*eXtreme*DB™
Kernel Mode Edition

Kernel-mode main memory storage solution for performance-critical applications.

> *"The Linux kernel is no place for 'self-expressive fancy.' "*
>
> *-- Andrew Morton,*
> *Open Source Development Labs*

*eXtreme*DB, the real-time embedded database for devices that are eXtremely innovative

Overview

The essential part of any operating system is the operating system "kernel". The kernel is responsible for resource allocation, scheduling, low-level hardware interfaces, networking, security, etc.

The Challenge

Increasingly, system software vendors try to put more application functionality into the operating system kernel. Running parts of the application in the kernel adds to the entire system's performance tremendously. At the same time, many modern application architectures require sophisticated data management functions such as concurrent data access, transaction control and complex searches. Putting the data management portion of the system into the kernel makes those functions faster and more robust.

For example, firewalls and other types of computer security software (access control applications, etc.) must run their policy engines as kernel components, to achieve needed performance. Their data structures are inherently complex, and lookups and updates must be blazingly fast.

Such data management is typically the job of database management systems (DBMS). Until recently, it was not possible to put a DBMS into the kernel – too much overhead introduced by available database engines would retard kernel performance, so developers had to "reinvent the wheel" of database capabilities, in limited fashion, for kernel tasks.

Introducing the Kernel-mode Database

McObject's *eXtreme*DB Kernel Mode (KM) is the first commercial database engine that deploys in kernel mode on most popular operating systems. From the start, the main principle behind the design of *eXtreme*DB has been to eliminate performance overhead while providing a predictable and reliable transaction model.

Representative applications performing computer security functions performed ***an order of magnitude faster*** with *eXtreme*DB-KM than with the standard user-mode database.

*eXtreme*DB-KM achieves this by providing ***direct access to kernel data***. The *eXtreme*DB-KM run-time maps its databases directly into the driver or kernel module address space, providing direct pointers to the data and eliminating expensive buffer management.

The *eXtreme*DB run-time code is directly linked with the module, so remote procedure calls are eliminated from the execution path. As a consequence, the execution path generally requires just a few CPU instructions. Concurrent access is coordinated by the database run-time. Kernel mode databases are also made available to user-mode applications through a set of public interfaces implemented via system calls.

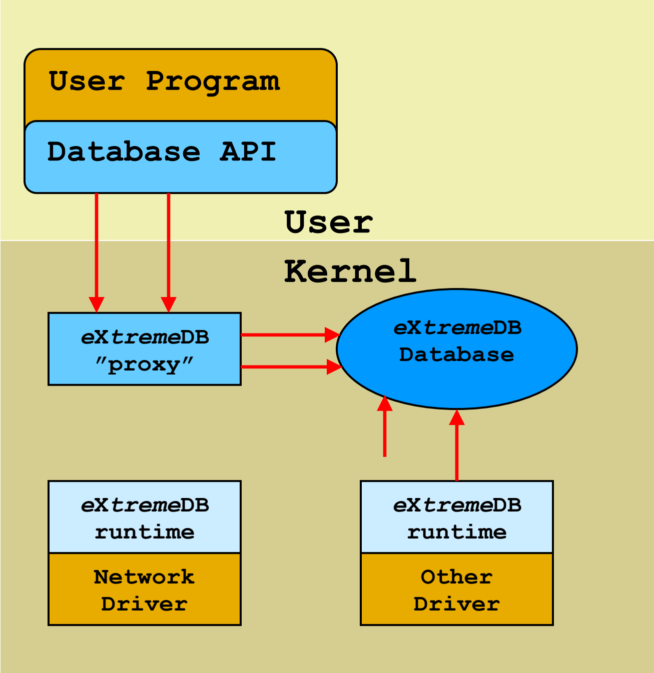

Packaging

*eXtreme*DB Kernel Mode Edition (*eXtreme*DB-KM) is distributed with complete runtime source code and includes a number of user-mode samples and kernel-mode modules.

Precision ***Data Management***

GAiLOGIC CORPORATION
ガイロジック株式会社
〒180-0005東京都武蔵野市御殿山1-6-8
Tel0422-26-8211 fax0422-26-8212
www.gailogic.co.jp
sales@gailogic.co.jp