SPE-163698-MS

Real-Time Streaming Data Management as a Platform for Large-Scale Mission Critical Sensor Network Applications

Ronnie Beggs, SQLstream Inc., and Victor H. Abadie III, Consulting Geologist

This paper was prepared for presentation at the 2013 SPE Digital Energy Conference and Exhibition held in The Woodlands, Texas, USA, 5–7 March 2013.

Abstract

Oil and gas operators must process, analyze, and react in real time to increasing volumes and rates of streaming data in order to improve safety, compliance, and profit. For example, real-time analysis of streaming data from drilling rig sensors, intelligent wells, and digital oilfield installations enables early detection of drilling hazards and pending equipment failures, thereby reducing rig time, intervention, and shut-ins.

Data problems include the need for real-time integration of operational systems and an inability to maintain accurate and current information across all systems and data warehouses. Poor integration leads to duplication of data and systems, and a lack of visibility across all monitored assets, resulting in the delayed identification of the root cause of problems.

This paper presents a streaming data architecture that is capable of processing and analyzing rig data and smart oilfield data in real time. Further, the architecture supports the continuous, streaming integration of any and all of an organization's sensor data, operational platforms, and data warehouses. The approach aggregates and analyzes live, streaming data on the fly, without the need to store the data first. Large volumes of high velocity streaming data in any format and from all sources can be processed continuously and delivered to existing operational systems and data warehouses.

The architecture increases efficiency by removing data silos and by ensuring that all systems and data warehouses are kept accurate and up to date in real time. This results in faster decisions that are better informed and based on consistent, accurate, and timely information.

Achieving new levels of operational efficiency

The oil and natural gas industry faces challenges to increase production, reserves, and profit while becoming safer and complying with ever more stringent environmental regulation. Global oil production capacity is forecast to grow to 102 million barrels a day by 2017, an annual average increase of 1.5 million barrels a day (Kent, 2012), and demand for liquefied natural gas has doubled in the last decade, and is forecast to double again over the next decade (The Economist, 2012). Meeting objectives such as these will require the industry to attain new levels of operational efficiency.

Obstacles to operational efficiency improvement

Operational processes execute in real-time (for operational monitoring and alerting) and in post mortem mode (for further analytics, modeling, and data mining); yet, there is concern that the operational processes and information technology (IT) systems are unable to support the levels of automation, control, and overall visibility required.

The issues hindering operational efficiency are similar to those of other industries, namely:

- Vertical siloed business units, where each has deployed and managed its own systems and operational infrastructure, resulting in duplication of systems and data across the organization, with little or no integration across the different business siloes. Vertical business siloes make it difficult, if not impossible, to fully automate end-to-end business processes and to share information effectively.
- Increasing data volume and velocity is degrading operational efficiency, particularly where data are captured and stored at a rate in excess of several terabytes per day.

This results in an inability to transform operational data into actionable information and and a failure to drive operational, management and corporate-level decisions. That is, the operational platforms must share the right data and information in order to support efficient and timely decisions (Abou-Sayed, 2012).

Obstacles to real-time business integration

Today's operational infrastructure is built on traditional systems and platforms, based on the technology advances of relational databases, visualization, and data mining tools. As a result, traditional data storage and mining practice discards or fails to make use of significant amounts of data.

The oil and gas industry is not alone. Until recently all industries were similar. However, recent advances in IT, wireless and sensor technology means increasing volume of high velocity data can be processed effectively and efficiently. It has been proposed that the adoption of real-time Big Data streaming technology can increase production rates significantly, particularly in older fields, and increase total field recovery by 2 – 15% (Abou-Sayed, 2012). Improvements arise as the result of two main functions: from the continuous integration of streaming data and from the ability to extract real-time and predictive analytics from the data as they stream past.

Streaming data integration and analytics: the driver for change

Streaming data integration can improve an organization's ability to transform raw data into actionable information and to improve operational efficiency. The challenge is to deliver a scalable real-time streaming integration architecture that integrates sensor data with existing siloed, offline, and manual systems, and to deliver real-time information that can be acted on with confidence across all operational divisions. Continuous integration also enables duplicated systems, databases, and data historians to be rationalized, and with fewer systems, organzations can operate with lower costs yet operate a more effective IT platform.

