Return to session menu

Dynamic Positioning Committee
Marine Technology Society

DYNAMIC POSITIONING CONFERENCE
October 9-10, 2012

SENSORS 1 SESSION

Information Aggregation on a Mobile Offshore Drilling Unit

By Trenton Martin, Donn Nguyen and Farooq Iqbal

Transocean

Information Aggregation on a Mobile Offshore Drilling Unit

Abstract

This paper is a consolidation of relevant points of interest and lessons learned when implementing information aggregation architectures on a mobile offshore drilling unit. Transocean is a leading international provider of offshore contract drilling services for oil and gas wells. The company specializes in technically demanding sectors of the global offshore drilling business with a particular focus on deepwater and harsh environment drilling services, and it operates one of the most versatile offshore drilling fleets in the world. Transocean's fleet has developed over many years through acquisitions and mergers. As a result this heterogeneity is reflected in the diversity of the controls systems being used onboard. In this paper we will review some of the challenges and propose best practices when tackling similar tasks.

The purpose of information aggregation within the context of this paper is to establish connectivity to all the vessel control systems so as to consolidate data and alarms, provide remote support capabilities as well as manage control system software. Leveraging mature, 'open' standards where applicable such as OLE for Process Control (OPC) is crucial to providing an extensible architecture that can leverage the latest tools without being tied into a single manufacturer's development roadmap. Control system aggregation and remote support capabilities start to blur the line between Control Systems Engineering and the Information Technology. Leveraging standards, skill sets and components from both of these specialties have provided an optimal architecture and support structure for this system.

Introduction

Transocean operates drilling rigs ranging from shallow water barges to high-specification deepwater drill ships. This paper is a consolidation of relevant points of interest and suggestions for implementing equivalent architectures on a mobile offshore drilling unit (MODU).

Our fleet has developed over many years through organic growth, acquisitions and mergers. This heterogeneity is reflected in the diversity of the controls systems originally installed and being used onboard. In the writing below, we will review some of the challenges and establish some best practices when tackling similar tasks.

Why Aggregation

There are at a minimum four primary control systems on an offshore drilling rig and these will be detailed later in the paper. Each control system has traditionally limited the responsibility of data logging and alarm management servers to their respective and specific system. Within the context of the drilling rig, our approach to information aggregation intends to consolidate data, alarms and other support capabilities to a single location onboard the rig. This would at a minimum consist of a server with some specialized software and network appliances, which we will discuss further.

The next level in aggregation is to digitally assemble data from multiple rigs back to shore-based consolidation at one or more locations. This would be considered enterprise aggregation as opposed to facility aggregation described in the paragraph above. This next step provides a unique challenge as the types of controls systems can vary significantly between rigs.

This consolidation of remote monitoring and support has some intuitively perceivable advantages. Once this data is in a single location, centralized or regional support and monitoring would be more effective; standardization of processes and tools can also be leveraged. In today's world where near real-time information is expected to be immediately accessible, these same expectations hold true for data from offshore drilling and production facilities. Our internal and external customers demand near real-time data and the ability to remotely support operations on-demand. Some of our control systems already support the ability to stream data back to shore and provide remote support of its equipment. Being able to provide this capability to all the control systems onboard all of our rigs, is the intent of this type of system.

Philosophy

At the onset of any similar project, it is recommended to first establish a philosophy. This could be in the form of a mission statement or a high-level presentation detailing the approach that is to be used. Our philosophy in fact evolved over time. The initial concept was abstract to a certain extent but as time went on and we acquired more experience and an understanding of what is required and expected, this became a tangible philosophy that we would refer to for guidance. For this application, our philosophy was more of an illustration than a mission statement for the project.

