

Documents to U / MNU Lending

University of Minnesota – Interlibrary Loan Lending

OCLC MNU & UMM * DOCLINE MNUMIN

15 Andersen Library, 222 21st Ave., S., University of Minnesota

Minneapolis, MN 55455-0439 USA

Phone: 612-624-4388, Fax: 612-624-4508, e-mail: docstou@umn.edu

This article is delivered directly from the collections of the University of Minnesota.

Thank you for using our service.

If you have problems with delivery, please contact us within 48 hours.

MNU 1161082 / 170370548

Title: Proceedings of the Winter 1994 USENIX Conference : January 17-21, 1994, San Francisco, California, USA /

Author: Glenn Fowler

ArticleTitle: A Flat File Database Query Languag, PLUS TITL AND CC PAGES AND PAGE WITH DATE STAMP. (SEE BORROWING NOTES)

ArticleAuthor: Glenn Fowler

Date: 1994

OCLC - 29862997; ISBN - 9781880446584;

Publisher: Berkeley, CA : The Association, ©1993.

Copyright: CCL

NOTICE CONCERNING COPYRIGHT RESTRICTIONS:

The copyright law of the United States [[Title 17, United StatesCode](#)] governs the making of photocopies or other reproductions of copyrighted materials.

Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specific conditions is that the photocopy is not to be "used for any purpose other than private study, scholarship, or research." If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of "fair use," that user may be liable for copyright infringement.

Symantec 1031

This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of that order would involve violation of copyright law.

For additional copies of these proceedings write:

USENIX Association
2560 Ninth Street, Suite 215
Berkeley, CA 94710 USA

The price is \$30 for members and \$39 for nonmembers.
Outside the U.S.A. and Canada, please add
\$20 per copy for postage (via air printed matter).

Past USENIX Technical Conferences


1993 Summer Cincinnati	1987 Summer Phoenix
1993 Winter San Diego	1987 Winter Washington, DC
1992 Summer San Antonio	1986 Summer Atlanta
1992 Winter San Francisco	1986 Winter Denver
1991 Summer Nashville	1985 Summer Portland
1991 Winter Dallas	1985 Winter Dallas
1990 Summer Anaheim	1984 Summer Salt Lake City
1990 Winter Washington, DC	1984 Winter Washington, DC
1989 Summer Baltimore	1983 Summer Toronto
1989 Winter San Diego	1983 Winter San Diego
1988 Summer San Francisco	
1988 Winter Dallas	

1994 © Copyright by The USENIX Association
All Rights Reserved.

ISBN 1-880446-58-8


This volume is published as a collective work.
Rights to individual papers remain with the author or the author's employer.

USENIX acknowledges all trademarks herein.

Printed in the United States of America on 50% recycled paper, 10-15% post consumer waste. 

USENIX Association

Proceedings of the Winter 1994 USENIX Conference



**January 17 - 21, 1994
San Francisco, California, USA**

cql – A Flat File Database Query Language

Glenn Fowler
gsf@research.att.com

AT&T Bell Laboratories
Murray Hill, New Jersey 07974

Abstract

cql is a UNIX* system tool that applies C style query expressions to flat file databases. In some respects it is yet another addition to the toolbox of programmable file filters: *grep* [Hume88], *sh* [Bour78][BK89], *awk* [AKW88], and *perl* [Wall]. However, by restricting its problem domain, *cql* takes advantage of optimizations not available to these more general purpose tools.

This paper describes the *cql* data description and query language, query optimizations, and provides comparisons with other tools.

1. Introduction

Flat file databases are common in UNIX system environments. They consist of newline terminated records with a single character that delimits fields within each record. Well known examples are */etc/passwd* and */etc/group*, and more recently the *sablime* [CF88] MR databases and *cia* [Chen89] abstraction databases.

There are two basic flat file database operations:

update – delete, add or modify records

query – scan for records based on field selection function

For the most part UNIX system tools make a clear distinction between these operations. Update is usually done by special purpose tools to avoid problems that arise from concurrency. Some of these tools are admittedly low-tech: *vipw* write locks the */etc/passwd* file and runs the *vi* editor on it; any other user running *vipw* concurrently will be locked out. On the other hand query tools usually assume that the input files are readonly or that they at least will not change during query access. *cql* falls into this category: it is strictly for queries and supports no update operations. Despite this restriction *cql* adequately fills the gap between *awk* and full featured database management systems.

In the simplest case a flat file database query is a pattern match that is applied to one or more fields in each record. The output is normally a list of all matched records. *grep*, *sh*, *awk*, and *perl* are well suited for such queries on small databases. These commands scan the database from the top, one record at a time, and apply the match expression to each record. Unfortunately, as the number of records and queries increases, the repeated linear scans required by these tools soon become an intolerable bottleneck. The bottleneck can be diminished by examining the queries to limit the number of records that must be scanned, but this requires some modifications, either to the database or to the scanning tools.

Some applications, such as *sablime*, ease the bottleneck by partitioning the database into several flat files based on one or more of the record fields. This speeds up queries that key on the partitioned fields, but hinders queries that must span the partition. Besides complicating the application query implementation, partitioning also

* UNIX is a registered trademark of USL.

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.