TECHNOLOGIES THAT ARE SHAPING THE OIL & GAS INDUSTRY

2024 Appalachian Gas Measurement Short Course (AGMSC)
Presented by Steve Hill
(originally authored by Julian Sanchez)
Emerson

TECHNOLOGIES THAT ARE SHAPING THE OIL & GAS INDUSTRY

The Oil & Gas industry is not known for its tendencies to embrace new technologies quickly. On the contrary, we continuously see multiple reports and world-wide industry analysis coming out of companies like McKinsey or Forrester that show how Oil & Gas is consistently the last industry vertical when it comes to "digital maturity". In a way this tendency can be understood since Oil & Gas is an industry very much focused on physical goods, and improvements in equipment, processes, and people skills typically yields a greater return than the adoption of new digital technology by itself. Unfortunately, this tendency to leave digital technologies for last has become a collective "badge of honor" rather than a recognized area of improvement. The default mode of operation has always been to not keep an eye on new technologies unless competitors implement that technology first. As a result, nobody ends up adopting new technology for long periods of time. This mindset may have worked for the last few decades, but we need to be prepared for an abrupt awakening.

Let's take a quick look at the pace in which new technology has been adopted throughout history. A good indicator to reflect the speed of adoption is to measure the time it takes for a new technology to reach 50 million users. It took 62 years for the automobile to reach 50 million people from the time of its invention in the 19th century. It took 50 years for the telephone to reach 50 million adopters while mobile phones only took 12 years to reach the same amount of people. The internet took only 7 years to reach 50 million users. Facebook took 4 years. The game "Pokemon Go" only needed 19 days to reach 50 million users.

Granted, these examples use commodity technologies and their adoption by the general public rather than having an industrial or enterprise lens. Nonetheless, it's a good representation of how new technologies are not only being released to the world quicker than ever before, but also how the pace of adoption of these new technologies is much quicker than before, and that pace is accelerating as time goes by. Even if the mindset within the Oil & Gas industry continues to be the slow adoption of new technologies, the definition of "slow" will certainly shift to something far faster than what we've seen before. The new slow will be very fast. Are we ready for it?

In this paper we are going to explore a few technologies that have begun reshaping the Oil & Gas industry, as well as peek at a few others that may have a greater impact in the near future. These are technologies that cannot be ignored any longer, and people within Oil & Gas must become versed with their purpose, applications, and business benefits.

THE CLOUD

It's very hard to classify the cloud as a new technology given that it's been around since 2007. Be that as it may, Oil & Gas is now just getting started with leveraging it. Up until now the bulk of the usage of the cloud has been limited to corporate applications (email, file sharing, virtual collaboration, etc.) but very soon, as we will see in the next section of this paper, this usage will extend to how we remotely monitor, control, and optimize our physical assets.

There seems to be a fair amount of confusion around what exactly is the cloud. Cloud is so ubiquitous in our personal lives, and cloud providers offer so many advanced features, that is difficult to discern how the cloud is different from previous technologies, and how the cloud brings tangible business value to the table.

The most common, and understandably unavoidable, comparison that many people make with the cloud is with server hosting companies (also known as server co-location or server co-hosting). Oftentimes I hear people say that the cloud is simply "a server hosted somewhere else". While there is some truth to that statement, it is vastly incomplete when it comes to what cloud technologies have to offer and how they work. Indeed, when you host something in the cloud it is inevitably hosted on a server that is managed by the cloud provider. Unlike traditional server hosting companies, however, the server that is hosting your application is neither dedicated to you nor it's explicitly assigned to you. The cloud involves multiple data centers geographically distributed around the world, and typically those data centers are linked by dedicated high-speed networks that operate separately from the internet, thus ensuring very high data transfer rates in between data centers. Each data center, as it could be expected, is swamped with server racks and network equipment. When you request the hosting of digital assets in the cloud (files, an application, etc.) the cloud provider will automatically and seamlessly allocate the necessary server and networking resources to host your digital assets. Back in the days of server co-location, the hosting company would "reserve" specific server and networking resources from their pool for your

own use. If those resources failed (a broken hard drive, for example) your digital assets would not be available until those physical resources were restored into working order again. In many cases it was up to you to detect the failure of those resources and alert the hosting company so they could replace the necessary parts. Hopefully you had backups! (that you were managing yourself). Also, if those resources were in a geographical region far away from where you are located, you would experience considerable slowdowns when accessing your digital assets.

