



## Vision for an Upstream Reference Architecture

---

*Business factors in the upstream oil and gas industry are driving the need for an information technology (IT) architecture that provides a common reliable environment that maximizes technology to improve the efficiency of upstream oil and gas analysis, operations, and business.*

*This paper describes the vision for the Microsoft Upstream Reference Architecture (MURA).*

*MURA is not prescriptive—that is, it does not lay out specifics of the architecture’s structure and function. Rather, it describes a set of foundational principles to which it must adhere. This descriptive approach provides an agreed-upon set of principles for establishing consistent performance, but also provides the flexibility for companies to innovate and establish competitive differences.*

*This agreed-upon reference architecture approach will encourage simplification and unification for all organizations in upstream oil and gas. System integrators and solution providers will benefit from an established, coherent environment within which to build solutions. Operators can be confident that application solutions will run and integrate into their IT environments.*

### Contents

Upstream Businesses Demand More from IT Architecture .....	2
Drivers for the Evolution of a More Efficient Architecture .....	5
Enabling the Evolution .....	6
Microsoft Upstream Reference Architecture: Approach and Principles .....	7
MURA Overview .....	10
Getting There .....	14

## Upstream Businesses Demand More from IT Architecture

An upstream reference architecture must support and respond to the functional activities of an upstream organization and provide the capabilities needed to effectively and efficiently run the business.

Oil and gas exploration and production (E&P) is a vast, complex, data-driven business, with data volumes growing exponentially. These upstream organizations 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, reservoir modeling, production, and other upstream activities. At the same time, large amounts of information pertaining to those same activities are generated in unstructured forms, such as emails and text messages, word processing documents, spreadsheets, voice recordings, and others.

Figure 1 shows the broad spectrum of structured and unstructured data upstream organizations use to orchestrate, automate, integrate, and execute integrated upstream operations and management activities.

