SPE 112134

Industry Standards and Improved Operations Integrity With a Model Framework—Reference Data Model

Adrian Chapman, Rik Evans, and Ron Montgomery, IBM Global Business Services

This paper was prepared for presentation at the 2008 SPE Intelligent Energy Conference and Exhibition held in Amsterdam, The Netherlands, 25–27 February 2008.

Abstract

Fundamental to the Intelligent Energy initiative are the acquisition and presentation of information across a distributed supply chain. This is for matters of integrity assurance and the wider elements of production optimisation, maintenance and normal operations. Thus, energy companies are seeking new solutions for the collection of real-time data from across their enterprise, from well-bore to export pipe and beyond, and the integration of this with other forms of information. The complexity which has prevented this from being realised until now arises in that the information required by the many disparate users is derived from multiple sources with differing protocols, data structures and suppliers, across the breadth of the asset's lifecycle.

This paper presents an enterprise scalable architecture that supports the activities of users in their searching for and recovery of information. It describes how it is possible to structure, partition and secure that data so that anyone can get access to the information when they need it, and yet remain in control so that it is possible to know where the information has come from, when it hasn't been delivered 'right' and when it needs to be updated to reflect process or technical changes.

This approach is significant in that it marks the drawing together for approval and publication of a number of widely used standards into a single data description and the deployment of this to meet a real-world problem that is affects all assets in operation or development.

Introduction

"In a five to ten years perspective, it is expected that Integrated operations (IO) by smart use of real-time monitoring, real-time control, visualisation and new work processes based on Information and Communication Technology (ICT), will constitute the single most important driver for increased efficiency in the Norwegian oil and gas industry" - Source: Norwegian Oil Industry Association OLF[1]

As an example, the economical potential of Integrated Operations on the Norwegian Continental Shelf (NCS) is estimated to be of the order of 250 billion NOK, against investments of 30-50 billion NOK over the next few years. The OLF report "eOperations on the Norwegian Continental Shelf"[2] shows that integrated operations can reduce operating costs by 20–30 percent, and increase production by 5–10 percent, an improvement on the benefits estimated by CERA under their Digital Oil Field of the Future (DOFF) study in 2004 (5-20 percent and 106 percent respectively)[3].

However, the complete realisation of these benefits is dependent on successfully addressing all of the components shown in the diagram below.

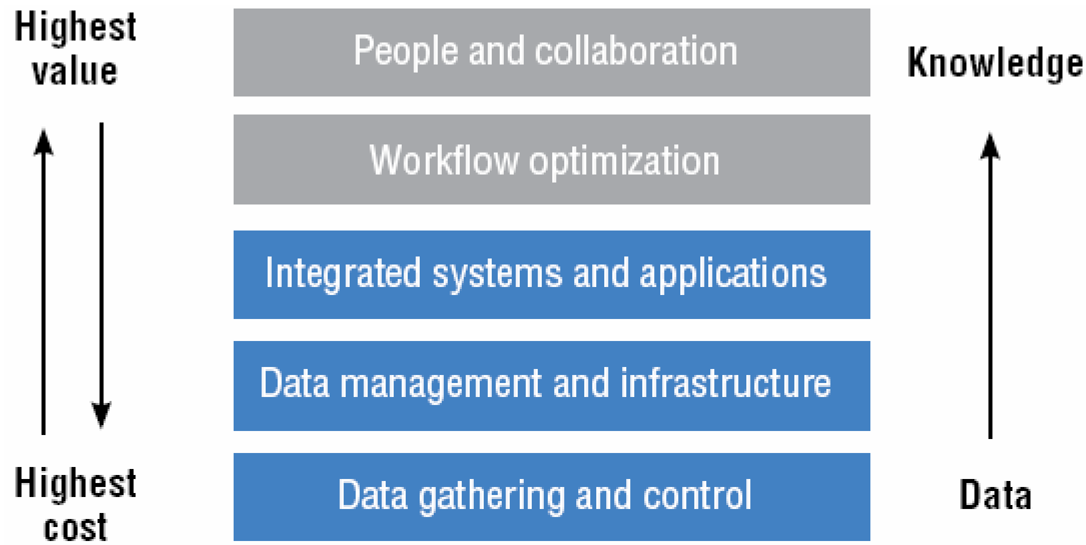

Figure 1 - A successful intelligent oilfield initiative relies on all components fully integrated for optimum results[4]

