

ORACLE'S TECHNICAL CONTRIBUTIONS TO LINUX

LINUX SUPPORT, COMMITMENT, AND LEADERSHIP

"Oracle's development work for the Linux kernel represents vital contributions to the open source community, which benefit anyone using Linux."

– **Andrew Morton,**
Linux Kernel Maintainer,
Google

"Oracle's significant Linux kernel development work with file systems, data integrity, and other contributions will prove beneficial to all Linux users."

– **Jim Zemlin,**
Executive Director,
Linux Foundation

With the release of the first commercial relational database on Linux more than a decade ago, Oracle has a long history of strong support and commitment to Linux, as evidenced by numerous, on-going [technical contributions to the Linux community](#). These contributions highlight Oracle's technical leadership and dedication to the success of Linux for organizations of all sizes, in all industries, worldwide. And with the Oracle Unbreakable Linux support program, Oracle is further committed to Linux by delivering true enterprise-class, low-cost support.

Long-Term Vision

Oracle's long-term vision is focused on continuing to enhance and extend the enterprise-class capabilities of Linux, and this vision is manifest through the many projects and code contributions that Oracle shares with the Linux community. Oracle continues to strengthen its involvement in the Linux community by providing enhancements that facilitate the development and deployment of enterprise Linux solutions. By developing enhanced capabilities and contributing code, Oracle's Linux engineering teams continue to make the Linux experience better for all.

This data sheet includes just some of Oracle's many technical contributions to the Linux community—contributions which benefit anyone using Linux. For additional information on Oracle's free and open source contributions, visit oss.oracle.com.

ASMLib

ASMLib is a support library for the Automatic Storage Management feature of Oracle Database 10g which simplifies database administration. [Oracle ASMLib](#)

Asynchronous IO (AIO) Kernel Subsystem

Oracle is working to replace the existing asynchronous I/O (AIO) interface in the kernel with a new AIO api that can be used to make system calls asynchronous in a generic fashion. The new kernel based implementation should provide a single access point allowing most system calls to become asynchronous, reducing complexity at both the kernel and application level. The new subsystem is expected to be faster for Oracle, and is intended to make it easier for other applications to benefit from asynchronous programming under any workload.

Btrfs

The Btrfs file system is designed to address the expanding scalability requirements of large storage subsystems, common in today's data centers. Btrfs delivers features missing from Linux today, such as enhanced scalability and simplified management for large storage configurations, and flexible snapshotting and fast incremental

backups of the data and metadata in the file system. Btrfs also maintains multiple copies of crucial file system structures. Currently, the Btrfs code is on track to be merged into the mainline Linux tree by the end of 2008. [Btrfs wiki on kernel.org](#)

Coherent Remote File System (CRFS)

CRFS is a network protocol that lets clients share access to a file system which is hosted on a server. Its main features include: general-purpose "POSIX" file system semantics, cache coherency, aggressive metadata performance, efficient and flexible snapshots, and robust data integrity protection. [CRFS Project Home Page](#)

Gold

Gold is a new GNU linker, and is released under the GPL as a part of binutils. Gold is much faster than the existing GNU linker, and eventually it will become an incremental linker. Gold supports the ELF file format only. Oracle has been contributing features and bugfixes in order to assure that Gold correctly builds the Linux kernel and the Oracle Database. Oracle also provides testing for Gold through the Labrat test system. [Gold Release Post](#)

Kernel I/O Subsystem Tuning

The development of storage technology is proceeding at a very fast pace. With the new generation of Solid State Disks (SSD), the transfer rate has increased by orders of magnitude. Oracle Linux kernel engineers are working on creating novel approaches in the area of block I/O, to fully exploit such higher disk speeds. [Block I/O Subsystem Maintainer Jens Axboe's Blog](#)

Libstdc++

Oracle is a major contributor and maintainer of the GNU standard C++ library, libstdc++. Bugfixes and features are constantly added as part of the C++0x standard implementation effort. [Libstdc++ Wiki](#)

NFS on IPv6

Oracle Linux kernel engineers are working to enable the Linux Network File System (NFS) to run natively on IPv6 networks. IPv6 is a network layer for packet-switched networks including the internet; IPv6 compatibility is a key requirement for software purchased by the U.S. federal government after 2008. The NFS IPV6 code has been integrated into mainline for the Linux client, and work on the server side is currently in progress. [IPv6 for Linux NFS Slides from Oracle's Chuck Lever \(PDF\)](#)

Oracle Cluster File System (OCFS) 2 v.1.4

Oracle Cluster File System 2 (OCFS2) is an open source, general-purpose, extent-based clustered filesystem which Oracle developed and contributed to the Linux community, and accepted into Linux kernel 2.6.16. [OCFS2 Project Home Page](#)

