we will outline the design, implementation, and management issues, as well as discuss methods to use and leverage the power available. Mapping applications onto the cluster and various methods to profile performance and tune may be covered if time allows.

**The instructor  
Chris Dwan**

*Chris is working as bioinformatics programmer at University of Minnesota and do consultant work on Linux clusters. He holds a M.Sc. in computer science. Chris has taught a number of classes on clusters, bioinformatics, and grid computing.*



**F6**

**Web Mining - Accomplishments & Future Directions**

13:30-17:00

Dr. Jaideep Srivastava

From its very beginning, the potential of extracting valuable knowledge from the Web has been quite evident. Web mining - i.e. the application of data mining techniques to extract knowledge from Web content, structure, and usage - is the collection of technologies to fulfill this potential. Interest in Web mining has grown rapidly in its short existence, both in the research and practitioner communities. A number of new concepts, e.g. PageRank, hubs & authorities, web communities, web interestingness measures, etc., and techniques to compute them have been developed. In addition, a wide variety of commercial enterprises regularly use Web mining in their daily operations, e.g. Amazon, Yahoo, Google, etc. This talk provides an overview of the accomplishments of the field - both in terms of technologies and applications - and outlines key future research directions.

**The instructor  
Dr. Jaideep Srivastava**

*Dr. Jaideep Srivastava is a professor on the faculty of the University of Minnesota. He directs research in the areas of databases, data mining, and multimedia computing. His lab has so far graduated 17 Ph.D. and 39 MS students, with whom he has co-authored over 150 papers in journals and conferences. Between 1999 and 2001 he took a two-year leave, during which he spent time at Amazon.com and at Yodlee Inc. This wide-ranging industry experience has provided him with a unique perspective on the application of various computer science technologies in various kinds of Web-based services. As a researcher, educator, consultant, and invited speaker in the areas of data mining, databases, artificial intelligence, and multimedia for over 15 years, Dr. Srivastava continues his active collaboration with the technology industry, both for research and technology transfer. An often-invited participant in technical and technology strategy forums, Dr. Srivastava has presented at a multitude of industry, academic and government meetings. The federal government has solicited his opinion on computer science research as an expert witness. He also served in an advisory role to the governments of India and Chile on various software technologies. Dr. Srivastava received his B.Tech. in Computer Science from the Indian Institute of Technology - Kanpur, and M.S. and Ph.D. in Computer Science from the University of California - Berkeley.*



*New architecture at the Copenhagen Waterfront*

Photographer: Ireneusz Cyranek

Tutorial Program

**F**



**F7****Introduction to Host Configuration and Maintenance with Cfengine**

13:30-17:00

Mark Burgess

Cfengine is a tool for setting up and maintaining a configuration across a network of hosts. It is sometimes called a tool for "Computer Immunology"--your computer's own immune system. You can think of cfengine as a very high level language, much higher-level than Perl or shell, together with a smart agent. The idea behind cfengine is to create a single "policy" or set of configuration files that describes the setup of every host on your network, without sacrificing their autonomy.

Cfengine runs on every host and makes sure that it is in a policy-conformant state; if necessary, any deviations from policy rules are fixed automatically. Unlike tools such as rdist, cfengine does not require hosts to open themselves to any central authority, nor to subscribe to a fixed image of files. It is a modern tool, supporting state-of-the-art encryption and IPv6 transport, that can handle distribution and customization of system resources in huge networks (tens of thousands of hosts). Cfengine runs on hundreds of thousands of computers all over the world.

**Who should attend**

System administrators with a minimal knowledge of a scripting language who wish to start using cfengine

to automate the maintenance and security of their systems. UNIX administrators will be most at home in this tutorial, but cfengine can also be used on Windows 2000 and above.

**Topics include:**

The components of cfengine and how they are used  
How to get the system running  
How to develop a suitable policy, step by step  
Security Examples  
How to customize cfengine for special tasks

**The instructor Mark Burgess**

Mark Burgess is a professor at Oslo University College and is the author of *cfengine*. He has been researching the principles of network



and system administration for over ten years and is the author of *Principles of Network and System Administration* (John Wiley & Sons). He is frequently invited to speak at conferences.

Tutorial Program

**F**

At Christmas the buildings and streets of Copenhagen get decorated. Here, the designer department store *Illums Bolighus* on *Amager Torv/Strøget*.



SA

Conference Program

Saturday, January 31  
08.30-17.00

KSA1

08:30-10:00


Keynote: Impact of Open Source Security Tools  
Wietse Venema

Wietse Venema is known for his software such as the TCP Wrapper and the POSTFIX mail system. He co-authored the SATAN network scanner and the Coroner's Toolkit for forensic analysis. Wietse received awards from the System Administrator's Guild (SAGE) and from the Netherlands UNIX User Group (NLUUG). He completed his two-year term as the chair of the international Forum of Incident Response and Security Teams (FIRST). Wietse has a Ph.D. in physics and is a research staff member at the IBM T.J.Watson research center in the USA.

**The Speaker**

**Wietse Venema**

*Wietse Venema discusses lessons learned from the software that he released over the years. This includes how the software came into being, the widely varying publicity that his work received, and the impact his work had on open source and security.*



SA1

10:15-11:00

OpenCM: Secure, Distributed Configuration Management  
Jonathan S. Shapiro, Ph.D.

Open Source software projects impose unique and challenging requirements on software development projects. Project teams are dispersed and highly mobile, crossing traditional organization boundaries, which raises both administrative and legal issues. With its new support for replication, and its existing facilities for inter-organizational collaboration, strong integrity checks, fine-grain access controls, and traceability, OpenCM now fully addresses these challenges.

In this talk we describe the new distributed development features of OpenCM, and for the first time provide a direct comparison against the feature sets of alternative tools. We will present the original motivating scenarios for the tool, and describe how the features of OpenCM address these scenarios. OpenCM 1.0 will be released at the conference, and CDs will be available.

**The Speaker Jonathan S. Shapiro, Ph.D.**

*Jonathan Shapiro, Ph.D. is an Assistant Professor at Johns Hopkins University. His current research focuses on operating system security and supporting tools. He is the creator of both the EROS*



*operating system and OpenCM. In previous lives, he has worked on compilers and software development tools, and was the primary architect and developer for what is now the SGI ProDev Workshop. His book, "A C++ Toolkit," was the first book on developing reusable code in C++. As an entrepreneur, he has founded three companies and served as CEO during the recovery of a fourth. He is presently juggling eight graduate students, one dog, a family, and far too many software projects.*

Conference Program

SA



**SA2****Windows to Linux migration Strategies**

10:15-11:00

Marcel Gagné

Desktop Linux is going mainstream. Governments, municipalities, non-profit organizations, and (yes) businesses are migrating to Linux. Should your organization be doing the same? Using recent implementations and case studies, Marcel Gagné, author of "Moving to Linux: Kiss the Blue Screen of Death Goodbye!", will look at the strategies behind successful migrations. He'll cover the challenges of saying goodbye to Microsoft from a technical standpoint and, perhaps more importantly, a human and "cultural" perspective. You'll learn about the benefits and, just as importantly, the costs. With the right tools and the right information, you'll discover how to make the transition with confidence and reap the benefits of moving to Linux.

Congratulations! You've managed to convince management that it's time to kiss the blue screen of death goodbye and migrate your organization to desktop Linux. While you may have those highly technical issues down pat, there are still the human issues to deal with. This tutorial will concentrate on several key issues to a successful Linux desktop migration starting with the technical side of the migration itself.

#### The Speaker Marcel Gagné

Marcel Gagné is probably best known as the award-winning author of the Linux Journal "Cooking with Linux Series", for which he received the Readers' Choice award for favorite



column two years in a row. His latest book is called "Moving to Linux: Kiss the Blue Screen of Death Goodbye!". His first book was the highly acclaimed "Linux System Administration: A User's Guide". One of the best known voices of the Linux world, he has written over 200 articles on Linux and open source

projects for various publications including Linux Journal, InformIT, Unix Review, SysAdmin magazine, and others. A long-time systems and network administrator, he now runs a systems and network consulting firm, is a published science fiction author and editor, a pilot, an avid science and astronomy buff, and a former top 40 disc jockey. He also folds a mean Origami T-Rex.