The intelligent oilfield is designed to help people (the highest valued component) work together more effectively, in order to reap the greatest value from the highest-cost component (data gathering and control) and its midlevel byproducts. Raw data (bottom right) migrates up through each component until it is converted into knowledge (upper right), which people use for improved decision making. Ultimately, all the technology components support the workflow of skilled personnel. Whilst these steps can be considered and even implemented independently in time, a successful initiative depends on fully integrating all five as part of a unifying initiative.

A critical component of this framework is Data Management and Infrastructure; not only is it a requirement in its own right, but it facilitates the transfer of data into knowledge. Indeed, the benefits of Integrated Operations can only be achieved if real-time and historical data can be captured and utilised within an enterprise-wide, integrated systems and applications that supports information harvesting (e.g. viewing and using the data within its context) and knowledge capture (e.g. applying the information to generate business benefit), thereby allowing for improved workflow and decision making.

- Having the data and information at your finger tips wherever you are and whenever you need it
- Being able to access data in context, without a priori knowledge of the asset topology
- Having a lifecycle data and information set across disciplines (Engineering Design, Production operations, Drilling and Completions, Maintenance, Projects, etc) allowing you to get the "complete picture"
- Understanding how your asset classes and / or specific equipment items are performing, against their design / build characteristics and in relation to other similar equipment items
- Tracing the root cause of changes and / or failures
- Having the confidence that the data and information is accurate and is single source of the truth

Information management challenges

The challenge arises in that the data and information often stems from multiple sources and from many disciplines, across the breadth of the asset lifecycle. Furthermore, because each discipline has its own specific data and information needs, and to a certain extent different suppliers, functional solutions tend to be developed in isolation with little consideration of the requirements of other disciplines and /or little opportunity for cross-referencing of information. For instance, during the projects' construction phase, data and documents about any one piece of equipment are developed and collated from a number of sources including, EPC, Operations team, ICSS vendor, equipment manufacturer, etc. Later, once the asset is Operational, information comes from additional locations, e.g. maintenance system, PSCM system, ICSS, historian simulations, etc. as well as all the legacy documents developed during, or even before, the MP phase. This means any future programme to optimise asset performance through reference or integration of data sources from these disparate silos will be compromised because the information cannot be accessed when and where you need it.

The traditional response of oil and gas companies has been to develop workaround, bespoke systems and / or separate databases/data warehouses. However, these solutions simply perpetuate the information access and connectivity challenges by building bigger information silos which are even less flexible and require more maintenance; which were precisely what the programmes were intended to avoid in the first place.

What is needed is a single information framework that provides the information to where it is needed, when it is needed, in a flexible, scaleable and extensible solution. This is the Integrated Information Framework.

Integrated Information Framework - IIF

The Integrated Information Framework is an integrated, asset lifecycle information management solution that consists of:

- A common and persistent data model based on leveraging industry standards and with enterprise wide ontology and consistent tag nomenclature
- Model-aware adapters, supporting automated configuration and model updates i.e. that relate and transform incoming data, providing contest to the information
- SOA-enabled services to enumerate, locate and refer to information irrespective of its location and originating application i.e. distinguishing business process logic from the data logic
- Embedded complex event processing and notification
- Multiple methods for data access, including independent and rich thin client SVG based visualization

A common and persistent data model

In conjunction with the following standards organisations (MIMOSA, OAGi, WBF, ISA88, ISA95, OLF) and organisations (OSISoft, DOW, SISCO); IBM have worked to unify the standards within a single Reference Semantic Model (RSM) – See Figure 2 overleaf - that can be used by manufacturing companies (process industries, oil & gas, automotive, aerospace, etc.) using a "harmonized" S95 and S88 throughout the enterprise

Importantly, it is not a data model and does not constrain the way applications implement the information contained within the model. The Reference Semantic Model facilitates the exchange of information, and is an implementable data exchange model.

