Improving Healthcare Provider Performance with Big Data

Architect's Guide and Reference Architecture Introduction

ORACLE ENTERPRISE ARCHITECTURE WHITE PAPER | FEBRUARY 2015

Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Table of Contents

Executive Summary

Healthcare Providers (including hospitals, clinics, rehabilitation centers, and extended care facilities) are facing a growing need to manage costs while improving quality of care and focusing on patient outcomes. In general, the industry's desire is to move towards evidence based medicine compared to trial and error approaches. In order to meet these goals, organizations are analyzing and managing vast volumes of clinical, genomic, financial and operational data while rapidly evolving their Information Architecture.

Health Information Organizations (HIO's) have long gathered information about patient encounters but it is only in the last few years that much of this information has entered the digitized world in the form of EMR (Electronic Medical Record) adoption which allows quick access to this data in a near real time basis. The data can help expose different treatments and their associated outcomes. Even though clinical data is fuelling initial set of analytics platforms, Provider organizations are looking beyond clinical information alone to provide superior care while reducing cost. Data sources that can be used together to gain a full picture of a patient can include:

- Clinical
  - EMR
  - Pharmacy and related
  - Radiology images
- Research
  - Genomic sequencing
  - Clinical trials
- Operational
  - Supply chain
  - Billing / Financial
- Customer Experience
  - Wearable devices
  - Customer sentiment and social

The rate that this data is generated is rapidly increasing leading to higher rates of consumption by the business analysts who crave such information. This increase in data velocity and sources naturally drives an increase in aggregate data volumes. Business analysts want more data to be ingested at higher rates, stored longer and want to analyze it faster. Emerging "Big Data" solutions are helping to HIO's to meet these requirements.

Other non-patient related data can also be critical for Providers to run their operations efficiently. Such data can assist in efficient energy management of buildings, proper and timely maintenance of equipment and facilities, and optimal management of supplies needed to run the facility and provide the best possible care.

This paper provides an overview for the adoption of Big Data and analytic capabilities as part of a “next-generation” architecture that can meet the needs of Healthcare Providers seeking to improve outcomes and efficiencies while increasing profitability.

This white paper also presents a reference architecture introduction. The approach and guidance offered is the byproduct of hundreds of customer projects and highlights the decisions that customers faced in the course of their architecture planning and implementations. Oracle’s advising architects work across many industries and government agencies and have developed standardized methodology based on enterprise architecture best practices. Oracle’s enterprise architecture approach and framework are articulated in the Oracle Architecture Development Process (OADP) and the Oracle Enterprise Architecture Framework (OEAF).

Key Business Challenges

Healthcare Providers historically used data warehouses and business intelligence tools to report on and analyze financial results, optimize operation of their facilities, and measure outcomes and quality of care. By deploying Big Data Management Systems that include data reservoirs (featuring Hadoop and / or NoSQL Databases), greater benefits can be reaped through advanced analytics. Typical goals these applications target include:

- Reduced hospital readmissions.
- Predicting high risk patients for Accountable Care Organizations (ACO).
- Simulation of patient reported outcomes to improve care quality and outcomes.
- Acceleration of time to market for new therapies with strategic portfolio modeling.
- Prediction of market access and optimized resource allocation for new therapies.
- Analysis of connected health consumers and recommended healthy behavior change.
- Quantification of health costs and recommended wellness interventions.
- Simulation of the financial risks and incentives of emerging reimbursement models.

Improving Patient Outcomes

Healthcare Providers seek to differentiate themselves by providing outstanding patient outcomes and patient experience, leading to referrals from patients and medical personnel. They must deliver services in a predictable and safe manner and offer transparency to payers and regulatory agencies. As the types of care and coverage models used by payers change over time, they must also intelligently forecast the likely impact of these changes and develop new business models that may overlap with services formerly provided in hospitals, rehabilitations centers, clinics, and / or extended care facilities. Hence, predictive analytics can have an important role in planning for such changes.

As sensors and associated applications become more prevalent and usable by home-bound patients, the number of data sources a Healthcare Provider will want access to will only grow. Such data volume will dwarf today’s data warehouses and require Big Data Management systems for processing and reporting.

Improving Efficiency and Managing Costs

Data warehouse solutions are often used today to understand staffing and facilities utilization trends, and other costs of running the business, including maintenance and supplies. Predictive analytics are sometimes used to predict whether current trends will continue and possible implications.

Sensor driven data will increasingly provide more information about the state of the facilities, their maintenance, and supplies in inventory. There is an opportunity to better understand when maintenance needs to occur affording greater efficiencies and cost savings. Predictive analytics solutions deployed across Big Data Management Systems (including Hadoop) will likely become a more common practice for such applications.

Where to Find Business Cases that Justify Projects

Many existing business capabilities can be enhanced when more and varied data becomes part of the Information Architecture. IT organizations at Healthcare Providers typically work with their lines of business to build solutions that deliver the following when defining Big Data projects:

1) Improved Quality of Care and Outcomes: Improving patient outcome and illness prevention is at the core of healthcare. The transition to a more holistic approach centered on the patient is accelerating. Now Providers are accountable for the cost and quality of care and outcomes. The reimbursement model is heading toward payment for performance. Hospital readmissions within a 30-day period for the same illness can incur penalties.

2) Improved Clinician Performance: Challenges related to cost, quality, and access to healthcare impact delivery of services and raise the question of performance of care givers and accountability at every level of care. Understanding whether doctors, and care givers are providing adequate care while meeting all standards and using that information to better train staff and optimize business performance is a growing focus area. Healthcare Providers must do so in a way that doesn't discourage highly competent employees and cause them to look elsewhere for employment.

3) More Flexible Patient Monitoring: The marriage of Electronic Medical Records (EMRs) with on-premise monitors promises more proactive and better diagnoses sooner with fewer errors compared to manual records of the past. A growing array of monitoring devices and sensors enable self-monitoring in homes by patients and the opportunity to leverage such data (and flexibility in location) by Healthcare Providers to further enable better quality of care while reducing cost. The consumer market featuring wearable devices performing health and wellness monitoring is transforming expectations and usage of technology. This also impacts healthcare organizations as care givers and patients alike expect instant access to and transfer of information regardless of location and device used. Accessing medical data at the point of care via mobile devices helps to reduce medical errors and improves productivity. The mobility in healthcare organizations is significantly increasing for better healthcare experience as a whole.

4) Better Management of Facilities, Staffing, Supplies, and Equipment: Healthcare Providers have a huge investment in facilities, staffing, supplies, and equipment. The use of data gathered from sensors and other sources of data should enable better utilization of facilities and equipment, but also eliminate waste and needless expense resulting in higher effectiveness.

5) IT operational efficiency: Not unique to Healthcare Providers and rarely driven from the lines of business (but a possible reason for embarking on extended architectures that include Hadoop) is the need to move data staging and transformation to a schema-less platform for more efficient processing and leveraging of IT resources. IT operational efficiency is often difficult to prove but is sometimes an initial justification that IT organizations gravitate toward when deploying these types of solutions.

