An Oracle White Paper
September 2009

Oracle Database 11g
DICOM Medical Image Support

ORACLE

Introduction

Governments, licensing authorities, payers, patients and healthcare providers agree that information technology investments will result in reduced costs, increases in healthcare efficiency and improvements in patient outcomes. Studies place cost saving estimates at over $80 billion annually. There are several initiatives underway to accelerate the broad adoption of electronic medical records.

Medical images are a key component of any electronic medical record. Oracle Database 11g brings the benefits of a modern, high performance relational database to the problem of managing vast quantities of data required for medical imaging.

- Build DICOM content enabled Electronic Health Records (EHRs) for secure and efficient sharing of images while sharing patient data
- Increase the security, integrity and reliability of medical image archives retained for legal and compliance reasons
- Improve the reliability and performance of PACS systems used for local storage of diagnostic images and other DICOM content
- Build reliable and efficient repositories of images used in clinical trials, research and education

Trends Affecting Imaging In The Healthcare Industry

The healthcare industry is one of the most regulated, resource-constrained, and scrutinized industries in the world. It is under intense pressure to deliver the highest quality of service as efficiently and securely as possible, while addressing the needs of the public at low cost.

The Growth of Information Technology in the Healthcare Industry

For many decades, the healthcare system operated on paper and film. Registration and billing were paper based and diagnostic imaging was film based. Film was the norm in the diagnostic arena, and light boxes were the "browser" of choice. Information technology was first introduced in administrative systems, not in diagnostic or medical records systems. Over the past decade, medical records systems were introduced to manage patient records electronically and PACS were introduced to manage the new digital radiology modalities. All of these first generation information technologies and digital techniques have been widely adopted by hospitals, medical centers, and managed care organizations, typically at the departmental level. There is now a strong requirement for sharing of information to improve quality and lower costs, and this includes the need to share images across organizations and make them part of electronic health records.

Image repository solutions are becoming very large, increasing rapidly, for four reasons:

- Clinics and hospitals are purchasing more and more devices that can collect digital images.
- More kinds of digital modalities are being offered by vendors and purchased by healthcare enterprises.
- Newer machines produce images with higher resolution and thus, larger image data size.
- Enterprise, regional, and even national archives of medical images and other health information are being designed and deployed.

Organization-wide networks and RHIOs (Regional Health Information Organizations) increasingly require the sharing of these images along with patient data. Sharing is facilitated by:

- Vendor-neutral storage and a non-proprietary data model for storing images.
- Storage of images with patient data to support image-enabled EMRs.
- Addressing privacy concerns while sharing data with consistent and uniform privacy policies that apply to images along with the rest of the patient data.

Oracle's Investment in Imaging

Storing images in a database is not a new idea. For more than 20 years, Oracle customers have used the power of the database to provide unfettered access to stored images. Originally, these

images were stored in files and described in relational tables. With the introduction of binary large objects (BLOBs), images and other unstructured data could be stored directly in the database. Later, Oracle included full support for image, video, and audio data types in the database using Oracle Multimedia technology.

Oracle Multimedia Enables the Database to Understand the Nature of Images

Oracle Multimedia (formerly Oracle *inter*Media*)* is in its seventh major release, spanning over a decade of development. Oracle Multimedia is built on the database kernel, and operates as a privileged component of the database. The advantages of using Oracle Multimedia to store media objects are:

- Both the descriptions of an image and the image itself can be stored using industry standard formats.
- The Oracle Multimedia object model and methods make application programming simple and application maintenance far easier.
- Support for standard streaming output technology enables convenient retrieval and easy delivery to media players.
- Extraction and indexing of all image metadata is greatly simplified.

STORING DICOM CONTENT IN ORACLE DATABASE 11g

The healthcare industry and the National Electrical Manufacturers Association (NEMA) defined the Digital Imaging and Communications in Medicine (DICOM) standard to aid the distribution and viewing of medical images and other content useful in diagnosis and treatment. Oracle Database 11g supports this industry standard will full understanding of DICOM content through special data types and methods directly in Oracle Database 11g.

The DICOM standard defines both an approved way to store DICOM content such as single-frame and mutiframe images, waveforms, video segments, and structured reports, as well as over 2,000 standard tags to describe the information (metadata) about the DICOM content, the patient, the study, and other relevant attributes of the content.