Detecting an event as it happens is important, but of greater importance is ability to predict in advance and take avoiding actions, such as avoiding kicks, lost circulation, stuck drill pipe, wellbore instability, failures of downhole and surface pumps and other equipment. Today's systems have the capacity to monitor operations spanning from the drill bit to the digital oilfield, midstream process control and transportation, to the retail fuel pump in real-time. However, taking the next step, to avoidance and real-time predictive analytics, requires a step change in approach and in the underlying supporting software systems.

Emerging Real-time and Big Data Technologies

The term "Big Data" is used to describe datasets whose volume, velocity, and variety are beyond the ability of traditional database and data management systems to capture, store, manage, and analyze (McKinsey, 2011). The emergence of streaming Big Data technology is focused on the challenge of managing high volume, high velocity streams of data, transforming these into actionable information, and responding in a timely, predictable, and reliable manner.

However, the emergence of Big Data technology is made possible by other factors such as Cloud computing and the availability of much cheaper server hardware and storage platforms (Bryant, 2008). Big Data is not just about volume and velocity. In fact, the volume of data created each year exceeds the world's global storage capacity. Furthermore, the rate of increase in data creation is faster than the rate at which storage capacity is expanding (IDC, 2012). The oil and natural gas industry is at the forefront of the next generation of sensor network management and the exploitation of Big Data. This raises two questions. First, how can the oil and gas industry exploit its Big Data? Secondly, if there is more data being created than available storage, how can all Big Data be exploited, when only a subset of the important information can be persisted?

The origins of Big Data and Stream Processing

The data management technologies underpinning Big Data and streaming data management are not new. As illustrated in **Fig. 1**, the initial model for data storage was the sequential data model, where data was stored as a sequence of data records with indexed access. The sequential model evolved into the hierarchical model for record databases, where complexity was managed by storing data in hierarchies, for example, IMS from IBM (McGee, 1977). Next came the major step forward in the form of the relational model, originated from IBM's System R project. This project also included the SQL language, a high level declarative model that has become the standard for data management. SQL is a specification for the problem to be solved that lets the underlying platform determine the more efficient way to compute the answers, to carry out automatic optimization, and to manage distributed processing (Chamberlain et al, 1981).

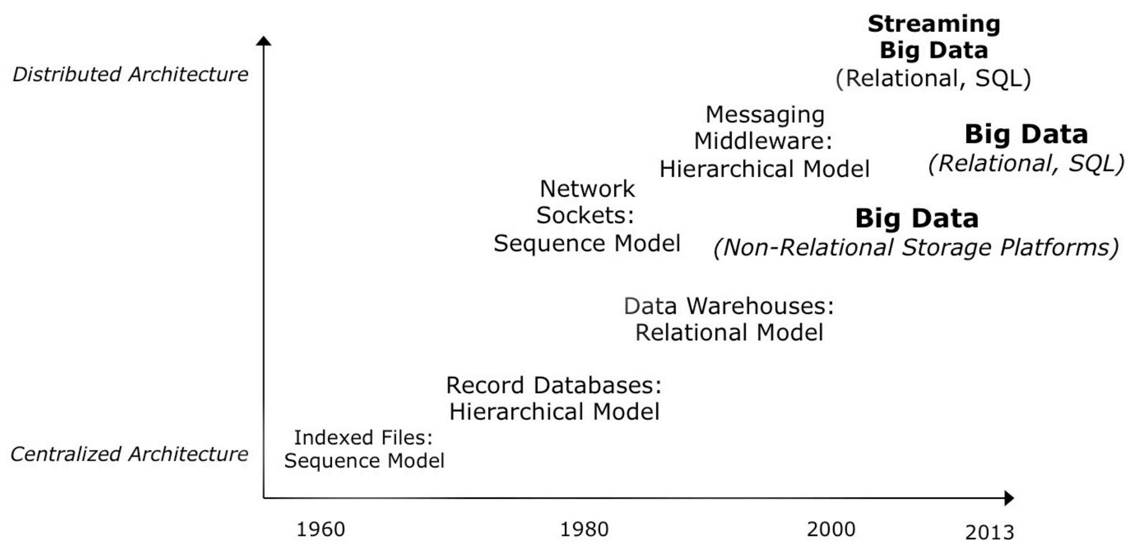

Fig. 1: The evolution of Big Data technology

So how has the history of data storage and management influenced Big Data technology? First, the original indexed file sequential model has returned in the form of name:value storage systems that underpin static Big Data storage platforms. Second, and more importantly, the SQL model is now being adopted as the de facto query language for Big Data management. The initial Big Data storage platforms were based on NoSQL technologies; however, SQL is now increasingly being added to the Big Data management portfolio (IDC 2012). The evolution of streaming Big Data has paralleled that of static data management from point to point pipes and sockets, through the hierarchical approach for managing complexity with messaging middleware, through to the adoption of SQL as a continuous query language for streaming data. The Big Data industry has now evolved to maximizing the capability of existing technologies for both streaming and static data analysis.