There are subtle but very important differences when the same scenario plays out when using a cloud provider. Your digital assets are not allocated to a specific physical server or geographical location. Hosting is now fully virtualized. Your digital assets may be stored in one server in the morning and moved to a different server in the afternoon. Maybe even a server on a different data center! But you wouldn't notice those things, only that your digital assets are always available when you need them. In addition to that, if you needed more resources to host more digital assets, all it would take is a few mouse clicks in order to make sure the cloud provider makes those resources available to you. Essentially, every resource made available to you by a cloud provider is *virtual*. The cloud provider is responsible to make sure the necessary physical resources are available so that your digital assets are available through corresponding virtual resources. And since all resources are virtual, you can expand/reduce the amount of resources you need at will while the cloud provider allocates the necessary physical resources from wherever they find most appropriate. It's the reason oftentimes people refer to the cloud as elastic, as there are no strict commitments to hardware or network resources.

That's how the cloud operates at its most fundamental level. What cloud providers have been able to do is build numerous services on top of that elastic functionality to empower customers with focusing on solving their specific business problem without having to worry about how the underlying technology is being managed. These services span data storage, application hosting, machine learning, robotic process automation, virtual networking, authentication and authorization, etc. Essentially what the cloud allows companies to do is to *treat IT as if it were a utility*. You simply pay for the services that you use whenever you need to use them. This is a huge game changer when compared to traditional approaches where every company was responsible for creating and maintaining their own data centers, disaster recovery sites, and expensive technologies needed to manage it all.

IOT

The Internet of Things has been a hot topic for the past 2-3 years, and we are already seeing multiple Oil & Gas companies actively seeking ways to adopt this technology given how promising it is. As the name suggests, IoT consists of connecting sensors and equipment securely to the internet. To date, IoT has made more advances on the consumer space but the industrial arena is quickly catching up. Oftentimes people will try to differentiate the consumer-related IoT with the industrial-level IoT by referring to it as IIoT (Industrial Internet of Things). Fundamentally they both use the same technologies and are going after the same goals. I personally refer to IoT for absolutely everything regardless of it coming from the consumer or industrial space.

In Oil & Gas, having sensors and equipment connected directly to the internet is the complete opposite to the widely accepted approach of isolating the entire control network from the rest of the world. You may know this approach as the "Purdue Model" or the "ISA 99" architecture. The reasons and motivations that called for that strict isolation have not been eliminated. Most industrial equipment relies on protocols that are intrinsically insecure. In most cases equipment and their protocols don't even support the concept of authentication or data encryption, therefore restricting who can connect to those devices and what network is used to transfer data back and forth becomes the most critical elements for keeping them secure.

Before we cover how the level of security is maintained while enabling direct internet connectivity, it's important to get out of the way a couple of important concepts around IoT. One of those concepts is what is known as an "edge device". In the broadest terms, edge computing is simply any type of computing that takes place as close as possible to the physical location where the data originates. It is very easy to see that, for example, a PLC or an RTU qualify as edge devices as you can have custom logic running inside them that will depend on the data captured by the device (and those devices are typically installed in close proximity to the sensors and equipment they are controlling). An edge device by itself, within this definition just given, is not necessarily an IoT device. As per the reasons described above, having these types of devices connected to the internet would not provide the sufficient level of security. Also, as it can be expected and as it is happening today, and edge device doesn't have to be connected to the internet at all in order to function and/or to be configured. It's just simply providing edge computing capabilities and its results are being consumed in multiple ways.

There is a second type of edge device that I like to call "IoT Edge Device" in order to differentiate them from the edge device defined above. The biggest difference between an IoT Edge Device and a more traditional edge device is that the IoT Edge device meets all the necessary criteria to fit within the realm of IoT. Namely:

- The IoT device always has direct access to the internet.
- The IoT device is the one that initiates the connection with the internet, not the other way around. In other words, an IoT device does not have to be *listening* for incoming connections.
- The IoT device has built-in hardware-based encryption capabilities that uniquely identifies each device and makes extremely difficult for someone else to tamper with. This hardware-based encryption is typically used

for secure communications as well to avoid a third-party intercepting and/or corrupting the data during its transmission.
- The IoT device is fully managed from the internet. You rarely perform any configuration directly on the IoT device, it's always done from the internet. This capability allows users to configure thousands of IoT Edge devices in a central location and have those configuration changes automatically propagated through the fleet of IoT Edge devices.

It's easy to see now that traditional equipment like PLCs and RTUs no longer qualify as an IoT Edge Device, as most of them lack at least one or more of the criteria listed above. In order to securely connect devices to the internet and securely manage those devices from the internet, all four criteria elements must be met. Some device vendors are working on adding the necessary features to their devices so that the four criteria above are part of their products. Other vendors and companies have chosen to have separate, dedicated devices that include those four criteria and can be deployed such that they sit in between the internet and traditional PLCs and RTUs. This is what people typically refer to as an IoT Gateway.