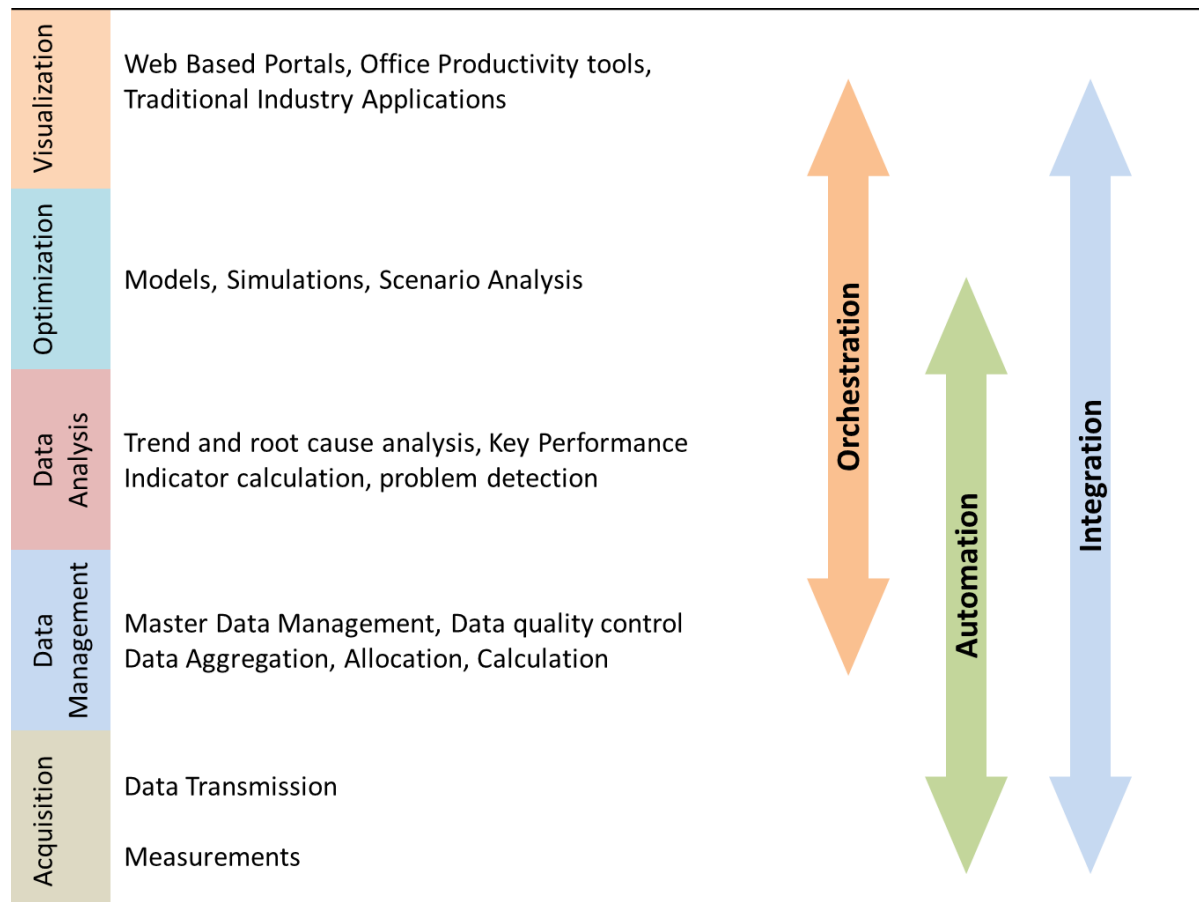


Figure 1 Upstream business activities use a broad range of structured and unstructured data.

Domain-oriented structured data is used for:

- Collaboration, including visualization, data fusion, decision tracking, and knowledge management.
- Optimization, including simulation, proxy models, decision selection, and implementation.
- Operations data analysis, such as trend- and root-cause analysis, solution evaluation, key performance indicators (KPI) and problem detection.
- Data management, which includes: quality control, validation, data storage and archiving, loss management, allocation and rate estimation and acquisition, including measurements and data transmission.

## A Day in the Life of an Upstream Organization

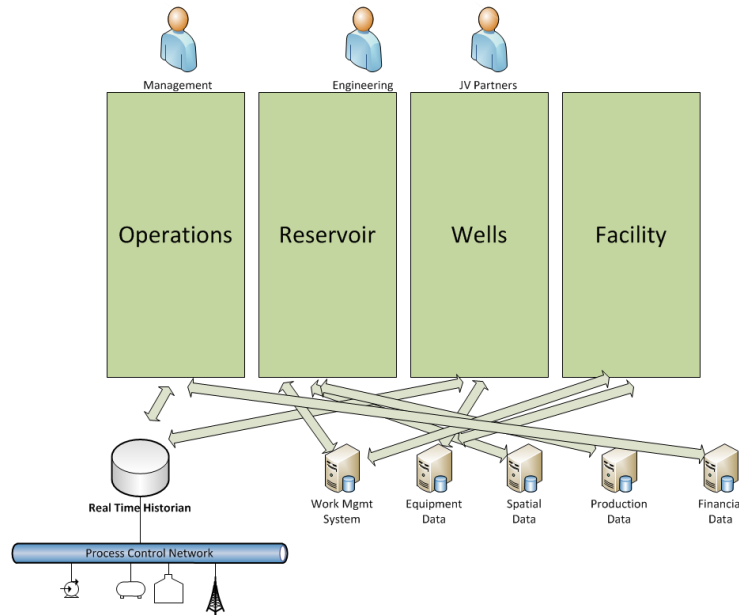
For a clearer understanding of the use of both structured and unstructured data, consider the following scenario.