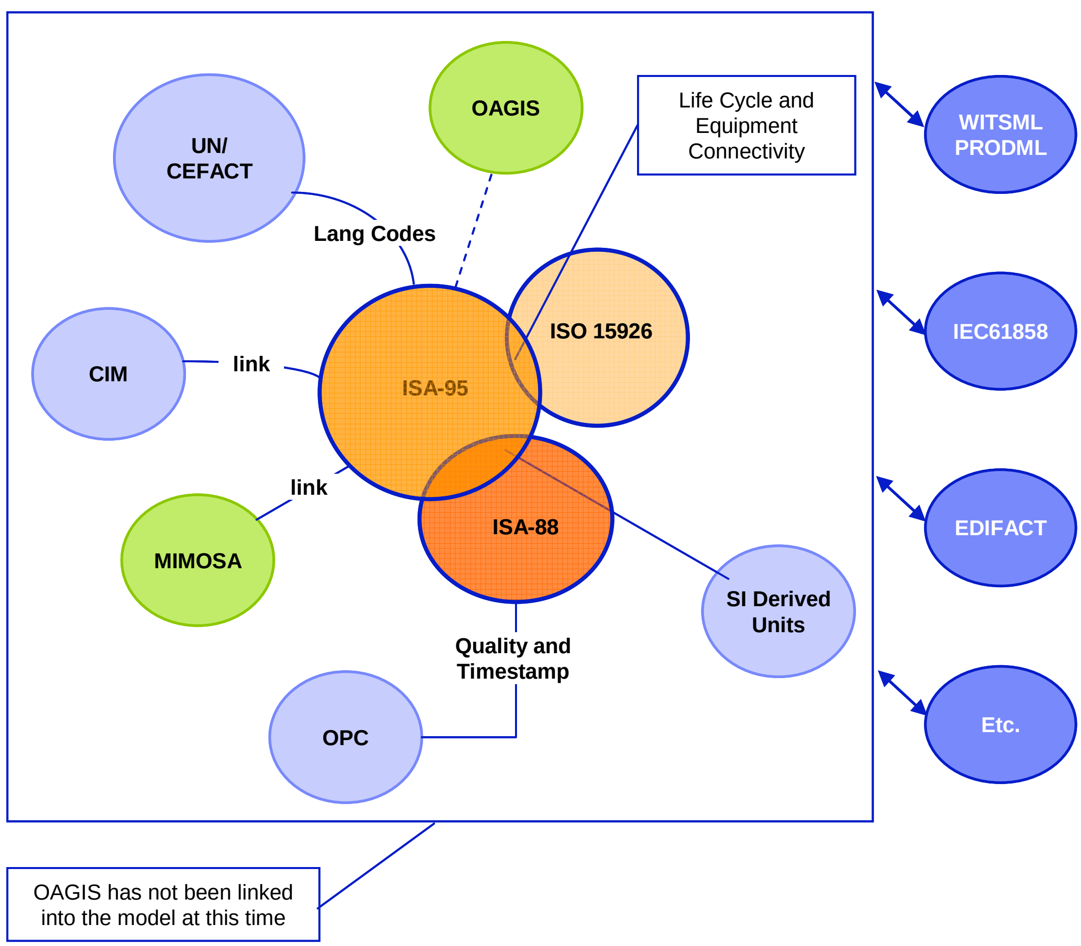

Figure 2 - The conceptual Reference Semantic Model (RSM)

Model Aware Adapters

Traditionally applications have been developed and deployed to support silos within the value-chain; meaning that as production or drilling operations have become more integrated, bespoke interfaces are developed to link the applications and ensure that decisions or results in one area can be properly accommodated in another. Couple this with the fact that each application has been developed to refer to its own unique reference and data model, results in a situation when the information that is interchanged between applications is conveyed through process tags containing little or no contextual information (e.g. location in the topology or asset hierarchy). The net result is that there is an overly high reliance on detailed engineering knowledge of the asset when interconnecting applications, and adding or additional or changing existing applications is time-consuming, costly and error-prone.

