

# Oracle Content Database: Infrastructure for Content Consolidation

*An Oracle White Paper*  
*July 2007*

# Oracle Content Database: Infrastructure for Content Consolidation

EXECUTIVE OVERVIEW .....	3
EVOLUTION OF THE CONTENT MANAGEMENT MARKET.....	4
ORACLE RESPONDS .....	4
THE NEXT STEP: ORACLE CONTENT DB.....	5
Enterprise-class File and Document Management.....	5
Web Services.....	7
Business Process Automation.....	7
Records Management.....	8
CONCLUSION .....	8
for more information .....	9

# Oracle Content Database: Infrastructure for Content Consolidation

## **EXECUTIVE OVERVIEW**

In the past few years, major changes have been taking place in the content management marketplace. Customers that previously considered content management to be a niche application focused on workgroup or departmental deployments for meeting specialized publishing-oriented requirements are increasingly requesting true enterprise deployments that deliver content management capabilities to all their users.

These changes have created a need for a scaleable, affordable, and highly usable solution that bridges the wide gap between limited capability file servers and specialized, expensive, and complex content management applications. Customers are looking for a solution that provides a common infrastructure for consolidating file servers, content repositories supporting mission-critical business applications and archiving applications. Ultimately, they are also looking to consolidate the number of vendors with which they have to interact.

Oracle refers to such a solution as *infrastructure for content consolidation*. This white paper discusses the drivers of market evolution and shows how Oracle Content Database (Content DB) delivers the solution that the changing market has been seeking.

## EVOLUTION OF THE CONTENT MANAGEMENT MARKET

The products that most people think of when they hear the term “content management” began to emerge about 15 years ago to support the creation, management, and publication of unstructured data (or content) for users who specialized in content production in highly regulated industries. However, the vast majority of content is actually used by those of us who are not content specialists (about 95% of all business users) and that content mostly resides on desktops and file servers—without any real management at all.

“The main business drivers for content management have historically been around productivity improvements within organizations for information workers and process workers. What has emerged over the past few years as a huge business driver is compliance, not just from the regulatory perspective, but also from responding to litigation requests, which can be extremely costly for organizations.”

**Jeetu Patel**  
**Executive Vice President**  
**Doculabs**

In the past few years, three conditions began driving major changes in the market and have given rise to a need for solutions that bridge the gap between limited capability file servers and specialized, expensive, and complex content management applications. First, the rapid and accelerating explosion of all types of content—documents, emails, instant messages, images, etc.—has driven a need for putting better content tools in front of all enterprise users to help them make sense of the flood of data or at least keep them from drowning in it. Second, a series of corporate scandals focusing on the discovery, alteration, or inappropriate destruction of documents and emails has dramatically increased awareness about the risks resulting from proliferation of poorly managed content. Finally, a plethora of government and industry regulations emerged—largely in response to these scandals—that dramatically increased the costs of compliance as well as the risks of non-compliance.

The net impact has been to transform content management from a niche market, catering to content specialists at the very top of the “user pyramid,” to a mainstream technology that serves virtually every user in an enterprise. The term “enterprise content management” or ECM was, until now, a misnomer since it originally meant a solution specifically for content production specialists that managed the different types of content in the enterprise— not one that met the needs of a broad range of users.

## ORACLE RESPONDS

“Historically the challenges we’ve seen that organizations have experienced for achieving widespread adoption of content management systems have primarily centered around cost and complexity - the cost being too high for content management systems on every single desktop within the enterprise, and the complexity requiring a level of training that can be prohibitive for every single user within the organization.”

**Jeetu Patel**  
**Executive Vice President**  
**Doculabs**

While this market evolution and expansion was unfolding, Oracle was steadily becoming more interested in the opportunity. It was also steadily improving its ability to eliminate the arbitrary distinction between structured (transactional) data and content. Unlike other content management products that use a complex conglomeration of database-managed metadata and a flat file-system for actual content storage, Oracle perfected the use of the database for storage of all content-related information—metadata, relationships, indices, system state—and the content itself. As a result, the reliability, security, and immense scalability of the Oracle Database, together with its powerful tools for business continuity, access control, search and query, auditing and tracking, etc. became available to content as well as transactional data. A new paradigm emerged where all information can be managed using the same tools and made available to all applications whenever and however it is needed. Oracle products such as the Internet File System (iFS), its successor the Content Management Software Developers Kit (CM SDK), and Oracle Files were prime examples of the paradigm shift.

By its second release in mid-2003, Oracle Files had become a robust and highly scalable solution for consolidated file management and sharing and lightweight document management. Oracle's own internal use of its content management technology is a powerful example of the product's capabilities. One application running on a single instance of the Oracle database met virtually all the file management and sharing needs of Oracle's then 50,000+ employees worldwide, managing over 22 million documents and 14 Terabytes of storage, which grew at a rate of over 30,000 new documents a day.