In Upstream Oil & Gas the use of IoT Gateways is increasing rapidly. It makes more sense to deploy an IoT Gateway that provides secure connectivity to the internet on one side and supports traditional (unsecure) protocols on the other side to interface with all field devices present at a well pad. This approach provides a greater level of flexibility as companies can opt to not install gateways on those places where power is at a premium or simply not needed while keeping PLCs/RTUs as lean low-powered devices specialized in I/O scanning and real-time equipment control. Another big reason why IoT Gateways are preferred is because they are all based on open standards and open source technologies.

IoT technologies also bring along a complete shift in the most fundamental paradigms that have reigned in traditional industrial control networks. We've already covered one of them, which is the complete shift from having all equipment in a network isolated from the rest of the world to have that equipment securely connected to the internet through IoT Gateways. Another significant paradigm shift is around the way data is handled. Up until now SCADA systems would rely on poll/response protocols and would interrogate every single field device in a round-robin fashion. This approach is executed mindlessly over and over. Detecting changes in data is up to the SCADA host as it compares the values just retrieved with values retrieved earlier. How fast and how often a SCADA system could retrieve data from field devices would depend on the speed of your network and the bandwidth available. With IoT, the devices themselves (or via an IoT Gateway) *publish* data onto a broker whenever that data has changed. This means that data is only transferred whenever it changes, and since it's the device itself acquiring and publishing the data we have a greater opportunity to see every data change at the exact moment the change takes place. Also, since most of the values managed by a given field device rarely change on a frequent basis (the density factor in a meter, the level of a tank, the position of a choke, etc.) it means that network bandwidth utilization is significantly reduced.

Having the ability to publish data, as described above, means having to use a different protocol than what has been used in the past. There are numerous protocols that support the concept of publish/subscribe, and most of those protocols are either open sourced or an open standard with multiple implementations and free applications available. The most prevalent protocol used these days is MQTT (particularly implementing the Sparkplug.B specification), but there are others like AMQP, NATS, OPC UA pubsub that are available. The advantage of using these protocols is not only the ability to only publish data whenever data changes but also the ability to "enrich" the message payload that it's published. With traditional protocols we had to deal with "naked" values (floats, integers, fixed-length strings, etc.) and it was up to the data consumer to figure out what was the meaning of that data value. With pub/sub protocols we can include rich metadata alongside the data value so that consumers have a full understanding of the message that is being received. When new data is made available, or existing data changes, all consumers have a much easier time adapting to the change. This is a huge time saver and efficiency boost when compared to more traditional protocols.

The last element to cover is the "internet" part. Of course, when we refer to the internet in general, we are really talking about the Cloud. We introduced the cloud earlier in this paper and highlighted its elastic and resilient nature, how you can simply use as much or little as you need without having to worry about the underlying technology that makes it all happen, and how you simply pay for those services you use. It's easy to see now that as we add IoT technologies into our facilities we also increase the amount of data that is being retrieved. More data not only means more storage space and computing capability (which the cloud can easily and automatically take care of) but also the need for more data analysis, the need to expose the data securely to a wider audience within the organization, the need to send this data securely to mobile devices. The built-in security, authentication, and scalability features of the cloud makes it the ideal vehicle to implement these types of systems. The low cost, pay-as-you-go approach of the cloud allows companies to marginally increment their costs while reaping all the benefits from having access to a much richer data set and analytics tools to analyze that data. Also, because IoT devices are managed from the cloud, we know have secured bi-directional interaction with our assets. This is another exciting game changer.

FREE OPEN SOURCE SOFTWARE (FOSS)

While not a technology per-se, free open source software (and hardware!) has been the unsung hero that has accelerated the creation and adoption of new

technologies. A practice that is already over 20 years old, but it's only recently that it's started to make inroads into Oil & Gas. And this will be the biggest game changer of all.

Back in the early 2000's the dominating mindset around the corporate IT community was that every piece of technology should be backed up by a large corporation. It was unthinkable to adopt any IT technology that didn't come from Microsoft, IBM, Oracle, or similar. Nobody would want to run Linux on their server rooms. The belief during those days was that hackers could get into a Linux system much easier than a Windows system because the source code for Linux was readily available on the internet, making it easier for people to find vulnerabilities. Also, since the source code was publicly available with no single corporation that owned and managed its direction, most companies didn't feel comfortable with not having a single company to call for support or hold accountable should something unexpected happened.