Model-aware adapters remove these constraints, because the information is mapped against an asset-model based on the RSM at the point where it is linked to the IIF. Because there is a single RSM-derived model for the enterprise, each information source must be unique within the context of the enterprise. Or put another way, since there is a unique piece of equipment in the physical world, which has unique physical attributes; there is a corresponding unique attribute associated with that unique object in information space. Thus, even if data comes from multiple sources to populate the objects' attribute-set, e.g. real-time temperatures, design data, maintenance records, etc. they can be mapped individually to the object at the interface. Thus when another application or visualisation refers to that object they can see all the attributes, irrespective of their originating application.

SOA-enabled

Historically, data has been retained across the asset in forms that are tightly coupled with the applications using the information, e.g. with multiple derivations or presentations of the same "truth". This results in the need for large amounts of manual entry and re-entry of information to transfer it from one system, or the development of complex interfaces between systems. At best this results a large maintenance burden to keep the system operational and inflexible implements: at worst this introduces errors into the process or introduces delay into work processes as people are forced to track down the right data.

SOA enables information to be presented as a service to those applications and processes using it. This means that the information is packaged identically, and promotes a greater understanding between the users. Because these is now a single source of truth, presented in a common way, maintenance is centralised, and there is greater flexibility to change any single element.

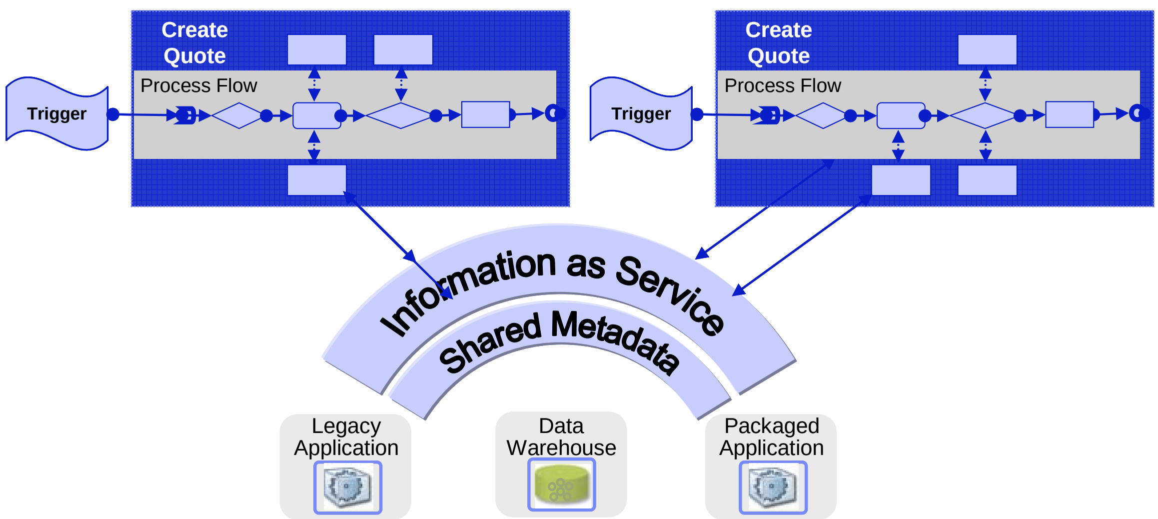

Figure 3 - Service Orientated Architectures improve flexibility

Embedded complex event processing

Existing solutions today allow complex logical or condition-based rules to be developed that can provide warning to control room operators, technicians or shore-based engineering personnel. Typically however these systems are limited by the data to with which they are fed, for instance, they will work on only maintenance related data (e.g. run-time between failures vs. design), production-data (e.g. actual, calculated or inferred values vs. design data), or historical-data (e.g. current performance vs. historical performance) and so on. Secondly, the configuration of these solutions is often overly complex, requiring a high degree of technical knowledge about the asset in order for the system to be built. They are also reliant on point to point interfaces to the information sources. Additionally, because of the need to have these links to the often disparate applications in place from one asset to another, it becomes an order of magnitude more complex to correlate events or performance from one asset to another, even when they are similar. Lastly, the structure of the solution often requires the user to 'know' not only the tag name by which the data-source is known, but also where, and thus when making changes he needs to make a manual correlation between the two.