Big Data Metrics

As previously stated, the volume, velocity and variety of sensor and other data in the oil and natural gas industry is sufficient to cause significant business and operational issues. Typical metrics given for Big Data problems include:

- **Volume**. The volume of data created in 2009 was 800 exabyes, forecast to grow by 40% per year (McKinsey, 2011).

- **Velocity**. Industry and sensor technology are at the forefront of Big Data velocity requirements. For example, a small scale V2I (Vehicle to Infrastructure) telematics application is in the order of 20 million events per second. Large scale GPS applications are now exceeding 1 million events per second; however, telecommunications and IP-based services (Internet Protocol) monitoring applications require the capacity to process many tens of millions of events per second.
- **Variety**. It is estimated that approximately 80% of all data in an organization is unstructured. This includes emails and documents (Cisco, 2012).

Big Data in the context of oil and natural gas

It is interesting to compare the scale of data processing capacities across different industries in order to better understand where Big Data technologies come into play. For example, data warehouses for exploration and production are typically of the order of:

- **Volume**. A large offshore field may deliver 0.75 terabytes of data weekly, a large refinery generates 1 terabyte of raw data per day, and super majors store up to 2 terabytes per day (Abou-Sayed, 2012).
- **Velocity**. Data rates for a typical offshore production platform are of the order of 4,000 to 10,000 events per second (Abou-Sayed, 2012).
- **Variety**. Unstructured and semi-structured data is increasing and poses a significanty greater challenge for data processing and integration than conventional structured and sensor data.

A relational streaming data management platform

The issues with operational management platforms can be summarized as a lack of business integration across operational siloes, increasing volume and velocity of data, and a lack of real-time analysis and forecasting over the live data. The function of a streaming data management platform is to address these issues by offloading the real-time analysis and continuous integration while retaining the existing operational systems.

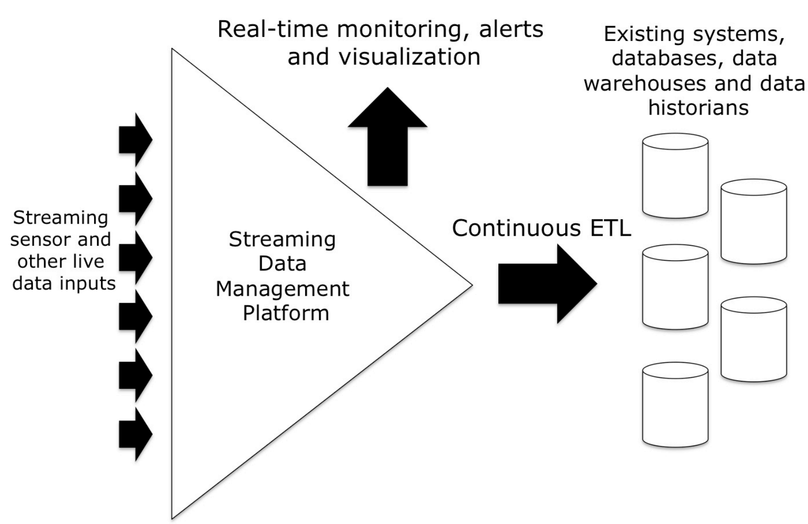

Fig. 2: Streaming data integration with Continuous ETL (Extract, Load, Transform)

Streaming data management is a paradigm for the continuous processing and transformation of real-time, dynamic data. A generic streaming data processing architecture is illustrated in **Fig. 2**. Sensors are the most common source of streaming, real-time data, but any static data source such as a log files and database can be instrumented using Change Data Capture adapters to transform new updates into real-time streams. Processed data streams are output to real-time dashboards, and also pushed to the existing database, data warehouse and data historian systems. The traditional approach of loading data into databases and data wareshouses is referred to as Extract, Load and Transform (ETL). With a streaming data platform, data are aggregated on the fly and delivered in real-time using continuous ETL operations. This eliminates the high latency, slow update overhead that is commonplace with traditional batch-based ETL solutions.