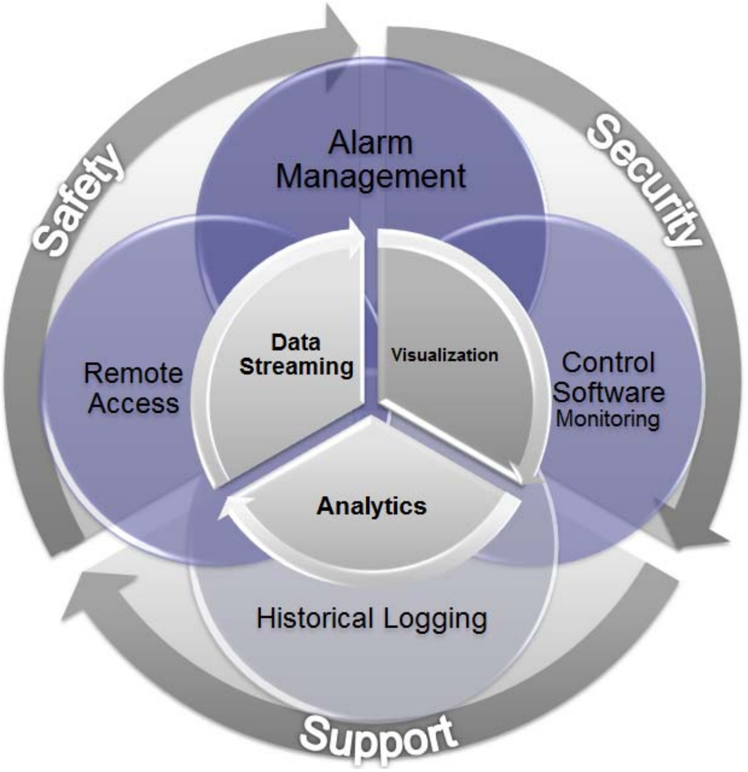

Figure 1: Areas of attention and concern for Control System Information Aggregation.

To clarify what this initiative entails, we further define some of these key philosophical and cultural components below:

Safety: Enabling subject matter experts with secure, password protected, virus-free, remote access to a specific control system, places that individual virtually onboard the vessel with the ability to troubleshoot and diagnose problems. As such, all applicable safety policies within the context of the task at hand, would be adhered to as if that individual were physically onboard, interacting directly with that control system. So any task specific safety pro forma or permits to work would need to be agreed to and signed off prior to commencement of the work. As usual, communications both by email and phone are crucial for these instances.

Security: Many of the systems that are being connected to were conceived and designed over a decade ago, when digital security for control systems was not one of the foremost concepts in mind. In addition digital security challenges are an ever-moving target. We will go into more detail about security issues in the context of these types of applications later on in this paper.

Support: Once a system is installed onboard the rig, it must have a support mechanism. There are various support models that can be leveraged for these types of systems. The approach to supporting these systems will be driven by as much by technology as by commercial issues. We would emphasize that support is paramount to the successful deployment of these types of systems

On every rig, there are typically three or four primary control systems that are critical to the rig's operation. When comparing Dynamically Positioned (DP) drill ships, semisubmersibles, jack-ups or moored vessels there will be variations in the systems. This variation can be due to differences in complexity or functional requirements. We categorize these systems in Table 1 below.

Table 1: Categorizing the 4 primary MODU control system types.

	DP Rig	Moored Rig	Jackup
Vessel and Power management	VMS/PMS	VMS/PMS	VMS/PMS
Station Keeping	DP System	Mooring Control	Jacking System
Drilling Control /Instrumentation	Drilling	Drilling	Drilling
Well Control System	MUX BOP control system	BOP multiplex control system / BOP discrete hydraulic control system	Surface BOP control system

VMS – Vessel Management System
PMS – Power Management System
MUX – Multiplexed Control System
BOP – Blow-Out Preventer

Next we established which fundamental roles our data aggregator will fulfill for the control systems.

Historical Logging: Data logging and aggregation would be one of the more common applications once the connectivity is established between the various control systems. Being able to bring together all the information from drilling, DP/VMS and subsea to a consolidated location is the first step. Once the data is aggregated, the next significant task is to normalize the tags so there is a common naming convention, common time base and standard set of engineering units. Then the data can be decimated and disseminated back to the shore-based enterprise historian.

Remote Access: One of the more rewarding benefits of such a system is remote access. This would predominantly be used when supporting a rig with expertise that is available onshore. Being able to provide remote access to drilling rigs tends to provoke positive and negative points of view. The reality is typically somewhere in between. Security is of paramount importance as soon as any degree of remote access to a vessel's control system is enabled.