A global asset team made up of geologists, geophysicists, and reservoir engineers, located in three different countries, works together to develop a field development plan to assess the economic potential for various tertiary recovery options on a key field with declining production. The team tunes the reservoir model in Petrel with all available G&G and production data, and plans several well options modeling fluid flow along streamlines and using reservoir simulation to assess the potential and impact of the placement and timing of those wells. The team works together on the shared models within the application software. However, the number of scenarios and the complexity of the analysis require that the work be an iterative, collaborative effort. So the team also discusses options and exchanges ideas using email and text messaging, and shares necessary documents through their secure team portal, which makes it possible for them to prepare multiple options in parallel for management and partner review.

## Current State: Overview and Challenges

The current state of IT infrastructure in most upstream businesses is unable to adequately support and respond to analysis, operations, and business needs.

In most organizations, the volume of information is increasing exponentially because digital sensors are deployed in more exploration and production plays, more data sources are connected to IT systems, and growing volumes of information are captured and stored in enterprise databases. Large volumes of domain-specific information are also embedded in various upstream applications. This data situation means it's difficult or impossible to use that data to quickly and efficiently get the information and answers needed.



### Current Architectures

Existing IT architectures in the upstream oil and gas sector 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, email, and other basic business applications.

**Figure 2. The current state of IT architectures for the upstream oil and gas sector.**

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

**Data Management.** That growing volume of data now typically resides in disparate source systems, such as Landmark’s SeisWorks® seismic interpretation software, or Schlumberger’s Petrel—or maybe even a combination of both. The Web-based tools used for viewing and collaborating on this 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. These one-off connections add time and cost, and cannot easily be shared or reused by other applications. Various industry solutions provide data- or application-integration frameworks—such as OpenSpirit—which create a common access layer to help address this integration problem.

**Integration.** Each discipline—petrophysics, geology, reservoir engineering, and others—tends to have and use its own analytic modeling systems, but currently little connectivity or interaction exists between those models. Therefore, changes in conclusions for one discipline are not always carried through to others, which can cause increase inaccuracy, errors and uncertainty.

**Collaboration.** With current IT infrastructure, collaboration is also difficult because there is no convenient, shared location where multiple internal and external partners can access information stored on the corporate network. For example, a seismic service company employee who works on prospects for multiple energy companies needs separate log-in locations and passwords for each of those collaborative partnerships. These same collaboration challenges also typically exist within individual oil and gas companies.

**Performance Management.** In the current state, KPIs, which are needed to understand and assess the current status and overall health of an organization, are often not readily available. The manual, time- and labor-intensive processes needed to gather and analyze KPIs means that managers and engineers waste valuable time waiting for answers, while data is collected, analyzed, and translated into the insights needed to understand and run the business.

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

## Drivers for the Evolution of a More Efficient Architecture

Powerful reasons are compelling oil and gas companies to seek a new and more efficient upstream IT architecture. Companies must have:

**Ability to deliver more with less.** In today's business and operational environment, companies must deliver more throughput with fewer resources and severely time-constrained work teams. To deliver better results faster, G&G and engineering workers must be able to spend more time doing domain-focused work—and less time searching for and preparing the data needed for that work. Workflows, data-driven events, and automated analysis should help drive their efforts to identify risks and help manage the exploration portfolio or production operations.

**Integrated views.** Workers also need integrated views that reveal all relevant data, both structured and unstructured, for a particular situation. For example, in an exploration scenario, that comprehensive perspective should include tornado charts that measure risk, analog well production histories, rock properties, log files, rig schedules, and other variables relating to the prospect in question.

**Easily accessible KPIs.** Management needs up-to-date KPIs to fully understand the current status and overall health of an organization. For example, ideally managers should be able to see a single screen showing the portfolio of current opportunities, which ones are drill ready, the 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 to the expected level of productivity.

**Plug-and-play technology.** The industry needs an architectural approach that allows upstream organizations 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 several months, to just a month or even weeks.

**Integration of structured and unstructured data.** 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 email 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 the Evolution

As noted, several emerging industry standards and technologies are now enabling the more flexible, integrated IT architecture needed in the upstream oil and gas.

