

2
no. 7
1998

Computer Communication Review



Special Interest
Group on Data
Communication

An ACM
SIGCOMM
Publication

Volume 28
Number 4
October 1998

Proceedings of ACM SIGCOMM'98 CONFERENCE



Applications, Technologies, Architectures,
and Protocols for Computer Communication

Computer Communication Review

a publication of the acm special interest group on data communication

Chairman:

Jon Crowcroft
Department of Computer Science
University College London
Gower Street, London
WC1E 6BT UK
Tel: +44 71 380 7296
j.crowcroft@cs.ucl.ac.uk

Vice Chairman:

Sally Floyd
Lawrence Berkeley National Laboratory
One Cyclotron Road
Berkeley, CA 94720
Tel: +1 510 486 7518
floyd@ee.lbl.gov

Secretary-Treasurer:

Patrick W. Dowd
Department of Electrical Engineering
A.V. Williams Building
University of Maryland
College Park, MD 20742
Tel: +1 301 405 6750
dowd@eng.umd.edu

Editor:

Martha Steenstrup
BBN Technologies
10 Moulton Street
Cambridge, MA 02138
Tel: +1 617 873 3192
Fax: +1 617 873 6091
msteenst@bbn.com

Information Services Director:

Greg Wetzel
AT&T Bell Laboratories
101 Crawfords Corner Road
P.O. Box 3030
Holmdel, NJ 07733-3030
Tel: +1 908 949 6630
Infodir_SIGCOMM@acm.org, or
G_F_Wetzel@att.com

SIG Program Director:

Pat McCarren
ACM
1515 Broadway, 17th Floor
New York, NY 10036
Tel: +1 212 626 0611
Fax: +1 212 302 5826
mccarren@acm.org

Associate Editor:

Lixia Zhang
4531G Boelter Hall
Computer Science Department
UCLA
Los Angeles, CA 90095
Tel: +1 310 825 2695
lixia@cs.ucla.edu

Award Committee Chairman:

David C. Wood
NATO C3 Agency
P.O. Box 174
2501 CD The Hague
The Netherlands
Tel: +31 70 3142420
Fax: +31 70 3142176
woodd@nc3a.nato.int

Executive Committee:

Jon Crowcroft
Sally Floyd
Patrick Dowd
A. Lyman Chapin
Martha Steenstrup
Greg Wetzel

ACM SIGCOMM Lecturers:

Raj Jain
Deepinder Sidhu
Ian Akyildiz

SIGCOMM Mission Statement:

SIGCOMM is ACM's professional forum for the discussion of topics in the field of communications and computer networks, including technical design and engineering, regulation and operations, and the social implications of computer networking. The SIG's

members are particularly interested in the systems engineering and architectural questions of communication.

Feedback on our mission is always welcome. Please email suggestions to any of the members of the SIGCOMM Executive Committee.

Advertising: ACM accepts recruitment advertising under the basic premise that the advertising employer does not discriminate on the basis of age, color, race, religion, gender, sexual preference, or national origin. ACM recognizes, however, that laws on such matters vary from country to country and contain exceptions, inconsistencies or contradictions. This is as true of the laws of the United States of America as it is of other countries. Thus ACM policy requires each advertising employer to state explicitly in the advertisement any employment restrictions that may apply with respect to age, color, race, religion, gender, sexual preference, or national origin. Observance of the legal retirement age in the employer's country is not considered discrimination under this policy.
For advertising information, please contact Walter Andrzejewski (Advertising Manager) at ACM, 1515 Broadway, New York, NY 10036 USA, tel: +1 212 869 7440, fax: +1 212 869 0481

COMPUTER COMMUNICATION REVIEW is a publication of the ACM Special Interest Group on Data Communication. Its scope of interest includes: data communication systems for computers; data communication technology for computers; reliability, security and integrity of data in data communication systems; problems of interfacing communication systems and

computer systems; computer communication system modelling and analysis.

Items attributed to persons will ordinarily be interpreted as personal rather than organizational opinions. Technical papers appearing in Computer Communication Review are informally reviewed.

COMPUTER COMMUNICATION REVIEW (ISSN 0146-4833) is published five times a year (January, April, July, and two issues in October) by the ACM, Inc., 1515 Broadway, New York, NY 10036. Annual subscription cost of \$14.44 is included in the member dues of \$50.00 (for students, the cost is included in \$24.00); the non-member annual subscription is \$37.00. Periodicals postage paid at New York, NY 10001, and at additional mailing offices.