Control Software Monitoring: Most, if not all, drilling companies have some form of a Management Of Change (MOC) policy or guideline. In addition there is most likely an upset recovery plan as well, so that re-installation of control system software can be accomplished onboard. Fortunately, there are tools in the control industry that aid in actively managing both of these areas.

Alarm Management: The automated reporting of Alarms and Events from the control systems is typically part of the day-to-day user interaction with these systems. Systems whose alarms have not been correctly rationalized tend to pose problems to the end users. Issues such as chattering alarms, incorrect categorization of alarms and events, and incorrect prioritization are examples of these problems. When control systems exist in isolation from one another, but still have interdependencies (e.g. power management and the drill floor), correlating events related to one or more systems are not immediately apparent. Including alarm management in our information aggregator was part of the system's design from its inception.

Having these four roles in place still may not justify deploying an aggregation system of this scale. When we start to leverage the information being gathered, there are certain minimum capabilities in this day and age that are expected.

Streaming: The current expectations of any information systems is the ability to be able to supply 'on-demand' or 'streaming' information to the users. One of the main limiting factors here is available bandwidth. Because of technology, cost or geographic location, the available bandwidth may vary significantly. As such, the ability to stream data must be scalable to accommodate a range of bandwidths.

Visualization: A clear deliverable from an information aggregator is typically found in the presentation layer. Here the real time and historical information would be presented to both rig and shore-based end users. How the data is presented is sometimes just as critical as what data is being presented. This effort is highly dependent on the Historical Logging components mentioned previously.

Analytics: Once we have the data and can display it, one of the final pieces of this puzzle is bringing value to what already exists. There is fantastic potential in this area once the data is co-located in a single spot. This could be where the return on investments realized in the long term. Analytics can be applied to the historical data just as easily as the alarm and events data.

Having these building blocks in place helps give us a defined path to move forward. This project did not happen in a single iteration, but as a continuous engineering development that was re-factored for improvement as both technology and software matured and changed.

Challenges

As with any engineering effort, challenges continued to present themselves. For the sake of brevity we will discuss some of the more interesting ones below.

Interdepartmental cooperation

In order to leverage information captured from a control system back in the office or provide remote support from shore based personnel, some secure connectivity must be established from the corporate intranet. Control Systems by design prioritize system availability above data confidentiality, whereas traditional IT systems prioritize confidentiality over system availability. The information aggregator is the system that needs to be in-between this dichotomy.

Leveraging best practices and standards from our IT group has provided a degree of scalability that was not originally envisioned. Leveraging the pre-existing corporate authentication and authorization as well as networking policies allows for granular access control to these remote rig based servers. By partnering with the IT department in this effort, we were able to leverage a common hardware platform, existing server images, existing patching regiments and existing server tools sets. Also, one of the most critical items that we could leverage to help support the installation and day-to-day support for the servers themselves was our own IT staff. A significant amount of cross-training occurred between Engineering and IT; both sides developed a better understanding of the technologies, processes and capabilities in this collaboration.

Human Factors

When a remote monitoring system is introduced in a facility, it may be met with some degree of skepticism and suspicion, if the motives for deployment are not understood. Once the system is installed, it must be supported both remotely and on the rig. From our experiences, we found noteworthy points in human factors that drove the success of this system:

Championing of the product by a team within the rig operations group was a critical step. This product must have a team who can support and promote its benefits to the operations departments across the organization. These types of projects are not trivial in terms of the initial cost of the system. A cost/benefit analysis is merited and once the commitment is made, support and direction from the executive group is critical to a successful implementation. Useful tools to aid in this process are use cases that must be developed from the perspective of the various users as well as those on board the rig. If an application which is directly coupled to the aggregation system has great benefit onboard the rig and is available for immediate use by the crew, the systems acceptance will most likely be accelerated.
From our experience, for this type of system deployment there is a reasonable amount of ‘internal marketing’ that needs to be done in advance.