There are numerous reasons why organizations contemplating large-scale storage of DICOM content should consider storing them in an Oracle database rather than in a separate file system.

Simpler Application Code

Oracle Database support for a specific type of content can include SQL language extensions, PL/SQL and JAVA and XML APIs and, in many cases, JSP Tag Libraries. The result is that developing applications that use this data involves substantially less code than without these language extensions. In addition, we frequently add algorithms that perform common or valuable operations through built in operators. By including operators that perform format

conversion or thumbnail generation of images, for example, the database greatly reduces the need for application logic.

Indexing and Searching

Oracle Multimedia populates an XML representation of all of the DICOM metadata, or a user-specified subset of the metadata, as an easily indexed and searched attribute of the DICOM object. The type of metadata search can be chosen based on the requirements: table column search using B-tree index for fast search of individual metadata attributes, XPath queries for flexible navigation of attribute hierarchy, full text search of all metadata, and linking in ontologies for capturing relationships between metadata terms and reconciling terms across multiple terminologies.

Flexibility

When DICOM content is stored in the database, it can be manipulated like other relational data. Sets of DICOM objects can be deleted, updated, or copied using simple SQL queries. Because the DICOM objects are presented to the application as columns in tables, they can be retrieved using relational joins or presented in views.

Improved Security and Auditing

The public desire for portable, electronic healthcare records including medical images is strong but it is conditional. People favor the establishment of regional and nationwide electronic information exchanges that would allow a patient's health information to be shared with authorized individuals, but want to ensure patient privacy and control over their own records. Oracle Database offers services that can assure both security and privacy of health information in the database instead of in an application.

Security

The U.S. Health Insurance Portability and Accountability Act (HIPAA) and similar regulations in other geographies require strict controls for access to medical information, including images. When DICOM content is stored as files, the application code must control access to these files. When DICOM content is stored in an Oracle database, the database can guarantee that unauthorized access to the content is forbidden. Oracle Database supports powerful role based security mechanisms to control access to all data in an archive. The Oracle Virtual Private Database (VPD) feature allows database administrators to strictly control the rows and columns visible to a database user and to mask information that should not be seen by that user.

Access to DICOM content can be controlled in file-based archives, but fine-grained access may have to be controlled by application code, adding complexity to the archive.

Auditing Access

DICOM content archives may be required to not only secure images and patient data, but also to keep an accurate record of all accesses to these repositories during their lifetime. Since data can be read, created, and changed using a variety of methods, implementing complete access auditing in application code is complex and often impractical. Database auditing records all access to DICOM content in an audit log allowing authorized personnel to monitor any event at the schema, row, statement, or content level of a column through fine-grained access controls.

Storage Reduction Techniques

With so much data being retained, another important consideration is how much space this stored data is using. Reducing storage requirements can provide performance improvements for data access as well as backup and recovery operations. Within Oracle Database, there are several techniques that can be used to reduce data storage requirements, including:

- Compressing DICOM content (when it is not already compressed)
- Removing unused space
- Shrinking data files

Encryption

If data, including DICOM content, must be secured even from the database administrators who manage the archive, Oracle's powerful encryption technology can be used to protect DICOM content from external access using secure encryption keys.

Transparent data encryption is a key-based access control system that enforces authorization by encrypting data with a secret key. No keys are stored in the database. They are stored in an Oracle wallet, which is part of the external security module.

Similarly, encrypting backup files can prevent access to DICOM content and patient data in misplaced backup tapes.

An Information Technology Platform for a New Generation of Healthcare Systems

The following sections describe how you can use Oracle Multimedia DICOM support to build powerful archives of images and other healthcare related data.

Oracle Multimedia DICOM Support Implementation

Oracle Database 11g adds a new data type for DICOM content: ORDDicom. This data type is similar in structure and function to a Java Class or a C++ object. Using this data type, any

column of any table can hold DICOM content. In relational terms, a DICOM object is an attribute of one instance of the entity that is stored in the table. Using a simple relational query on a table containing patient information, you can also retrieve the associated DICOM content.

The ORDDicom object used for DICOM content includes the following components:

- A set of metadata attributes
- The DICOM content, identical to the DICOM Part 10 file or bit stream delivered to the database
- A set of methods used to create, validate, store, anonymize, manipulate, or retrieve DICOM content

DICOM Metadata

For DICOM content, application metadata includes patient information, physician information, modality type, and series and study identifiers. DICOM metadata is mapped to an XML representation that is populated with the other object attributes when the DICOM content is stored in the database.

