

Microsoft - Oil and Gas

Upstream IT Reference Architecture

Introduction

Today's oil and gas industry needs a common information technology (IT) reference architecture for use by upstream organizations, system integrators (SIs) and solution providers. Microsoft is working to develop a coherent reference architecture that will enable better integration of upstream applications through the composition of common application components.

A single architectural approach will encourage simplification and unification in the upstream sector. It will also give SIs and solution providers an established environment within which to design and build solutions, while allowing upstream operators to design architectures as a requirement for potential vendors.

Upstream Business Needs

Any upstream reference architecture must respond to the functional activities of the organization and the capabilities needed to support the business.

A workable architectural approach should recognize that oil and gas exploration and production organizations (much like organizations in virtually every industrial sector) work simultaneously with both structured and unstructured data. Structured data is handled in the domain-specific applications used to manage surveying, processing and imaging, exploration planning, production and other upstream activities. At the same time, large amounts of information pertaining to those same activities are generated in unstructured form, such as e-mails and text messages, word processing documents, spreadsheets, voice recordings and other forms.

As shown in Figure 1, upstream organizations use a broad spectrum of structured and unstructured data to orchestrate, automate, integrate and engage in integrated upstream management activities.

Domain-oriented structured data is used for: collaboration, including visualization, data fusion, decision tracking and knowledge management; optimization, including simulation, proxy models, and decision selection and implementation; operations data analysis, such as trend and root cause analysis, solution evaluation, key performance indicators (KPIs) and problem detection; data management, including quality control, validation, data storage and archiving, loss management, allocation and rate estimation; and acquisition, including measurements and data transmission.

Of course, personnel in any upstream organization also generate and use a wide variety of unstructured data in the form of e-mails and text messages, electronic documents, voice messages and other information.