**SA3****Creating virtual "soft" devices with User-mode-Linux**

10:15-11:00


Sapan Bhatia &amp; Laurent Reveillere

Developing device drivers can be highly tedious as it entails direct interaction with hardware devices, which are difficult to analyse in trying to find the cause of unexpected behaviour. User-mode Linux simplifies this task by allowing developers to test and debug their device drivers in user-space.

In this paper, we describe a systematic approach to create virtual "soft" devices for the purpose of testing device drivers while they are still in the stage of development. Soft devices run as user-space processes and can have a GUI interface. We have used the existing emulation capabilities of User-mode Linux and extended them by adding some of our own. We have designed a language named "Saint" to specify soft devices, and implemented a virtual coffee-machine soft device as a proof of concept.

#### The Speakers

Sapan Bhatia is doing a PHD jointly with the COMPOSE group of INRIA and the University of Bordeaux I. His current research interests lie in the areas of Operating Systems, Network Systems and Program Specialization.



Laurent Réveillère received his PhD from the University of Rennes I in 2001. Currently, he holds the position of Associate Professor in the department of Telecommunications of Enseirb, an engineering school located in Bordeaux (France). He is also a research scientist in the COMPOSE group at INRIA. Laurent Réveillère's research interests lie in the areas of programming languages, program analysis and transformation, software engineering and operating systems.

Conference  
Program

SA



**SA4**

11:15-12:00

**Libxml and libxslt, a portable XML and XSLT toolkit**  
Daniel Veillard

This presentation will present the libxml2 and libxslt libraries which provide a relatively complete and portable toolkit for the processing of XML resources. After a tour of the different specifications supported by the toolkit, it will show some simple examples of code using those libraries. We will conclude with some performance and deployment data and some guidelines on the expected future development done for the toolkit.

**The Speaker**

**Daniel Veillard**

*Daniel Veillard is ex co-chair of the W3C XML Linking WG, member of the XML Core WG, main implementer of the libxml and libxslt C and Python libraries implementation. He is also implementer and maintainer*



*of the rpmfind services. Veillard is a Member of the Board of the Gnome Project. PhD in distributed system, Grenoble 96.*

**SA5**


11:15-12:00

**Enhancements to the fast Filesystem to Support Multi-Terrabyte Storage Systems**  
Dr. Marshall Kirk McKusick

This talk describes a new version of the fast filesystem, UFS2, designed to run on multi-terabyte storage systems. It gives the motivation behind coming up with a new on-disk format rather than trying to continue enhancing the existing fast-filesystem format. It describes the new features and capabilities in UFS2 including extended attributes, new and higher resolution time stamps, dynamically allocated inodes, and an expanded boot block area. It also describes the features and capabilities that were considered but rejected giving the reasons for their rejection. Next it describes changes that were made to the soft update code to support the new capabilities and to enable it to work more smoothly with existing filesystems. The talk covers enhancements made to support live dumps and changes made to filesystem snapshots needed to avoid deadlocks and to enable them to work efficiently with multi-terabyte filesystems. Similarly, it describes changes that needed to be made to the filesystem check program to work with large filesystems. The talk concludes with some comments about performance, and describes areas for future work including an extent-based allocation mechanism and indexed directory structures.

**The Speaker Kirk McKusick**

*Dr. Marshall Kirk McKusick writes books and articles, consults, and teaches classes on UNIX- and BSD-related subjects. While at the University of California at Berkeley, he implemented the 4.2BSD fast filesystem and was the Research Computer Scientist at the Berkeley Computer Systems Research Group*



*(CSRG) overseeing the development and release of 4.3BSD and 4.4BSD. His particular areas of interest are the virtual-memory system and the filesystem. One day, he hopes to see them merged seamlessly. He earned his undergraduate degree in Electrical Engineering from Cornell University and did his graduate work at the University of California at Berkeley, where he received Master's degrees in Computer Science and Business Administration and a doctoral degree in Computer Science. He is a past president of the Usenix Association, and is a member of ACM and IEEE.*

*The transportation system in the Greater Copenhagen Area operates with frequent fares and is highly efficient and modern.*



Photographer: WICo

Conference Program  
**SA**



**SA6**

11:15-12:00

**A system for managing sensitive machine configuration information on a large-scale**  
Vladimir Bahyl

This paper discusses the design and implementation of a system for a large-scale management of sensitive machine configuration information. By sensitive machine configuration information we mean password and configuration files as well as key-pairs.

The paper doesn't talk about yet another account management system nor it describes ready to be installed tool. It rather inspects our unique approach in building the trusted relationship with bare nodes in a large cluster and distributing sensitive information to them.

We explain the problem in depth, we mention constraints we had to face, we describe tools we used for our solution and we also dedicate a significant part of this paper to implementation details. As a part of the conclusion we mention possible enhancements for the future.

### The Speaker Vladimir Bahyl

Vladimir Bahyl (Age: 29 Nationality: Slovak) After obtaining a Masters Degree in Computer Science from the Czech Technical University in Prague in 1998, Vladimir continued his UNIX career at UNI-C in Copenhagen. There he extended his knowledge to administer 8 different UNIX flavors but also participated in several projects that required Perl skills. In 2001 he moved to CERN in Geneva. Here his work allowed him to refocus on 1 operating system, but on a large scale. He is currently a member the team that designs the architecture of many interactive, batch and other special clusters in the CERN's computer center. In total, these clusters consist of around 1500 nodes.



Wall Signs are pretty common at the gable ends of the old buildings in Copenhagen and add to the city atmosphere.

**SA7**

13:15-14:00


**A physicist's take on version control**

David Roundy

What happens when a physicist becomes interested in version control? He develops a theory for how patches interact, of course! This talk will introduce the theory of patches which underlies darcs. It will then show how this theory is used to create a flexible revision control system. The design decisions behind darcs will be discussed, in particular where those decisions are (and are not) dictated by the underlying theory.

### The Speaker David Roundy

David Roundy obtained his PhD from U.C. Berkeley in 2001 in the field of computational condensed matter theory, and has since worked as a postdoctoral researcher at MIT in the field of photonics. He writes free software during his spare time. He also knits, plays classical guitar and plays bridge at the local bridge club.

Conference  
Program

SA



**SA8**

13:15-14:00

**Handling Multiple User Databases on a single machine**  
Volker Lendecke

One area where Unix has room for improvement is the handling of the user databases in large environments. For modern Unices you have a 32-bit space for numeric UIDs and GIDs. This is more than enough for a single box. If you plan well, this can handle even the largest installations across many locations. However, trouble starts when you have to merge user databases from different machines or companies. Then you will inevitably have to face and resolve conflicts.

This is an area where the Unix world might learn from the way Windows does it: Windows allocates IDs out of a 128-bit space. 32 bits of those are comparable to the Unix ID space. The other 96 bits are an identifier of the issuing user database. This way in all the places where User and Group IDs are used, especially in the file system, any number of user databases can be used concurrently on a NT system.

This talk will show the underlying mechanics of Windows, the Samba way to map the 128 Bit space onto the 32 Bits and the problems Samba has to face when mapping the large into a small number space.

**The Speaker Volker Lendecke**

*Volker Lendecke is a Samba-Team member for many years now. One of his early contributions to Samba was writing the Linux client smbfs for it. The course notes for the many trainings he has given have evolved into a German book on Samba, published by the dpunkt Verlag.*



*Late contributions of Samba have been in the area of Kerberos Authentication and AFS integration. Volker is co-founder of the Service Network GmbH, a Network and Security consulting company in Göttingen, Germany. SerNet hosts the yearly Samba eXPerience, the only international conference focused completely on Samba.*

**SA9**

13:15-14:00

**What's New in OpenLDAP 2.2**

Howard Chu

There are a number of significant new features in OpenLDAP 2.2 including a new database backend, new replication mechanism, new plugin support, and more. This talk will present these new features along with a look at the internal design changes that made the new features possible.