DICOM Object Content

To retain the integrity of the original DICOM content, it is preserved in a BLOB exactly as it was presented to the database in a transaction. When the DICOM content is retrieved from the database, Oracle guarantees that the exact same content will be delivered to the user. The ORDDicom object in the database contains the following attributes, easily accessible from an application:

- SOP_INSTANCE_UID: the SOP instance UID of the embedded DICOM content.
- SOP_CLASS_UID: the SOP class UID of the embedded DICOM content.
- STUDY_INSTANCE_UID: the study instance UID of the embedded DICOM content.
- SERIES_INSTANCE_UID: the series instance UID of the embedded DICOM content.
- source: the source for the embedded DICOM content.
- metadata: the XML representation of the DICOM metadata
- contentLength: the length of the embedded DICOM content, in number of bytes

Indexes can be built on these attributes for rapid access to specific ORDDicom objects.

DICOM Methods

Methods are functions or procedures that can perform operations on the DICOM content stored in the ORDDicom data type. The most important methods are these:

- *import()* – This method transfers DICOM content from a DICOM Part 10 file in the database server's file system to the database.
- *setProperties()*– This method parses the DICOM content and populates the ORDDicom object attributes.
- *extractMetadata()* – This method extracts a user-selected subset of the DICOM metadata and makes it available as an XML document.
- *processCopy()* – This method makes it possible to copy and convert DICOM content into a JPEG image or other popular supported formats for display in a browser or publication in a document. It can also perform other image processing operations such as scaling and cropping.
- *makeAnonymous()* – This method creates a new ORDDicom object from an existing one, removing or overwriting any patient identifying information. The DICOM attributes to be overwritten or removed are identified using an anonymity document specified by the user.
- *isConformanceValid()* – This method tests an ORDDicom object to determine if it complies with local requirements and a set of customized rules specified in a conformance document. Conformance documents can contain complex rules, invoke other conformance documents and can contain macros. The rules can check for conformance with the DICOM Standard or with local rules or policies.
- *export()* – This method produces a copy of the DICOM content in an ORDDicom object to an external file on the file system.

The DICOM Data Model Repository

All aspects of Oracle Multimedia DICOM support are controlled and managed by a series of XML documents stored in the database in a DICOM data model repository. A fully functional data model repository is installed when Oracle Multimedia is installed so users need only access the data model repository to modify default behavior.

The repository contains a series of XML documents:

- Dictionaries of public and private DICOM data elements
- Mapping documents controlling how DICOM attributes are mapped to XML representations of the metadata
- Anonymity documents that control how personally identifying information is removed from a copy of an ORDDicom object

- Constraint documents that contain rules of acceptability for various kinds of ORDDicom objects. These also contain the action to be taken when a constraint rule is violated including the logging messages used to describe the reason that an object violates a particular constraint rule
- A preference document that controls run-time preferences for Oracle Multimedia DICOM features
- A UID document that lists and categorized the registry of DICOM unique identifiers as defined by the DICOM standard

Update access to the data model repository is controlled by the new ORDADMIN database role. Administrators can reserve the data model to make changes and commit or roll back all changes made in one session. Users control when they see any changes committed by administrators using the setDataModel() procedure.

Loading DICOM Content in an Oracle Database

Oracle Multimedia supports loading DICOM content to an Oracle Database directly from files or from streams of data delivered by an application that implements the DICOM Communication Protocol. Loading DICOM content from files or creating DICOM format files from ORDDicom objects can be performed using simple PL/SQL procedures. In addition, bulk quantities of DICOM content can be loaded in a database using PL/SQL or Java procedures, SQL*Loader, or DataPump.

ENHANCEMENTS IN ORACLE DATABASE 11g Release 2

Performance Enhancements in Metadata Extraction

Performance of DICOM metadata extraction has been improved in two ways.

- A subset of attributes can now be specified for extraction. DICOM images can contain hundreds of metadata attributes, but often only a few are required for indexing, searching, and partitioning. Extraction of only a subset of attributes improves performance, and the specification of this subset is completely flexible.
- Metadata extraction can now be performed outside the database, by a client tool or in the middle-tier. This can improve performance and also satisfies requirements that need the metadata to be extracted before the images are loaded into the database. For example, to partition DICOM data efficiently, the partitioning key must be populated before the data is loaded into the database. If DICOM metadata is part of the partitioning key, it can now be extracted outside the database so DICOM data can be loaded directly into the correct partition.

