



An Oracle White Paper  
September 2010

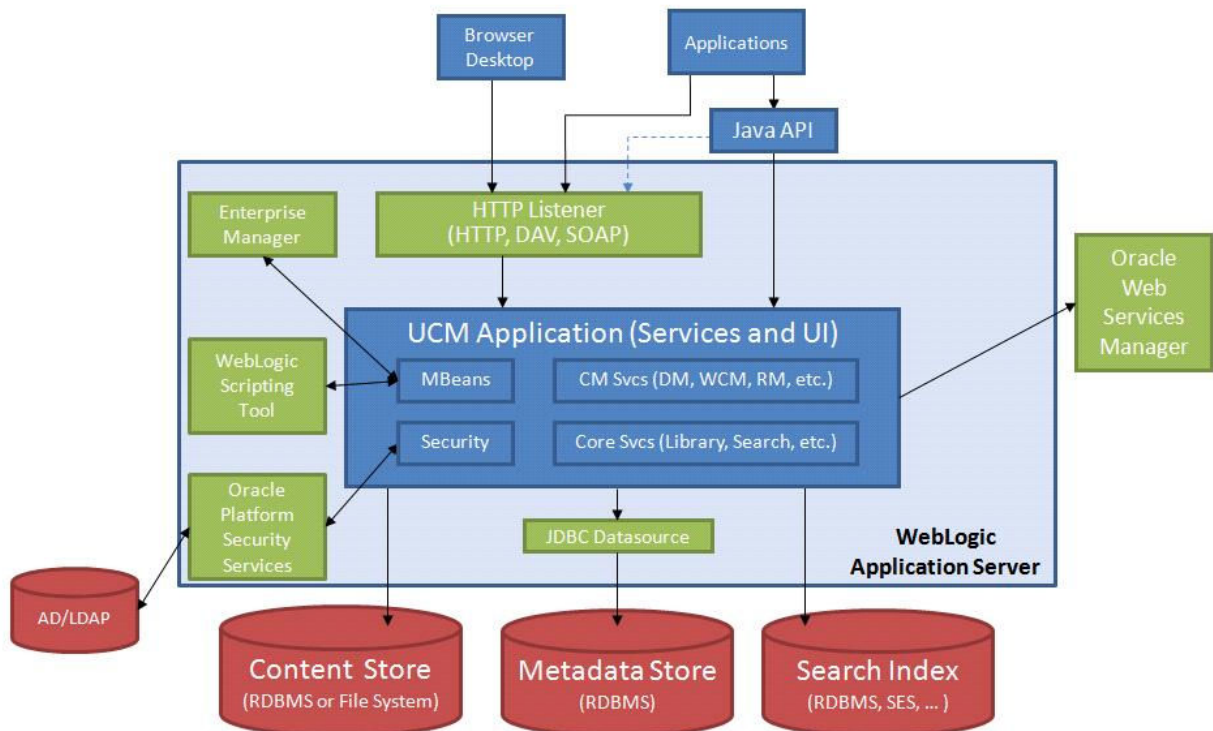
## Oracle Enterprise Content Management Suite Extreme Performance, Extreme Scalability

## Introduction

Oracle Enterprise Content Management Suite is a highly scalable, high performance enterprise content management system. All components of the suite, including Oracle Universal Content Management (UCM), Imaging and Process Management (I/PM), and Universal Records Management (URM), run on top of the Universal Content Management repository. The UCM repository is designed to support a wide range of content management solutions including high ingestion claims processing applications as well as high-consumption Web sites. The 11g version of UCM was enhanced to enable exceptional performance even when dealing with extremely large volumes of content, including ingestion rates of over one hundred million content items per day.

## Oracle Universal Content Management Architecture

The following diagram provides a high level view of the Oracle UCM 11g architecture:



Oracle Universal Content Management 11g is deployed as a managed server in Oracle WebLogic Application Server. UCM 11g is integrated out-of-the-box with various Oracle Fusion Middleware

Technology products such as Oracle Identity Management, Oracle Enterprise Manager and Oracle Web Services Manager. Because of these enterprise integrations, proven scalability, and content management features, UCM is the ideal strategic repository for all your unstructured content needs.

## High Volume Content & Imaging Application Characteristics

In many large organizations, critical business applications depend on handling large volumes of content. Accounts payable, accounts receivable, and human resources deal with extremely large volumes of paper documents that need to be scanned and contributed to a repository where they can be managed. Very large companies may actually need to manage millions of new documents per day.

In certain industries, such as insurance, you may deal with high volumes of documents even if you are a mid-sized organization. While explanation of benefits, explanations of payments, and claims processing may generate a great deal of content, the number of consumers may be very low (perhaps only a few hundred). Furthermore, content search functionality is streamlined so that content can be quickly retrieved using well known metadata (e.g. claim number, or customer ID) as opposed to ad-hoc keyword and full-text search. In addition, in these high volume imaging applications, content conversion (to HTML or PDF) is not required. Content is natively generated in TIFF or PDF format and is consumed in that format as well.

## High Performance Web Access / Delivery Characteristics

Many organizations use Oracle Universal Content Management to deliver high-performance Web sites. These sites may have a few contributors (in the case of an internet site) or many contributors (a corporate intranet). They may provide a strictly informational site, or combine that information with structured information as part of a catalog or custom Web application.