Installation Challenges

The most significant component of the cross-disciplinary engineering effort lies with the system’s preparation and deployment. We found that the field service and the engineering teams significantly benefit in having the following skill sets:

- Diverse knowledge of interconnectivity with the control systems. This includes the control network, field bus protocols, HMI, and microcontrollers (such as PLCs).
- Strong software troubleshooting skills related to various main stream operating systems and data stores used in the process control industries.
- Comprehensive knowledge of networking to include managed switches, firewalls, VLANs, TCP and UDP traffic. Traditional copper Ethernet connectivity as well as Fiber Optic.
- One of the most important deliverables at the completion of the installation is a documentation package that includes updated topologies representing the physical and logical connectivity.

Those planning or executing the installation and configuration of the historian, alarm management and software monitoring application must be expert users with those software components. As the majority of the effort with this type of system may be spent collecting needed data and configuring the software, the following experience would be crucial:

- Management and mapping of incoming tags from the control systems to a normalized set that will be consumed onshore. Drillships can contain well beyond 50k tags in the combined control systems.
- Historian installation and configuration as well as in-depth knowledge of the suite used to capture the real-time data from the control systems (OPC-DA/UA or other software connectivity).
- Most control system software management packages have a fairly intuitive interface with similar concepts to mainstream software revision management systems (check in/out, lock, compare). The time spent in this area will predominantly be on data collection. Once again, excellent organization and documentation skills are needed here.
- The alarm management software may not collect as much data as the historian however its connectivity to control systems tends to be more challenging. With the exception of OPC A&E, which is rarely seen, there is typically no standard method or format that alarms and events are provided in for external use.

Once the aggregator is deployed, plans must be in place to repeatedly upgrade the system with the security appliances and tools available. It is inevitable at some future date that some utility or component will be upgraded or replaced with a better widget. Patching for the aggregators operating system will be part of the regularly scheduled tasks.

A survey of the rig prior to ordering software and equipment is critical to the systems success. The installation surveyors must be aware of the different controls systems and methods that will be used to connect. To ensure the correct hardware and software to connect is captured; the server installation should be aligned with established industry best practices. These standards may be driven by the engineering team or IT department (or ideally both) within the organization.

Architecture, Design and Topology

Because of the variety of control systems we have integrated on our rigs, an agnostic approach to aggregation was established. We then looked at handling each connection in its own unique fashion, abstracting the overall architecture from the connection.

Consider the physical topology where the core strength to implementing a similar system is leveraging the existing Ethernet networks of the controls systems. There are exceptions to the rule, as some control systems do not support Ethernet and may require that data be taken from the field bus networks. For the initial design we focused exclusively on Ethernet-based data transport mechanisms. Figure 2 below, indicates the logical connections from the main control systems to a local aggregator on the rig. Although the firewall is depicted as solely a connection interface, it is in fact applying the firewall functionality between the control system and the aggregator as well.