On the next page, we show a table that summarizes several typical business challenges in Healthcare Providers and illustrates the opportunity for new or enhanced business capability when adding new analytic capabilities.

TABLE 1 – HEALTHCARE PROVIDER FUNCTIONAL AREAS, BUSINESS CHALLENGES & OPPORTUNITIES

FUNCTIONAL AREA	BUSINESS CHALLENGE	OPPORTUNITY
Care Planning and Quality of Care	Provide the right care to achieve the best outcomes at the lowest cost	Improve care by doctors & staff Increase patient satisfaction, referrals Identify early high risk patients Gain timely payer reimbursement Decrease litigation cost Move patients to lower cost facilities when viable
Clinician Performance Management	Understand doctors, nurses, and other healthcare givers ability to provide proper quality of care	Decrease potential litigation Educate underperformers on how to improve Improve quality of care Decrease cost of care
Outpatient and In-Patient Monitoring	Centralized monitoring of patients	Decrease cost of staffing, travel to home-bound / off site patients Improve quality of care through more timely detection or prevention Improve patient satisfaction
Financial Administration & Human Resources	Optimal management of hospital utilization and staffing	Grow / gain profitability Hire and maintain right mix of staff in the right place at the right time Predict services demand Take facilities / floors offline when underutilized
Supply Chain Management	Optimal management of drugs and medicines, treatment supplies, food, and other facilities supplies	Lower waste and warehousing Provide better quality of care Detect potential fraudulent activities
Promotions and marketing	Optimal spending on promotions to drive profitable business results	Gain and maintain patients Cross-sell / up-sell to existing Improve community profile
Maintenance	Optimal maintenance scheduling for equipment and facilities	Minimize breakdowns Minimize potential liability Extend life of equipment

Establishing an Architectural Pattern

The following Figure illustrates key components in a typical Information Architecture. Data is acquired and organized as appropriate and then analyzed to make meaningful business decisions. A variety of underlying platforms provide critical roles. Management, security and governance are critical throughout and are always top of mind in Healthcare Providers. These components are further described in the "Information Architecture and Big Data" whitepaper posted at http://www.oracle.com/goto/ea.

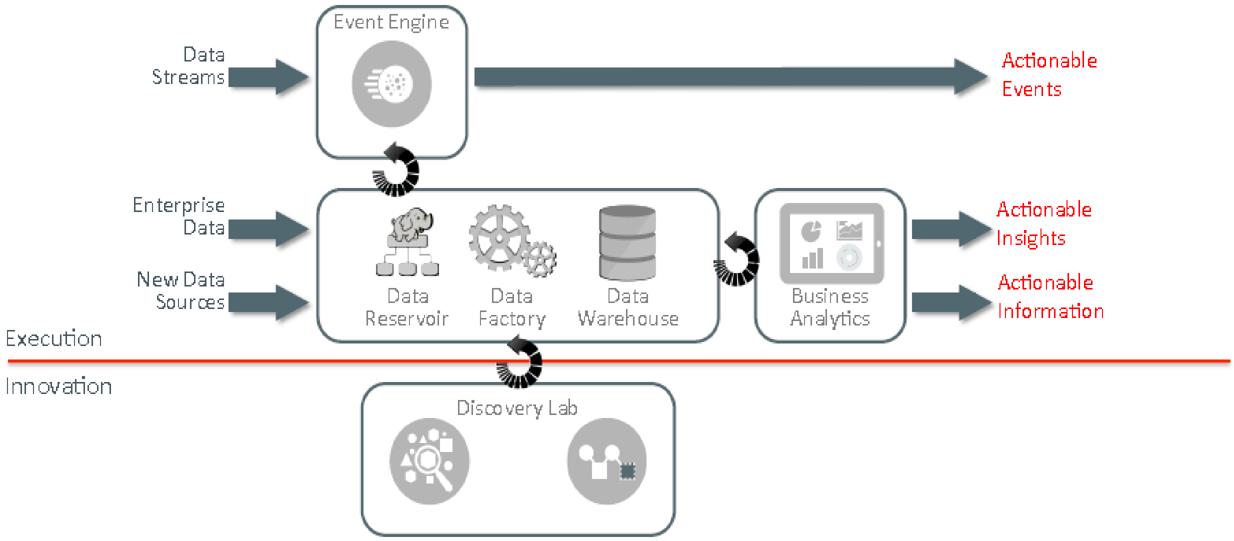

Figure 1: Key Information Architecture Components

How do we determine which of these components should be part of the architecture to meet the needs of a specific organization or company? If we create an Information Architecture diagram and trace the data flow from sources to the applications (end-user), we can build a logical configuration of the components to support the functions.

The first step in defining a future state architecture is documenting the current state, its capabilities and any functional gaps. Typically a current state data warehouse environment might look something like Figure 2.

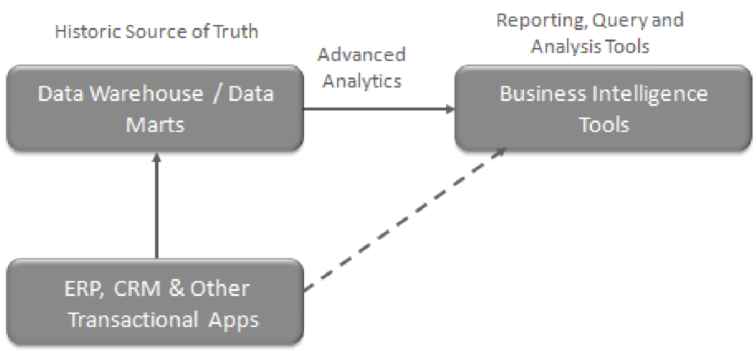

Figure 2: Typical Current State Data Warehouse

The first gap that typically has to be closed is a need to provide a more agile reporting and analysis environment where new data and ad-hoc reports are needed on an ongoing basis. Information and data discovery engines can provide this type of capability. When information discovery is incorporated into the architecture it would look something like the illustration in Figure 3.