**The Speaker Howard Chu**

*Howard is an OpenLDAP core team member and has been working with OpenLDAP since 1998. He specializes in security and distributed computing technologies and has been a contributing developer on many projects/technologies over the years including*



*GCC, GNUmake, Kerberos, OpenSSL, Cyrus SASL, and others. Howard's contributions to OpenLDAP are sponsored by his company Symas Corporation, which was founded on the principles of open standards and open source. In addition to being Chief Architect at Symas, Howard is also an accomplished musician.*

*Copenhagen's colorful old buildings bring out historic remembrance.*



Photographer: Cees van Reedin

Conference Program

SA



**SA10****DocBook**

14:15-15:00

Norman Walsh

The description of this speak was not available when the layout finished. Please find the information on the webpage.

**SA11****The GNU Hurd**

14:15-15:00

Marcus Brinkmann

The GNU Hurd, the GNU project's replacement for the UNIX kernel, is the only Free Software project that aims to develop a POSIX compatible operating system based on the multi-server philosophy.

We think that the monolithic UNIX kernel design used in the BSD and Linux kernels not only limits flexibility, scalability and security but unnecessarily complicates maintenance. The Hurd servers work together in user space to provide the commonly expected POSIX and UNIX interfaces. Since the interfaces are not system calls but RPCs to different user space programs, any user can override the default system policy by choosing to use a different, possibly specialized, server without the intervention of the system administrator.

Experience with the current GNU Hurd system suggests some major improvements can be made with respect to performance. One theory is that the Mach kernel dictates too much policy especially with respect to IPC and memory management. Currently, we are exploring an alternative microkernel, L4.


Over the course of my presentation I will give an overview of the design of the Hurd system, the status of its development and a roadmap for possible future developments.

### The Speaker Marcus Brinkmann

*Marcus Brinkmann was born in and grew up in Dortmund, a town located in northwest Germany, a congested region currently undergoing a change from an industrial to a service based economy. Presently, he is a student at the Ruhr-Universität Bochum studying mathematics and physics.*

*At the age of seven, Marcus discovered computers and programming soon became a hobby. It was not until the Internet replaced the German Btx, however, that he became part of the world-wide community of free software developers. He considers this essential to his maturity as a programmer: it not only permitted him to easily retrieve and study large amounts of real world source code under a free software*

*license, but made it possible to communicate and eventually collaborate with the programmers. After becoming familiar with the Debian GNU/Linux distribution, Marcus realized its value and decided to become a contributor. In 1997, after a long time user of GNU/Linux systems, Marcus joined the Debian project to help develop their GNU/Linux distribution. The following year, he discovered the GNU Hurd and refocused his efforts to work on a binary distribution of the GNU system. Since then, he has joined the GNU project and now works on several of its software packages, principally the GNU Hurd. He is now, together with Neal H Walfield, leading an effort to redesign the GNU Hurd and run it on the L4 microkernel.*



*With its cafes and fishing boats, the picturesque canal wharf of Nyhavn is one of Copenhagen's most popular outdoor meeting places.*

Conference  
Program**SA**

Photographer: W&amp;G




**SA12** Courier Mail Server  
14:15-15:00 Sam Varshavchik

Sam Varshavchik is the author of the popular Courier-IMAP server, as well as it's big brother, the Courier mail server (<http://www.courier-mta.org>), which supplements the IMAP server component with SMTP, webmail, and mail filtering. Recently he also released the Cone mail client (<http://www.courier-mta.org/cone/index.html>)

**The Speaker Sam Varshavchik**  
*Born in Russia, he emigrated with his family to the United States at age 10; he presently works as a programmer/analyst consultant in New York; and supports Courier and his other software in his spare time.*



*The House of the Society of Danish Engineers (IDA) is situated at the Copenhagen Harbour and is frequently hosting different meetings.*



Photographer: Cees van Rooijen

**SA13** Network Telescopes: Assessing Global Internet Security Threats  
15:15-16:00 David Moore and Colleen Shannon

Network telescopes provide the unique ability to see large-scale globally-dispersed network security events, such as denial-of-service attacks and the spread of Internet worms. A network telescope is a portion of routed IP address space with little or no legitimate traffic. By monitoring unexpected traffic arriving at a telescope, we can determine remote victims of DoS or hosts infected by a worm. More than 100 distributed denial-of-service attacks are occurring on average every minute of every day. Highly infectious Internet worms have become prevalent: in August 2001, CodeRed infected 360,000 machines in 10 hours. In January 2003, Sapphire/SQL Slammer infected over 75,000 machines in ten minutes. This talk covers trends in DoS attacks and victims over the past 2 years, as well as the spread dynamics of the Code-Red, CodeRedII and SQL Slammer/Sapphire worms.

**The Speaker David Moore**

*David Moore is a principal investigator and assistant director of the Cooperative Association for Internet Data Analysis (CAIDA). He is also a computer science PhD candidate at the University of California, San Diego. His research interests are high-speed network monitoring, denial-of-service attacks and infrastructure security, and Internet traffic characterization.*



**The Speaker Colleen Shannon**

*Colleen Shannon is a staff researcher at the Cooperative Association for Internet Data Analysis (CAIDA). Her research interests focus on network security, including traffic measurement, characterization, and network application identification.*

*She recently completed a study of the prevalence and characteristics of IP packet fragmentation, the results of which were published in IEEE Transactions on Networking in December 2002.*





**SA14****DarwinPorts as a modern way to provide ported Software**

15:15-16:00

Ole Guldberg Jensen and Felix Kronlage

Distributing ported software is a challenge, to which quite a few solutions exist. Unlike others, DarwinPorts did not choose an existing framework as the base, but instead was designed from scratch to overcome and avoid known problems, while meeting the needs for a modern ports tree. DarwinPorts, a project of OpenDarwin.org, was started by Landon Fuller, Kevin van Vechten and Jordan K. Hubbard in 2002 and has improved significantly since then. Initially intended for Apple's Mac OS X, the DarwinPorts infrastructure does work on a large variety of platforms, mostly due to the effort to stay as portable as possible. In the course of the talk, an introduction to DarwinPorts will be given, covering installation, maintenance and the basic structure of Portfiles, the core of each port. Furthermore the internals of DarwinPorts will be explained, accompanied by coverage of the APIs and reasoning for design decisions. Last, but not least there will be a quick glance at the OpenDarwin project.

**The Speaker****Ole Guldberg Jensen**

*(olegb@opendarwin.org)* Comitter at the DarwinPorts project and has ported much of gnome 2.2 along with other ports. Since his first encounter with Linux in '94, he has been interested in unix and \*BSD. As part of that interest he has been involved in arranging BSD user group meetings in his local area.

**The Speaker Felix Kronlage**

*(fkr@opendarwin.org)* Felix Kronlage has been part of the DarwinPorts team almost from the beginning. Being part of the project lead, he has done lots of work in various areas of the project. Besides working on DarwinPorts, being a student at the University of Oldenburg and running a company his time is mostly spent working on unix technology, mainly around OpenDarwin and the various BSDs.

**SA15****What's the State of Spam?**

15:15-16:00

Eric Allman

No one needs to be told that e-mail spam is a serious problem, but some people don't truly understand how serious it is. The speaker now gets about 500 spams every day, a great many of them in character sets he can't even render, and is seeing a doubling rate of about four months. Many solutions have been proposed, falling primarily into two areas: legislative and technological.

The current state of spam will be reviewed, including some about the current legislative climate (and whether legislation has any chance of doing any good) and quite a bit about the various technologies that are being discussed and deployed. Although opinions will be offered, no conclusions will (or can) be drawn in an environment changing as quickly as we are seeing with e-mail today.

**The instructor****Eric Allman**

*Eric Allman is the Chief Technology Officer and co-founder of Sendmail, Inc. He authored sendmail while at the University of California at Berkeley, and continues to spearhead sendmail.org, the global team of volunteers that maintain and support the sendmail Open Source platform.*

*Eric was the Chief Programmer on the INGRES database management project at U.C. Berkeley and an early contributor to the Unix effort, authoring syslog, tset, the -me troff macros, and trek. He designed database user and application interfaces at Britton Lee (later Sharebase), contributed to the Ring Array Processor project for neural-network-based speech recognition at the International Computer Science Institute, and was lead developer and provided a large-scale research software infrastructure at the Mammoth project at U.C. Berkeley. He was also Chief Technical Officer at InReference, Inc. He co-authored the "C Advisor" column for Unix Review magazine for several years, is a former member of the Board of Directors of USENIX Association, and is a member of the Editorial Review Board for ACM Queue magazine. Eric received his M.S. degree in Computer Science from the University of California at Berkeley in 1980.*

Conference  
Program**SA**



**KSA2**  
16:15-17:45

**UNIX, the Internet and Linux: Birthdays and Crises**  
Dr. Peter H. Salus

UNIX and the Internet were both "born" in 1969 - they are now 35. While Linux was created in 1991, the 1.0 kernel is now 10 years old. Linux is thus just a teenager, or just a decade old.