The IIF solution solves these problems by deploying the event detection within the same architecture as the information bus. Consequently any user can configure the system simply by locating the data point within the asset/enterprise topology rather than having to know what it is called. The framework itself locates and routes the data to the event detection engine, meaning that the user does not need to know what information source its coming from or what it is called within that application. Again, because the information is presented within the context of the information bus, information from multiple sources can be used in the same event, with no need for additional application interfaces to be linked, hence complex events are much more easily configured and detected.

Support for multiple visualisations

The IIF in itself provides a visualisation portal based on open-standards GUI including direct capability for Scalable Vector Graphics as shown in the diagram below.

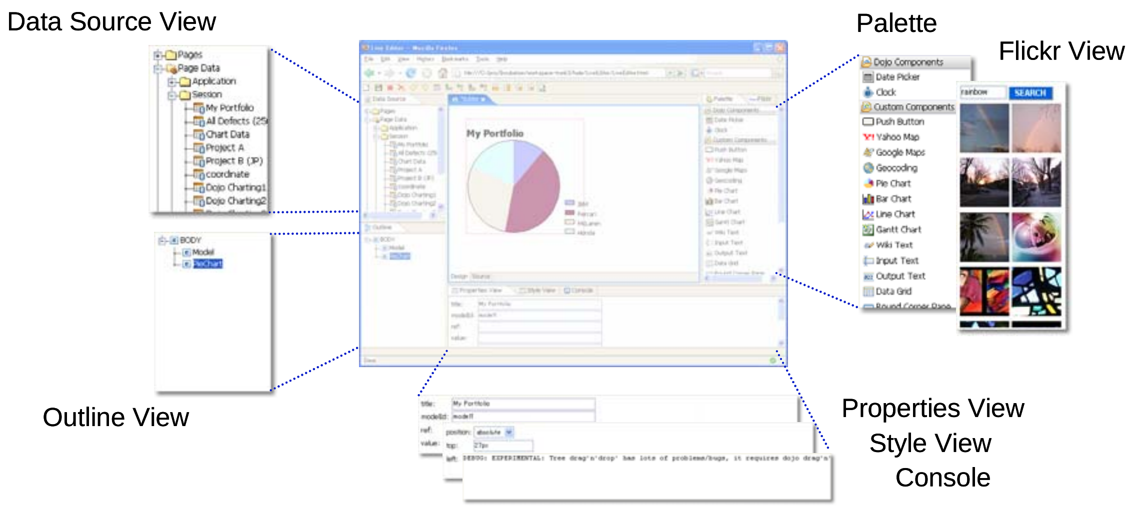

Figure 4 - Embedded views

In truth, however, most implementations will utilise existing portals and visualisation solutions, for which the IIF provides bi-directional connectivity through model-aware interfaces irrespective of the use that it is being put to. This means that your existing portals can be used without need for expensive re-engineering or reconfiguration.