As illustrated in **Fig. 3,** a streaming data management platform enables multiple applications to be deployed on a single core platform. Applications include data cleansing, monitoring and alerting, integration and visualization. Each application has access to any or all of the arriving data streams simultaneously. The platform assembles the streams that each application has requested, effectively sharing all the streams across any number of applications. Applications process the data streams using relational Views, where each View is generated by a continuously executing SQL query. Streaming SQL queries are active queries that execute over the live streaming data while the data are still in flight, without having to store the data first. Data and information is pushed out continuously to external systems as streams of results and processed data.

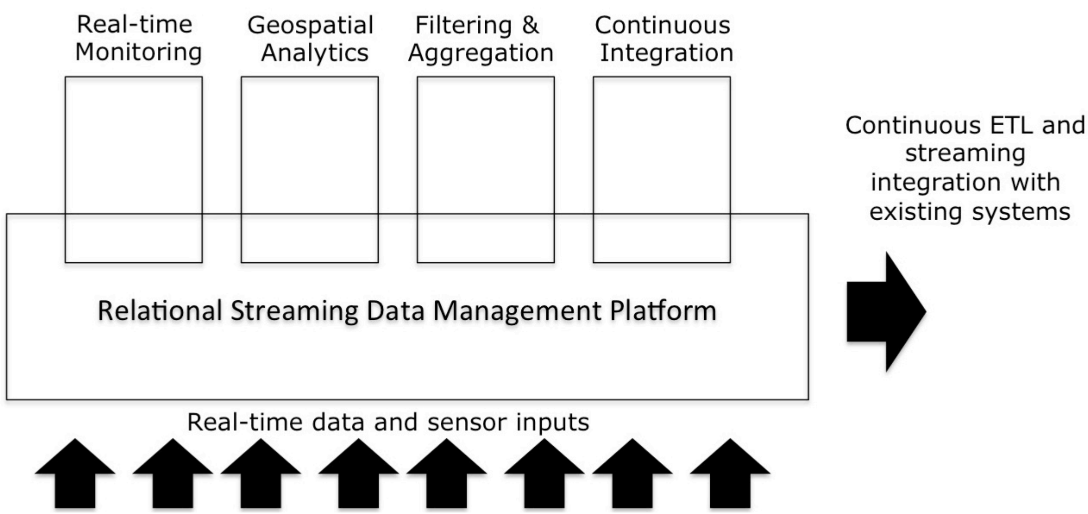

Fig. 3: Relational Data Stream Management Platform

Streaming SQL and continuous queries

The architectural paradigm for streaming data management has both similarities with and significant differences from traditional Relational Database Management Systems (RDBMS). Both RDBMS and streaming data management platforms are based on the industry standard SQL language; however, a traditional RDBMS must first store the data (data are persisted) before the data can be queried and analyzed. The primary differences between the two paradigms are described in **Table 1**.

	Relational Streaming Platform	RDBMS
Query Duration	Queries execute continuously	Queries complete and exit
Query Scope	Queries over arriving data	Ad-hoc queries over stored data
Query Federation	Stream processing distributed in-memory over many server nodes	Processing executed centrally over a single in-memory or disk-based repository

Table 1: Streaming and Traditional RDBMS SQL query comparison

In summary, in a streaming data management platform, the arriving data streams are processed in-memory before the data are stored. The same RDBMS SQL queries can be deployed as streaming queries on a streaming data management platform.

However, unlike an RDBMS query that always completes and returns a fixed data set, a streaming data SQL query is continuous and executes forever.

A Basic Streaming SQL Query Example

SQL is a declarative language that has been proven to offer a natural and powerful approach for data processing and manipulation. It also supports the rapid creation of new processing applications without having to write low level code. Streaming SQL queries are scoped over explicit time windows based on business rules. The business rules specify time windows measured from milliseconds to months, depending on the application.