As an operating system, UNIX was recognized as a valuable tool for the Net early on. And without the Net, Linux would never have arisen.

But UNIX, as it grew, faced the problems of licensure and lawsuits. The Internet confronts daily efforts on the part of global telecommunications companies and governments to limit its spread and use.

Linux, too, confronts the problems of corporate greed and code control.

This talk will concern this as well as making clear the futility of the current SCO efforts.

**The Speaker**


**Dr. Peter H. Salus**

*Dr. Peter H. Salus has been Executive Director of the USENIX Association and of the Sun User Group and Vice President of the Free Software Foundation, following 20 years' experience in academia and a stint at IBM's T.J. Watson Research Center.*



*Dr. Salus has written or edited over a dozen books, including A Quarter Century of UNIX, Casting the Net, the four-volume Handbook of Programming Languages, and the Big Book of IPv6 Addressing RFCs. He has been "The Bookworm" in ;login: for over a decade.*

*Amalienborg Palace is the residence of The Royal Danish Family and Copenhagen's most popular royal attraction.*



Photographer: Ireneusz Cymak

**P3**

**Cryptar: Secure, Untrusting, Differencing Backup.**  
Jeff Abrahamson and Adam J. O'Donnell

*P3 is an extra paper presentation that will replace the first cancelled speaker at the conference, or be held as BoF*

We present an algorithm and its implementation to perform secure incremental backups to untrusted servers. The algorithm avoids unencrypted data leaving the local host, but is nonetheless able to compute a set of differences between the encrypted data on the remote server and the changed local files. Significant motivation comes from the rsync algorithm and program, and our program is conceptually similar to running rsync with the remote image encrypted.

**The Speaker Jeff Abrahamson and co-author Adam J. O'Donnell**

*Jeff Abrahamson is a PhD student in Computer Science at Drexel University in Philadelphia. He has an SB in mathematics from MIT and a masters in theoretical mathematics from the University of Pennsylvania. He also spent several years working in industry. He currently works in structural pattern recognition, and with further research interests in image matching, graph theory, and the theory of algorithms.*



*Adam J. O'Donnell is an NSF Graduate Research Fellow pursuing a Ph.D. in Electrical Engineering at Drexel University. He graduated Summa Cum Laude from Drexel with a Bachelor of Science in Electrical Engineering with a concentration in Digital Signal Processing. Adam has optimized RF Amplifier subsystems at Lucent Technologies, where he was awarded a patent for his work, and has held a research position at Guardent, Inc. He has worked on several books, serving as the technical editor and contributor to "Building Open Source Network Security Tools", a contributing author on "Hacker's Challenge", and co-author of "Hacker's Challenge 2". His current research interests are in networking, computer security, and distributed systems.*



Conference Program

SA



SU

Conference  
Program

# Sunday, February 1

## 08.30-17.00

SU1

09:00-09:15

MySQL: Past, Present and Future

David Axmark

MySQL is a different from other FOSS (Free and Open-Source) projects since it was started as commercial and OpenSource from day one developed by a company (MySQL AB). Most other Free Software/Open-Source databases was either introduced by 'normal' proprietary commercial companies at the end of their commercial life or the result of university research. The main benefits of using MySQL are Speed, Robustness and Practicality/Usability.

The talk will start with a history of MySQL and then continue with an overview of the current/developing MySQL functionality (version 4.0/4.1 & 5.0).

MySQL has with these new features started to enter the Enterprise Database arena. MySQL are already used in mission critical applications at companies like Yahoo!, Cisco and Google.

MySQL AB uses a dual licensing scheme where the SAME source code is available under both a GPL and a non GPL commercial license. This unique way of mixing a normal software business with Free Software will be covered in the talk.

With over 33000 server downloads per day only through MySQL.com and an estimated 4 millions installations the MySQL Database system (TM) is one of the world's most used SQL databases.

### The Speaker David Axmark

*Joe is the founder of Scalable Informatics - a small company focusing on Linux clusters, bioinformatics and scientific computing. Joe has a Ph.D. in computational physics, and has been working with UNIX and Linux for many years.*



SU2

09:15-10:00

### Harnessing distributed computing power using Open source tools

Carlos Justiniano

This presentation discusses how to use open source tools to tap into the vast computing potential of distributed machines throughout the Internet. The presentation examines the topic using the ChessBrain project as a demonstration of what's possible. The lecture will be structured into three parts; the first is a brief historical look at distributed computation, the second is an introduction to building an online community of members who will support your project by donating the use of their machines. An understanding of community building is vital to helping to attract and retain members. The third portion of the lecture focuses on technical considerations. Attendees will gain insight into important topics such as server design, cross platform development, open standards and security.

Open source tools will be examined throughout the lecture: Apache, PHP and Perl, GCC, CVS etc...

### The lecture will feature topics covered in two published articles:

The ChessBrain Project A Global Effort To Build The World's Largest Chess SuperComputer Carlos Justiniano and Dr. Colin M. Frayn Published in the Journal of the International Computer Games Association ICGA Journal Vol. 26, No. 2, pp. 132-138. <http://www.chessbrain.net/chessbrainproject.html>

ChessBrain: a Linux based distributed computing experiment Carlos Justiniano Published in the Linux Journal September 2003 Issue 113, pp. 66-70. <http://www.chessbrain.net/linuxjournal200309.html>

### Sample sub-topics:

#### Part one: Community building

- You can't sell something you don't believe in.
- Determine your needs and prepare to convince others that your project is worth their time.
- Establish an online presence to communicate project status and to solicit the involvement of volunteers.
- Treating your members as valued customers.

Conference  
Program

SU

Continues next page



**Part two: Network building**

- Hardware considerations. Primary servers.
- Bandwidth considerations
- Cross platform software development
- Communication protocols, a look at HTTP, XML, SOAP
- Securing your network, a look at security using intrusion detection systems and protocol security using encryption.

**The Speaker Carlos Justiniano**

*Carlos Justiniano has developed software for the past 20 years. He has consulted for companies such as Disney Online and IBM, and has served both technical and managerial roles including Dir. of Software Development and Chief Software Architect for leading software companies.*



*Mr. Justiniano founded the ChessBrain project in the summer of 2001 to explore distributed computation and the game of Chess. Today, ChessBrain consist of 500-600 machines from over 50 different countries.*

**SU3**

10:15-11:00

**Gnuskole**

Jens Karsø and Christian Hansen

The GnuSkole project would like to tell about the "hands on experience" that have been collected:

- Where do we stand with Linux in schools?
- Where are the possibilities on the server side, thin clients and traditional clients.
- The reality for a schools administrator today!
- What are the needs for open source educational programs?
- What are the economical possibilities for the use of open source in schools?
- How can we help the distribution of Linux to schools?

**The Speakers Jens Karsø and Christian Hansen**

*Jens is teaching 5 grade at Hjørring Private School, and his subjects are mathematics and science. Moreover, Jens is the system administrator at the school which has 4 servers and a number of thin clients - all running Linux. In his spare time he plays quite a lot with Linux, and he is actively contributing to the GNU-Skole project.*



**SU4**

10:15-11:00

**Achieving Maximum Cluster Efficiency with HPC Middleware**

Henry Gabb

- Off-the-shelf packages that facilitate cluster configuration and management
- Improving I/O performance with parallel file systems
- Configuration options for optimum application performance
- Scheduling options for effective resource sharing and maximum resource utilization

Middleware is the software layer between the operating system and user-space applications. It allows users, administrators, and applications to interact with the cluster; in this case, a cluster designed for high-performance computing. Correct choice and optimum configuration of middleware can improve overall cluster efficiency.