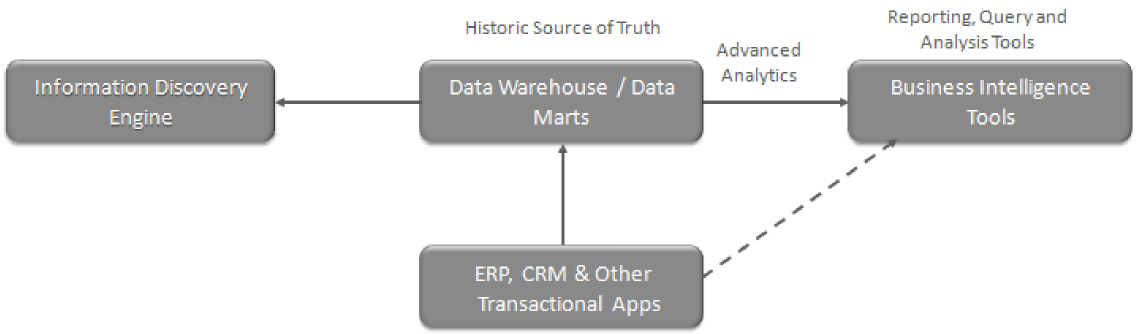

Figure 3: Typical Introduction of Information Discovery

Now that we're better able to analyze the data we have, the next step would be to explore bringing in new data and new data tapes. These data sets might be internal, 3rd party, structured, unstructured or of unknown structure. When storing data of unknown structure, the most efficient way to store data sets is often in a Hadoop-based data reservoir. Initially, such projects are often considered experimental in organizations and therefore they might be independent efforts separated from the traditional environments, as illustrated in Figure 4.

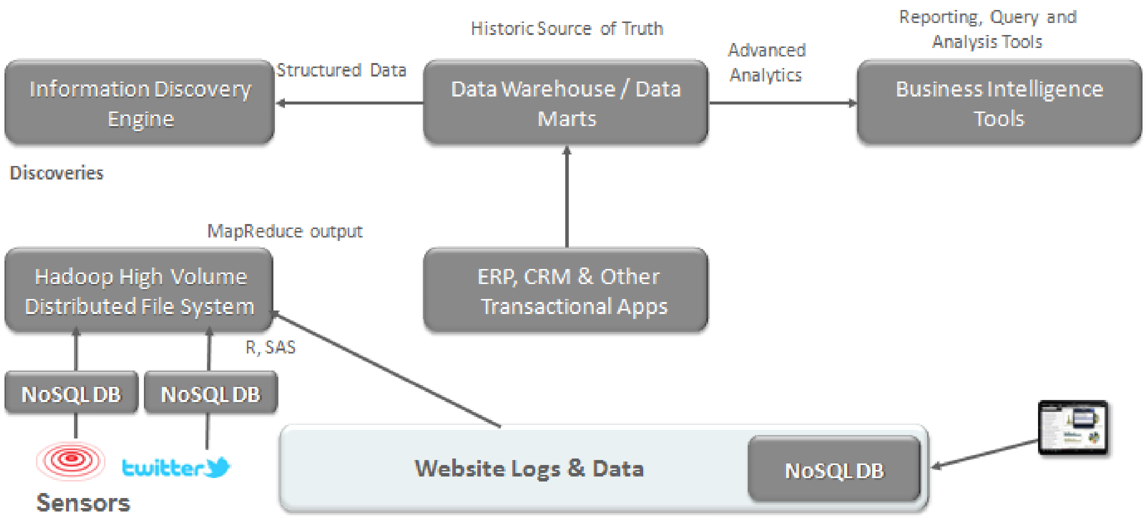

Figure 4: Typical Early Hadoop Environment separate from the Data Warehouse

The profile of the data such as how it is acquired, how it should be formatted, the frequency of updates and quality of the data will help us put the right technology in place best suited for the particular situation. We need to understand whether real-time or batch processing is appropriate. We should understand the periodicity of processing required based on data availability. Below is a partial list of the characteristics that should be considered:

» Processing Method – prediction, analytics, query, ad-hoc reports
» Format and Frequency – external data feeds, real-time, continuous or periodic on-demand
» Data Type – web/social media, machine generated, human generated, biometric, legacy or internal, transactional
» Consumer Application – Web Browser, Intermediate processes, Enterprise Application

When business value is found in analyzing data in a Hadoop-based data reservoir, lines of business generally begin to see a need to link data there to historical data stored in their data warehouse. For example, a business analyst might want to compare historical transactions for a shipment stored in the data warehouse to sensor data tracking that shipment in the data reservoir. Various linkages are often established as pictured in Figure 5.

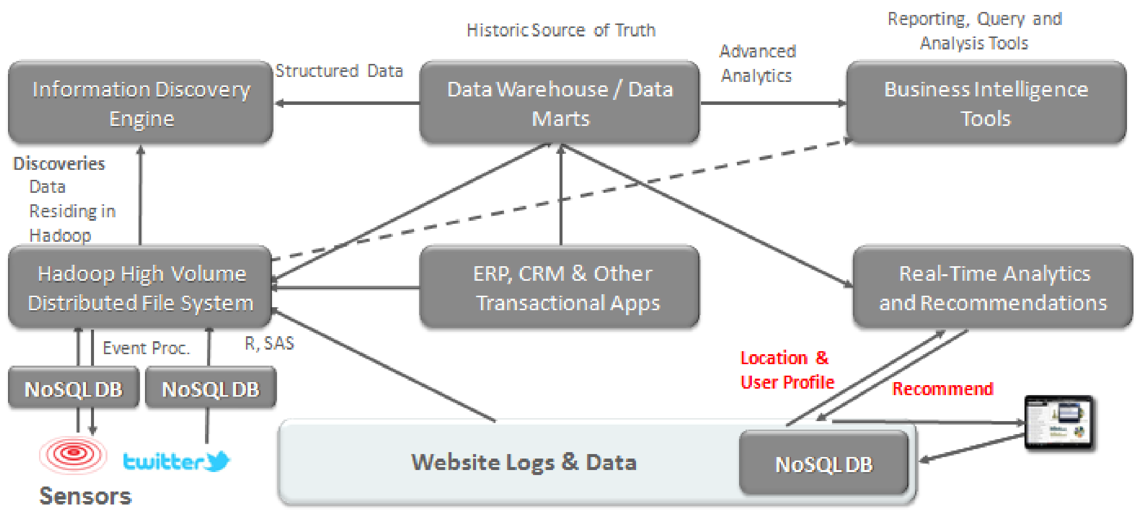

Figure 5: Integration of Hadoop Infrastructure and Data Warehouse

We also added something new to Figure 5, a real-time analytics and recommendation engine. In many situations, the latency inherent in the data movement pictured above means that the recommendation from analysis would come too late to take action in near real-time. A way around this is to perform periodic advanced analytics in the data reservoir and / or data warehouse and provide updates to a real-time recommendation engine that becomes more fine-tuned through self-learning over time.

IT Operational ETL Efficiency

In Figure 5, you might have noticed a line pointing from the transactional sources to the Hadoop cluster. This is to illustrate a popular ETL alternative, leveraging Hadoop as a data transformation engine.

Let's now consider the type of data typically stored in today's data warehouse. Such warehouses are typically based on traditional relational databases using a "schema on write" data model. The data sources can vary, but the structure of the data is determined before the data in imported into the data warehouse. In the example below there are two data sources. These two data sources go through an ETL process to prepare the data to be loaded into the warehouse.