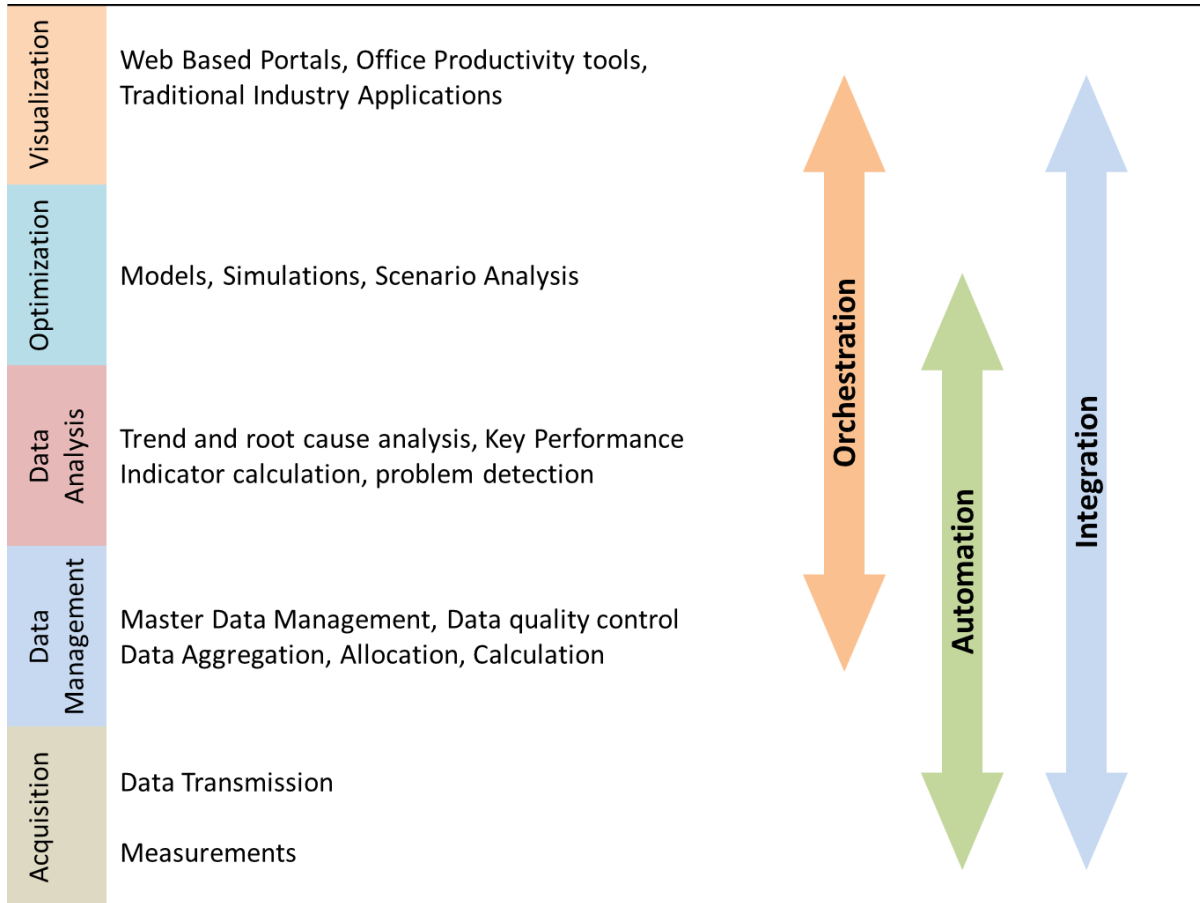


Figure 1 – Upstream Business Activities

A day in the life of an upstream organization

To better appreciate the use of both structured and unstructured data, consider the daily activities of three typical oil and gas professionals.


Reservoir engineer. A geophysicist named Jill might perform streamline simulations, manage uncertainty and plan future wells using a Petrel simulation workflow, then leverage advanced upscaling techniques to recreate geologically accurate models for full reservoir simulation. At the same time, Jill uses corporate e-mail, text messaging and a collaborative work sharing program to communicate with colleagues about key reservoir activities.

Production engineer. A subsurface engineer named Robert might use production surveillance, analysis and forecasting software to manage reservoir/well interfaces such as perforations, sand control and downhole flow, while leveraging specialized applications for oil and gas workflows, lift management, and well and network optimization. Robert also uses word processing, spreadsheets and other familiar business applications to manage vital unstructured project data on a daily basis.

Drilling engineer. Tasked with managing the technical aspects of drilling exploratory, production and injection wells, a drilling engineer named Elliot uses sophisticated drilling software to reduce well-related cost and risk, and to handle casing and tubing design, well surveillance, completions and interventions. Like his counterparts in production and reservoir management, Elliot also relies on a variety of non-technical applications – such as word processors, telephony and collaboration software – to store and manage important but less-structured, project-related information.

The current state

It may also help to examine the current state IT infrastructure in most upstream petroleum businesses. In most organizations, the volume of information is increasing exponentially as digital sensors are deployed in more exploration plays, as more data sources are connected to IT systems, and as growing volumes of information are captured and stored in enterprise databases. Large volumes of domain-specific information are also embedded in various upstream applications.



Current Architectures

Existing IT architectures in the oil and gas production space are often limited by siloed applications, poor integration and barriers to collaboration. Paradoxically, the most common activities across all of these domains are word processing, spreadsheet, e-mail and other basic business applications.

Figure 2 - The Current State of IT Architectures

A few basic issues define the requirements of an upstream IT architecture.

Data Management. That growing volume of data, however, now typically resides in disparate source systems, such as the Landmark SeisWorks® seismic interpretation software, and/or Schlumberger’s Petrel. The web-based tools for viewing and collaborating on such information are not fully integrated. That means when a geologist is reviewing seismic data for a prospect and he or she needs to cross check core samples, that information can typically be accessed only through an inconvenient and time-consuming search of the different systems, rather than from one common interface.

When integration does exist, it is usually through point-to-point connections or intermediary database tables. Those one-off connections add time and cost and cannot be easily shared or reused by other applications. There are various industry solutions that provide data- or application-integration frameworks – such as OpenSpirit – that create a common access layer to help address this integration problem.

Integration. Each geophysical discipline—petrophysics, geology and reservoir engineering—tends to have and use its own analytic modeling systems, but there is currently little connectivity or interaction between those models. Changes in conclusions for one discipline are not always carried through to other specialists, causing overall reservoir descriptions to be inaccurate.

Collaboration. The current state also makes it more difficult to collaborate, because there is no convenient, shared location where multiple partners can access information stored on the corporate network. A seismic service company employee, for example, who works on prospects for five or six energy companies, may need separate log-in locations and passwords for each of those collaborative partnerships. These same collaboration challenges typically exist within individual oil and gas companies as well.