Cluster efficiency refers to everything from ease-of-use to application performance to maximum resource utilization. The communication layer has a direct affect on application performance. For I/O-limited applications, a parallel file system can improve performance. The configuration options for these two middleware components will be discussed.

Several off-the-shelf software packages are available to facilitate cluster configuration, management, and resource utilization. The advantages and disadvantages of these packages will be discussed in terms of the cluster's production environment and user requirements. The presentation ends with an illustration of how job scheduling policies affect resource utilization.

**The Speaker Henry Gabb**

*Henry Gabb is a parallel applications engineer at the Intel Parallel Applications Center. Henry works with a team of engineers dedicated to improving the performance of ISV applications through parallel computing. Prior to joining Intel, Henry was Director of Scientific Computing at the U.S. Army Engineer Research and Development Center, Major Shared Resource Center, a DoD high-performance computing facility. He holds a PhD in biochemistry and molecular genetics and has several years of parallel programming experience, primarily in applications pertaining to the life sciences.*





**SU5** **The International OpenOffice.org:  
How it Works**  
11:15-12:00 Louis Suarez-Potts

From its beginning OpenOffice.org has been an international project. In this, it is not unique; all major open-source projects are international, or more accurately, transnational, with Internet-enabled collaboration freely crossing national borders. But OpenOffice.org differs from others in the structure of its international emphasis. Whereas, as in other projects, development is in English, in OpenOffice.org much of the work of support, testing, marketing, and, increasingly, localization, takes place in a user's native language. At present, two years after the first such project, virtually every major European and Asian language is represented by a native-language project. This paper examines the way in which the native-language projects work and at what the native-language projects have accomplished and the challenges they face.

**The Speaker**  
**Louis Suarez-Potts**

*The Community Manager of OpenOffice.org since its inception, Louis Suarez-Potts holds a Ph.D. in English from the University of California, Berkeley. He is currently working on an in-depth examination of the difference corporate-sponsored open source projects have made to the logic and practice of open source.*



**SU6** **The Grid - the new way to use the Internet**  
11:15-12:00 Brian Vinter

Around 10 years ago the Internet became a part the vocabulary of many non-science people, today the net is an integrated part of most peoples every day lives. When saying Internet most people actually mean the world-wide-web and now the replacement of the web is slowly being introduced - the Grid. 'The Grid' is a pure name and is meant as an analogy to the power-grid. The idea is to make access to computing-resources as simple as access to electricity. As consumers we do not care who is producing the electricity we are using as long as delivery is in compliance with the protocol, stable and cheap. 'The Grid' is meant to provide access to all possible computing resources, contrarily to today's WWW model that allows access to information only. 'The Grid' will give us access to applications, processing power, storage, instruments and many other types of resources. The ultimate purpose of the Grid is to view all computers and their accessories as one large virtual machine. The Computer Science challenges in achieving this goal are tremendous and the talk will introduce some of the many problems we are facing to achieve this end. We will also discuss fundamental issues such as the fact that the speed of light seems horribly slow in a Grid world.

**The Speaker Brian Vinter**

*Brian Vinter, MSCE, PhD, Associate Professor of Computer Science at the University of Southern Denmark, with supercomputers as his field of expertise. He is currently the director of the Nordic DataGrid Facility.*



Conference Program

SU



*The skating rink at the Tivoli Christmas Market is a popular winter activity.*



**SU7**

13:15-14:00

**Sun a11y**

Bill Haneman

Accessibility, or "Universal Access", means removing barriers to meaningful life activities. This includes not only access to the graphical desktop by blind and low-vision users and people with who cannot use a keyboard, but improved usability by people may not identify themselves as disabled (for instance older users). Accessibility is vital not only as a means of realizing the egalitarian potential of technology, but for meeting a range of legal requirements which are coming into force across the US, Europe, and globally.

Until recently, accessibility of graphical Linux and Unix systems has been quite limited in comparison to other desktops; however this situation is rapidly changing. The GNOME Accessibility Project has the goal of providing a full-featured free-and-open-source desktop which is fully accessible to people with disabilities. We will:

- present the current state of this project;
- mention relevant standards and requirements,
- demonstrate the first bundled free-and-open-source assistive technologies
- explain what platform and application developers need to do to make their systems fully accessible.

**The Speaker**

**Bill Haneman**

*Bill Haneman is architect and technical lead of the GNOME Accessibility Project (an effort which is building a freely-redistributable, fully accessible graphical computing environment for Linux and Unix) and sits on the GNOME Foundation Board of Directors.*



*He works for Sun Microsystems in Dublin, Ireland. (<http://www.gnome.org>), (<http://developer.gnome.org/projects/gap>) (<http://www.sun.com/access>)*

*Like many other squares in Copenhagen Højbro Plads is often full of people enjoying a day out.*

**SU8**

13:15-14:00

**Commodity High Performance and Throughput Computing in Research and Business**

Joseph Landman

As engineers, and scientists move toward more accurate and detailed models, a number of trends have emerged with regards to their processing needs. Data appears to grow monotonically as model accuracy and data set size scales, and as seen in the life science computing and bioinformatics community, the growth rates are sometimes exponential. More data could mean more accuracy, only if you are able to use and comprehend the data. This requires that you are able to process it in a reasonable time. If you are fighting data growth rates that exceed the Moores law power increases, you might be fighting a losing battle. To give you a better chance at solving your problems, you might look to a high performance computing (HPC) or high throughput computing (HTC) system. For a properly designed solution, these systems give you the opportunity to leverage significant power. Issues will be discussed that are relevant to solution design considerations and implementation, including best practices. Specific examples of leveraging HPC/ HTC in a business as well as research environment will be discussed.

**The Speaker**

**Joseph Landman**

*Joe is the founder of Scalable Informatics - a small company focusing on Linux clusters, bioinformatics and scientific computing. Joe has a Ph.D. in computational physics, and has been working with UNIX and Linux for many years.*



Photographer: Henning Ojannik

Conference Program

**SU**




**SU9**  
14:15-15:00  
**Calorie restricted OpenBSD - running production systems on min. hardware**  
Wim Vandeputte

In this talk, Wim takes a survey of what possibilities and restrictions there are running OpenBSD production systems on limited hardware, like the Soekris hardware platform or small VIA motherboards booting from Compact Flash or booting over PXE. When dealing with embedded hardware, several issues pop up:

- 0. overview of platforms and hardware choices. Small, smaller, smallest.
- 1. installing and pruning the OS to fit on a small medium like a compact flash
- 2. dealing with boot loader issues
- 3. performance and tuning possibilities, dealing with limited resources
- 4. crypto accelerators to the rescue
- 5. storing logfiles and tmp files on CF
- 6. remotely monitoring the system
- 7. upgrading possibilities

**The instructor**  
**Wim Vandeputte**

*Wim studied computer sciences and business management at the university of Ghent, Belgium. He has been babysitting UNIX systems since 1994 and became an OpenBSD developer in 1998. He's currently the OpenBSD booth bunny for Europe, taking care of most of the logistics for the project and organizing presence at way too many events.*




**SU10**  
14:15-15:00  
**Systems Security & Ethics: Who's interests do sysadmins represent?**  
Hugh Daniel

Being a systems admin is getting more complex as computers both become more a part of the everyday infrastructure of society and pawns in the legal & political spheres. How we as architects and admins design and operate our systems must be affected by this changing landscape.

What are the new questions, problems, challenges and opportunity's for both our users and our selves in the near future? More importantly what directions and responsibility are we as a community willing to take on in a complex world of multiple society's but a single(?) Internet?

**The instructor**  
**Hugh Daniel**

*Mr. Daniel has been a general systems consultant on and off for over 20 years during which he worked on The Well, Project Xanadu and helped start the Cypherpunks. He has put over five years into managing the Linux FreeS/WAN project to develop ubiquitous Internet privacy technology. He is now also the CTO of Xelerance.com, a commercial FreeS/WAN custom development and deployment company.*



*The Black Diamond  
Night-time view at  
The Royal Library.*

Conference  
Program

SU



**KSU1**  
15:15-16:45

**Content Wars:  
Fighting the Gatekeepers**  
J.D. "Illiad" Frazer

J.D. uses UserFriendly.org as a case study of not only why, but how, the traditional content distributors should be fought tooth and nail, given the advent of the Net. He explains the harsh realities of doing business as a content creator and covers both the upside and downside of the medium that we rely on for much of our information and entertainment.