ACM



SIGCOMM'98

Proceedings



SCIENCE & ENGINEERING
LIBRARY

FEB 05 1999

EMS COLLECTION
UCLA

**Applications, Technologies, Architectures,
and Protocols for Computer Communication**

<http://www.acm.org/sigcomm/sigcomm98>

Vancouver, British Columbia, Canada
September 2 to September 4, 1998
(Tutorials August 31 and September 1)



Bellcore



SIGCOMM 1998 in Vancouver, B.C Canada

**The Association for Computing Machinery, Inc.
1515 Broadway
New York, New York 10036**

Copyright © 1998 by the Association for Computing Machinery, Inc. (ACM). Permission to make digital or hard copies of portions of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted.

To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permission to republish from: Publications Dept. ACM, Inc. Fax +1-212-869-0481 or E-mail permissions@acm.org.

For other copying of articles that carry a code at the bottom of the first or last page, copying is permitted provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

ACM ISBN: 1-58113-003-1

See also: <http://www.acm.org/sigcomm/sigcomm98>

You may order additional copies — prepaid — from:

ACM Order Department

P.O. BOX 11414 Phone: 1-800-342-6626
Church Street Station (U.S.A. and Canada)
New York, NY 10286 +1-212-626-0500
(All other countries)
Fax: +1-212-944-1318
E-mail: acmpubs@acm.org

ACM Order Number: 533980
Printed in the U.S.A.

Router Plugins

A Software Architecture for Next Generation Routers

Dan Decasper¹, Zubin Dittia², Guru Parulkar², Bernhard Plattner¹

[dan|plattner]@tik.ee.ethz.ch, [zubin|guru]@arl.wustl.edu

¹Computer Engineering and Networks Laboratory, ETH Zurich, Switzerland

Phone: +41-1-632 7019 Fax: +41-1-632 1035

²Applied Research Laboratory, Washington University, St. Louis, USA

Phone: +1-314-935 4586 Fax: +1-314-935 7302

1. ABSTRACT

Present day routers typically employ monolithic operating systems which are not easily upgradable and extensible. With the rapid rate of protocol development it is becoming increasingly important to dynamically upgrade router software in an incremental fashion. We have designed and implemented a high performance, modular, extended integrated services router software architecture in the NetBSD operating system kernel. This architecture allows code modules, called *plugins*, to be dynamically added and configured at run time. One of the novel features of our design is the ability to bind different plugins to individual flows; this allows for distinct plugin implementations to seamlessly coexist in the same runtime environment. High performance is achieved through a carefully designed modular architecture; an innovative packet classification algorithm that is both powerful and highly efficient; and by caching that exploits the flow-like characteristics of Internet traffic. Compared to a monolithic best-effort kernel, our implementation requires an average increase in packet processing overhead of only 8%, or 500 cycles/2.1ms per packet when running on a P6/233.

1.1 Keywords

High performance integrated services routing, modular router architecture, router plugins

2. INTRODUCTION

New network protocols and extensions to existing protocols are being deployed on the Internet. New functionality is being added to modern IP routers at an increasingly rapid pace. In the past, the main task of a router was to simply forward packets based on a destination address lookup. Modern routers, however, incorporate several new services:

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. SIGCOMM '98 Vancouver, B.C. © 1998 ACM 1-58113-003-1/98/0008...\$5.00




Figure 1. : Best Effort vs Extended Integrated Services Router (EISR)

- Integrated/differentiated Services
- Enhanced routing functionality (level 3 and level 4 routing and switching, QoS routing, multicast)
- Security algorithms (e.g. to implement virtual private networks (VPN))
- Enhancements to existing protocols (e.g. Random Early Detection (RED))
- New core protocols (e.g. IPv6 [8])

Figure 1 contrasts the software architecture of our proposed Extended Integrated Services Router (EISR) with that of a conventional best-effort router. A typical EISR kernel features the following important additional components: a packet scheduler, a packet classifier, security mechanisms, and QoS-based routing/Level 4 switching. Various algorithms and implementations of each component offer specific advantages in terms of performance, feature sets, and cost. Most of these algorithms undergo a constant evolution and are replaced and upgraded frequently. Such networking subsystem components are characterized by a relatively "fluid" implementation, and should be distinguished from the small part of the network subsystem code that remains relatively stable. The stable part (called the core) is mainly responsible for interacting with the network hardware and for demultiplexing packets to specific modules. Different implementations of the EISR components outside of the core often need to coexist. For example, we might want to use one kind of packet scheduling on one interface, and a different kind on another.

In this paper, we propose a software architecture and present an implementation which addresses these requirements. The specific goals of our framework are:

- **Modularity:** Implementation of specific algorithms come in the form of modules called *plugins*¹.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.