Improving Oil & Gas Performance with Big Data

Architect's Guide and Reference Architecture Introduction

ORACLE ENTERPRISE ARCHITECTURE WHITE PAPER | APRIL 2015

ORACLE

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

The ability to access, analyze, and manage vast volumes of data while rapidly evolving the Information Architecture is increasingly critical to oil and gas exploration and delivery companies. These companies are looking to improve business efficiency and performance while facing a number of challenges including the uncertain and volatile pricing of oil and gas, changing energy policies, environmental concerns (such as from global warming and dealing with the aftermath of shale extraction), emerging competition from new sources of energy, and ongoing operational management costs and inefficiencies.

Faced with these challenges, many see Big Data and sensors that can provide such data as important sources of information needed to optimize exploration, drilling, production and delivery. New data sources such as social media can provide important insight into the sentiment of local communities impacted by facilities and pipelines. Well managed real estate assets are also critical to maintaining a positive reputation when facing evaluation of financial assets.

Today, oil and gas companies analyze data from a variety of sources. These data sources can include:

- Data from sensors during oil and gas drilling exploration, production, transportation, and refining
- Traditional enterprise data from operational systems
- Social Media
- Web browsing patterns (on informational web sites)
- Response data from job postings
- Demographic data
- Historical oil & gas exploration, delivery, and pricing data

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. “Big Data” solutions help to enable oil and gas companies to meet these requirements.

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 in the dynamic oil and gas market.

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

Oil and gas companies exist in a rapidly changing marketplace. New methods for extracting energy and alternative forms of energy can enter the market and create an oversupply. Political events can create shortages. Understanding the direction of market pricing and demand is crucial. For example, refiners might need to explore other markets (in new niches or other geographies) when local supply outstrips demand.

While facing this uncertainly, many continue to place an emphasis on better asset management and control. There also remains a need to focus on efficient and effective oil and gas exploration whenever it is undertaken. Maintaining a favorable environmental track record and good public and government relations remains top of mind as good business but also to gain government approval for new exploration and production when needed.

The types of data used in these analyses can vary widely with much of it coming from sensors and other streaming data sources. By deploying Big Data Management Systems that include traditional data warehouses and newer data reservoirs (featuring Hadoop and / or NoSQL Databases), broader types of data can be analyzed to ensure that the business becomes more agile. A key challenge remains making sense of this growing mountain of data in meaningful ways that will impact the business.

Exploration and Refinement Effectiveness

Oil and gas companies have long focused on improving exploration effectiveness through advanced analytics applied to a variety of data. Seismic mapping provides clues on where to drill. In today's challenging environment, cost of new exploration must not increase while success rates need to improve to maintain profitability. Similarly, increasing reservoir capacity as part of refinement and containing cost are major factors in maintaining profitability.

Such data has traditionally been analyzed in data warehouses consisting of relational databases. Today, some of this activity is being moved to Hadoop clusters, partly due to the relatively low price point of Hadoop clusters and partly due to its "schema-less" file system being ideal for predictive analytics workloads.

Operational Efficiency

The need for gaining operational efficiencies at drilling sites, pipelines, and in refineries is well understood. Increasingly, this data is gathered from sensors providing real-time information on the state of operations. There is an opportunity to better understand when maintenance needs to occur sooner and, by monitoring the changing state of key components, affording even greater efficiencies and cost savings. Predictive analytics solutions deployed across Big Data Management Systems (including Hadoop) will likely become common practice to increase overall safety, reliability, and reduce cost. Such analysis can also point to potential safety issues or identify environmental risks before expensive accidents occur.

Predictive analytics might also be deployed in data warehouse solutions today, such as when optimizing the routing of vehicles, crew and supplies to exploration and production facilities. Such routing is fundamental to providing the lowest possible cost of delivery while maintaining margins. It is also critical for delivering parts on time thus possibly avoiding major maintenance and safety issues.

Understanding a Changing Market

Oil and gas prices can be quite volatile, yet their value during the life of a project can determine whether exploration, production, and refining make financial sense. Predictive analytics can play a critical role here in gaining an understanding as to the likely direction of pricing and identifying the right exploration and production level choices. Such data can also provide guidance as to the right real estate investments.

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 oil and gas companies typically work with their lines of business to build solutions that deliver the following when defining Big Data projects:

1) Improved Exploration Return on Investment: Oil and gas exploration requires extensive seismic, environmental, and cost of production analysis in order to determine whether yields might be profitable and achievable. Understanding changing market conditions is also required.

2) Improved Production Efficiency (e.g. greater production at lower cost): Production efficiencies are ensured by ongoing analysis of drilling operations and timely maintenance of equipment. Remote monitoring and analysis are also critical to determining the conditions of pipelines and refineries. The analysis also plays an important role in helping to avoid environmental and safety issues that can be a by-product of less optimally run operations.