**The Speaker**  
J.D. "Illiad"  
Frazer


*By the turn of his third decade, J.D. "Illiad" Frazer discovered that his years working as a project manager*



*left very little in the way of sanity. This left him qualified for two occupations only: politician or cartoonist. Since he still had a solid grip on his ethics, the latter choice was his only option.*

*Rated as one of the top 50 most influential people in the Linux community in the world by Linux Magazine, he quickly changed his name and appearance as people began to seek him out to answer questions regarding "channel binding" and "SYN packets." The only thing that J.D. knows about binding involves his underwear, and he readily admits to being a synner, although he'll never get used to the American spelling.*

*Often found at speaking engagements across Europe and North America, J.D. remains quite incredulous at this sudden change in his life. He thinks everyone he's met at the conferences are really nice folks, although he does wonder a bit at the people who follow him around in llama costumes. His favourite color is green but only on trucks and walls. He stopped drooling after he quit work as a project manager, and thank you for asking.*



*Next to "The Black Diamond" - the latest addition to the Royal Library - lies the impressive Søren Kierkegaard Plads with excellent views of the waterfront.*

Conference  
Program

SU

Photographer: Cees van Ierland



## G

The sixth NordU/  
USENIX Conference

# General Information

## Website

For the latest information about the conference, please visit the conference website <http://www.nordu.org/NordU2004/>

## Currency

The official currency in Denmark is Danish kroner (DKK). International Credit Cards are widely used in Denmark, including most taxis etc. Please notice that only hard currency can be used in public transportation. It is recommended to have 1-200 DKK in cash to small expenses. Some shops especially turist-shops do accept EUROS.

## Timezone

The timezone in Denmark is GMT +1 hour.

## How to get to Symbion

### Address

Symbion Science Park  
Fruebjergvej 3  
2100 Copenhagen Ø (East)  
<http://www.symbion.dk/>

### DKUUG office Address

DKUUG  
Danish UNIX-Systems User  
Group  
Fruebjergvej 3  
DK 2100 Copenhagen Ø (East)  
Telephone: +45 3917 9944  
Fax: +45 3920 8948  
e-mail: [sek@dkuug.dk](mailto:sek@dkuug.dk)  
<http://www.dkuug.dk/>

### Public transportation

Symbion can be reached by public transportation (S-train or bus). The busses 42 or 43 stops near Symbion. You can also take the S-train to the nearest station (Emdrup or Ryparken) and walk the rest of the way (10 minute walk). You can read more about getting to Symbion at the Nordu website. The webpage for public transportation in Denmark is: <http://www.rejseplanen.dk> (Travelplanner)

## Electricity

Electrical current in Denmark is 220V/50Hz. Round, European-style two-pin plugs are used. Appliances designed to operate on 100/120 Volts need a voltage converter and a plug adapter.

## Climate and dress

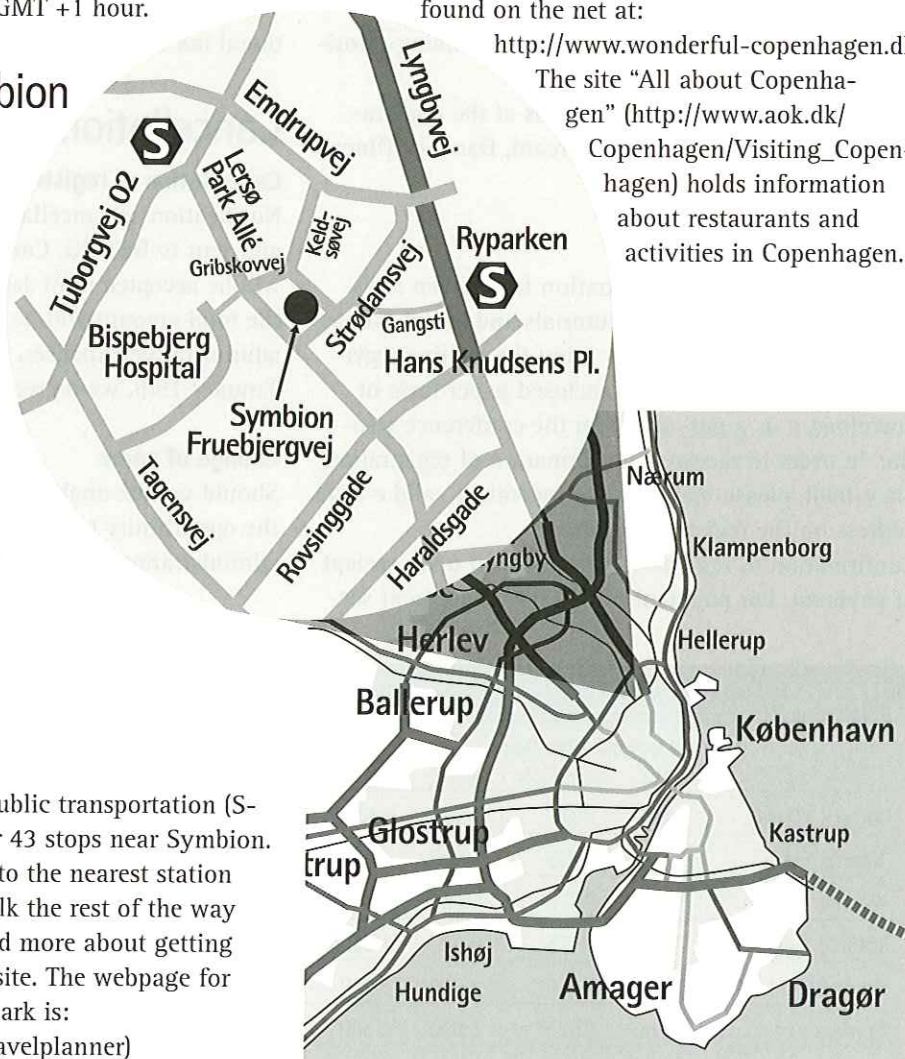
The weather in Denmark in January is usually cold with temperatures around 0-5 degrees celcius, and it might be snowing. Temperatures below zero degrees celcius can happen.

## Tourist information

Denmarks official tourist information services can be found on the net at:

<http://www.wonderful-copenhagen.dk>

The site "All about Copenhagen" ([http://www.aok.dk/Copenhagen/Visiting\\_Copenhagen](http://www.aok.dk/Copenhagen/Visiting_Copenhagen)) holds information about restaurants and activities in Copenhagen.





## Registration

### Registration on-line

We prefer booking via the on-line booking system at the conference website: <http://www.nordu.org/NordU2004/>

### On-site registration at Symbion

On-site registration will start on January the 28th at 08.00. The registration desk and conference Secretariat is located at Symbion Science Park and will be open during the following hours:

Wednesday the 28th .....	08.00 - 17.00
Thursday the 29th.....	08.00 - 17.00
Friday the 30th.....	08.00 - 20.00
Saturday the 31th .....	08.00 - 18.00
Sunday the 1st .....	09.00 - 16.00
Monday the 2nd.....	08.00 - 09.00

## Payment

Online registration and payment is strongly encouraged. Invoice can be sent if no on-line payment is possible. Please allow 10 days for processing.

All amounts due must be paid before attending tutorials or conference.

We accept the following creditcards at the registration desk: Visa, Eurocard/Mastercard, Dankort, Diners Club, JCB Card

### Registration guidelines

Please use the on-line registration form when registering for the Conference, tutorials and social events. Should you have difficulties using the on-line registration form please use the enclosed paper form or download it as a pdf- file from the conference website. In order to receive a confirmation of registration via e-mail, please make sure to indicate a valid e-mail address on the registration form.

Confirmation of registration will be sent upon receipt of payment. For payment details please see next section.

Registration for events which are included in the registration fee must also be indicated on the form, in order to obtain a ticket.

The registration fee for the Tutorial Sessions, January 28 - 30th, includes admission to the ordered tutorials and documentation, daily tea/coffee and lunches.

The registration fee for the Conference, January 31 - February 1st, includes admission to the Conference Sessions, the exhibition, daily tea/coffee, lunches, Welcome reception and Conference Dinner.