## THE NEXT STEP: ORACLE CONTENT DB

Oracle Content DB is the next generation of Oracle's content management technology - building on the powerful CM SDK infrastructure with a new level of functionality, usability, and expandability. Delivered with a standards-based service oriented architecture (SOA), Oracle Content DB provides precisely the tools and facilities customers are looking for, seamlessly surfaced in familiar environments when and where users need it and without changing the way they work. Oracle Content DB provides key customer benefits including:

- Increased productivity by making information throughout the enterprise easier to manage, find, and share
- Reduced risk through better control of information and establishment of consistent information policies and processes
- Facilitate compliance with government and industry regulations through records management
- Lower costs through consolidation of server hardware, software licenses, and administrative support

The capabilities of Content DB can be summarized as follows:

### Enterprise-class File and Document Management


Oracle Content DB provides a number of significant file and document management capabilities in three areas: rich user interfaces, flexible access control and user management, and policy-based management behaviors.

Oracle Content DB delivers a rich experience for both Web and Windows desktop users. A new highly interactive Web interface allows content to be accessed easily anytime, anywhere from any Web browser. This interface provides a user experience scarcely distinguishable from thick client applications and familiar desktop tools via a dynamic tree view, right-click and pull-down menus, and drag and drop capabilities.

Advanced integration with Microsoft Windows through the Oracle Drive Windows plug-in permits users to seamlessly access Oracle Content DB content and functionalities directly from Windows Explorer. Finally, Oracle Drive features seamless offline content management capabilities, allowing users to access and manage content easily when they are disconnected from the network.

" Oracle Content DB marries the traditional file system look and feel with the power of the Oracle 10g Database. This enables us to manage our documents effectively while we focus on our business transactions."

**Kyle Lambert**  
Vice President, Information Systems  
**John I. Haas**



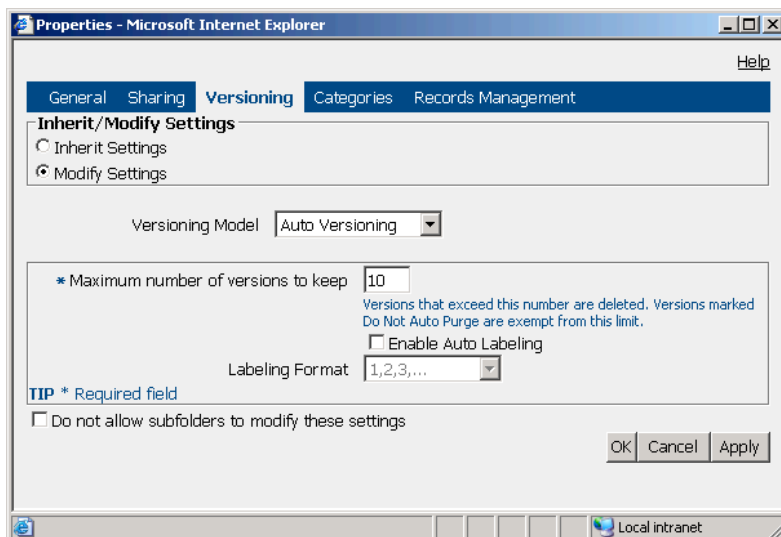
< Oracle Content DB Web UI

people *don't* have to do. There is no time and effort wasted looking for the right piece of information"

**Sean O'Donoghue**  
Vice-President of Business  
Applications and Technology,  
The Macerich Company

The security model within Content DB is extraordinarily flexible, with fine-grained permissions and access specified at the folder and document level. Role-based access control includes a rich set of out-of-the-box roles and also permits customer-defined roles to be built up from the basic permission bits. Finally, access management is available through groups as well as individual users.

Behaviors controlling management of the information lifecycle such as versioning, attribution, and records management can be specified through policies established on a folder-by-folder basis. Versioning can be applied to content automatically based on the occurrence of a specific event such as check-in, applied manually by users, or disabled. The maximum number of versions to retain before archiving can also be specified.



Category attributes can similarly be applied manually or automatically. Individual Category attributes can be given default values (set as user-editable or not), and mandatory user entry of values can be specified. Records management (detailed below) is also managed through policies on a folder-by-folder basis. While Content DB policies will normally be managed by a small subset of users with specific administrative permissions, the tools for setting up policies are straightforward and consistent, allowing use of a delegated policy management model as dictated by the needs of the business.

By utilizing automated policies instead of requiring the user to remember to perform a lot of extra steps, these important lifecycle management capabilities become truly useful and the risk of user rejection often associated with content management deployments is minimized.

### **Web Services, SOA, and Extensibility**

Oracle Content DB provides complete access to application functionality programmatically via J2EE- and Microsoft.NET-compatible Web Services APIs. These APIs are being used to integrate Content DB with other Oracle products including E-Business Suite applications and Oracle Portal and can also be leveraged by ISVs and implementation partners to integrate Content DB with other applications, as well as to automate and extend application functionalities.