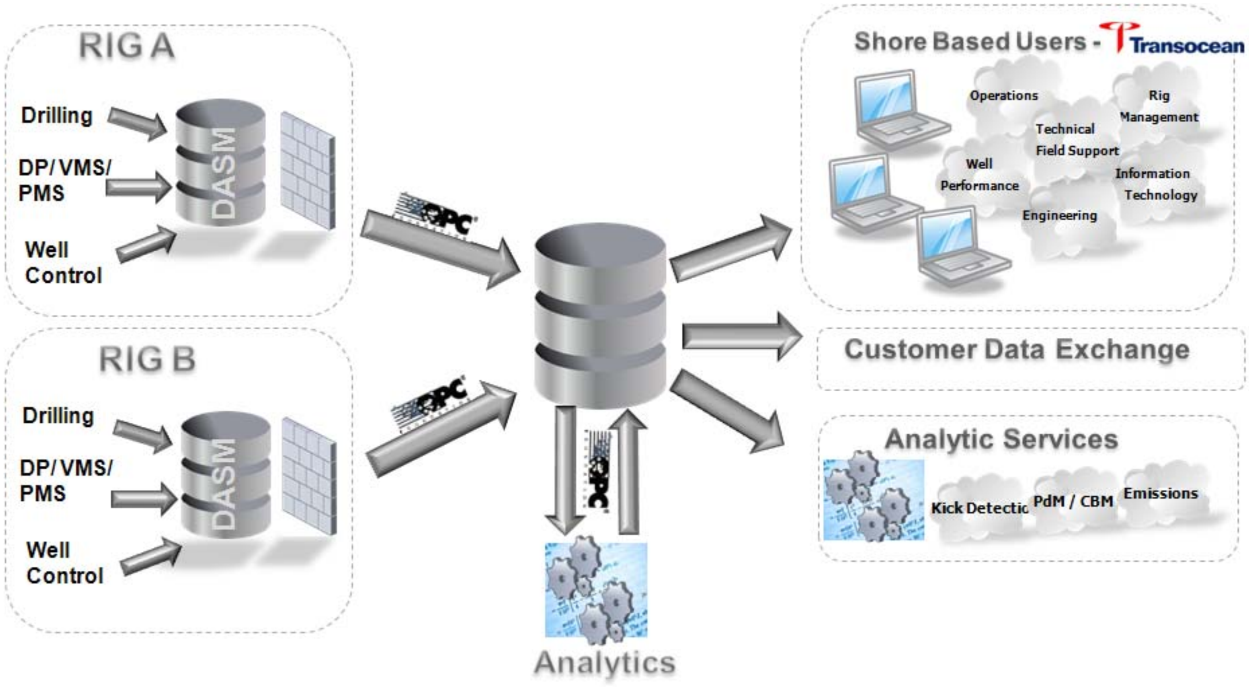

Figure 2: Generalized view of Rig to Shore aggregation framework.

OLE for Process Control

Example OPC architecture is shown below in Figure 3. There are two separate interfaces of the OPC standards being utilized here.

OPC DA - OLE for Process Control (OPC) Data Access (DA). This can be considered as the streaming data coming from the control systems. Each data channel, also called a *tag* has quality and time information associated with the updated value. In addition metadata such as *descriptions*, *units* and customizable entry can be associated with a specific tag.

OPC HDA -OPC Historical Data Access. This DCOM (Distributed Component Object Model) interface provides the ability to query data sets from the past. This functionality will require a data store in the background from which the query will pull the data. Many historians can provide this interface and an OPC-HDA server. OPC-compliant client tools will leverage this interface. For example, if the user wanted to plot a trend of what the Top Drive RPM was doing for the past hour, the trending tool would use its OPC HDA client and connect to the historian's OPC HDA Server interface. There is yet another use of the OPC HDA interface that is leveraged in the distributed architecture offshore. Consider a period of satellite communications outage. If the shore-based enterprise historian just used OPC DA, it would have a gap in its logged data that could not be recovered. However, if the shore-based historian was leveraging OPC had, it could have the facility to recover the missed data by querying past data sets. During our project, we selected historians and utilities that would provide this ability to recover from communication outages. It does, however, add a layer of complexity as illustrated in Figure 2 below.

To be able to benefit from the logged data onshore from multiple rigs across various controls system with a common set of tools, we recognized that the systems tag lists would need to be normalized both in name and engineering units. Mapping the vendor's control system tag name to a generic form had to be done. In addition, the data from the vendor needed to be converted to a standard set of units, in our case SI.