Performance Management. In traditional IT situations, the KPIs needed to assess the health of a business are not readily available. The manual processes needed to gather and analyze KPIs are time- and labor-intensive. In all too many oil and gas firms, managers and engineers wait up to a week as KPIs are collected, analyzed and translated into the insights needed to understand and run the business.

In this siloed environment, it can also be difficult to locate information and to ensure the timeliness and quality of that data. Three or four different systems may compile lists of available field data, for example, but the organization may lack a single, comprehensive and up-to-date list of those crucial subsurface assets.

Evolving to a more efficient architecture

There are powerful reasons why oil and gas companies should seek a new and more efficient IT architecture.

In today's business and operational environment, companies must deliver more throughput with leaner staffs and fewer resources. To do that, geoscientists should spend more time in the interpretation activity—and less time searching for and quality controlling data. They should rely more on workflows, data-driven events and automated analysis to identify risks and help manage the exploration portfolio.

They also need integrated views that reveal all relevant data, both structured and unstructured, for a particular situation. That comprehensive perspective should include tornado charts measuring risk, analog well production histories, rock properties, log files, rig schedules and other variables relating to the prospect in question.

Management needs up-to-date KPIs to fully understand the health of the business. Ideally, managers should be able to see a single screen showing, for example, the portfolio of current opportunities, which ones are drill ready, available rigs and the prospect peer reviews that are scheduled for the next two weeks. With appropriate drill down details, managers can focus their time on the under-performing evaluation teams to quickly take remedial action to bring them back the expected level of productivity.

The industry needs an architectural approach that allows exploration companies to use more flexible and cost-efficient plug-and-play business logic. If a technology supplier comes up with a better seismic viewer, the architecture should allow that solution to be deployed quickly and economically. This approach reduces the constraints on IT, gives companies access to best-of-breed solutions, and can reduce the time needed to deploy new solutions from years or months, to just months or even weeks.

Lastly, upstream organizations also need the ability to connect and integrate the large volumes of unstructured data generated and used by non-domain-specific sources, such as word processing and e-mail programs, unified communications and collaborative applications. This requirement recognizes that much of the information needed to manage upstream projects is in fact hosted in non-domain applications and environments.

Enabling that evolution

As noted, several emerging industry standards can now enable the more flexible, integrated IT architecture needed in the oil and gas exploration sector.

New XML standard-based technologies (including the WITSML, PRODML and RESQML standards refined and supported by Energistics and other organizations¹), provide the common data interfaces needed to ensure plug-and-play access to best in class hardware and software solutions. If for example a company currently has a wellbore viewer that uses WITSML, with that interface in place, the company can in the future deploy any WITSML-based wellbore viewer solution.

¹ <http://www.energistics.org>

Solution Architecture and Approach

Approach

This Reference Architecture is based upon five foundational “pillars”, or principles – Performance Oriented IT Infrastructure, Enhanced User Experience, Optimized Domain-Specific Infrastructure, Rich Application Platform and Comprehensive Interoperability. These foundational pillars are described in Figure 3 and the following sections.