Please follow the payment instructions carefully.

### Who pays prices with VAT

Persons and companies in Denmark must always pay amounts including VAT.

Other EU VAT. Registered Businesses pay amounts without VAT, if the VAT number is entered on the registration formular / on-line system. Everybody without a VAT number pays the amount including VAT.

Outside EU companies must pay amounts including VAT, and can get refunds according to international tax rules. Ask your local tax advisor.

## Cancellations

### Cancellation of registration

Notification of cancellation must be made in writing and sent to DKUUG. Cancellations of registrations will be accepted until January 15th., up to which date the total amount will be refunded less 500 DKK for administrative expenses. For cancellations made after January 15th. we regret that no refunds can be made.

### Change of name

Should you be unable to attend, you will be given the opportunity to send a colleague in your place. An administration fee of DKK 100,- will then be charged.

## Prices

	Paid before 22th Dec.		Normal		Student	
	Excl VAT	Incl VAT 25%	Excl VAT	Incl VAT 25%	Excl VAT	Incl VAT 25%
Tutorial 1/2 day	1 500	1 875	1 500	1 875	750	937,5
Tutorial 1 day	3 000	3 750	3 000	3 750	1 500	1 875
Tutorial 2 days	5 500	6 875	6 000	7 500	3 000	3 750
Tutorial 3 days	7 500	9 375	9 000	11 250	4 500	5 625
Conference 1 Day	1 000	1 250	1 000	1 250	500	625
Conference 2 Days	1 500	1 875	2 000	2 500	500	625

Early registration prices apply if you book online or DKUUG has received your registration form on December 22. 2003.

To acquire student prices requires that you send a copy of your valid student card with photo and date, which confirms that you are a student of an official recognized educational institution when NordU2004 happens. Send this via mail to DKUUG ([sek@dkuug.dk](mailto:sek@dkuug.dk)) or to the DKUUG fax: +45 39208948.



### Cancellation of hotel reservation

The general scheme is that you book your hotel rooms directly, see below or the NordU2004 webpage, for rebates. Speakers and invited people book during the DKUUG sekretariat sek@dkuug.dk.

For hotels booked via DKUUG, contact DKUUG as soon as possible.

### Scientific information

Please note that the program in this announcement is preliminary and may be subject to changes. Changes will occur at our webpage as soon as possible. Check the latest news section.

### Language

The official language of NordU 2004 is English. No simultaneous translation will be provided.

### Badges

Each participant will receive a name badge upon registration. For security reasons all participants are requested to wear their badge during all the conference activities and social events. The cost for replacing the badge is DKK 100.

### Meals

Coffee, lunch and social events are included in the registration fee and will be served daily.

## Hotel Information

Hotel booking must be done directly at the hotels. DKUUG have made a company discount deal with Choice Hotels Scandinavia, which covers the following hotel brands: Comfort Hotels, Quality Hotels and Clarion Hotels. The discount agreement covers all Choice hotels, but we recommend three hotels, near the Conference site "Copenhagen Science Park Symbion." The hotels are listed in order of ascending distance from the conference site.

To get the DKUUG company discount, please refer to the DKUUG company discount when you book. All Choice Hotels has Wireless Internet access in all rooms, at the following rates:

2 Hours DKK 50	6 hours DKK 100	24 hours DKK 120 (only some hotels)
-------------------	--------------------	---

Please bring your own Wireless LAN card.

## Hotels and prices

The official conference hotel is:

### Choice Hotel Østerport

Oslo Plads 5  
DK 2100 Copenhagen East

Tel.: +45 3311 2266

Fax: +45 3312 2555

Email: info.oesterport@  
comfort.choicehotels.dk

Located at the Østerport  
Railway station.

Prices:

incl DKUUG Rebate: All days

Double room DKK 1020

Single room DKK 860

Ask for the DKUUG  
company discount.

Other Choice  
Hotels recommended:

### Comfort Hotel Esplanaden

Bredgade 78  
DK-1260 København

Tel.: +45 33481000

Fax: +45 33481066

Email: info.esplanaden@  
comfort.choicehotels.dk

Close to the Royal Palace

Prices

incl. DKUUG Rebate: All-days

Double room DKK 1210

Single room DKK 1035

Ask for the DKUUG  
company discount.

### Comfort Hotel Excelsior

Colbjørnsensgade 6  
DK-1652 København

Tel.: +45 3324 5085

Fax: +45 3324 5087

Email: info.excelsior@  
comfort.choicehotels.dk

Very close to main railway  
station

Has special offer: Stay 3  
nights - pay for 2

Double room DKK 1880

Single room DKK 1480

Regular prices

incl. DKUUG Rebate: All days

Double room DKK 945

Single room DKK 1120

Ask for the DKUUG  
company discount.

Youth Hostel:

### DANHOSTEL Copenhagen Amager

Vejlandsallé 200  
2300 København S

Tel.: +45 3252 2908

Fax: +45 3252 2708

Hosts: Dorte & Torben Borch

Email:  
copenhagen@danhostel.dk  
www.copenhagenyouthhostel.dk

Cheapest prices:

per night per person:

DKK 95 (dormitory rooms)

Per night per room:

2 persons rooms DKK 300

3 Persons rooms DKK 390

4 Persons rooms DKK 460

5 persons rooms DKK 475

Breakfast DKK 45



**BoFs**

A BoFs (birds-of-a-Feather session) is an opportunity to the conference participants to chair a discussion of their own choice. The program committee encourages any conference attendee wishing to chair a BoFs to submit a short description to nordu2004@dkuug.dk. Scheduled BoFs will be displayed on the conference message board during the conference.

**Terminal room**

There will be a room with a limited number of Internet terminals (Secure Shell, web browser) at the conference. Moreover, you can connect your laptop to the conference LAN and the wireless network.

**Wireless LAN**

NordU 2004 is pleased to offer Internet connectivity during the conference and tutorials. A wireless network, with Internet access, will provide DHCP addressing for your laptop. Any DSSS card which is IEEE 802.11 compliant will work with this network. A limited number of PCMCIA cards will be available for check-out (with a credit card guarantee) at the registration desk. Drivers for Linux, FreeBSD, OpenBSD, Mac OS, and Microsoft Windows are available on the Internet.

**Conference proceedings**

The conference proceedings will be available at our website. The password is handed out at the registration desk.

**On demand print of proceedings**

Hewlett-Packard will provide print-on-demand of papers during the conference. You can obtain a hard-copy of any paper - just ask at the HP booth!

**Commercial exhibition**

A commercial exhibition will be arranged in conjunction with the conference. Please contact the organizing committee or the DKUUG secreteriat, if you are interested in participating at the exhibition (see the address elsewhere). The exhibition will be open during the following hours:

Friday the 30th. . . . . 18.00 - 20.00  
Saturday the 31th. . . . . 09.00 - 18.00  
Sunday the 1st. . . . . 09.00 - 15.15


**Social events**

- Welcome reception (Friday, January 30th, 2004, 18.00) The welcome reception will take place in the exhibition area. Snacks and beverages will be provided.
- Chess match (Friday, January 30th, 2004 at 18.00) The ChessBrain project <<http://www.chessbrain.net>> will play against the grand master Peter Heine.
- Conference dinner (Saturday, January 31st, 19.00-?) To be announced at the webpage.

**Disclaimer**

The Organizing committee and DKUUG accept no liability for any injuries/losses incurred by participants and/or accompanying persons, nor loss of, or damage to, any luggage and/or personal belongings. We are providing the information about the conference as correct as possible, but can not be held liable for any losses incurred from the use of this information.

*The Marble Church with its dome is located just behind Amalienborg Palace and "Amaliehaven".*







Kenneth Christiansen, a 24 year old computer science, mathematics and language student at the University of Copenhagen. A long time member of the GNOME open source desktop project as well as other smaller opensource projects.