Data are presented as a dynamic stream of records or messages, with each representing one row in a logically infinite stream of rows. The following simple example query is a streaming SQL query that operates upon an `ORDERS` message stream, returning only orders that originate in New York.

```
SELECT STREAM * FROM ORDERS WHERE CITY = 'NEW YORK'
```

The difference between this query and a traditional RDBMS query is that its results are delivered continuously and in real-time as soon as the input data arrive. A traditional RDBMS SQL query returns all `ORDERS` that are currently stored. The streaming SQL query executes continuously, publishing each New York message in the `ORDERS` stream immediately as it arrives.

Streaming data is transformed over time `WINDOW`s using standard relational operations such as filtering, aggregation, and correlation. For example:

```
SELECT STREAM *
  FROM Orders OVER sla
  JOIN Shipments OVER sla
  ON Orders.id = Shipments.orderid
  WHERE city = 'New York'
  WINDOW sla AS (RANGE INTERVAL '1' HOUR PRECEDING)
```

This query delivers all orders from New York that are shipped within the time window of their service-level agreement (in this example, one hour). It is only necessary to execute this query once. The query operates continuously, automatically delivering all orders that meet their service-level agreement to all interested subscribers as the Order and Shipment information arrive.

Streaming is the complement to traditional operational systems

Streaming data management is a complement to traditional RDBMS-based solutions. Both share the concept of a data model centered on processing relational rows, queries, and views. Both share common data manipulation and definition languages standardized as SQL. They are able to share a common security model and application programming interfaces, such as JDBC (Java DataBase Connectivity), and a common representation of metadata.

A streaming data management platform is based on predetermined queries executing continuously over arriving data, while an RDBMS is used for ad hoc queries over historical, stored data, processing each query until it terminates. Transactional processing is support in both an RDBMS and a streaming data management platform. In an RDBMS, transactions mark the start and end of an update operation; in a relational streaming platform, the transactions delineate the arrival and delivery of data.

The next generation of data processing architectures must utilize the strengths of different data management technologies. Examples are traditional RDBMS for data warehousing and Master Data Management (MDM), Big Data and NoSQL technologies for offline, batch-based preprocessing, and streaming data integration with in-memory analytics for the real-time, intelligent integration fabric across all operational systems.

A Streaming Blueprint for the Oil and Gas Industry

Planning an ideal architecture is not just about applying new technology. It is also important to understand how to integrate new technology with existing platforms and systems. This includes augmenting what is working already while offloading the real-time performance and streaming Big Data management bottleneck.

Requirements

The ideal architecture is a streaming data management platform that is capable of high volume, high velocity data acquisition, continuous integration of data across all existing operational siloes, with in-memory fast analytics, geospatial analysis, and predictive alerts over the data as it streams past into the offline systems. In summary:

- **Low latency**. Eliminate latency at all stages: data acquisition, real-time analysis, real-time forecasting, and predictive analytics.
- **Eliminate data siloes**. Streaming integration of data between data historians, applications, and databases enables wider visibility and better automation of key operational processes.
- **Combine real-time and historical trend information for optimum real-time decision-making**. Databases, data warehouses, and data historians contain mined data and trend information that can be kept up to date continuously and in real-time by streaming in data as they arrive. Trend information can also be streamed out and joined with the real-time arriving data in order to separate ‘business as usual’ events from real business exceptions.
- **Collect all data sources**. Augment existing information by streaming and joining data of all types, including sensors, log file updates, financial data, unstructured, and semi-structured data.

Solution Architecture for Real-time Operations

It is currently difficult to share data sources across different applications in real-time when the data is held in many different vertical business and operational siloes. Each acquires and processes its own data separately, resulting in redundant integration and data processing. This becomes a significant issue as data volume and velocity increase.