3) Cost Effective & Timely Supply Chain and Logistics Management: Delivery of parts, supplies, equipment, and personnel in a timely manner are key when optimally maintaining production facilities and ensuring minimal down-time. Effective logistics management can also ensure a more cost effective supply chain.

4) Better Market Analysis Driving Investments: Market analysis of current pricing and trends is a critical factor in determining when to explore, when to go into production, when to buy and sell real estate, and when to change the investment strategy.

5) Improved Public & Government Relations: Understanding public opinion and responding rapidly to requests from the public and government agencies are extremely important when establishing trust with broader communities.

6) IT Operational Efficiency: Not unique to oil and gas companies 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 Oil and Gas companies and the opportunity for new or enhanced business capability when adding new analytic capabilities.

TABLE 1 – OIL AND GAS COMPANIES, BUSINESS CHALLENGES & OPPORTUNITIES

FUNCTIONAL AREA	BUSINESS CHALLENGE	OPPORTUNITY
Exploration	Understand the most viable areas to explore with potentially the greatest return (subject to market conditions)	Improved geologic analysis through fault modeling and seismic analysis leading to more finds, more predictable outcomes, and better well planning Improved market forecasting to better determine when exploration and production could be viable
Drilling and Production	Maximize operational efficiency while meeting environmental standards	Better predictive maintenance enabling higher equipment availability Sensors on drills indicating abnormal pressure readings helping to predict and avoid risk of catastrophic failure Better energy management enabling decreased cost of drilling Better asset and personnel management enabling lower cost of production Remote operational control and monitoring of pipelines and equipment for lower cost, lower environmental risk, and better safety
Supply Chain and Logistics Management	Optimal time for delivery of parts, supplies, and personnel needed at exploration and production sites	Predictable time to delivery of critical parts, equipment, and qualified personnel Minimal down-time Optimal fuel utilization and cost of delivery
Real Estate Management	Retire and sell properties if no longer believed to be financially viable	Sell properties while the properties maintain reasonable value Use property sales to fund new exploration in more lucrative areas Track environmental regulations in locales and understand potential clean-up costs
Public and Government Relations	Maintain good working relationship with community and government officials	Faster determination of community sentiment toward exploration, drilling and production More proactive PR where sentiment suddenly changes negatively Fast response to requests for environmental impact and other drilling production records
Market Pricing Forecasting	Understand market pricing direction and future impact on profitability of sites being explored and in production	Make smarter decisions sooner on where and when to explore for oil and gas Better align distribution of personnel, equipment, and supplies to demand Better manage real estate holdings

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 oil and gas companies. These components are further described in the "Information Architecture and Big Data" whitepaper posted at http://www.oracle.com/goto/ea.

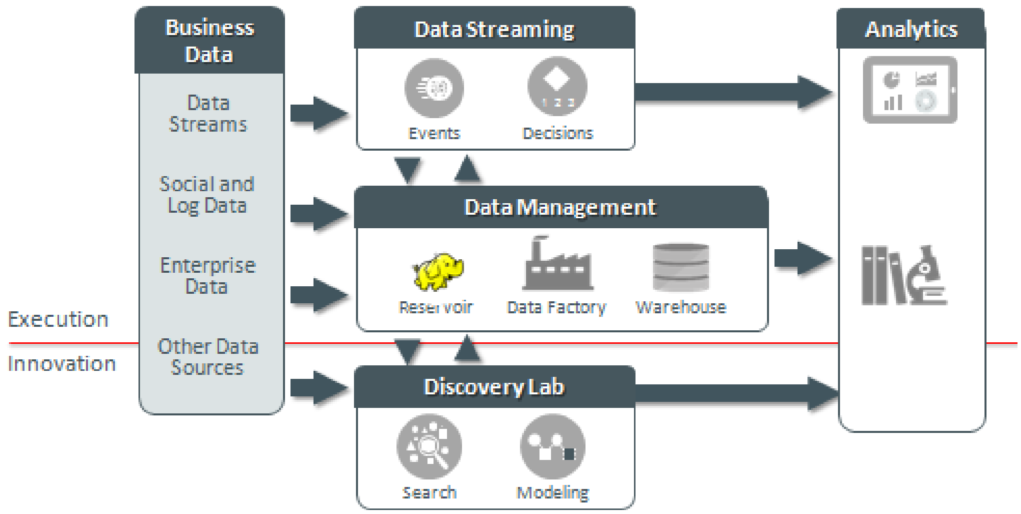

Figure 1: Key Information Architecture Components

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 the sources to the application (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.