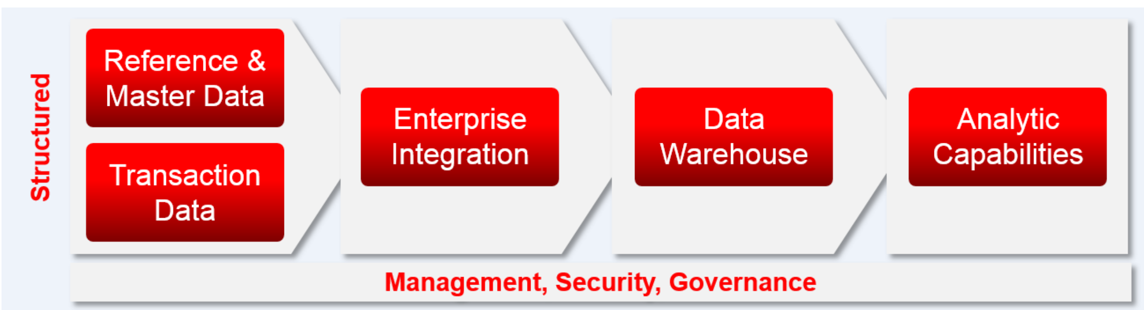

Figure 6: Structured Data and the Data Warehouse

Extending the architecture can enable a more agile workflow by incorporating data sets for which there is not rigid structure. This data model is best defined as "schema on read". That is, we store the data without the traditional ETL processing, as we don't know exactly how we want to access the data. In the example below we are using multiple data sources with varying structures.

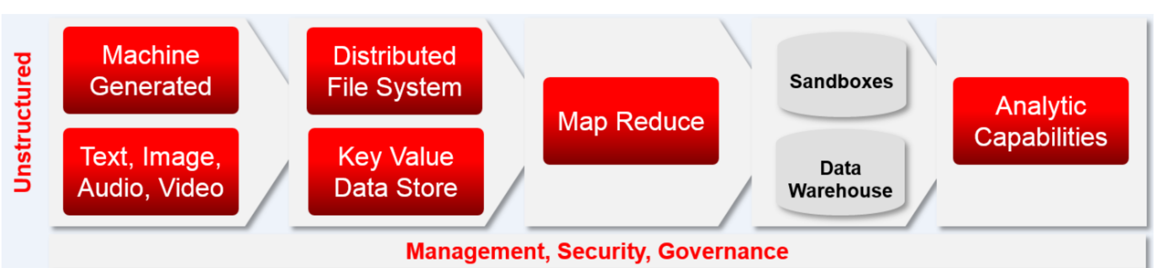

Figure 7: Unstructured Data, Distributed File Systems and Key Value Data Stores

These two environments should not be separate and unique. Building an integrated Information Architecture that can handle data sets of known structure as well as unknown structure enables us to augment the capabilities of existing warehouses as well as leverage data center best practices that are already in place.

Oracle Products in the Information Architecture

In Figure 8, we illustrate how key Oracle products could fit in the generic architecture diagram previously shown.

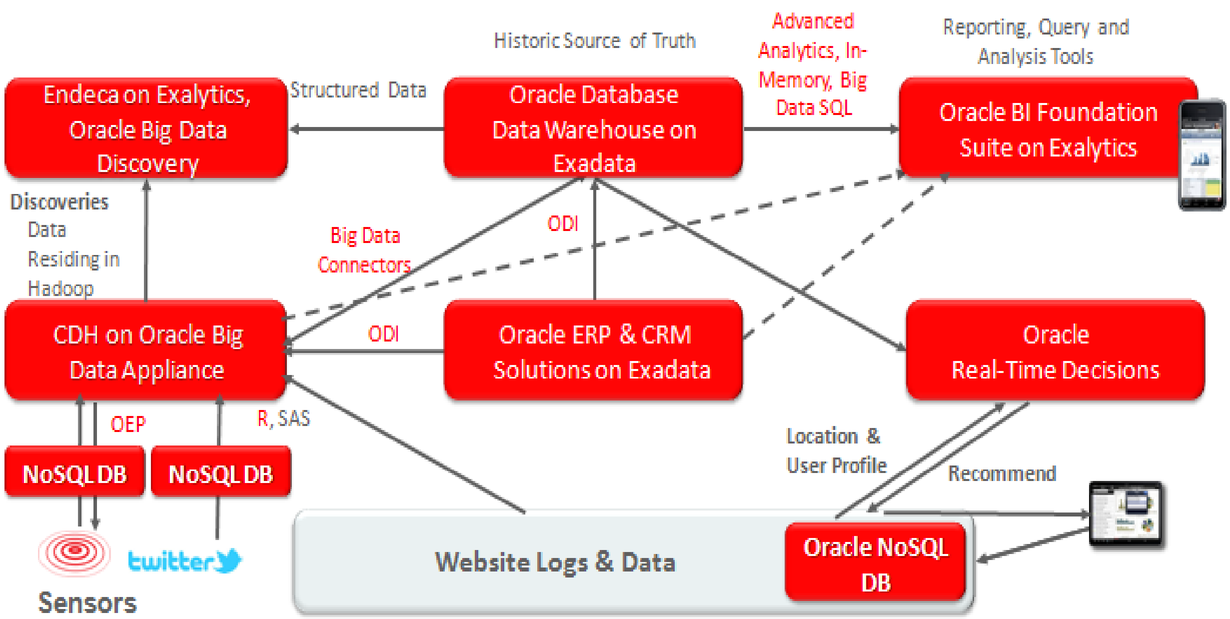

Figure 8: How Key Oracle Products Fit in the Generic Architecture

While Oracle can provide a more complete integrated solution, many organizations mix and match products from a variety of vendors. Therefore, such architecture diagrams often show such a mixture of products from Oracle and other vendors.

Defining an Information Architecture is all about linking it to a specific use case. For example, a footprint that includes various operational sources and the Oracle Healthcare Data Repository and Healthcare Data Model might look like Figure 9:

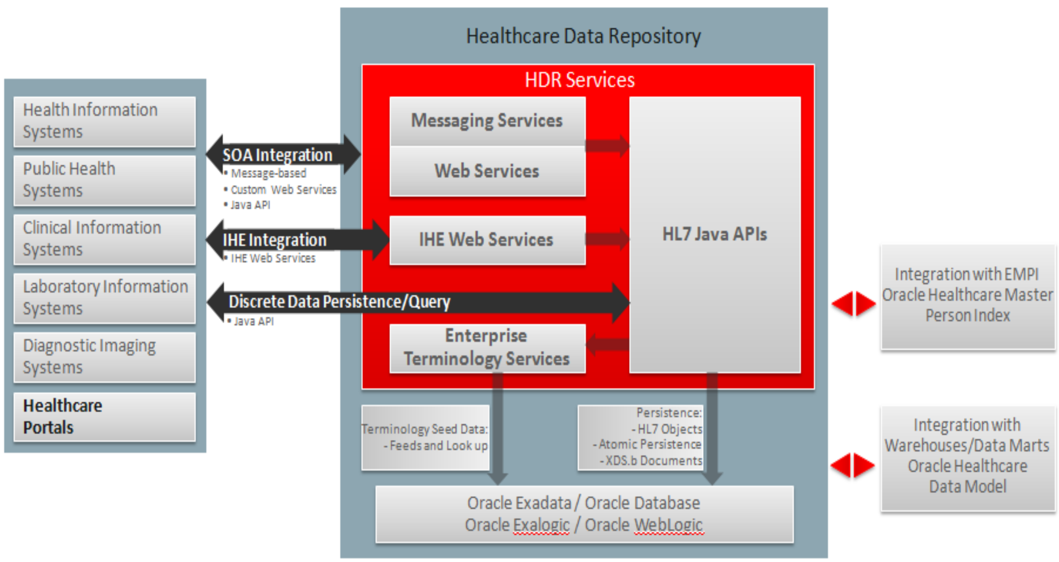

Figure 9: Typical Provider Data Repository

The various software capabilities required in a typical architecture might include these Oracle components:

» Oracle Relational Database Management System (RDBMS): Oracle Database 12c Enterprise Edition is designed for performance and availability, security and compliance, data warehousing and analytics, and manageability. Key data warehousing options often include In-Memory, OLAP, the Advanced Analytics Option, and Partitioning.

» Oracle Business Intelligence Enterprise Edition (OBIEE): A business intelligence platform that delivers a full range of capabilities - including interactive dashboards, ad hoc queries, notifications and alerts, enterprise and financial reporting, scorecard and strategy management, business process invocation, search and collaboration, mobile, integrated systems management and more.

» Oracle Real-time Decisions: A real-time recommendation engine.

» Hadoop Distributed File System (HDFS): A scalable, distributed, Java based file system that is the data storage layer of Hadoop. Ideal for storing large volumes of unstructured data.

» Flume: A framework for populating Hadoop with data via agents on web servers, application servers, and mobile devices.

» Oracle Data Loader for Hadoop: A connectivity toolset for moving data between the Oracle RDBMS and the Hadoop environment.

» ODI: Oracle Data Integrator is a comprehensive data integration platform that covers all data integration requirements: from high-volume, high-performance batch loads, to event-driven, trickle-feed integration processes, to SOA-enabled data services.

» Oracle Enterprise Metadata Management: Data governance and metadata management tool providing lineage and impact analysis, and model versioning for business and technical metadata from databases, Hadoop, business intelligence tools, and ETL tools.

» Endeca: An information discovery tool and engine.

» Oracle Big Data Discovery: A Hadoop-based information discovery tool.

» Oracle Big Data SQL: An optimal solution for querying an Oracle Database on Exadata and combining the results with data that also answers the query and resides on Oracle's Big Data Appliance.

» ORE: Oracle R Enterprise enables analysts and statisticians to run existing R applications and use the R client directly against data stored in Oracle Database (Oracle Advanced Analytics Option) and Hadoop environments

- Oracle Enterprise Manager: An integrated enterprise platform management single tool used to manage both the Oracle structured and unstructured data environments and Oracle BI tools.
- Oracle Essbase: An OLAP (Online Analytical Processing) Server that provides an environment for deploying pre-packaged applications or developing custom analytic and enterprise performance management applications.

The software products listed above can be deployed in an integrated environment leveraging these engineered systems:

- Big Data Appliance (BDA): Eliminates the time needed to install and configure the complex infrastructure associated with build-out of a Hadoop environment by integrating the optimal server, storage and networking infrastructure in a rack.
- Exadata: Streamlines implementation and management while improving performance and time to value for Oracle relational database workloads by integrating the optimal server, storage and networking infrastructure.
- Exalytics: Provides an in-memory server platform for Oracle Business Intelligence Foundation Suite, Endeca Information Discovery, and Oracle Essbase.

Obviously, many variations are possible. Healthcare Providers generally rely on a mixture of EMR and ERP applications from vendors and highly customized data sources.

Additional Data Management System Considerations

In defining the Information Architecture, it is important to align the data processing problem with the most appropriate technology.

When considering the choices you have in database management systems to include in an Information Architecture, you might consider if the form of the incoming data or ACID properties or fast data availability is most important. Other considerations should include manageability, interoperability, scalability, and availability. Of course, you should also consider the skills present in your organization.

Some of the various data management technologies in a typical architecture include:

Relational Databases

Typically already in use at most companies, RDBMS' are ideal for managing structured data in predefined schema. Historically they excel when production queries are predictable. Support of dimensional models makes them ideal for many business intelligence and analytics workloads. They frequently house cleansed data of known quality processed through ETL workloads. Relational databases also excel at transactional (OLTP) workloads where read / write latency, fast response time, and support of ACID properties are important to the business.

These databases can usually scale vertically via large SMP servers. These databases can also scale horizontally with clustering software.

Example RDBMS Product: Oracle Relational Database

MOLAP Databases

Typically used for highly structured data, MOLAP databases are ideal when you know what queries will be asked (e.g. facts and dimensions are predefined and non-changing) and performance is critical. These databases excel at certain business intelligence and analytics workloads.

Example MOLAP Product: Oracle Essbase, Oracle Database OLAP Option

NoSQL Databases

NoSQL databases are without schema and are designed for very fast writes. Often, they are used to support high ingestion workloads. Horizontal scale is most often provided via sharding. Java and Java scripting (JSON) are commonly used for access in many of the commercial varieties.

NoSQL databases are sometimes described as coming in different varieties:

Key Value Pairs: These databases hold keys and a value or set of values. They are often used for very lightweight transactions (where ACID properties may not be required), and where the number of values tied to a key change over time.

Column-based: These databases are collections of one or more key value pairs, sometimes described as two dimensional arrays, and are used to represent records. Queries return entire records.

Document-based: Similar to column-based NoSQL databases, these databases also support deep nesting and enable complex structures to be built such that documents can be stored within documents.

Graph-based: Instead of structures like the previous types, these databases use tree-like structures with nodes and edges connecting via relations.

Example NoSQL Database Product: Oracle NoSQL Database

Distributed File System

Not a database per se as the name would indicate, highly distributed file systems have the advantage of extreme scalability as nodes are added and frequently serve as a data landing zones or data reservoirs for all sorts of data. Read performance is typically limited by the individual node of the “system” when accessing data confined to that node, however scalability to a huge number of nodes is possible driving massive parallelism. Write performance scales well as data objects can be striped across nodes.

The most popular distributed file system used today is Hadoop. Given its role as a data reservoir, it is increasingly a location for performing predictive analytics. SQL access is available via a variety of interfaces though various levels of standards support are offered.

Example Distributed File System Product: Cloudera Hadoop Distribution (featuring the Cloudera Hadoop Distributed File System and other features)

Big Table Inspired Databases

There is an emerging class column-oriented data stores inspired by Google’s BigTable paper. These feature tunable parameters around consistency, availability and partitioning that can be adjusted to prefer either consistency or availability (given these are rather operationally intensive.