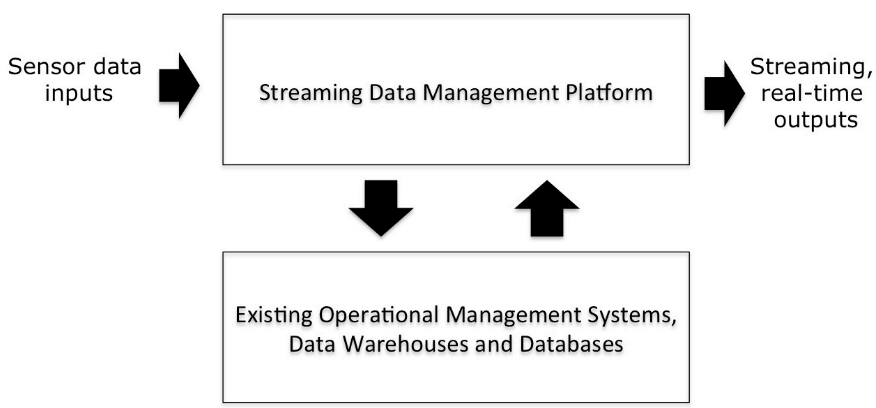

Fig. 4: Real-time Streaming Integration Infrastructure

The relational streaming data management platform (**Fig. 4**) is responsible for real-time collection, processing and sharing of all data. All data sources can be collected on a single, scalable platform in order to support an integrated data acquisition strategy. Multiple data management applications can be deployed on the platform where each processes data streams between any source and data sink.

Information and aggregated data is streamed (pushed) out continuously to existing operational systems. For example, use of continuous ETL integration with existing data warehouses and databases, providing output feeds for external systems, other operational management centers and partner organizations.

The result is a very efficient mechanism for collecting, analyzing, and distributing streaming data across an organization. New data sources and data sinks can be plugged in while the system is running, meaning that there is zero downtime for adding new data sources and applications.

Hence a streaming data management platform can be used to bridge the gaps between operational siloes, integrate all data in real-time across the organization, filter and aggregate raw data delivering only meaningful information, and with zero downtime, is suitable for Big Data mission critical applications.

Hierarchical Operational Architectures

The ideal solution architecture for real-time operations is driven by technology advances but also by other factors such as data security, network bandwidth, and autonomous local operations. This tends to drive an enterprise architecture consisting of a centralized operations center with a potentially large number of local or remote on-site operations.

As illustrated in **Fig. 5,** the role of a relational streaming data management platform is to connect all operations centers continuously and in real-time. Intelligent integration is key, along with the ability to control the flow of information from the local centers, distributing only the required data back to the central operations center.

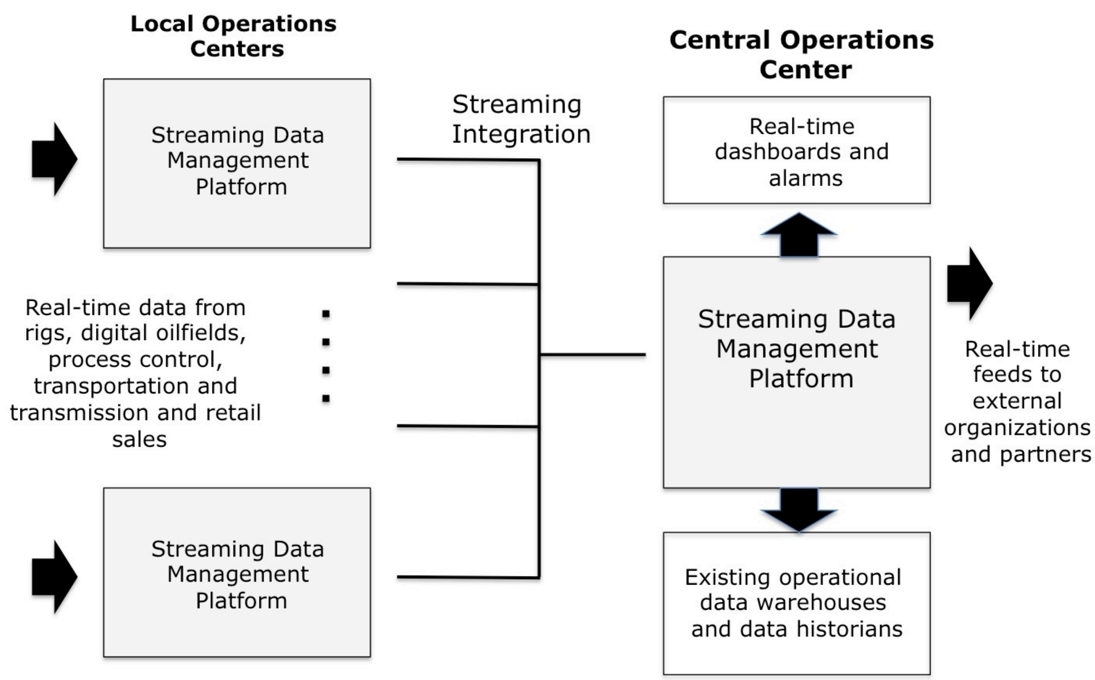

Fig. 5: Real-time Streaming Integration Infrastructure

The potential for streaming improvement

The benefits of integrated streaming data management can be illustrated through a number of specific oil and gas industry use cases. These will vary from one organization to another depending on the area of business expertise and the level of integration with existing operational platforms. For example:

- Integrate live data streams and analytics of digital oilfield data with intelligent agents, which are well-suited for large amounts of data from physically distributed sources where response times are short, especially in production control and optimization via distributed control systems (Bravo et al., 2012; Mikkelsen and Jørgensen, 2012).
- Perform multiple, independent calculations of kick tolerance simultaneously and in real-time and compare the results to drive real-time automated responses or human intervention (Santos et al., 2011; Reitsma, 2011).

- Integrate high velocity, high variety rig data in real time with WITSML-based (Wellsite Information Transfer Standard Markup Language) systems in order to predict and avoid drilling hazards and assist in rapid decision-making (Braswell; 2011, Pickering et al., 2009, 2012).
- Integrate streaming digital oilfield data in real time, in particular distributed fiber optic sensor data, in order to increase production, operational efficiency, safety, and compliance and reduce intervention. Distributed fiber optic sensor data volumes are increasing, and measured quantities are expanding from temperature, pressure, strain, vibration, and acoustics to chemical sensing, resistivity, and magnetic fields (Powers, 2012). The data are real-time; performing the processing and analytics on the live, streaming data enhances real-time results, optimization of production, and reduction of costly intervention.
- Compare data from downhole fiber optic distributed temperature sensors against modeled zonal flow in real time to determine and alter zonal allocation of flow from producing wells or into injector wells (Mehtiyev et al., 2012).

Other use cases include the ability to optimize enhanced oil recovery in real time from water flood, steam flood, and CO_2 flood (Angelo and Mershon, 2012), to monitor and control downhole equipment such as submersible pumps in real time (Medina et al., 2012), and to fully automate in real-time the curation of digital oilfield data (Chelmis, 2012).

Conclusions

The oil and natural gas industry is striving to achieve new levels of operational efficiency. Significant improvement at reasonable cost is being hindered by a number of IT and operational factors:

- Operational business units structured as vertical IT siloes, each with its own data acquisition, monitoring and business automation platforms. Siloed business operation makes it difficult and expensive to integrate information across an organization, and to deliver automated end-to-end business processes.
- Inability to deliver real-time operations as data volume and velocity increases beyond the capability of existing operational infrastructure. Data volume is increasing as a result of new sensor technology such as RFID (radio frequency identification) and wireless sensors, and technologies such as distributed fiber optic sensors that increase both the volume and velocity of data acquisition.

Real-time, streaming Big Data technology provides the infrastructure required to manage both increasing data volume and the continuous integration of acquired data with existing systems. Real-time streaming technology offers continuous integration of streaming data with the ability to extract real-time and predictive analytics from the data as they stream past.

A relational streaming data management platform is a complement to, rather than replacement for, existing operational systems. The relational streaming platform offloads the real-time performance and integration bottleneck from existing systems while maintaining the currency of those systems through continuous ETL operations. However, this will also enable the rationalization of existing IT systems and the retirement of duplicated databases, data warehouses, and data historians.

What does this mean in real terms for the oil and gas industry? Addressing the business issues faced by the industry today will improve operational efficiency, enhance real-time processes, support wider information sharing, and increased process automation. For example:

- Improved operataional safety, reliability, and compliance.
- Automating business processes to address the skills gap, for example, a lack of skills in new technologies or bridging the gap where experienced personnel have retired.
- Integrating the latest RFID and wireless sensor technology into existing operational platforms.

In summary, real-time and streaming Big Data will be the core pillars of the next generation of operational platforms. Utilizing the latest IT technology, integrating business siloes, and sharing information in real-time will be key in attaining the required new levels of operational efficiency. Ultimately, real-time streaming operations will enable the industry to meet future production targets and drive greater profitability while achieving IT and operational cost reduction.