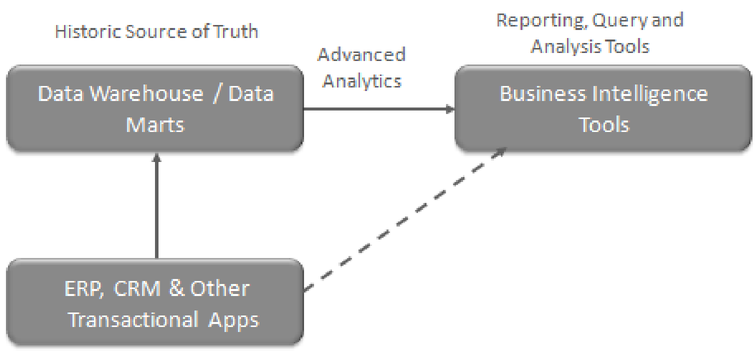

Figure 2: Typical Current State Data Warehouse

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.

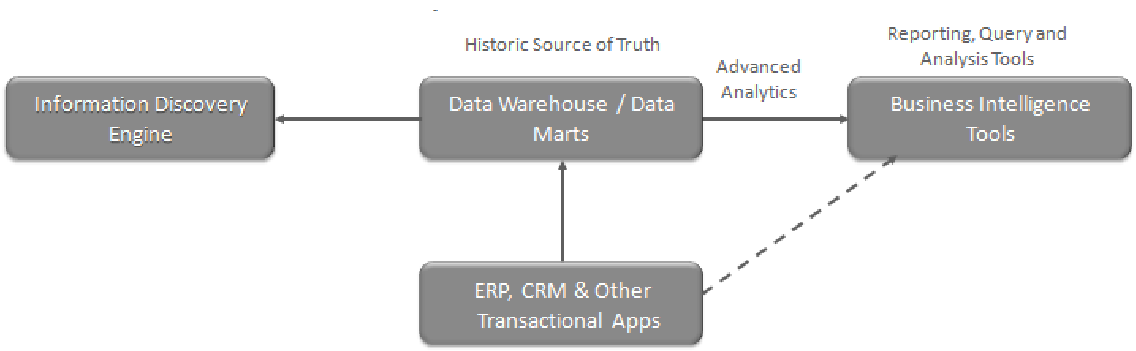

Figure 3: Typical Introduction of Information Discovery

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.

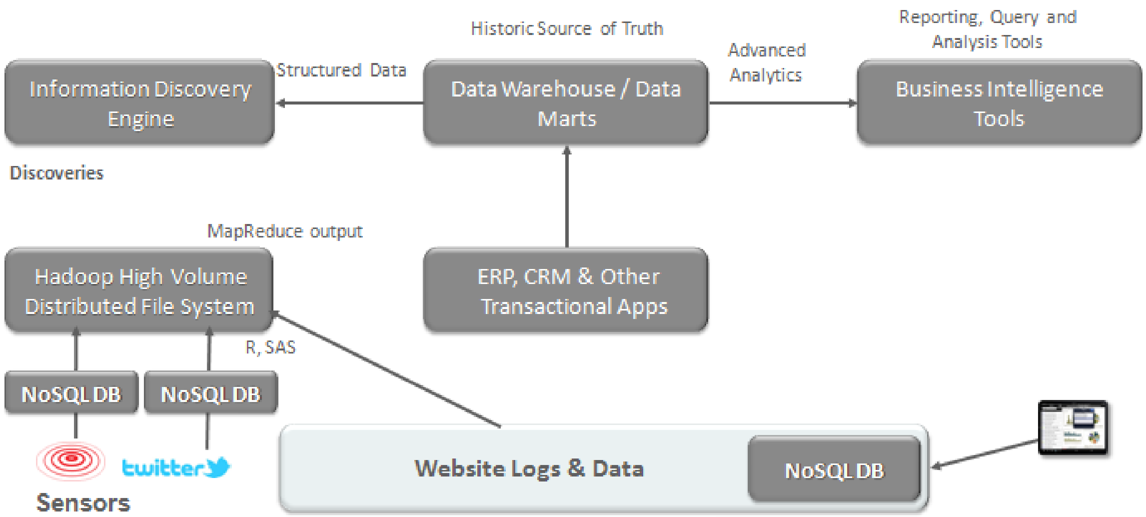

Figure 4: Typical Early Hadoop Environment separate from the Data Warehouse

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.