Oracle Linux Test (OLT) Kit

Available as open source under the GPL and Artistic licenses, the Oracle Linux Test Kit, derived from the Oracle Validated Configurations program, is designed to verify Linux kernel functionality and stability essential for the Oracle Database. The Oracle Linux Test Kit provides an automated mechanism to define, execute and analyze tests. [OLT Kit Project Home Page](#)

Oracle Validated Configurations for Linux and Virtualization

Oracle Validated Configurations for Linux and Virtualization are pre-tested, validated architectures with software, hardware, storage, and networking components with included documented best practices for deployment. [Oracle Validated Configurations Project Home Page](#)

PHP

The Oracle Linux engineering team devotes resources to the improvement and maintenance of PHP and its Oracle-specific extensions. The team publishes and updates a book, "The Underground PHP and Oracle Manual," about PHP and its usage with Oracle, and newer PHP packages in .rpm format are available for free download to Linux users. The Oracle Call Interface (OCI8) database driver for PHP provides new, improved integration between PHP and Oracle Database 11g, to allow a server-side connection pool shareable across web servers and languages, significantly enhancing the scalability of web-based systems. [PHP Developer Center](#)

RDS Reliable Datagram Sockets (RDS)

RDS is an effort to provide a socket API which is uniquely suited to the way Oracle does network Inter Process Communication (IPC). The Oracle Linux kernel development team created an open source implementation of the API for the Linux kernel. The code is now integrated into the OpenFabrics Enterprise Distribution (OFED) stack. OFED aims to deliver a unified, cross-platform, transport-independent software stack for RDMA, including a range of standard protocols. [RDS Project Home Page](#)

T10-DIF Data Integrity

Oracle, in collaboration with Emulex, is implementing a leading, first of its kind initiative to bring enterprise-class data integrity to the Linux platform. An open source interface is being implemented by Oracle to expose the T10 Protection Information Model (also known as DIF) standard to the Linux kernel and end-user applications.

The Oracle code contribution, which has been accepted into the 2.6.27 Linux kernel, includes generic support for data integrity at the block and file system layers, as well as support for the T10 Protection Information Model and the Data Integrity Extensions. The Data Integrity Extensions, developed in collaboration with Emulex, augment the T10 standard by allowing protection information to be transferred to and from host memory, enabling comprehensive data integrity protection.

The Oracle-developed code is the first implementation of the T10 Protection Information Model standard for an operating system, and should allow applications and kernel subsystems to take advantage of these crucial data integrity features, enabling reduced system downtime and cost savings for end-users. [Linux Data Integrity Project Home Page](#)

Testing of Open Source Projects

Testing the mainline kernel is essential so the Linux community can get a long-term regression picture of how the kernel performs and works. Mainline kernel testing and quality assurance (QA) benefit the entire community. In addition, The Oracle Linux Test (OLT) Kit is a test framework designed to verify Linux kernel

functionality and stability; see additional information above. Other projects Oracle tests include Systemtap and GCC, with the Labrat build and test system. [Oracle Kerneltest Results](#)

Yet Another Setup Tool (YaST)

Yet Another Setup Tool (YaST) helps make system administration easier by providing a single utility for configuring and maintaining Linux systems. Available under GPL, this code can be freely accessed by anyone. The Oracle Linux Engineering team ported the Yast tool to OEL from SUSE. Oracle Unbreakable Linux support customers have access to the YaST functionality integrated with the Oracle Management Pack for Linux, thereby gaining a fully supported, easy to use, feature-rich system management environment for their Linux servers. [YaST Project Home Page](#)

Xen Hypervisor

Oracle VM is free, next generation, scalable server virtualization software that supports Oracle and non-Oracle applications. Oracle VM provides an easy-to-use graphical interface for creating and managing virtual server pools running on x86 and x86-64-based systems across an enterprise; both Linux and Windows guests are supported.

Part of Oracle VM includes the Xen hypervisor, and Oracle's engineering team contributes heavily to feature development of Xen mainline software. The most significant contributions are in the area of Windows Paravirtualized drivers, hardware virtualized timers, guest and hypervisor debugger and bugfixes, transcendental memory, SSL live migration and xend locking, as well as participation in XenAPI changes.

The Oracle QA team also provides stabilization efforts through testing Xen configurations with Oracle workloads and OEL kernels. Oracle is also a member of the Xen Advisory Board which serves in an advisory capacity to the Xen project leader for all community and development activities as well as management for the Xen trademark. [Xen Home Page](#), [Oracle VM](#)

Contact Us

For more information about Oracle's technical contributions to Linux, please visit oracle.com/linux. And for additional information on Oracle's free and open source contributions, visit oss.oracle.com.



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2009, Oracle and/or its affiliates. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners. 0109