## Organizing Committee

Organizing Chair: Kristen Nielsen, DKUUG/TDC A/S  
 Tutorial Chair: Sidsel Jensen, DKUUG BSD-DK/DIKU  
 Programme Chair: Kenneth Geisshirt, DKUUG/Silex Science ApS  
 Webmaster: Benny Kjærgaard, DKUUG/DTU  
 Preceedings: Flemming Kraglund, DKUUG  
 Papers coordinator: Martin Wahlen Europen.SE  
 Financial Officer: Hanne Schmidt Vilmann, DKUUG  
 Jan Säll, Europen.SE/Irial  
 Seppo Kauppinen FUUG  
 Kenneth Christiansen Gnome/DIKU  
 Alexander Funcke Europen.SE  
 Ulla Sandberg Europen.SE  
 Peter Håkansson Europen.SE  
 Anita Nilsson Rojning Europen.SE  
 Håkan Carlsson Europen.SE



Alexander Funcke

Jan Säll is the chairman of EurOpen.SE. Apart from being the chairman he is running his own company (Irial) in sweden and uk. Jan focus his work on internet, security and open source.



Kristen is 40 years old. He learned his first computer language BASIC by reading the manual of the Commodore PET 2001 computer around 1977, and has been programming computers since then. He learned Comal on an alpha-LSI from CA. Kristen took his first education as an Electronic Technician in 1984, where he booted UNIX for the first time on a Philips PMDS system. In 1986 he taught his first unix classes on UNIX System III produced by Christian Rovsing A/S. Besides being a teacher of computer classes for 8 years, he has worked with IP Networking and UNIX as his main workarea for 9 years. Besides his full-time employment Kristen is a student at the Computer Science dept. at University of Copenhagen, where he studies Computer Science and Electronic engineering. Kristen is the vice chair of DKUUG and has been a member of the DKUUG board and head of the DKUUG networking group for 9 years. He has been in the NordU programme committee since 1999. Kristen has created the opensource computer summer camp thecamp.dk together with Benny Kjærgaard.



Benny Kjærgaard is a 29 years graduate student in Bachelor IT engineering, at the Technical University of Denmark, he has been a programmer for many years, Besides that he has been a member of DKUUG and their networking group for one year. He is also one of the founders of the computer summer camp "TheCamp.dk", that focuses on FreeBSD and Open Source. This year he has done the web pages of the NordU 2004 Conference. His first education was as a Repair and Maintenance Mechanic at Danfoss A/S, where he did programming of PLC's (Programmable Logic Controllers) as well as regulation processes. He did planned maintenance under the Total Productive Maintenance concept, build machines, and did documentation with respect to the ISO-9001 quality standard. At Danfoss A/S his interest in programming started. He also was the "Sysop" of his own BBS system, back in 1994.



Sidsel Jensen is a computer science graduate student at DIKU, where she also works in the IT-department. Apart from being a member of DKUUG, she is boardmember of the Danish BSD user group BSD-DK and one of the organizers of the Danish LinuxForum conference.

### Reviewers of papers

We wish to thank all our anonymous reviewers for their work, reviewing the papers for this conference.



Martin Wahlen is a founding member of SSLUG the Skåne Sjælland Linux User Group. He is currently working on his Ph.D. thesis at Lund University in the algorithm theory group. He serves on the board of EurOpen.se.



Kenneth Geisshirt has been working with UNIX for almost 15 years. Today he primarily uses Linux, and his interests include clusters, scientific computing, and non-technical issues of open source software. He is member of DKUUG, and has organized a number of technical conferences. Currently, Kenneth is working as principal scientist at Silex Science - a consultant company focusing on Linux clusters and bioinformatics.



Flemming Kraglund



Seppo Kauppinen



Hanne Schmidt Vilmann. Financial officer at NordU 2004. Commanding officer of the DKUUG general office. Organizer of the Danish LinuxForum conference.







**W** Wednesday 28th.  
January 2004

**T** Thursday 29th.  
January 2004

**F** Friday 30th.  
January 2004

08:30-12:30

**WTF1** page 4  
**FreeBSD 4.8 Kernel Internals: Data Structures, Algorithms, and Networking (3 days)**  
Dr. Marchall Kirk McKusick

**WTF1** page 4  
**FreeBSD 4.8 Kernel Internals: Data Structures, Algorithms, and Networking (3 days)**  
Dr. Marchall Kirk McKusick

**WTF1** page 4  
**FreeBSD 4.8 Kernel Internals: Data Structures, Algorithms, and Networking (3 days)**  
Dr. Marchall Kirk McKusick

**W2** page 5  
**FreeS/WAN: Getting more network security out of less sysadmin work!**  
Hugh Daniel (Full day)

**T2** page 8  
**PostgreSQL**  
Bruce Momjian (Full day)

**F2** page 10  
**Logging & Security: Building an Enterprise Logging Infrastructure**  
Tina Bird (Full day)

**W3** page 6  
**New Features in MySQL.**  
David Axmark (½ day)

**T3** page 8  
**Python from the Ground Up**  
Fredrik Lundh (Full day)

**F3** page 11  
**New Technology In Sendmail**  
Eric Allman (Full day)

**W4** page 6  
**Introduction to Perl 6**  
Allison Randal (½ day)

**T4** page 9  
**Migrating Samba 2.2 server to Samba 3.0**  
Volker Lendecke (Full day)

**F4** page 11  
**Linux Migration Challenges for the Systems Administrator**  
Marcel Gagné (Full day)

**F5** page 12  
**Building Linux Clusters**  
Chris Dwan (½ day)

12:30-13:30

**Lunch**

**Lunch**

**Lunch**

13:30-17:00

**W5** page 7  
**Introduction to Parrot (PERL 6 Virtual Machine)**  
Dan Sugalski (½ day)

**T2** page 8  
**2nd part**

**F6** page 12  
**Web Mining - Accomplishments & Future Directions**  
Jaideep Srivastava (½ day)

**W2** page 5  
**2nd part**

**T3** page 8  
**2nd part**

**F7** page 13  
**Introduction to Host Configuration and Maintenance with CFEngine**  
Mark Burgess (½ day)

**T4** page 9  
**2nd part**

**F2** page 10  
**2nd part**

**F3** page 11  
**2nd part**

**F4** page 11  
**2nd part**

18:00-20:00

page 36  
**Welcome Reception**  
See Backcover

**M** Monday 2nd.  
February 2004

**T** Tuesday 3rd.  
February 2004

08:00-08:30  
**Registration**

08:00-08:30  
**Registration**

08:30-17:00

**MT1** page 5  
**FreeBSD 5.X: An Intensive Code Walkthrough (2days)**  
Dr. Marchall Kirk McKusick


**MT1** page 5  
**FreeBSD 5.X: An Intensive Code Walkthrough (2days)**  
Dr. Marchall Kirk McKusick



## Keynotes

**KSA1**  
Page 14

**Keynote: Impact of Open Source Security Tools**  
Wietse Venema



**KSU1**  
Page 28

**Content Wars: Fighting the Gatekeepers**  
J.D. "Illiad" Frazer



**KSA2**  
Page 22

**UNIX, the Internet and Linux: Birthdays and Crises**  
Dr. Peter H. Salus



## Welcome Reception

Is held Friday 30th of January, 2004 at 18.00 - 20.00  
Meet old and new friends  
Snacks and beverages  
Exhibition is open  
Distributed Chess Computing World Record Attempt

## Distributed Chess Computing World Record



"ChessBrain (with the help of hundreds of international participants) will attempt to establish the first world record in distributed computation involving a game. The key here is that no such record actually exists and we'll make the very first attempt at establishing a new world record." ChessBrain will be play against the Danish

grandmaster Peter Heine Nielsen who rates as high as 2626. After the World Record attempt the conference participant will be able to play against ChessBrain. Carlos Justiniano will talk about the ChessBrain Project on Sunday, please see SU2, page 23. To participate in the World Record attempt, see: <http://www.nordu.org/NordU2004/>

**The Chess Grandmaster Peter Heine Nielsen**  
*30 years old. Grandmaster in chess and a history student. 4 times and current Danish champion. Ranked 68 in the world. Twice participated in the world championship and came to the second round, losing to defending Champion Anand in Moscow 2002. Sole danish qualifier for the next world championships. Nordic player of the year 2003. World under 18 bronze medalist. Olympic bronze medalist. Winner of this years edition of the worlds oldest chess congress Hasting international.*

**ChessBrain**



[www.chessbrain.net](http://www.chessbrain.net)