References

Abou-Sayed, A., "Data Mining Applications in the Oil and Gas Industry, 2012, JPT Vol.64 No.10, p.88-95

Angelo, C., Mershon, K., "Creating digital Thermal EOR Projects Using Wireless Technology and New Data Integration solutions", paper SPE 153921, presented at the 2012 SPE Western Regional Meeting, Bakersfield, California, 19-23 March

Braswell, G. "Open Standard Developer Tool Streamlines Upstream Data Integration", JPT Online, 27 Dec. 2011

Bravo, C., Saputelli, L., Rivas, F., Perez, A., Nikalaou, M., Anngi, G., de Guzman, N., Mohaghegh, S., and Nunez, G., "State of the Art of Artificial; Intelligence and Predictive Analytics in the E&P Industry: A Technology Survey", paper SPE 150314, presented at the 2012 SPE Western Regional Meeting, Bakersfield, California, 19-23 March

Bryant, E., Randal, Katz, H., Randy, "Big-Data Computing: Creating revolutionary breakthroughs in commerce, science, and society", Computing Research Association, 22 Dec. 2008

Chamberlain, Donald D. et al, "A History and Evaluation of System R", Communications of the ACM, Vol. 24 No. 10 p.633-646 1981

Chelmis, C., "Semiautomatic, Semantic Assistance to Manual Curation of Data in Smart Oil fields", paper SPE 153271, presented at the 2012 SPE Western Regional Meeting, Bakersfield, California, 19-23 March

Cisco Systems Inc., "For Big Data Analytics There's No Such Thing as Too Big", March 2012

The Economist, "A Liquid Market", 14 July 2012, p.15-17

IDC, "Managing Big Data: SQL or NoSQL", Jan 2012

IDC, "The expanding didigtal universe", March 2007, updated in May 2009 and May 2010.

Kent, S., "IEA Sees Relief for Energy Customers", The Wall Street Journal (On-line), 12 Oct 2012

McKinsey Global Institure, "Big data: The next frontier for innovation, competition and productivity", Report, 2011

McGee, W. C., "The information management system IMS/VS", IBM Systems Journal, Vol.16 No.2 p.84, 1977

Medina, M., Torres, C., Sanchez, J., Boida, L., Leon, A., Jones, J., and Yicon, C., "Real-Time Downhole Monitoring of Electrical Submersible Pumps Rated to 250°C Using Fiber Optic Sensors: Case Study and Data Value in the Leismer SAGD Project", paper SPE 153984, presented at the 2012 SPE Western Regional Meeting, Bakersfield, California, 19-23 March

Mehtiyev, D., Bourgoyne, D., and Rahman, M., "Determination of Water Injection Zonal Allocation from Distributed Temperature Sensing Data", paper SPE 153809, presented at the 2012 SPE Western Regional Meeting, Bakersfield, California, 19-23 March

Mikkelsen, L., and Jørgensen, B. "Towards Intelligent Optimization of Offshore Oil and Gas Production Using Multi-agent Software Systems", paper SPE 153815, presented at the 2012 SPE Western Regional Meeting, Bakersfield, California, 19-23 March

Pickering, J., Kreijger, J., Grøvik, L., Franssens, D., Deeks, N., Doniger, A., and Schey, J., "WITSML 'Comes of Age' for the Global Drilling & Completions Industry", paper SPE 124347, presented at the 2009 SPE Annual Technical Conference and Exhibition, New Orleans, Louisiana, 4-7 October

Pickering, J., Shields, J., Gidh, Y., Johnson, D., AlShery, M., Al-Khudin, M., Farnan, M., and Schey, J., "WITSML: Laying the Foundation for Increasing Efficiency of Intelligent Sellsite Communications", paper SPE 150278, presented at the 2012 SPE Intelligent International, Utrecht, The Netherlands, 27-29 March

Powers, G., "Downhole Fiber Optics Are Changing the Well's Footprint", 2012, JPT Vol.64 No.6, p.18-21

Pritchard, D., Roye, J., Cunha. J., "Trends in Monitoring: How to Use Real-Time Data Effectively", 2012, JPT Vol.64 No.1, p48-52.

Reitsma, D., “Development of an Automated System for the Rapid Detection of Drilling Anomalies using Standpipe and Discharge Pressure”, paper SPE/IADC 140255, presented at the 2011 SPE/IADC Drilling Conference and Exhibition, Amsterdam, The Netherlands, 1-3 March

Santos, H., Catak, E., and Valluri, S., “Kick Tolerance Misconceptions and Consequences to Well Design”, paper SPE/IADC 140113, presented at the 2011 SPE/IADC Drilling Conference and Exhibition, Amsterdam, The Netherlands, 1-3 March