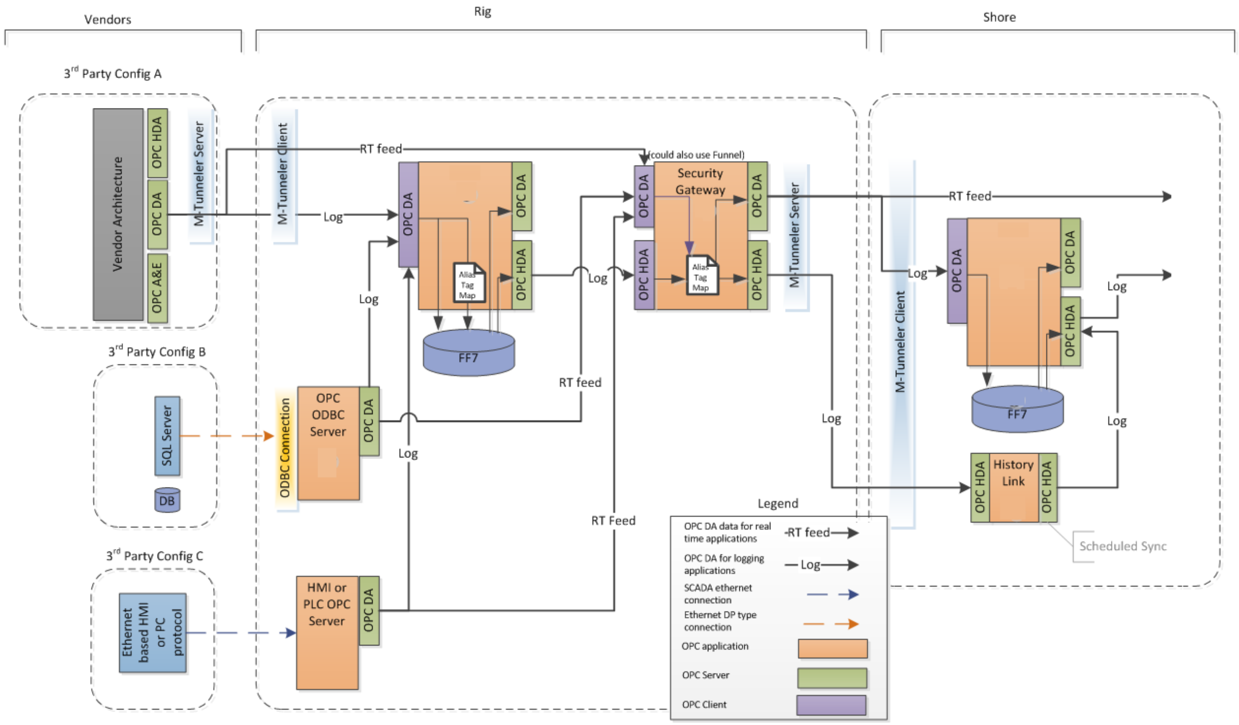

Figure 3: Example of an OPC DA and HDA synchronization architecture with tunneling and aliasing.

There are a few other OPC protocols that deserve mention. They are not used as extensively, but we do encounter them in certain cases.

OPC A&E: This is OPC for Alarms and Events. Alarms are derived from triggers that are actions by measured values exceeding a defined limit, and are by nature discrete. Events are to a certain extent expected. Data from this type of OPC source would not necessarily be captured by an OPC HDA historian. OPC A&E would be handled by the data store that is specifically design to capture Alarms and Event management data.

OPC Security: We would like to bring the reader's attention to this protocol only to provide awareness about avoiding it. In order to benefit from it both client and server applications must support it. There are very few if any mainstream client applications that support OPC Security (1.0).

OPC UA: Although this protocol uses the OPC (OLE for Process Control) acronym, it is in fact not at all dependent on the Microsoft® DCOM (Distributed Component Object Model) protocol. The UA stands for "unified architecture". It encompasses functionality provided by DA, HDA and A&E. The specification for OPC UA has been available since 2006 and is slowly gaining momentum, however the existing OPC DA and HDA still appear to be the dominant protocols.

As mentioned above, when using OPC DA, HDA or A&E, some additional challenges are introduced with Microsoft® DCOM. Challenges with network Firewalls as wells as Microsoft® Windows Security become immediately apparent when trying to share data between two PCs from separate control systems. As a result, many solutions to tunnel the DCOM communications between PCs are available that support all three OPC protocols.

The following graph is a sample breakdown for average traffic that was generated on a specific installation. The majority of the data is related to the tunneled and compressed OPC-DA and HDA data being transmitted from offshore to the onshore enterprise historian.