Oil and Gas Reference Architecture’s Foundational Principles

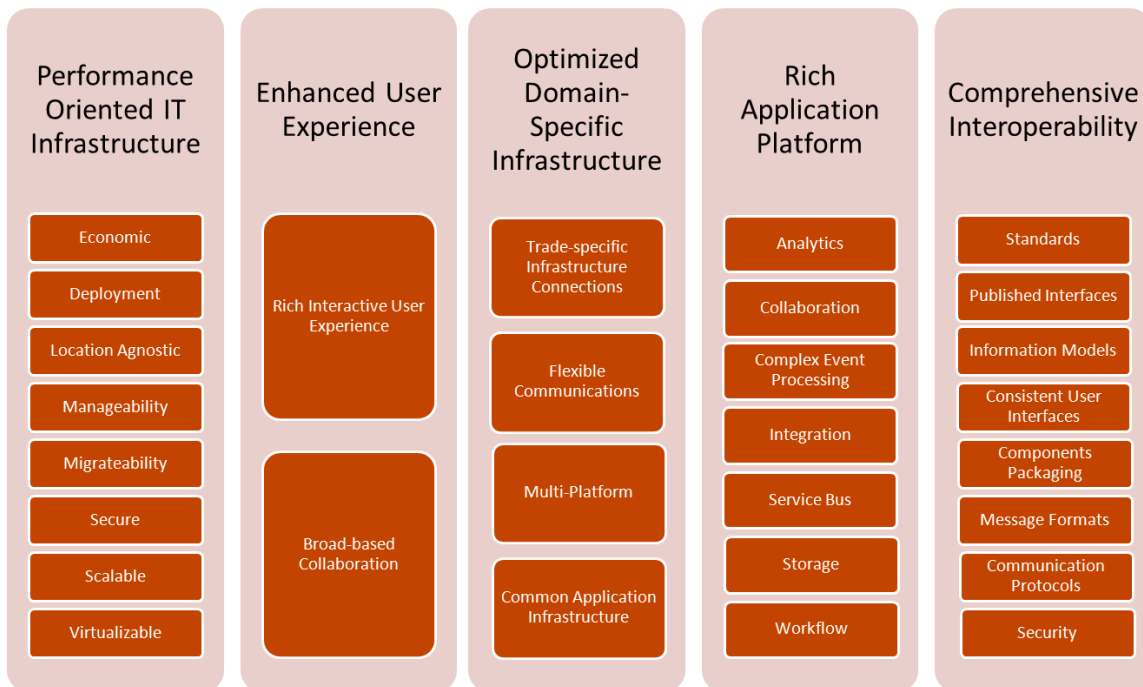


Figure 3 – Foundational Principles of the Microsoft Upstream Reference Architecture

Performance Oriented IT Infrastructure

A Performance Oriented Infrastructure includes those features that make an architecture complete and appropriate to business needs. These include:

- **Economic:** The infrastructure must provide cost effective means to deploy and integrate new functionality.
- **Deployment:** Components have to consider flexibility in how and where they can be deployed.

- **Location agnostic:** Services are designed so that they can be deployed on-premise or in the cloud. Users and software components have access to platforms and services wherever they are located.
- **Manageability:** Infrastructure components can be efficiently deployed, managed and monitored.
- **Migrateability:** Functionality and information can be migrated easily from one version of underlying infrastructure components to another with minimal interruption or intervention.
- **Secure:** Deployed components, functionality and associated information are protected from unauthorized access or malicious attacks.
- **Scalable:** Support for more users, larger models, increased transaction volumes, etc. can be accommodated through increasing hardware performance (scale-up) or the linear addition of hardware and network resources (scale-out).
- **Virtualization:** Components can be deployed in a manner that optimizes the use of hardware resources.

Enhanced User Experience

An Enhanced User Experience enables all qualified upstream employees and partners to view the architecture from the perspective of other participants.

To Microsoft, this equates to ensuring that the host oil and gas company understands how participants experience the world and how technology fits into that experience. A technology architecture that facilitates this comprehensive perspective will then necessarily consist of:

- **A rich, interactive user experience** for field workers, operations staff, control center personnel, and for use at home and on the road
- **Broad-based collaboration** using both thick and thin clients, across a variety of devices, and leveraging today's most advanced collaborative tools through corporate portals and services

Optimized Domain-Specific Infrastructure

The Microsoft Oil and Gas Upstream Reference Architecture connects domain-specific equipment to the operational network. An optimized domain-specific infrastructure incorporates:

- **Trade-specific infrastructure connections:** Leveraging unified communications to manage compliant devices, from downhole sensors and well heads, to pumps, intelligent bits, compressors and other plant equipment, then flowing that data into appropriate operational systems.
- **Flexible communications:** Deployments can leverage a variety of communications paths and technologies and are easily reconfigured minimizing the time required to make new information available to users.

- **Desktop, server, embedded mobile operating systems and Internet:** Operating systems (OS) can be effectively employed leveraging the right OS at the right level for the right role with the right performance.
- **Common Application Infrastructure:** This is the architecture for applications infrastructure and services for commonly used capabilities so developers can focus on domain-specific functionality optimizing speed to market and reliability of solutions.