A typical use case might be where consistency and write performance are needed with huge horizontal scaling. HBase (deployed on a Hadoop Distributed File System) in particular has been deployed to 1,000 node configurations in production.

Example Big Table inspired Product: Cloudera Hadoop Distribution (Cloudera HBase)

Extending the Architecture to the Internet of Things

Thus far, we've focused on the analytics and reporting and related data management pieces of the Information Architecture. Where sensors are providing key input, the architecture for data capture, security, and linkage to the rest of the Information Architecture can require additional consideration. The following illustrates what is often described as an Internet of Things footprint for connected healthcare in the Provider space:

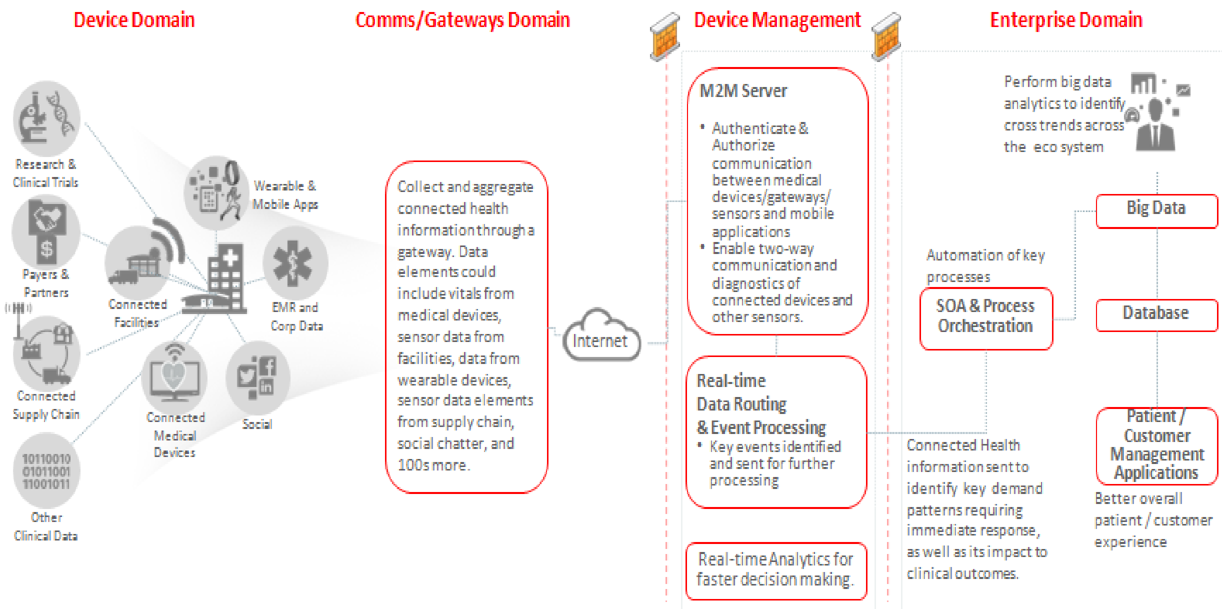

Figure 10: Connected Devices and Healthcare Providers

Items to the far right of Figure 10 have largely been previously discussed in this paper. Many of the other items pictured are what Oracle typically describes as Fusion Middleware components. For example, much of the sensor programming today takes place using Java. Security is extremely important since most would not want unidentified third parties intercepting the data provided by the sensors. Applications closer to the sensors themselves are often written using Event Processing engines to take immediate action based on pre-defined rules. There are also various message routing, provisioning, and management aspects of such a solution.

Figure 11 illustrates a typical capability map of this architecture for connected healthcare among Providers:

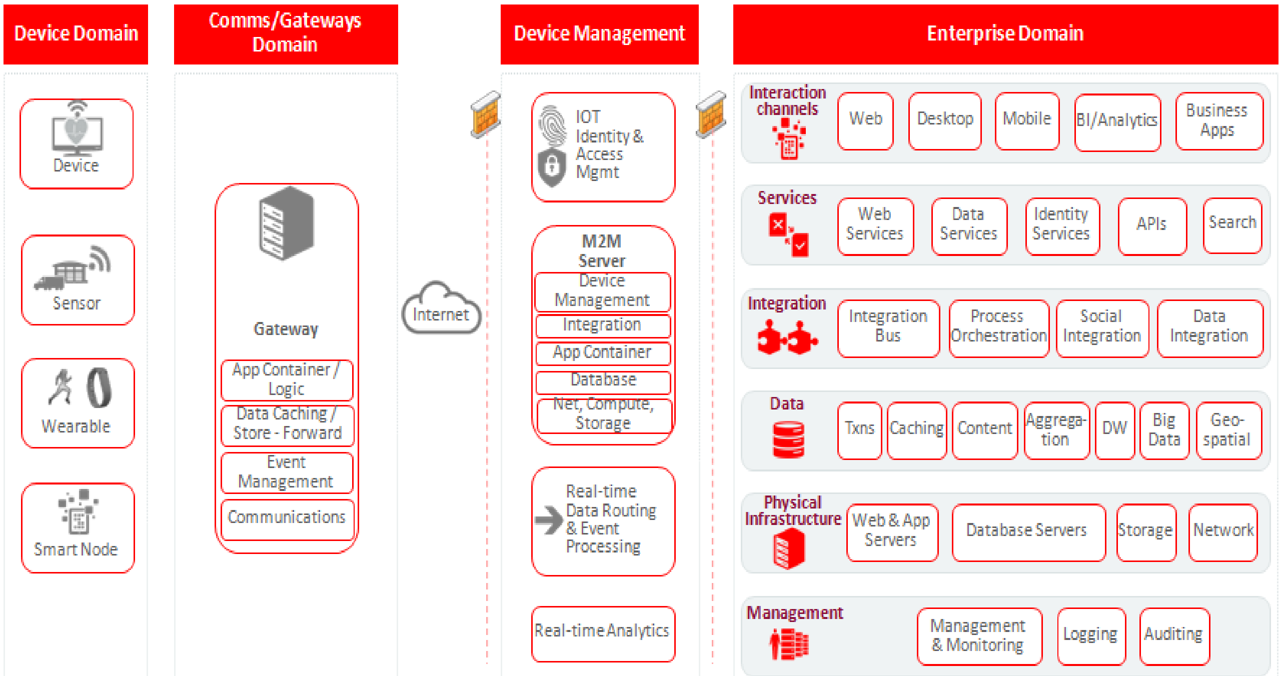

Figure 11: Connected Devices Capability Map for Healthcare Providers

Sensors are increasingly providing critical data regarding the health of individuals. They are also used to monitor energy consumption in facilities, supplies, and equipment state. This data will continue to grow and enable Healthcare Providers to better determine the status of and manage people, equipment, and services that are being offered.