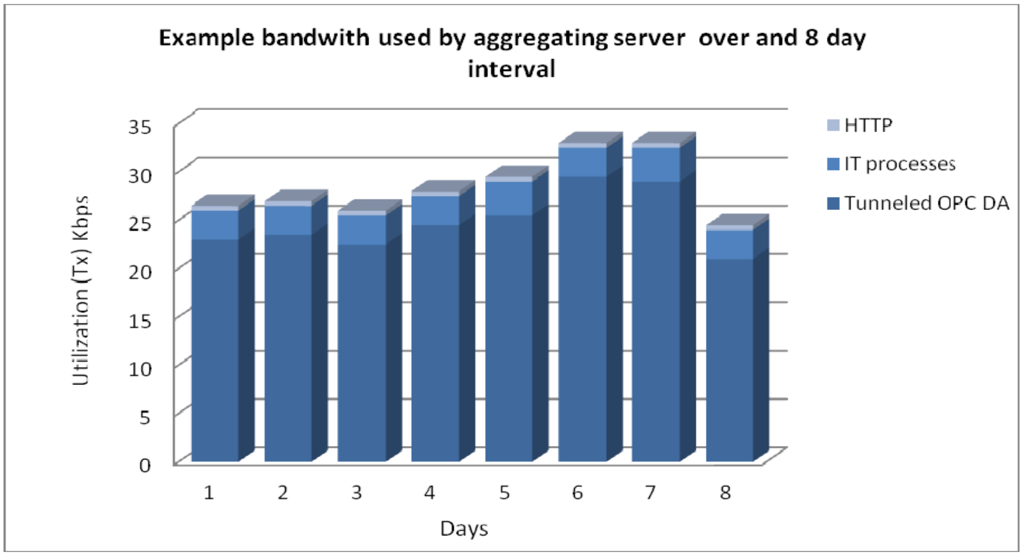

Figure 3: Example bandwidth utilization (transmitting to shore) for rig based aggregating server.

Alarm Management and Aggregation

In terms of data collection, alarm management follows an almost identical topology to Data Aggregation. Alarm and Event collection differs sometimes not only in format but also in the mechanism. Typically we will find this information being captured in at least one of the following mechanisms for a particular control system

OPC A&E: this is not necessarily one of the more common methods; however it is definitely preferred from our perspective as it simplified our interfacing and parsing tasks.

Alarms written to a file: This is one of the more simple mechanism alarms; and event history is stored in a system.

Output to a printer: This is typically a difficult mechanism to access if it is a physical printer but there are access methods using virtual printer ports.

Dump to an ODBC compliant data base: This provides a reasonable mechanism to access the appropriate data with the given permissions.

Alarm management and aggregation tools need to support the above protocols in order to access these sources. Just as with Data aggregation, Alarms and Events need to be normalized into a common format. A few of the mainstream Alarm and Events capture frameworks leverage a SQL database as its storage mechanism. Another capability these Alarm management tools must provide is a configurable parsing tool or method. During the parsing process, it will be important to identify and differentiate between what is considered an alarm and what is considered an event.

Most of the mainstream alarm management packages for process control systems have built-in visualization and analytical tools that can be leveraged ‘out of the box.’ The benefit of these tools and how they can help isolate problem areas for a control system become apparent once the system starts to accumulate alarm and events in the database.

Remote Connections

Remote connections through a robust and secure connection would provide the capability to remotely view designated screens. Some obvious uses would be for remote support, troubleshooting and collaboration. Most modern operating systems used for HMI/MMI applications have some sort of built-in or installable remote viewing capability or utility. Microsoft® Windows-based operating systems have built-in support for what is called Remote Desktop Protocol. Other operating systems may have support for what is known as Virtual Network Computing (VNC) protocol. These software technologies allow a user on one PC to log in and view remotely applications on another computer across a network. This type of remote connection can be configured to require some authentication before allowing the connection to be established. There are also some 3rd party remote access tools that share the same desktop context, which is being displayed on the local console. This way, when the remote user is operating on the remote terminal, all parties can see the operations being performed.

It is paramount that the security layers are well established and configured on the aggregating proxy server. So for any remote desktop type connections to a control system, there are two remote sessions that are established. The first is from shore to the proxy server. The second is from the server to the target control system in question.

Configuration Management Aggregation

The fourth component and arguably the most critical is configuration management and monitoring. These very specialized software suites have some specific tasks. The purpose of this package is tri-fold.