Rich Application Platform

The Microsoft Oil and Gas Upstream Reference Architecture acknowledges from the outset that no one vendor will be able to provide all of the application functionality needed to implement this approach. This Reference Architecture seeks to offer a rich platform that makes it easy for partners to develop and deploy their applications. Notable aspects of the applications platform include services for:

- **Analytics:** Rich statistical and analysis packages for data mining, discovery and reporting for diverse information consumers
- **Collaboration:** Tools, services, and applications enabling interaction between users and between equipment
- **Complex Event Processing:** Stream processing engines which can detect and filter events
- **Integration:** Messaging and database technology for linking together workflow, processes and data optimization
- **Service Bus:** Services and components for communication of device and equipment data
- **Storage:** Repositories for capture and enabling analysis of utility operational and business data
- **Workflow:** Services for managing the automation of applications as well as business processes

By providing these services to developers, Microsoft partners will only need to worry about using their expertise for the solution of domain-specific problems, leaving the platform to provide the common capabilities needed across many vertical domains. As a result, customers will find it easy to leverage those application offerings.

Comprehensive Interoperability

The Microsoft Oil and Gas Reference Architecture must enable interoperability in order for the ecosystem to develop in a cost-effective manner. Otherwise, the vision for the ecosystem will go unfulfilled. Pragmatic integration approaches will need to be considered and the Reference Architecture should be flexible to allow deploying new components without custom integration.

Interoperability considerations include:

- **Standards** that define a consistent industry-wide interface to allow new component deployment

- **Published interfaces** that are transparently publicized for open-industry use even if a standard is not available and also satisfy important interoperability needs. All the elements of an interface are well defined so that applications can be independently developed to leverage the interface
- **Information Models:** Consistent ontology for referring to equipment and assets to enable exchange of information throughout the enterprise and the value chain
- **Consistent User Interfaces:** Consistent content and behavior in presentation of information and interaction with the user
- **Components Packaging:** Well defined sets of functionality packed for developer and integrator reuse
- **Message Formats:** Key construct of Service-Oriented Architecture (SOA)² defining format and content that enable services exchange messages using the defined format (eg Publish – Subscribe pattern)
- **Communication Protocols:** Format, content and exchange mechanism so applications can be written to transfer information using the protocol definition
- **Security:** Definition of the security implementation including authentication, authorization, identity lifecycle management, certificates, claims and threat models to enable secure interoperable design and deployment.

² “In [computing](#), **service-oriented architecture (SOA)** provides a set of principles of governing concepts used during phases of [systems development](#) and [integration](#). Such an architecture will package functionality as [interoperable services](#): software modules provided as a service can be integrated or used by several organizations, even if their respective [client](#) systems are substantially different.” [Wikipedia.org](#)

Architecture

The oil and gas IT community is now adopting Service-Oriented Architectures (SOA) as a more flexible and responsive alternative to traditional hard connections between applications and source data.

The industry also is now embracing the Enterprise 2.0 concept and social media technologies such as status updates and notifications used by social networks, messages, blogs and wikis. As engineers begin to use these technologies to manage their personal connections, the industry is adapting those network-based capabilities to foster cross-discipline collaboration and to better understand and manage the oil and gas upstream operations environment.

Reference Architecture

Microsoft recommends the following reference architecture for the oil and gas upstream environment. The architecture reflects a service-oriented computing environment that encompasses domain applications, business productivity tools and back office applications in an integrated manner.

By following the service-oriented architecture approach for interactions between components on different layers of the architecture, we will be able to realize the full benefits of this environment. As shown in Figure 4, to function as planned and to eliminate the need for additional integration work, this Reference Architecture must provide inherent, seamless integration to connect all layers and applications. This built-in integration enables frictionless communication and information flow up and down the layers, and provides the ultimate “plug-and-play” efficiency needed by oil and gas companies, technology vendors and other partners.

We will walk through the architecture and discuss the interactions that take place between the components. We will provide sample scenarios along the way and how we envision a solution can be delivered leveraging the various components of this architecture.