High performance public Web sites are characterized by relatively few contributors (tens or hundreds) but lots of viewers (thousands or millions). These sites typically provide strong metadata organization and full text searching capabilities. They tend to require robust content conversion capabilities – converting images (converting Photoshop images to multiple size renditions in JPEG format) or text (Word documents to HTML for viewing on laptops as well as mobile devices).

Applications	Input Volume	Viewers	Search	Conversion
Imaging/High Volume Ingestion	Extremely high	Few	A few fields	Little to none
Web/High Volume Delivery	Low to medium	Thousands to millions	Full-text, possibly combined with metadata	Extensive

## Content Ingestion & Web Delivery Performance with Commodity Hardware

Oracle conducted tests to measure how much content can be ingested into a single node UCM Content Server at a time. The tests used configurations similar to what would be expected of a claims processing or accounts payable application, running on commodity hardware (2 CPU 2.33 GHz Xeon®, 16 GB RAM).

The tests were conducted using a variety of file sizes (4 KB to 250 KB) and types (Text, MS Office, PDF). The results of these tests were as follows:

TABLE 2. ORACLE UCM 11g HIGH-VOLUME INGESTION TEST WITH VARIOUS FILE SIZES					
FILE SIZE	4 KB	20 KB	40 KB	100 KB	200 KB
Content Checked In Per Second	270.00	187.00	162.00	149.00	128.00
Content Checked In Per 24 Hours	23,328,000	16,156,800	13,996,800	12,873,600	11,059,200
UCM Server CPU Usage (%)	62.38	44.71	41.0	42.19	43.33
Database Server CPU Usage (%)	51.07	58.53	57.19	57.20	58.54

As the above results table shows, **a single node of Oracle UCM 11g** can ingest anywhere from **11 Million to 23 Million** content items per day on generic server hardware.

In addition to high-volume contribution, the same setup can provide high-speed delivery of dynamic Web sites, including corporate intranet sites, extranets, and public facing dot-com sites. Starting with UCM 11g, Oracle now offers a high-volume Web content management architecture called Site Studio for External Applications. The consumption performance metrics are likewise impressive on commodity hardware:

TABLE 3. ORACLE UCM 11g WEB CONTENT MANAGEMENT PAGE ACCESS / DELIVERY PERFORMANCE	
Page Access Per Second	124
Page Access Per Hour	446,400
UCM Server CPU Usage (%)	89%

One key point to note is that according to the CPU usage percentages, the software could process even higher rates of consumption with different infrastructure. This is evident because neither the UCM server nor the database had close to 100% CPU utilization rates. These kinds of applications are called

I/O limited, which means that overall performance is restricted by the speed of your hard disks, and the bandwidth in your network.

## Performance with Oracle Exadata Database Machine

In order to take performance to the next level, you will need specialized hardware to address the I/O limitations, such as the Oracle Exadata Database Machine. The Oracle Exadata Database Machine is the world's fastest for any type of database workload, and the only database machine that runs transaction processing applications. At the heart of this system is the Oracle Exadata Storage Servers, which has smart storage software built in. The smart storage software offloads data-intensive query processing from Oracle Database 11g servers and brings it closer to the data.

Oracle performed additional tests with 100KB files, which were run using multiple configurations of the Oracle Exadata Database Machine: quarter-rack and half-rack systems. The results were as follows:

TABLE 4. ORACLE UCM 11g INGESTION TEST WITH ORACLE EXADATA DATABASE MACHINE		
	SINGLE UCM SERVER WITH QUARTER RACK SYSTEM	TWO NODES UCM SERVER WITH HALF RACK SYSTEM
Content Checked In Per Second	1060	2070
Content Checked In Per 24 Hours	91,584,000	178,848,000

As can be seen from the results above, **a single node of Oracle UCM 11g** can ingest over **91 Million files per day** with a Oracle Exadata Database Machine Quarter Rack (which includes Oracle Database 11g Release 2 and Oracle Exadata Storage Server software), and almost **179 Million files per day** with Oracle Exadata Database Machine Half Rack.

An important point to note here is that once hardware and network bottlenecks are eliminated, **Oracle UCM 11g can deliver extremely high performance content management solutions at about a 95% scalability rate.**

## Conclusion

Oracle provides the most complete, open, and unified enterprise content management (ECM) platform – that enables you to build in a single repository both high volume imaging applications such as accounts payable and claims processes as well as high performance delivery applications, such as consumer-based Web sites. The solution is simple, smart and highly scalable – offering organizations extreme content ingestion in the tens of millions per day and only stopping at hardware limitations. Combining it with the Oracle Exadata Database Machine, the world's fastest for any type of database workload, Oracle Universal Content Management scales to hundreds of millions of items per day with only a two node UCM server configuration and Exadata Database Machine half-rack configuration.



Oracle Enterprise Content Management Suite  
Extreme Performance, Extreme Scalability  
Authors: Vijay Ramanathan, Michelle Huff  
September 2010

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  
oracle.com



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

Copyright © 2010, 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 and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0410

**SOFTWARE. HARDWARE. COMPLETE.**