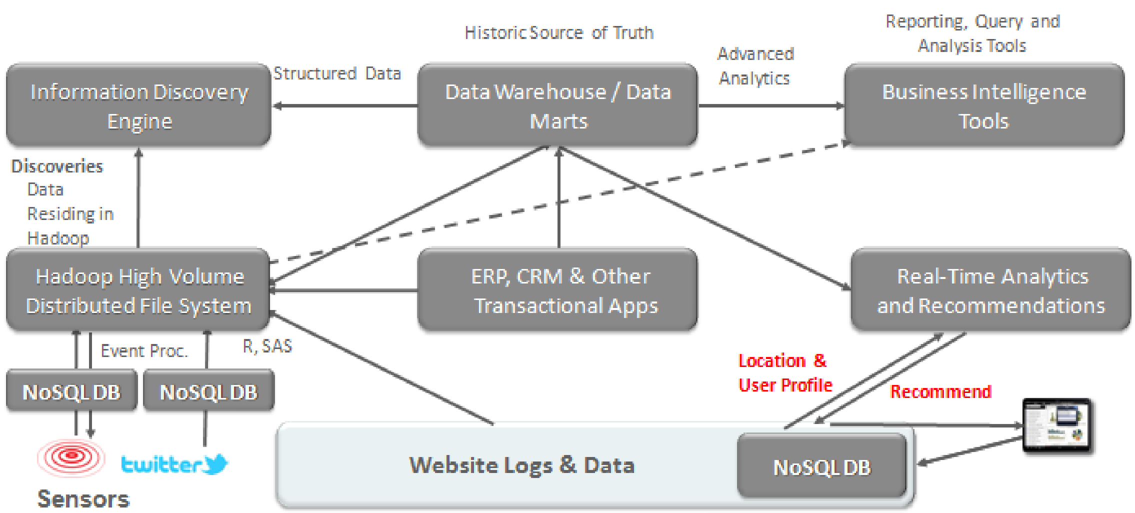

Figure 5: Integration of Hadoop Infrastructure and Data Warehouse

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.

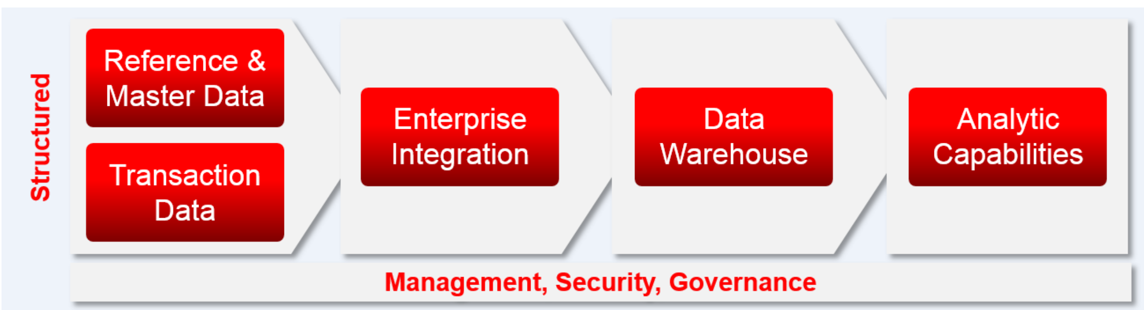

Figure 6: Structured Data and the Data Warehouse

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.

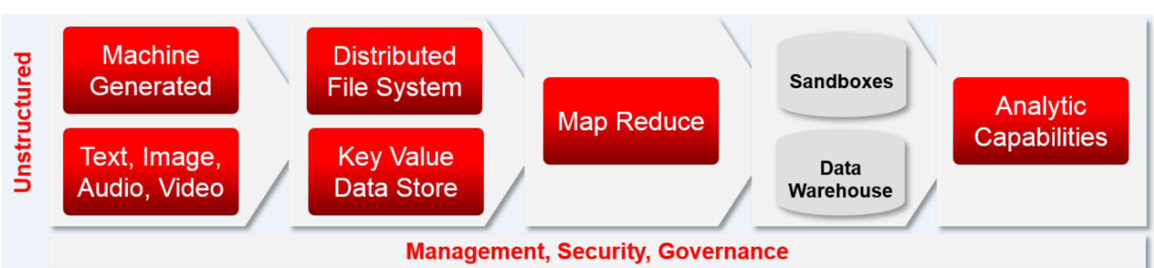

Figure 7: Unstructured Data, Distributed File Systems and Key Value Data Stores

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.