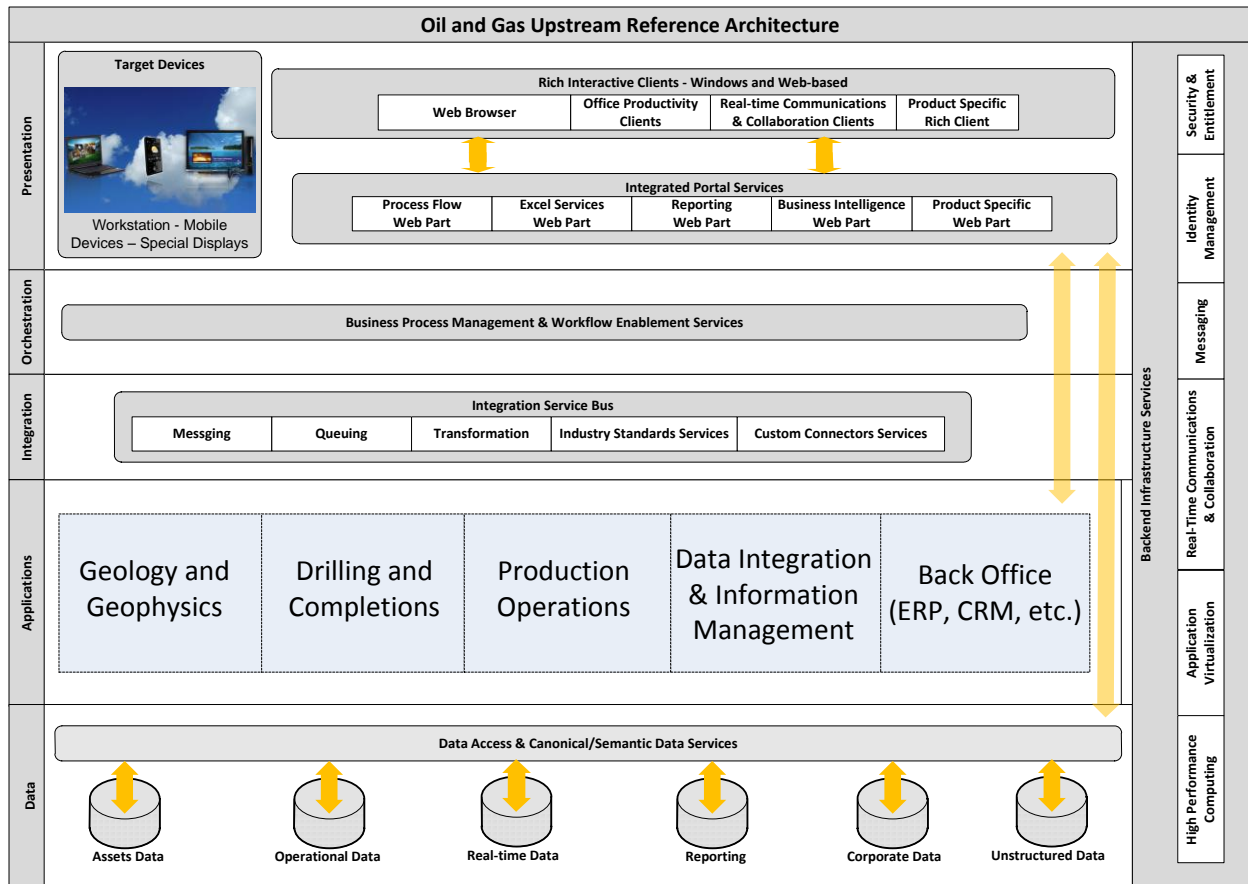


Figure 4 – Upstream IT Reference Architecture

Integrated Portal

The Integrated Portal functions as a common platform where geoscientists, engineers and managers access IT-based work processes of all kinds. The Portal establishes a single location where authorized employees find and use a wide range of data, including dashboard-based KPIs, applications for the interpretation of seismic data, log files, field studies, scout reports, rig activity and business intelligence systems.

This web-based portal also provides mechanisms that support blogs, wikis and social networks used to establish and maintain cross-domain collaborative systems. Rather than logging on to a specific system (such as SAP to access a work order), users simply log on to the Integrated Exploration Portal to access work orders, KPIs, analytic and other exploration-related systems.

This portal-based approach allows all disciplines and managers to focus on drilling assets, technical resources and reserve replacement ratios —rather than on finding data in various siloed applications. If, for example, a problem arises with a drilling rig, a drilling engineer or operational staff member can

quickly and easily use the tools available through the Portal to drill down and see all pertinent data relating to that rig to analyze the problem and make a timely operational decision.