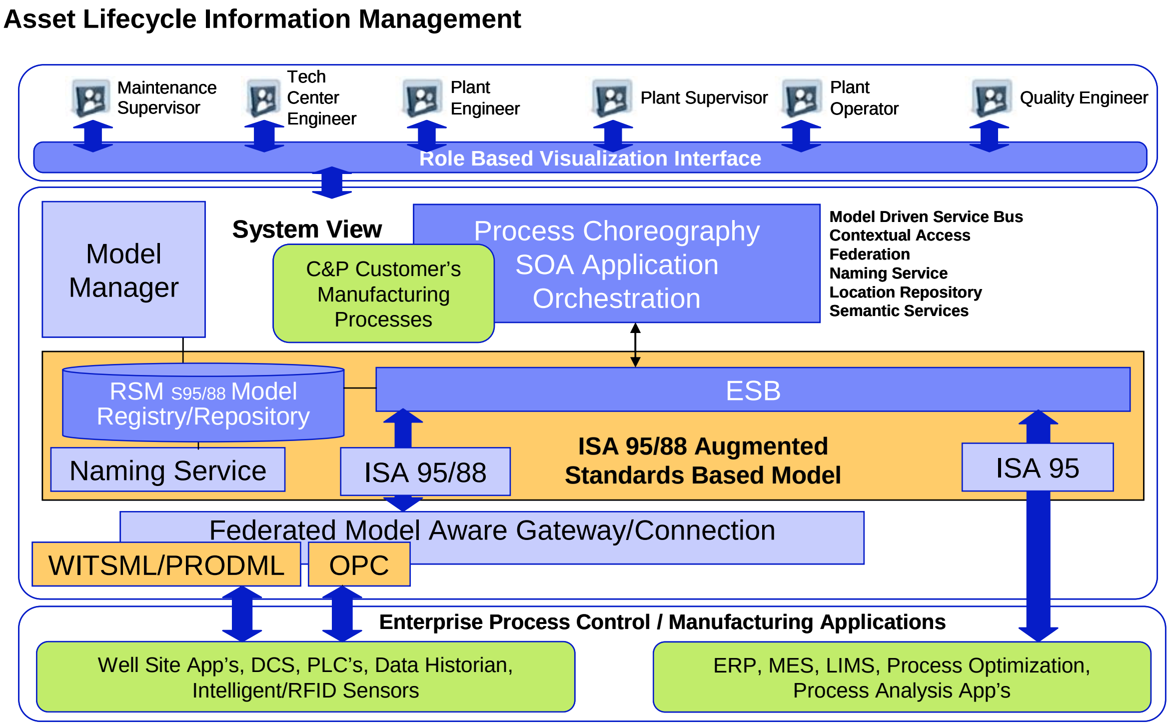

Figure 5 - Integrated Information Framework (IIF) - Architecture

The hub described in this paper can marshal the real-time drilling information from all well sites, can access real-time and historical information from process control systems and historians and interlace it with information garnered from other applications such as ERP, MES, engineering information systems, MRO solutions. This can then be provided it back to the enterprise for other applications to consume as required. As information is delivered from the application polling interface, a common "model-based" representation of the information is made and then delivered to applications for analysis and decision making. This is represented in Figure 5.

Furthermore consistent application design using these methodologies does not require you to redevelop any of the current applications or integration programmes; they provide a wrapper, or model-based interface definition on which further integration between applications across the breadth of an asset and along its lifecycle can be made. Least of all, this is not a suggestion that the solution is to create ever larger and more complex data warehouses to collect information from multiple sources into a single place, merely that the relationships between the existing databases be federated so that a user or application requesting data can find it irrespective of its originating location and without the need for extensive interfaces and complex data structures.

Benefits of IIF

IIF delivers significant business and IT advantage. At it's core, IIF establishes the ability to quickly and cost effectively view and analyze "asset" i.e. equipment part performance (e.g. condition based monitoring), transactions and events (alerts and event recognition) in the context of actual equipment connections. Given the reference model against which all equipment data items in all systems are referenced, tag data can be viewed and understood, without understanding the context of the tag and / or the data, thereby allowing for more confidence in decision making.

Given the model is consistent across the business (in terms of business units and assets), asset performance can be compared and contrasted across the enterprise. This combined with the fact that IIF allows for learnings and / or actions to be recorded against the "asset" means that best practices can be more effectively understood, learnt and leveraged across the organisation.

The learning capability, combined with the greater transparency provided by the reference model means that human resources and subject matter experts can be used more effectively and efficiently in operations, thereby alleviating some of the challenges of the "big crew change".

The reference model provides consistency of information access across disciplines (surface, subsurface, reservoir, wells, export etc) . This means that all information relative to a specific asset can be viewed through a single source, allowing for multiple dimensions of performance to be considered and, more importantly providing the ability to optimise across the value chain.

Finally, the reference model allows for consistent information access across the value chain (projects, operations). This means that information for an equipment iteam can be viewed across that equipment item's lifecycle. A typical scenario might be where a production engineer wishes to compare real time equipment performance against as designed or as built information and commentary as well as against the maintenance records. In summary, the IIF provides greater information transparency as well as greater information integrity allowing for improved decision making and ultimately asset integrity.