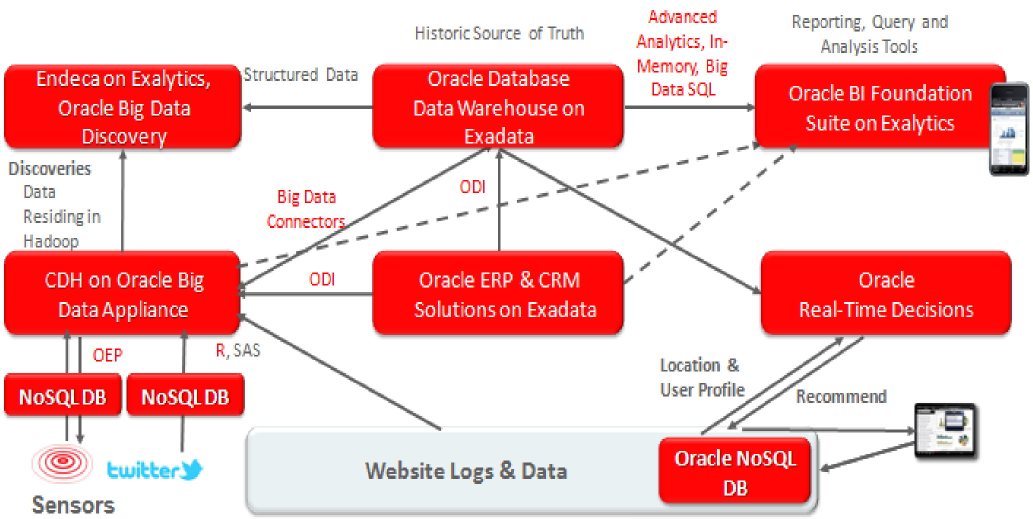

Figure 8: How Key Oracle Products Fit in the Generic Architecture

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, an Information Architecture diagram that focuses on analysis of data from drilling rigs might look like Figure 9:

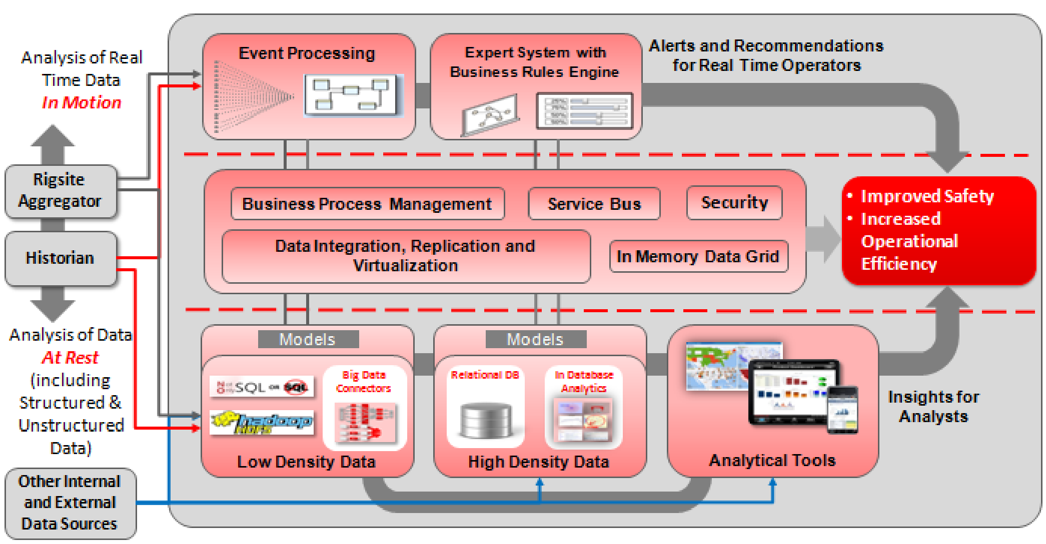

Figure 9: Analysis of data from Drilling Rigs

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 in Oracle products and other products are possible when defining and deploying your Information Architecture.

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 oil and gas companies monitor and take action on data from sensors, the architecture discussion is extended to the "Internet of Things". This extended architecture for data capture, security, and linkage to the rest of the Information Architecture can require additional consideration. The following illustrates a typical footprint:

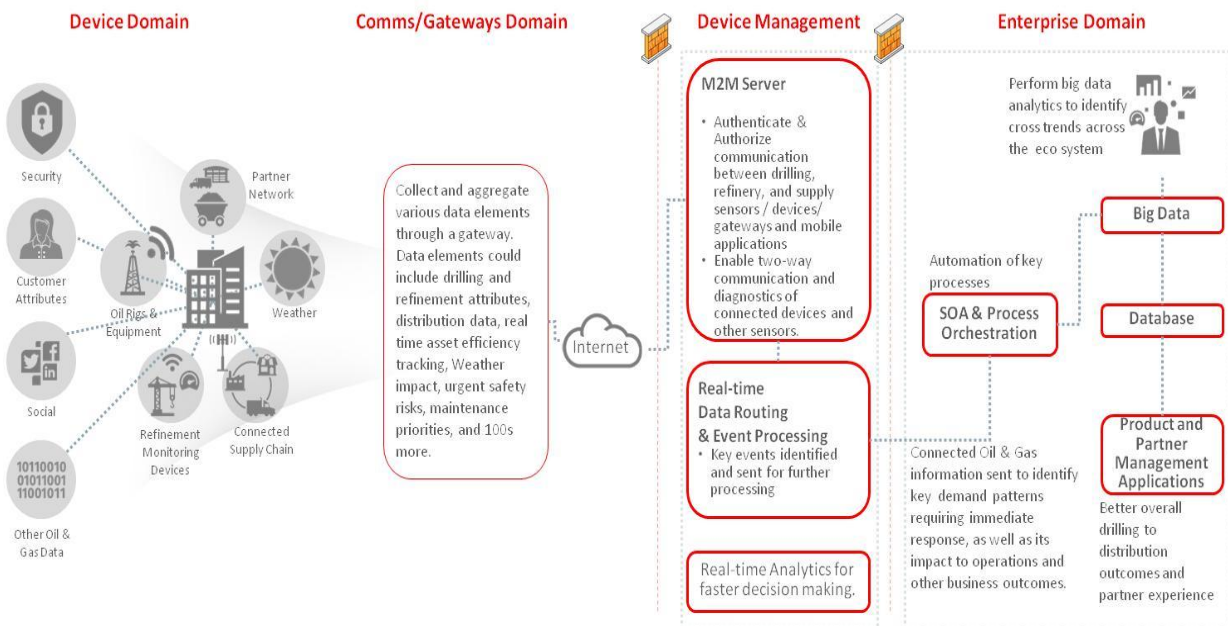

Figure 10: Connected Devices in an Internet of Things Footprint

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 oil and gas companies:

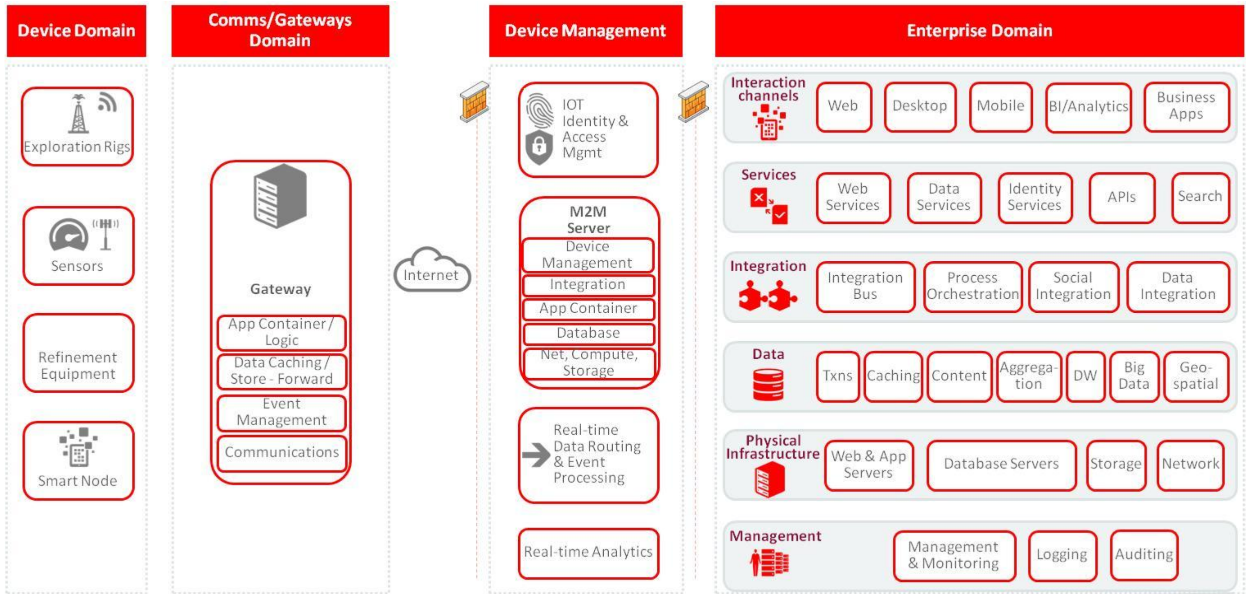

Figure 11: Connected Devices Capability Map

Sensors are increasingly used to monitor the state of exploration, production, transportation, and refining facilities and equipment. The real-time monitoring enables proactive measures to be taken sooner enabling great efficiencies and reducing potential environmental and safety risk.