Data integration

The Data Integration and Business Process Management components are in many ways the heart of a more effective IT architecture, providing a central mechanism for the movement of data between systems, equipment and other elements in the IT infrastructure.

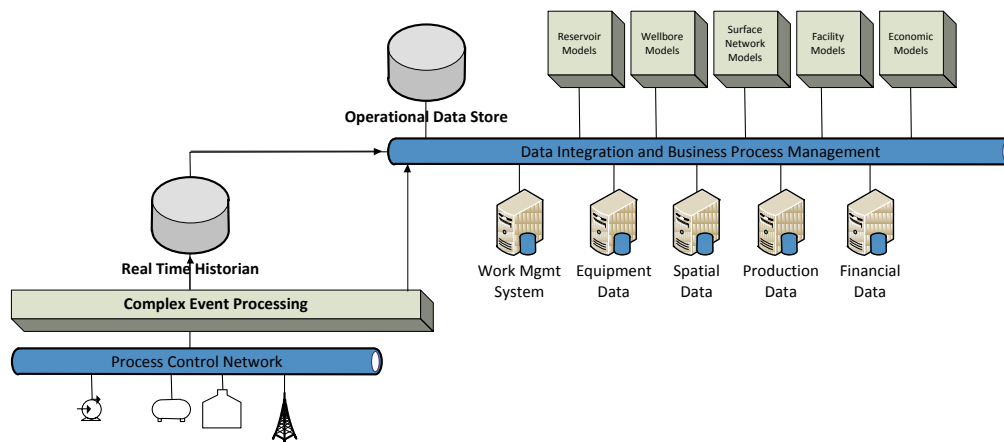


Figure 5 – Data Integration and Business Process Management

As illustrated in Figure 5, the Data Integration and Business Process Management element provides a centralized repository for incoming data from log analysis and from work management, spatial, production and financial systems. Using defined business rules, this element orchestrates the movement of data between various systems, including the Seismic Data Store, Operational Data Store and the economic evaluation applications used to model the expected ROI of lease opportunities.

The Data Integration and Business Process Management components also serve to synchronize hierarchical and metadata across systems to ensure operational coordination and reporting accuracy. This capability addresses a number of common issues in a dynamic oil and gas upstream operations.

When a lease evaluation or other type of work process is initiated, for example, the intelligence built into this element identifies the correct source of all data needed for that activity. That data is then collected and packaged into an XML file or other standardized format and forwarded to the application being used to handle the economic evaluation or other work process.

By establishing a common path for all relevant field data, the Data Integration and Business Process Management components also measurably reduce the time, cost and complexity of deploying new applications.

Industry Standard Interfaces (such as WITSML, PRODML, and RESQML) ensure interoperability and the smooth flows of information throughout the entire architecture.

A Canonical/Semantic Data Services element serves to translate and present data in domain-oriented terms that make sense to geologists, geophysicists, and petroleum or facilities engineers. Engineers, for example, do not think in terms of “SAP work orders”, but rather as “sub-surface work orders”. The Canonical/Semantic Data Services element presents data in a form that is logical and familiar.

The Orchestration Layer delivers data to the simulators used to run various models, as well as to the visualization tools, reporting systems, and business intelligence systems accessed through the Integrated Portal.

This Orchestration Layer also coordinates activities in a business process workflow. When a work process—say, a lease review—is initiated, this service pulls a list of wells, all relevant work orders, production data, well tests and other required information. This element then packages that data and serves that information up to the appropriate application.

Security and Entitlement Services allows organizations to provide convenient, role-appropriate access to enterprise data. A particular engineer might be given access to rock properties, well production histories and rig schedules relating to a specific set of wells but not to reservoir data that is proprietary to the business. This service is used to manage security and entitlement for employees, partners, and others in the exploration environment.

Forecast model synchronization

The ability to update and synchronize forecast modeling applications is critical to the success of a dynamic oil and gas exploration operation. This is a key capability within the Data Integration and Business Process Management Components of this architecture, ensuring synchronization of data for reservoir, wellbore, surface network, facility and economic modeling applications.

Here is how model synchronization works in this architecture: If a geologist makes a change in an earth model that reflects new rock properties from core data, that employee then re-runs the reservoir model to update the reserves estimates —and model synchronization then automatically pushes those changes to all other affected models. In this example of a reservoir change model, synchronization ensures that updated and accurate data is pushed to the pipeline and facilities models, and the economic model. Thus, users of all those modeling systems are working from a single, updated set of forecast data.