Conclusion

In terms of technology benefits, IIF offers an accelerated data management solution to the business that significantly reduces data management expenditure. As the IIF solution is based on industry standards, is effectively a data model "translator" (contains no data, but rather the semantic reference model) and is SOA based, it is easily scaleable and extensible across an entire enterprise. It there reduces individual project data management set up costs, new partner / supplier incorporation and accelerates delivery.

As IIF leverages existing databases and repositories (e.g. maintenance management and information management systems, process control system historians, SAP and other ERP systems, and so on), it enhances the return on existing technology investment. Furthermore, it removes the need for any business, regional or corporate data warehouses as it pulls the data from existing sources through the model. Hence, any application purely has to reference the IIF rather than go through a data warehouse.

Finally, as more and more assets / businesses are incorporated onto the IIF solution, IIF facilitates increased standardisation and integration across the enterprise. Tied to the fact that the solution is future proof (as it is based on industry standards), IIF therefore enables reduced longer term data management investment.

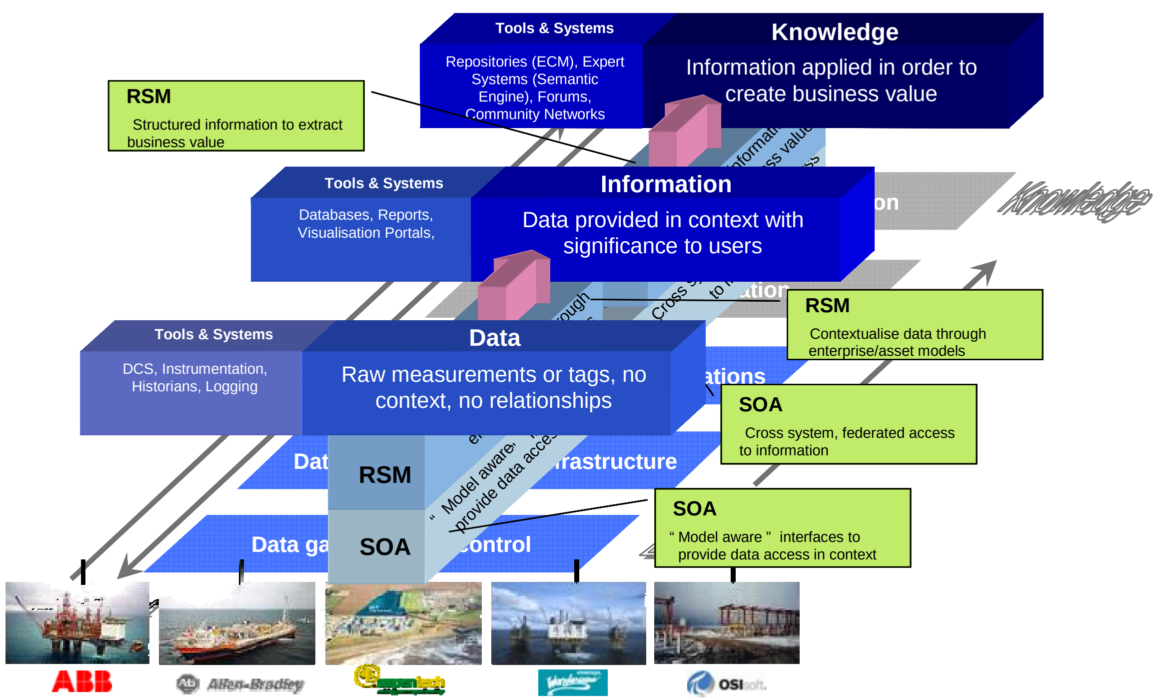

Figure 6 - IIF - Transforming Data into Knowledge

References

1. OLF report – "Integrated Work Processes: Future work processes on the Norwegian Continental Shelf" – Autumn 2005
2. OLF report - "The Shelf in 2005 – Integrated Operations" http://www.npd.no/English/Emner/E-drift/Sokkelaret_2005_edrift_120106.htm
3. CERA report "Making the leap toward DOFF adoption", January 2005
4. "The intelligent oilfield: meeting the challenges of today's oil and gas exploration and production industry", IBM Global Business Services, 2006