Other metadata extraction enhancements include extraction of a portion of a metadata attribute such as the patient's last name.

DICOM Content Processing

DICOM vidoes can be converted to AVI and MPEG videos for viewing in any web browser. Like the conversion of DICOM images into JPEG for viewing for non-diagnostic purposes DICOM videos also can be converted into web friendly formats.

More complex constraint definitions can now be specified which is useful for testing conformance of DICOM structured reports.

Use Cases

This section examines several possible uses for Oracle Multimedia DICOM support.

DICOM-enabled EMRs

Oracle Multimedia facilitates inclusion of DICOM content in Electronic Medical Records with minimal development effort. EMR systems connect to different sources of patient data and provide unified access to the data in a centralized store. This enables clinicians across the hospital to access and share patient data and also to share the data with referring physicians outside the hospital. DICOM support in Oracle Multimedia allows the EMR systems to include DICOM content in the patient record for a complete and 360-degree view of all patient data. Since the management of DICOM data is native in Oracle Database, EMR applications can easily be extended to include DICOM content without being burdened by the infrastructure for managing DICOM content. EMR systems can include DICOM metadata in search applications and provide a web friendly view of DICOM content by converting images to JPEG and video to AVI or MPEG formats.

Managing DICOM content in an Imaging Modality

Oracle Multimedia can simplify the application code needed to create an image and assemble the DICOM metadata used in diagnostic medical devices. Modern modalities sometimes store DICOM content for weeks or months after transferring the DICOM content to a PACS or DICOM archive. Oracle Database can make that storage secure and reliable with high performance storage and retrieval.

PACS / RIS

Many current PACS systems store information about patients, schedules, and DICOM content in an Oracle database and store the DICOM content separately as files. PACS vendors can reduce

their software maintenance costs and improve security, transaction integrity, auditing, indexing, retrieval, and management by allowing Oracle Database to manage the storage and retrieval of DICOM content. This allows PACS vendors to devote more resources to application functionality, visualization tools and ease of use rather than on platform data management services.

By storing DICOM content as part of the database transactions, synchronization between patient and study information and related DICOM content is enforced, even if the information has to be restored after a hardware failure. Hospitals and clinics that purchase these PACS systems save on operational costs because a single procedure is used to back up both the DICOM content and the data.

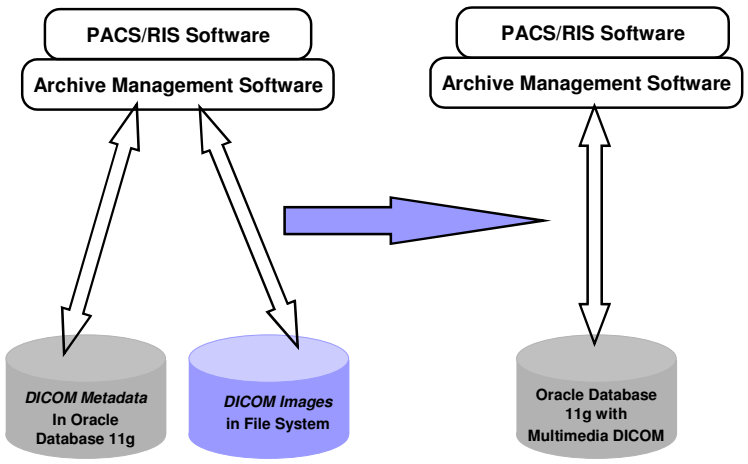

FIGURE 1: USING ORACLE MULTIMEDIA 11G DICOM TECHNOLOGY WITH A PACS SYSTEM

Enterprise DICOM Archives

Large healthcare organizations may select PACS and modalities based on department and specialty needs but often rely upon corporate IT to manage long term storage of the DICOM content stored in these departmental PACS. This results in several expensive problems:

- PACS require professional management of large quantities of storage in each local department.
- As PACS storage needs grow, resources must be provisioned throughout the enterprise in various storage system environments.

- DICOM content stored in one PACS is not readily transportable to another department. Physicians may order expensive new studies rather than use those from another department.