Managing mountains of data

The exploration process is the most data intensive of all the E&P processes. Seismic data constitutes the largest data type by volume, and one of the fastest growing data types. Online storage requirements are approaching multiple petabyte (1 petabyte = 1000 terabytes) for most large oil and gas companies.

Offline storage requirements are even larger. In order to continue accessing data sets greater than 10TB in size and make informed investment and management decisions concerning exploration opportunities, the oil and gas industry needs high-performance solutions for its computational and data-driven problems, along with the capability to leverage the expertise of geoscience experts worldwide and from its vendor community. Microsoft's Windows High Performance Compute (HPC) cluster server environments provide an easy-to-access compute platform that makes the analysis of large data sets easy for the non-expert user to schedule and run. With Windows HPC, users can deploy, manage, monitor and maintain a Windows Server-based cluster using the same tools already used to manage existing Windows Server systems.

Operational data store

The Operational Data Store serves to aggregate hierarchical data and metadata across the entire architecture. Here, data are also optimized for different purposes and synchronized for use in a wide range of exploration-oriented applications. Well lists, for example, are maintained in the Operational Data Store, and when a change of any kind is made to that list, the change is pushed out to every system that includes or relies on an accurate, up-to-date well list.

In this architectural approach, data is stored in a domain-based data model, rather than being identified and organized by the system that originates the data. As noted elsewhere in this paper, this domain-oriented method of handling data allows engineers and other employees to relate data more directly to everyday oil and gas exploration activities.

It should also be noted that this Operation Data Store does not create another system of record. Information is stored here only for reporting purposes or to be presented through the Integrated Portal or is metadata needed for data management purposes.

Securely collaborating with partners

This architecture also provides for highly secure and convenient collaboration between the operator and various partners. An external network, or cloud, establishes a secure location where partners can log in, store data, share and collaborate as needed.

Federated security capabilities allow producers and partners to establish secure user groups and identifications. In this federated security environment, an employee from Contoso Oil Corporation would log in using a Contoso identity, while personnel from Frabikam would log in to the same external network using their own corporate security credentials.

The result is a convenient yet secure external network designed specifically to serve collaborative partnerships in the oil and gas exploration sector.

Getting there

How can oil and gas firms best realize the IT architecture described in this paper? In the Microsoft view, companies should consider the following factors when seeking to create a more flexible and proactive IT infrastructure.

First, don't try to "boil the ocean". Instead of trying to do everything at once, pick a process (such as well reviews) and build the infrastructure, connectivity and processes needed to accomplish that process within the integrated, service-oriented environment described here. Recognize that this high-level architecture is an objective and can at the very least serve as a guideline for service providers seeking to develop and deliver compatible applications for a more common architecture.

Focus on business processes and work to incorporate robust data management into those processes. If a process includes data on a well to be drilled in the future, formulate a solution that transfers that updated information to any application that uses data on well counts, production volumes or other relevant metrics. If the process is exploration focused, concentrate on the search, discovery and collaboration aspects of the process to enable a robust discussion and consideration of the various insights and innovation each discipline, partner or vendor brings to the dialogue.

Use the vocabulary of your business when creating a solution. Build data models that use and present information in ways that engineers and geoscientists understand—not based on the systems that generate that data.

Maintain a "system of record" for data. It makes little sense to create additional databases or to gather all data into a temporary data warehouse to support daily work processes. Rather, seek to establish an integrated, service-oriented data model that recognizes where information is stored, easily accesses needed data and serves that information efficiently up to systems, employees and managers. Create data warehouses only when they can deliver and improve the speed and performance of the end user experience.

To deliver faster insights and make better decisions, IT solutions must enable visibility and collaboration. Microsoft has created a robust partner ecosystem designed to help oil and gas firms transcend traditional barriers, accelerate decision making and drive efficiencies. Our partners' E&P software can help companies gain the full benefits of a more effective IT architecture.

The reference architecture described here and the process of transitioning to this more efficient future state can be applied to address the real-world needs of oil and gas upstream operations. To learn more about Microsoft's oil and gas industry solutions, please visit www.microsoft/oilandgas.