### Standards

New XML standards-based technologies (including the WITSML and PRODML standards refined and supported by Energistics and other organizations, [www.energistics.org](http://www.energistics.org)) provide the common data interfaces needed to ensure plug-and-play access to best-in-class hardware and software solutions. For example, if a company currently has a wellbore viewer that is WITSML compliant, with that interface in place, the company can deploy any WITSML-based wellbore viewer solution.

### Technology

**SOA.** The oil and gas IT community is now adopting service-oriented architecture (SOA) as a more flexible and responsive alternative to traditional hard connections between applications and source data. SOA is a collection of connected services that communicate with one another. These services may communicate one-to-one, passing data back and forth, or may communicate among several services, which includes the ability to make applications that might use the services “aware” of their existence.

**Cloud computing** refers to remote centers for storing and accessing data and applications using the Internet, which are designed to save businesses money, in part by reducing the need to build major on-premise computing infrastructure.

The cloud approach is ideal for complex upstream operations, with its multi-vendor, multi-partner environment and huge volumes of data that require a combination of strict security and easy sharing with appropriate partners.

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

## Microsoft Upstream Reference Architecture: Approach and Principles

The Microsoft Upstream Reference Architecture (MURA) is not prescriptive—that is, it does not lay out specifics of the architecture’s structure and function. Rather, the MURA describes a set of foundational “pillars,” or principles, that govern it. This descriptive approach provides an agreed-upon set of principles for establishing consistent performance, but also provides the flexibility for companies to innovate and establish competitive differences.

Figure 3 shows the five principles (at the top of each box), which are also described in the following sections.

# Oil and Gas Reference Architecture’s Foundational Principles

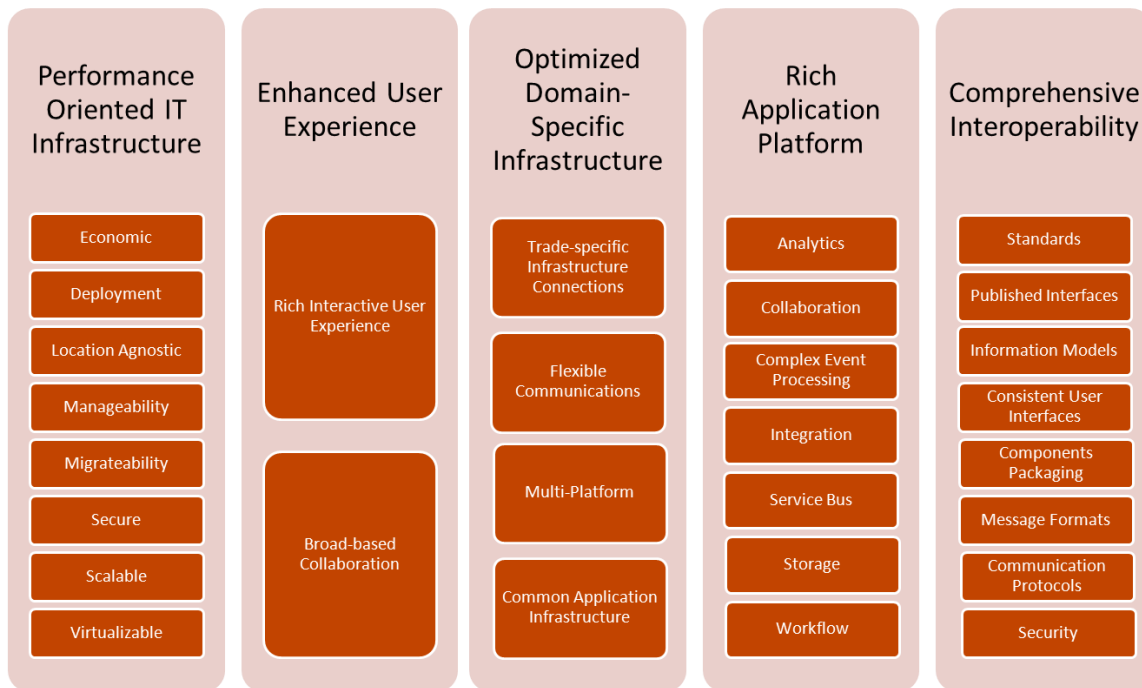


Figure 3. Foundational Principles of the Microsoft Upstream Reference Architecture (MURA).