Management of Change: Putting policies and procedures in place to establish a particular company’s directive in regards to change management is an effective first step. An automated and auditable software package that provides a check and balance to this effort would help further these policies. These software packages can automate the audit process programmatically, which can be cost effective when increasing control system software audit frequency. These tools have automated notification via email and summarize the current state of control system software on the rig.

Upset Recovery: In the event of a PC or microcontroller failing, even if the spare parts to replace the component are on hand, the logic/code/application must be available to load back on to the device. Having an archive of the entire mission-critical control software available in a consolidated location has its obvious benefits.

Version Management: Most of these types of Change Management (CM) applications have some roots in the traditional software version control realm. These CM aggregations will have the traditional 'check-in' / 'check-out' concept as well as an ability to lock a tree node and roll back to previous working configurations. .

In our experience, this is an excellent tool that we have introduced both in mid-life and during the vessel's construction, when the controls systems are first being installed and integrated.

Network Security

The design of the rig's network infrastructure to support data aggregation is a key factor for the success of the system. A main design goal was to connect to each control network in a manner that would require no changes to the network devices yet still extract the information needed. Most devices on the control networks do not have default gateways configured and therefore cannot talk to any device that is not directly connected to the control network. This clearly presents a challenge for aggregator system that must reside at a central location and collect data from each of the control networks. Each control system may have 2, 3 or more networks. There is also a need to keep these networks' traffic from intermingling. The aggregator network access must be designed with that purpose in mind, also, which is the next challenge to be met. Some requirements of design goals that should be considered are summarized in the list below:

- Allow the aggregator to access any control system subnets. This may seem apparent, however scalability is key.
- Prohibit any control system subnets to access the aggregator or other control system subnets
- The solution must allow the aggregator to reach different control system subnets even if they are using overlapping IP address spaces. This occurs more frequently than one would expect, primarily on some of the older drilling rigs.
- Allow restricted access to aggregator data by client and control systems
- Include the ability to monitor and report network configuration changes in the aggregation network and possibly the control system network as well.

It is a modest challenge to meet these goals and to keep things relatively standard and scalable in the network design while keeping complexity low. The first step is to ensure the physical and logical separation between the three zones of activities. The first zone is the corporate zone, such as an intranet. The second zone would be for the information aggregator and the third zone is for the control systems.

The corporate zone has its own set of switches with virtual LANs (virtual local area networks, VLAN) and so does the aggregator zone. The control system zones have their own switches but do not typically leverage VLANs for its purposes. Instead, each control system network may be using the default VLAN preconfigured in the switch. In most cases, switches allow devices in the same VLAN to freely communicate with any member in that VLAN on layer 2 of the Open System Interconnection (OSI) model. Some mainstream switches have a "protected port" feature that prevents devices from communicating to each other at layer 2, even if they are on the same VLAN.

The "protected port" feature forces the devices to route traffic via layer 3 devices such as a firewall or a router. To keep the zones logically separated, we need a demarcation device such as a router or a firewall. Here the term firewall is used to describe both these devices for the sake of brevity. Access restrictions can be implemented simply in either direction for any IP address and/or port combination to/from any of the zones using access rules. Slight complexity is added when network address translation is introduced. Network address translation makes the aggregator traffic appear to be coming from the directly connected VLAN/subnet of any control system network. Control system networks respond to the information aggregator just as if it were another device on their directly connected VLAN/subnet. If unable to change one of those subnets, one option available is called virtualization. On a router, it is called virtual routing and forwarding, and on a firewall it is referred to as virtual firewalls or contexts.

Virtualization allows each 3rd party subnet to be treated as a virtual router/firewall and therefore overlapping IP addressing space becomes a non-issue. This is a useful feature to be considered during hardware selection.

Summary and Conclusions

Consolidating all of the control systems for the purposes of aggregating real-time data, and alarms, managing software and remote support for just a single rig is a significant engineering effort and accomplishment. Designing the system for scalability to a fleet of drilling rigs increases the degree of difficulty significantly. Leveraging open standard protocols, solid engineering design, interdepartmental synergies, cross-disciplinary skills and a bit of salesmanship are tools a team would benefit from when deploying a similar information aggregation system on mobile offshore drilling units.