Figure 12 illustrates some of the Oracle products aligned to the previously shown capability map:

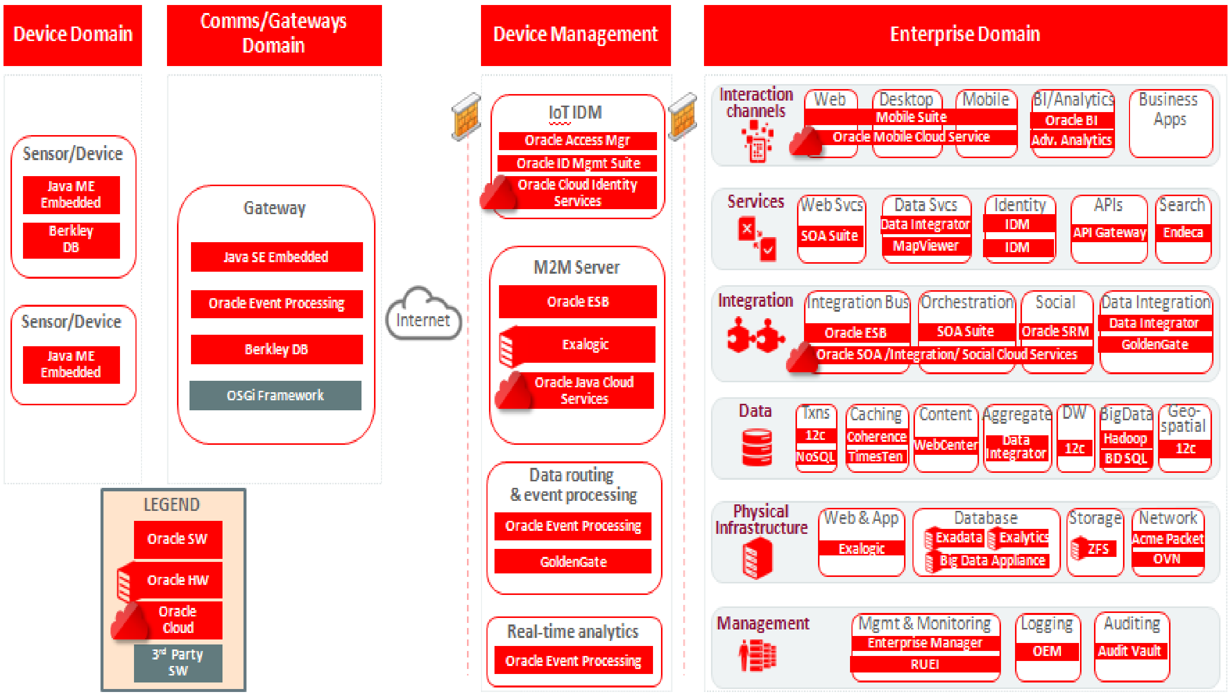

Figure 12: Oracle Products aligned to Capability Map

Keys to Success

One of the most significant keys to success in a large project undertaking is to gain alignment between the business needs and goals and with the IT architecture design and deployment plans. Key business sponsors must be engaged and active in all phases.

Methodologies based on phased approaches are almost always the most successful.   To start, you'll need to understand the current state and its gaps so that you can better understand how to build towards the future state. You will need to modify the architecture as business needs change. Therefore, a common method to help assure success is to deploy quickly in well scoped increments in order to claim success along the way and adjust the plan as needed. A complete Information Architecture is never built overnight, but is developed over years from continued refinement.

Figure 13 illustrates such an approach, beginning with defining an initial vision, then understanding critical success factors and key measures tied to use cases, defining business information maps based on output required, linking the requirements to a Technical Information Architecture, defining a Roadmap (including phases, costs, and potential benefits), and then implementing. Of course, an implementation leads to a new vision and requirements and the process continues to repeat. Pictured in the Figure are some of the artifacts Oracle often helps deliver during Enterprise Architecture engagements and Information Architecture Workshops.

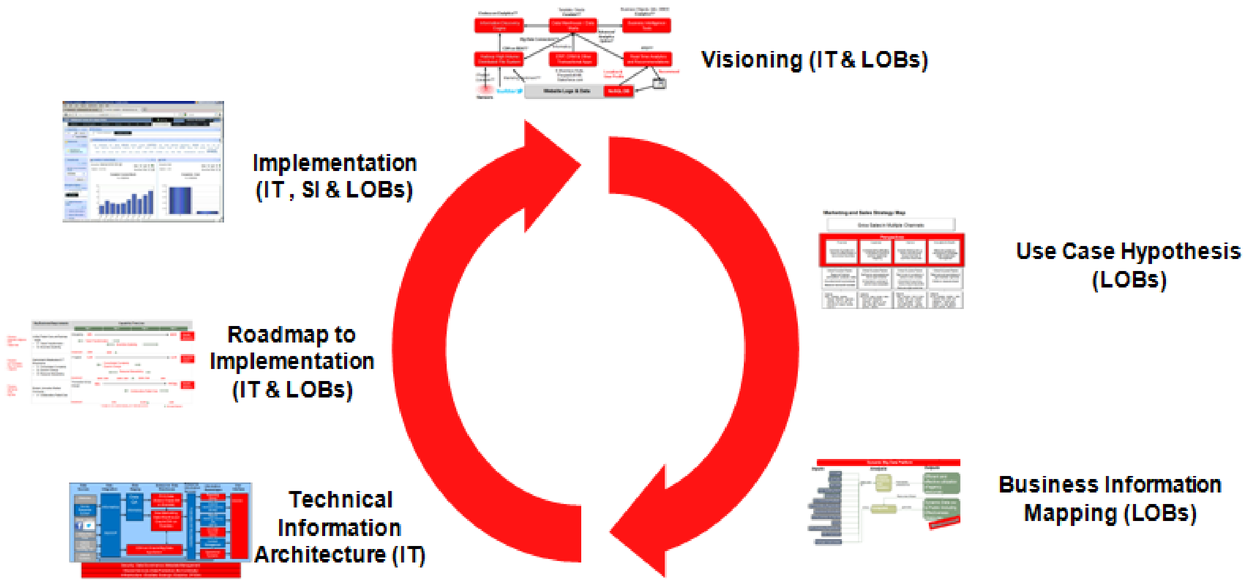

Figure 13: Typical Methodology for Information Architecture Projects

Usability needs will drive many of your decisions. Business analysts will likely have a variety of business requirements and possess a variety of analysis and technical skills. They could require solutions ranging from simple reporting to ad-hoc query capability to predictive analytics. You'll need to match the right tools and capabilities to the right users. One size does not usually fit all.   While new features in the data management platforms can provide more flexibility as to where you host the data for such solutions, the data types, volumes and usage will usually determine the most optimal technology to deploy. A common best practice is to eliminate as much movement of data as possible to reduce latency.