Corporate IT organizations seek to reduce these costs by implementing a centralized archive for long-term storage of the DICOM content created in the various departments. After images are created and interpreted in the departmental PACS, they are transferred to the central archive. If they are needed again, they can be easily accessed from the central archive or copied back to a departmental PACS.

Oracle Database can ease the development and management of enterprise archives by providing atomic transactions, simple application programming tools, backup and disaster recovery solutions and powerful and economical grids for middleware, database servers and storage.

Research Data Warehouses

The manner in which research on drugs and clinical treatments is conducted by academic and commercial organizations has changed substantially. The world's largest drug companies frequently contract with outside organizations for clinical trials as part of their development process. Once satisfied with reports from these contractors, many research organizations now expect to collect DICOM content from these external partners. In addition, the purchase of medical records, which have been made anonymous from regional and national health delivery organizations, is being used to build massive medical information warehouses to support discovery and research. Longitudinal studies of medical data and DICOM content from these information warehouses promise new insights into drugs and treatments as well as greater patient safety.

Using Oracle Multimedia DICOM technology, all biomedical data can be gathered into the same database and mined together. For example, in a disease discovery research scenario, a newborn baby's medical record could be combined with ultrasound images, videos, and gene sequence data in an Oracle database. That database could then be searched to find correlations among genetic anomalies and image content. This process could lead to the discovery of new disease screening methods with input from multiple modalities.

National or Regional Medical Image Repository

Many governments and other organizations are building regional, and even national repositories of medical records. These repositories are expected to provide a number of benefits, including these:

- Lower costs, by reducing the number of duplicate tests.
- Better care, by offering each physician complete medical records, including DICOM content, for each patient encounter.
- More efficiency, by supporting remote reading and diagnosis of medical images.

- Portability of medical records for patients who seek treatment at any clinic or hospital in their area.

In this model, local clinics and hospitals control DICOM content in local Picture Archive and Communication Systems (PACS). When appropriate, the DICOM content and the responsibility for maintaining it is transferred to the central repository. If any physician or clinician needs the DICOM content again, it is available as part of the patient's medical record. By establishing a central repository of record for DICOM content incorporated with the patient's record, clinics and hospitals can remove the practice of patients and staff carrying film and CDs from practitioner to practitioner.

National and regional projects that follow this model are underway in the United Kingdom, Canada, Sweden, the United States, and many other countries.

Clinical Trials Image Management

When pharmaceutical companies contract with outside organizations for clinical trails their control over the clinical research image data is limited. Their access to data is indirect, expensive, not timely, and insecure. When the data is transferred from the contract research organizations to the pharmaceutical companies the process is ftp or media/courier based, which is disjointed, manual, and not audit trailed. Lack of image anonymization tools makes it hard to meet patient privacy regulations. Finally the data quality is highly variable with no tools to monitor quality of image data coming in from the external organizations.

Oracle Multimedia provides the infrastructure for a solution that is centered on the needs of the pharmaceutical company. Using partner solutions on this infrastructure, pharmaceutical companies can have direct, immediate, cost-effective access to clinical research image data. Data quality standards can be consistently applied, enhancing data quality and increasing transparency and accountability of the external contract organizations. There is more control over the data for improved review, reporting, analysis, and reanalysis. Security, access control, and regulatory compliance needs are addressed effectively by the rich security support of the database.

Conclusion

The support that Oracle Multimedia provides for safe, secure, and auditable storage of DICOM content in large-scale databases offers the potential to improve the quality and reduce the cost of healthcare. By enabling the delivery of historical DICOM content throughout an enterprise or a region, healthcare enterprises can benefit from the reduced need to repeat expensive tests, simplified delivery of remote radiology services, easier access to DICOM content by remote physicians, lower IT management costs and easier, auditable compliance with privacy and data security regulations.

Patients will benefit from accessible repositories of DICOM content when they can go to any hospital or clinic in a region and have all their historical medical records, including images, securely accessible at that location.

DICOM support in Oracle Database 11g can ease the development effort for vendors of next generation PACS products and regional health archives by eliminating the development and maintenance effort required for tasks such as format conversion, indexing of metadata, and data security.

Managers of these systems will achieve cost savings because they will use the same tools to load, manage, and back up traditional relational data tables and DICOM content.

Oracle Database 11g DICOM Medical Image Support
September 2009
Author: Melliyal Annamalai
Contributing Authors: Dongbai Guo, Susan Mavris, Jim Steiner

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

0109