### Performance-Oriented IT Infrastructure

A performance-oriented infrastructure includes those features that make an architecture complete and appropriate to business needs, such as:

- **Economic.** The infrastructure must provide cost-effective means to deploy and integrate new functionality.
- **Deployment.** Components must 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.
- **Migrate-ability.** 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 way 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 MURA 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, and then flowing that data into appropriate operational systems.
- **Flexible communications.** Deployments can leverage a variety of communications paths and technologies, and are easily reconfigured. This minimizes 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 by leveraging the appropriate OS for the right workload to deliver optimal performance.
- **Common application infrastructure.** Provide the applications infrastructure and services for commonly used capabilities so developers can focus on domain-specific functionality, thereby optimizing speed to market and reliability of solutions.



## Rich Application Platform

Inherent in the MURA approach is the notion that no one vendor can provide all of the application functionality needed to implement it. The purpose of MURA is 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 that 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 that enable analysis of operational and business data.
- **Workflow.** Services for managing the automation of applications and business processes.

By providing these services to developers as part of the architecture, Microsoft partners can focus their expertise on solving domain-specific problems, also making it easier for customers to leverage those applications.

## Comprehensive Interoperability

For the MURA to successfully deliver the cost-effective, integrative benefits for which it was conceived, it must enable comprehensive interoperability. Pragmatic integration approaches must be considered, and the MURA must 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 that 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.
- **Consistent User Interfaces.** Consistent content and behavior in presentation of information and interaction with the user.
- **Information Models.** Consistent ontology (naming system) for referring to equipment and assets to enable exchange of information throughout the enterprise and the value chain.
- **Components Packaging.** Well defined sets of functionality packaged for developer and integrator reuse.
- **Message Formats.** Key construct of SOA defining format and content that enable services exchange messages using the defined format (e.g., 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.

## MURA Overview

Figure 4 shows the MURA, which reflects a service-oriented computing environment that encompasses the integration of domain applications, business productivity tools, and back office applications.

By following the service-oriented architecture (SOA) approach for interactions between components on different layers of the architecture, upstream businesses, technology vendors, system integrators, and other Microsoft partners can realize the full benefits of this environment. To function as planned and to eliminate the need for additional integration work, the MURA must provide inherent, seamless integration to connect all layers and applications. Built-in integration enables frictionless communication and information flow, up and down the layers, and provides the ultimate plug-and-play efficiency needed.

The next sections walk through some key aspects of the MURA and discuss the interactions that take place between the components. Scenarios serve as examples and show how a solution can be delivered by leveraging the various components of this architecture.