Underlying the Web services is a comprehensive set of Java APIs known as the *FDK*. These APIs provide essentially the same functionality as the Web services, but in an in-process environment that provides application developers with some additional functionality as well as lower processing overhead that might be required for some higher performance applications.

Oracle Content DB also provides a rich facility for adding custom processing and handling of virtually every event that occurs in the product. Easily developed Java applets can be added which, when triggered by events, can utilize the FDK Java APIs to provide a broad range of additional functionality to the product.

### **Business Process Automation**

Oracle Content DB supports custom workflows through Oracle BPEL Process Manager. Workflow templates that can be customized or extended as required are delivered out-of-the-box with Content DB. New workflows can also be built with BPEL Process Manager and made available to Content DB.

Also included are powerful capabilities that allow workflows to be associated with specific folders and automatically triggered when events such as document check-in or deletion occur in those folders. Workflows can be specified as blocking (the triggering event remains in a pending state until the workflow completes) or non-blocking (the triggering event completes and the workflow is launched as a side-effect.) In addition, workflow template parameters can be set up with defaults and/or user entry enforcement in much the same way as Category attributes.

These custom workflows can be used to drive review and approval cycles, notify users that a new document version has been checked in, or prevent the deletion of a document without a manager's permission. In addition, workflows can call Content DB Web services for performing an extensive set of management actions, enabling the automation of a wide range of business processes.

## **Records Management**

With regulatory compliance issues becoming increasingly important, records management capabilities are also becoming an essential requirement for many organizations. Oracle Records Database (Records DB) works in concert with Oracle Content DB, to provide the ability to specify that a document or other content entity is to be retained for a certain period of time, prevents or controls changes to the document during the retention period and disposes of the document in a prescribed way once the retention period has expired.

Oracle Records DB provides a dynamic records management Web UI with secure login, and supports record organization based on a File Plan, record search, and flexible management of record classification, retention and disposition policies. Records may be declared manually by users or through automated policy-based record declaration applied on a folder-by-folder basis. Records DB includes the ability to place specific content under legal hold for litigation and related purposes. The product provides comprehensive storage management capabilities that enable customers to specify what types of online, near-line, and archival storage are to be used for various record types at different stages of their life cycles.

With the acquisition of Stellent, Oracle obtained a powerful new records policy engine that offers significant advantages over the original Records DB release. This engine, known as Universal Records Manager or URM, not only provides full DoD 5015 certification (including Chapter 4 for government classified content), it also utilizes a federated approach, which will allow it to manage records and retention policies not only for information in Content DB but also in products such as MS/SharePoint, file servers, third-party archives, and other content management applications. Given these superior capabilities, the URM/Content DB integration will supersede Records DB during calendar 2007.

## **ORACLE CONTENT DB: DIRECTION**

In the next few calendar quarters, the rich services and APIs, robust, highly scalable infrastructure, and process automation capabilities of Content DB will be utilized to host an expanding array of purpose-built content applications. These will include the Imaging and Business Process Management application (Oracle I/PM) obtained in the Stellent acquisition as well as an array of services permitting Content DB to be utilized as an *active archive* for a wide range of content. This later application will be described more fully in the upcoming white paper *Oracle Content Database: The Universal Active Archive*.

"Oracle Content DB and Records DB enables POWER Engineers to more easily manage the lifecycle of approved documents – from document and project completion to the recording of all relevant project documents. Oracle provides us a robust content and records management environment that addresses our compliance needs, integrates into our existing business processes, and is easy for our people to use."

**Loren Dugan**  
IT Director  
Power Engineers

## CONCLUSION

The content management market is evolving from specialized vertical applications to true enterprise deployments. These changes will expand the market substantially and require solutions that can bridge the huge gap between ubiquitous file servers and traditional high-end ECM products. Such solutions must provide horizontal functionality across the entire enterprise and complement existing departmental content management focused on specific vertical applications.

To meet the requirements for true enterprise deployments, a solution must deliver content management for all enterprise users, everywhere they want it, and for every business processes that requires it, it must provide users precisely the functionality needed without changing the way they work, and it must scale to support the largest enterprise— at costs that finally make enterprise deployment affordable.

Oracle Content DB meets these requirements, providing customers with increased productivity, reduced risk, enhanced compliance support, and lower cost, and making it possible for customers to benefit from consolidating their file servers, application content repositories, and content archives across their enterprise.

## FOR MORE INFORMATION

For more information about Oracle Content DB and Oracle Records DB, please visit:  
[www.oracle.com/database/contentdb.html](http://www.oracle.com/database/contentdb.html)



Oracle Content DB: Infrastructure for Content Consolidation

April 2007

Author: Rich Buchheim

Contributing Authors: Katherine McMillan

Oracle Corporation

World Headquarters

500 Oracle Parkway

Redwood Shores, CA 94065

U.S.A.

Worldwide Inquiries:

Phone: +1.650.506.7000

Fax: +1.650.506.7200

[www.oracle.com](http://www.oracle.com)

Copyright © 2006, Oracle. 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.