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

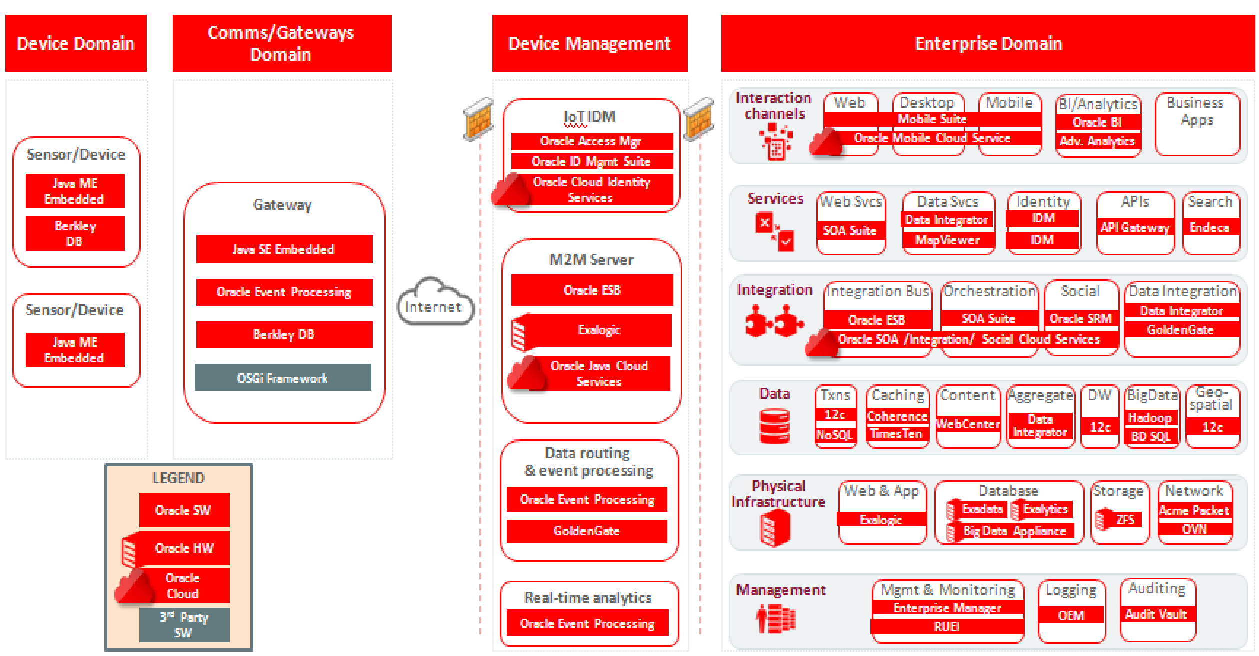

Figure 12: Oracle Products aligned to Capability Map

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.

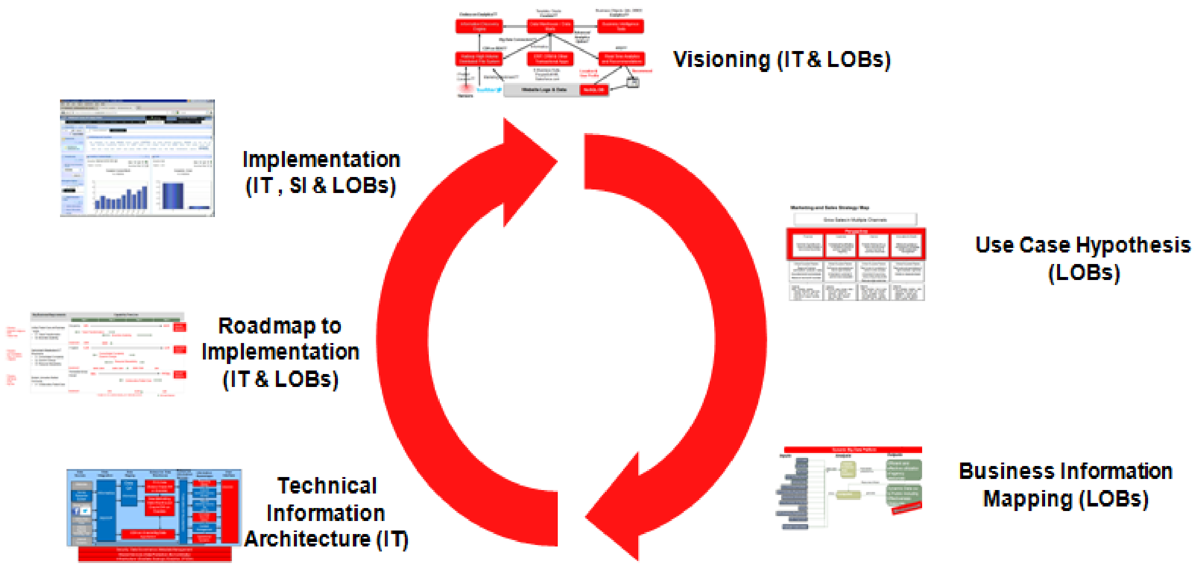

Figure 13: Typical Methodology for Information Architecture Projects

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. Oil and gas companies want data to remain private unless they specifically agree to share it. So 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 oil and gas companies. 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:

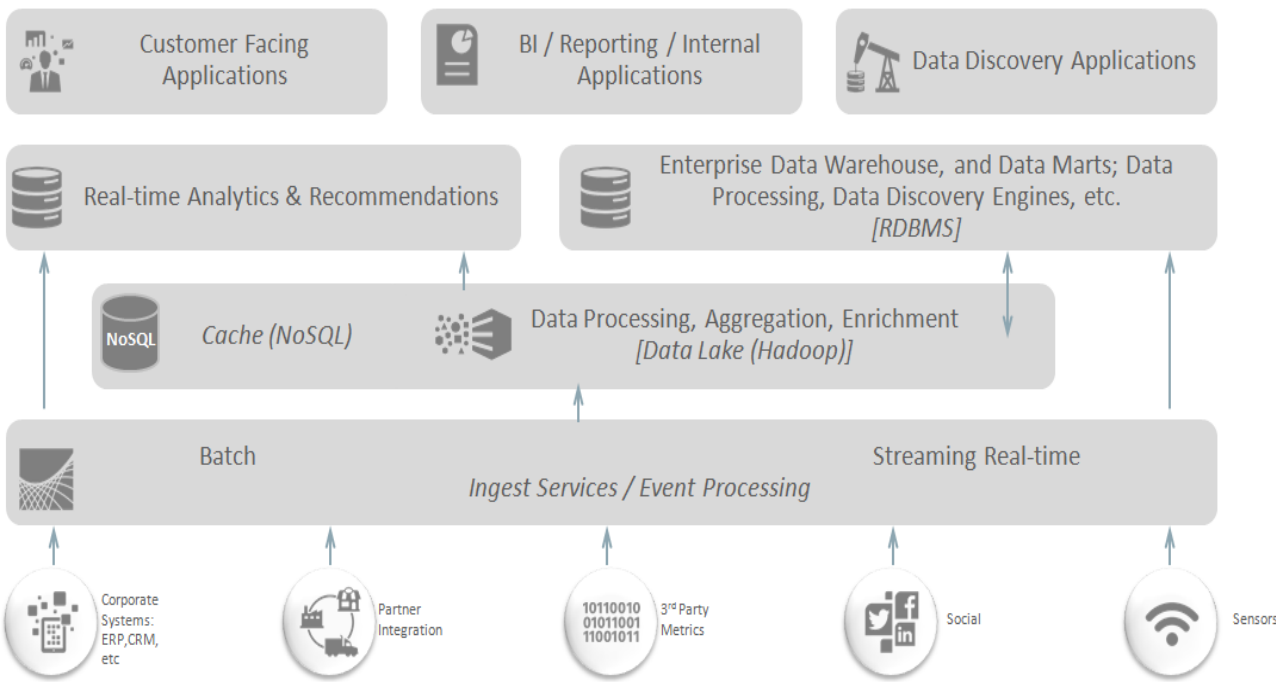

Figure 14: Typical Conceptual Future State Diagram

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.

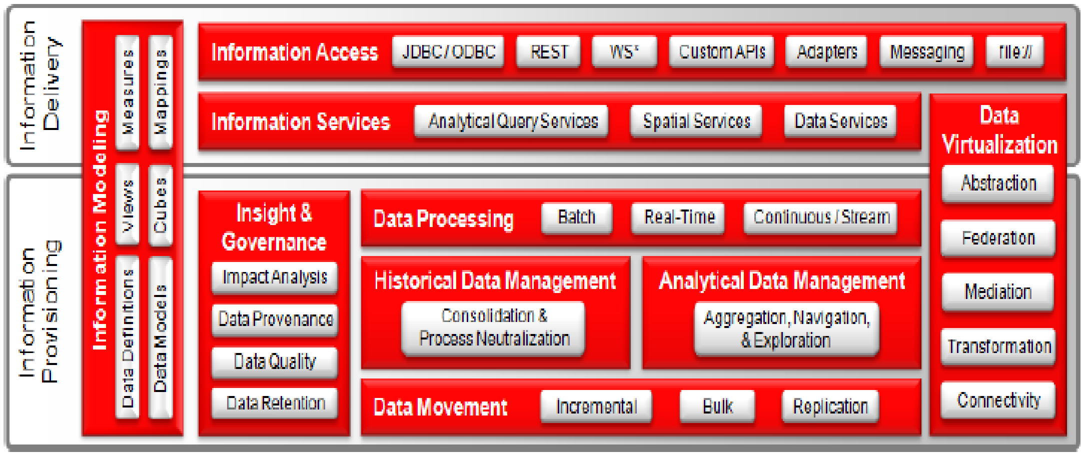

Figure 15: A more detailed Reference Architecture Diagram for Information Delivery and Provisioning

The architecture discussion can also lead to consideration on where to host and analyze the data (e.g. in the cloud versus on-premise). Most oil and gas companies choose to host data in the location where the data initially lands with an eye on minimizing network data traffic while securing the data at rest and in motion. Once the data lands, reporting and predictive analytics often take place in that data management system. Where some of the data is stored in the cloud and some is on-premise, there should be careful consideration of the impact of network bandwidth on analysis performance where data from both locations is required.

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 competitive nature of oil and gas companies 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.

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

April 2015
Oracle Enterprise Architecture White Paper – Improving Oil & Gas Performance with Big Data
Author: Robert Stackowiak, Venu Mantha, Alan Manewitz, Art Licht.