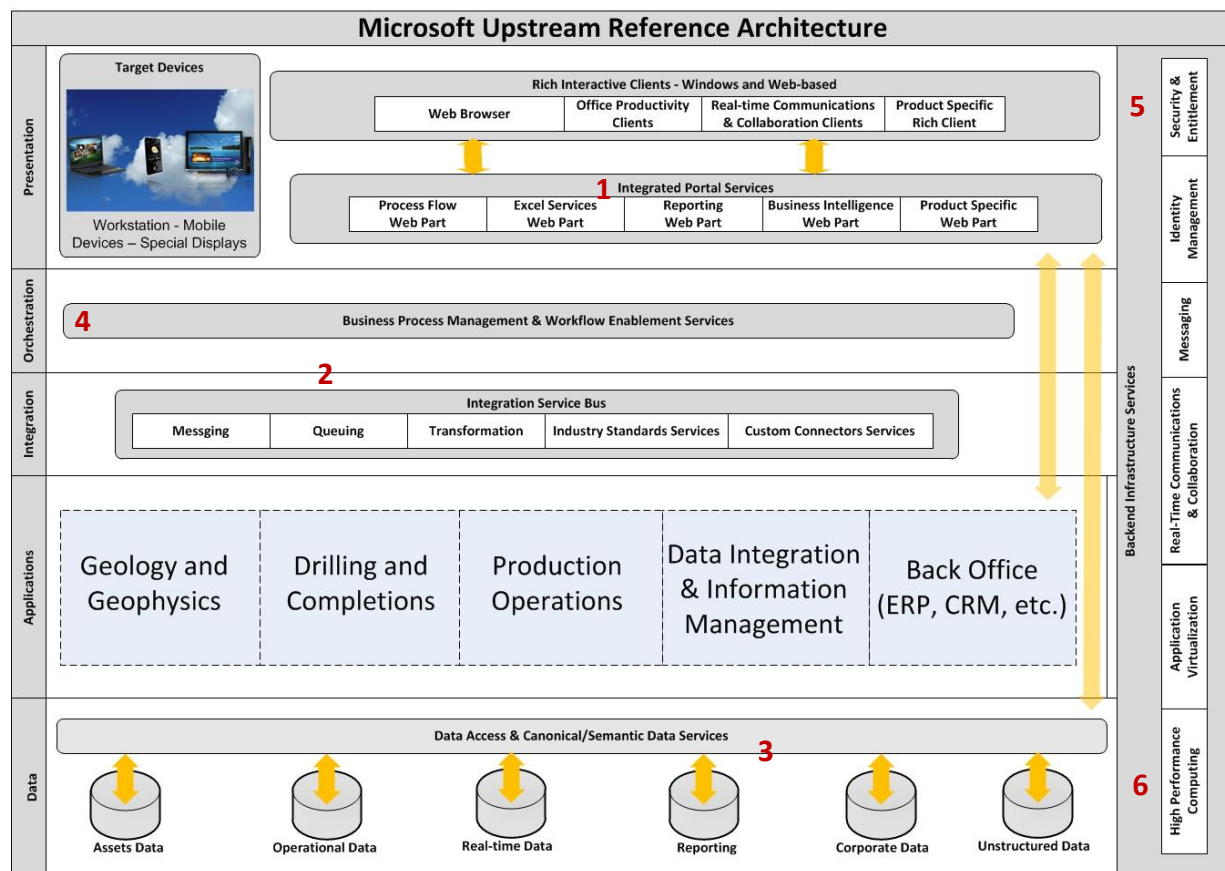


Figure 4 Microsoft Upstream Reference Architecture (Red numbers on this figure refer to references in text below.)

## Integrated Portal

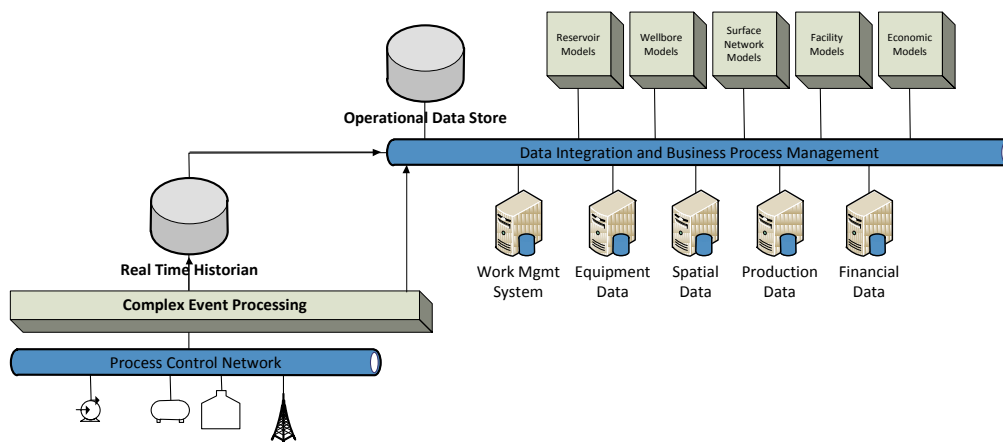
The Integrated Portal (1) functions as a common platform where geoscientists, engineers, and managers access IT-based domain work processes of all kinds. The Portal establishes a single location where authorized employees find and use a wide range of data, including dashboard-based KPIs, technical applications, such as those used 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 Portal to access work orders, KPIs, analytic, and other exploration- or production-related systems.