Data security and governance are also a key consideration. Healthcare Providers gather sensitive data and face strict HIPAA regulations. Therefore securing access to the data, regardless of data management platforms, tools, and data transmission methods used, is critical. Data governance needs regarding the meaning of data as well as its accuracy and quality will often require close coordination with and among multiple lines of business.

Finally, as fast time to implementation important to the success of any business driven initiative, you will want to leverage reference architectures, data models and appliance-like configurations where possible. These can speed up the design and deployment and reduce the risk of incomplete solutions and severe integration challenges. Leveraging engineered systems and appliances where possible can simplify the architecture, reduce time to value and improve architecture reliability.

Final Considerations

This paper is intended to provide an introduction to applying Information Architecture techniques for Healthcare Providers. These techniques guide the extension of current architecture patterns to meet new and varied data sources that are becoming part of the information landscape. Oracle has very specific views regarding this type of information architecture and can provide even more of the individual components than were described in this paper.

The following diagram provides a conceptual future state that can encompass all types of data from various facets of the enterprise:

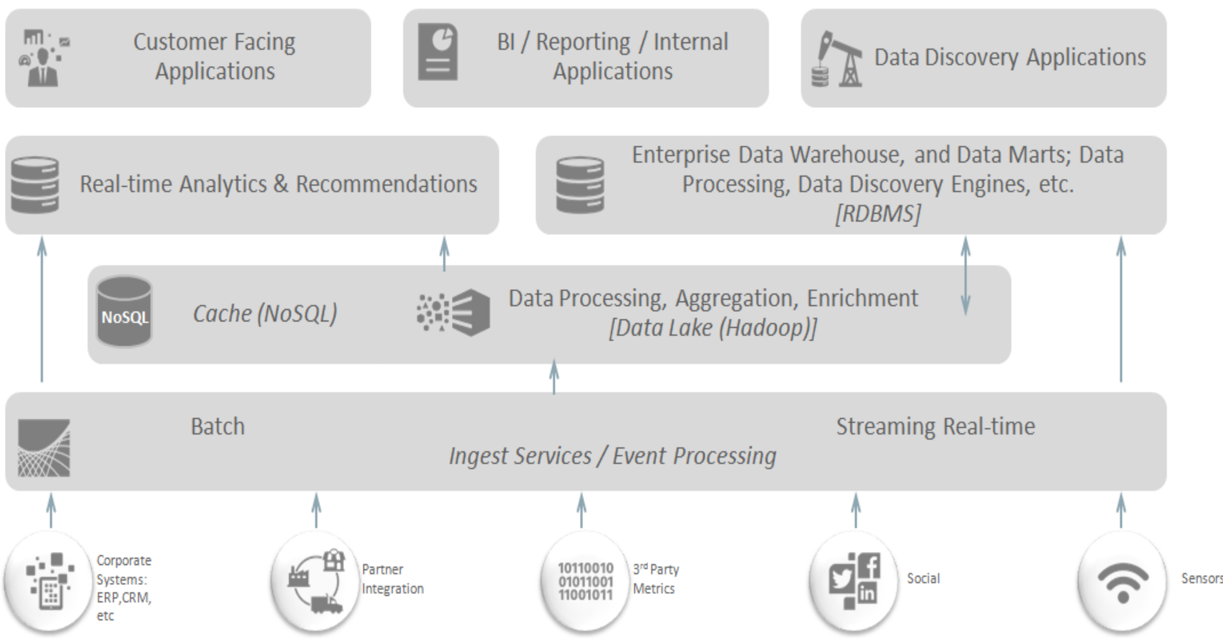

Figure 14: Typical Conceptual Future State Diagram

A more detailed look at "Business Analytics" reference architectures appears in documents posted to the Oracle Enterprise Architecture web site at http://www.oracle.com/goto/ITStrategies.

The following is a figure from one of the just referenced documents to give an idea as to the level of detail that might be considered around information delivery and provisioning.

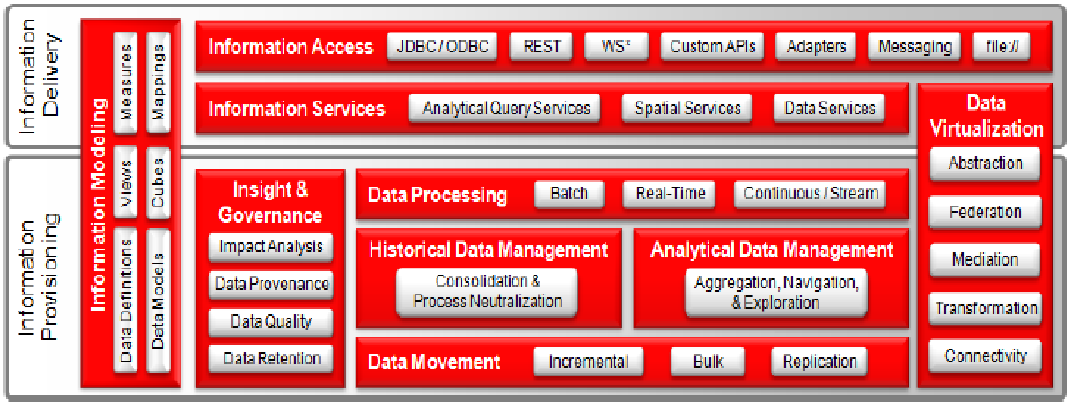

Figure 15: A more detailed Reference Architecture Diagram for Information Delivery and Provisioning

Often, the architecture discussion also leads to consideration on where to host and analyze the data (e.g. in the cloud versus on-premise). Aside from security considerations, most Healthcare Providers come to the conclusion that another motivating factor to storing the data on-premise is the volume of data being produced and a desire to minimize network data traffic. In other words, most organizations are coming to the conclusion that it makes sense to analyze the data where it lands. And once it lands, reporting and predictive analytics often take place in the data management system holding the data.

An additional consideration not addressed in this paper is the availability of skills needed by the business analysts and the IT organization. A future state architecture evaluation should include an understanding as to the degree of difficulty that a future state might create and the ability of the organization to overcome it.

The growing competitive nature of Healthcare Providers as well as the regulations and payer oversight they face will assure that those that take advantage of these new data sources to augment what they know about their business will continue to be leaders. They will continue to invent new and better business processes and efficiencies and they will do so by evolving their Information Architecture in an impactful manner. Some will likely leverage their advanced footprints to offer data subscriber networks, thereby going into competition with data aggregators and further monetizing their IT investments.

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

blogs.oracle/enterprisearchitecture

facebook.com/OracleEA

twitter.com/oracleEAs

oracle.com/EA

Hardware and Software, Engineered to Work Together

Copyright © 2015, 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.

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. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0215

February 2015
Oracle Enterprise Architecture White Paper – Improving Healthcare Provider Performance with Big Data
Author: Robert Stackowiak, Art Licht, Sean Keane, Venu Mantha

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