This portal-based approach allows all disciplines and managers to focus on drilling assets, technical resources, and reserve replacement ratios, instead of working to find data in various siloed applications. For example, if a problem arises with a drilling rig, a drilling engineer or operations personnel 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 (2) 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**

Figure 5 shows how the Data Integration and Business Process Management component provides a centralized repository for incoming data from log analysis and from work management, spatial, production, and financial systems. Using defined business rules, this component 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 dynamic upstream operations.

For example, when a lease evaluation or other type of work process is initiated, the intelligence built into this component identifies the correct source of all data needed for that activity. The 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 and PRODML) ensure interoperability and the smooth flow of information throughout the entire architecture.

A Canonical/Semantic Data Services component (3) serves to translate and present data in logical and familiar, domain-oriented terms that make sense to G&G, engineering, and other upstream professionals.

The Orchestration Layer (4) 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. For example, when a work process, like 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 component then packages that data and serves that information up to the appropriate application(s).

Security and Entitlement Services (5) allow organizations to provide convenient, role-appropriate access to enterprise data. For example, a particular engineer might be given access to rock properties, well production histories, and rig schedules related 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 involved in a project or venture.

### **Forecast Model Synchronization**

The ability to update and synchronize forecast modeling applications is crucial to the success of a dynamic upstream oil and gas 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 (refer to Figure 5), synchronization ensures that updated and accurate data is pushed to the surface network, facilities, and

economic models. Thus, users of all those modeling systems are working from a single, updated set of forecast data.

## Managing and Processing Mountains of Data

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

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 computing (HPC) solutions for its computational and data-driven problems. Ideally, the high-performance computing solutions should be integrated with the MURA and provide the capability to leverage the expertise of geoscience experts worldwide and from its vendor community.

High performance computing (HPC) cluster server environments powered by Windows® HPC Server provide an easy-to-access HPC platform that makes the analysis of large data sets easy for the non-expert user to schedule and run. With Windows HPC Server, users can deploy, manage, monitor, and maintain an HPC server-based cluster using the same tools already used to manage existing Windows Server systems.

## Operational Data Store

The Operational Data Store (Figure 5) serves to aggregate hierarchical data and metadata across the entire architecture. Here data is also optimized for different purposes and synchronized for use in a wide range of exploration- or production-oriented applications. For example, well lists 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 and production activities.

Also important to note: this Operational 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 as metadata needed for data management purposes.

## Securely Collaborating with Partners

This architecture also provides for highly secure and convenient collaboration between an 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 operators and partners to establish secure user groups and identifications. For example, employees from Contoso Oil Corporation log in using their Contoso

identities, while employees from partner Fabrikam 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 upstream oil and gas industry.

## Getting There

How can oil and gas businesses best realize the MURA described in this paper? Microsoft urges companies to consider the following factors when seeking to create a more flexible and proactive IT infrastructure.

**Start small.** Instead of trying to do everything at once, pick a domain 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 up that information efficiently 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 businesses transcend traditional barriers, accelerate decision making, and drive efficiencies. Our partners’ upstream software solutions can help companies gain the full benefits of a more effective IT architecture.

The MURA 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](http://www.microsoft/oilandgas).