The next ten years witnessed a slow but gradual change of mindset. People slowly started to realize that community-backed software was even more robust than private software, as there were more eyes looking at the code and suggesting changes for improvement. It was also the time when new business models emerged around free open source. Companies like RedHat, who became one of the first corporations to provide enterprise-level support and services for Linux, began to appear around the world, further helping the adoption of open source technologies in the enterprise.

Open Source Software, contrary to popular belief, is highly regulated. There are numerous legal license models that most pieces of open source software will adhere to. The type of license offered by a specific piece of open source software will be determined by its original authors. Some open source software is licensed such that it can only be used for non-commercial purposes. Other licenses allow you to use that software for your own commercial purposes for as long as you provide the source code every time you sell your software. Other licenses force you to publish your own source code enhancements done to the original source code you started from. The possible combinations are endless, but they have also provided the necessary legal basis for companies, individuals, and communities to develop software in a more open and collaborative environment such that there is a clear expectation for all parties involved.

Off all those combinations, there are two that have been the most prevalent. They could be described as "community-led, corporate-backed" and "corporate-led, corporate-baked", respectively.

A "community-led, corporate-backed" open source technology is one such that it originates within an open community. This could be academia, a group of friends, or a group of people that have met virtually through the internet and have collectively created a new piece of technology that is released as open source. At that point, if the technology becomes successful, one or more companies will be created with the purpose of offering official support and consulting services around this new piece of open source software. The company itself will have staff that actively contributes to the open source project. There will be an open community of individuals and students that will also continue contributing to the project, therefore maturing and strengthening the solution. In most cases, those companies offering consulting and support services around this open source project will also create a number of private premium features and offer those premium features at a higher price to customers. A great example of this approach is Linux itself. Linux was born in an online community, and companies were formed at a later time (RedHat, Canonical, etc.) that provided support and consulting services for Linux, as well as created premium features on top of the publicly available Linux source code as further means of differentiation with their competitors.

A "corporate-led, corporate-backed" open source technology is very similar to the approach described above. The biggest difference is that the open source technology originates inside a corporation, and it's that corporation that decides to make that technology open source for other people around the world to contribute to the project, as well as use that technology for free. Examples of this approach includes Node-RED, which was created by IBM employees and was released as open source later on. Google released their Tensorflow library for machine learning, and even companies like The Washington Post, The Guardian, and other newspapers have open sourced the technologies they use internally for managing content and their main news website.

Nowadays we see a landscape that is the complete opposite to what we had in the early 2000's. The IT community demands the use of open source software. Giants like Microsoft, IBM, etc. have open sourced their own technologies in order for their customers to feel comfortable with continue using those technologies. Most of the products and services that we use today are either based on open source or they are fully open sourced for the whole world to see, copy, and enhance at their own will. It's this philosophy of open source that has truly accelerated the creation of new technologies. The bulk of technologies around data analytics, machine learning, cloud management, data storage, communications, etc. have been made possible thanks to open source and the open collaboration of worldwide communities that include individuals as well as large corporations.

Unsurprisingly, *every* technology that powers *every* cloud service, every IoT feature, comes from open source. The application of open source mindset is what has accelerated the development and adoption of new technologies, and that pace will continue to accelerate! As we explored at the beginning of this paper.

Such wide adoption of open source results in a less than obvious outcome. By having so many powerful

and sophisticated technologies available for free, the worldwide community is essentially "raising the bar" with regards to minimum capabilities available for people to use at no charge. You no longer need to pay money to store data efficiently, or securely encrypt communications, for example. As a consequence, you can focus on the part of the technology that will be truly differentiating for you. You can focus on the business value aspects of technology without having to reinvent the wheel or having to pay extra for it. If you, in turn, open source your differentiator (or parts of it), you will be a contributor of that bar raising that is taking place, allowing other people to leverage what you've created while only having to focus on their own differentiating piece.

This collaborative bar raising is a very stark contrast with how technology is being handled within Oil & Gas. Imagine what could happen if a lot of technologies that are essential in Oil & Gas today were open sourced. If we were able to raise the bar for Oil & Gas by making a lot of technologies free open source, we could collectively accelerate the creation and adoption of new and differentiating technology. What would be the next set of challenges to be solved in Oil & Gas if a lot of the foundational pieces were available as free open source? It's up to us to find out.

DIGITAL TRANSFORMATION

Digital Transformation can be defined as a company's unique and personalized approach to improve their own competitiveness in an increasingly digital world. It's very important to touch on digital transformation in a paper where we cover so many technologies, as one should not think that any given company can simply pick cloud, IoT, and/or open source technologies and expect that organization to successfully leverage that new piece of technology.

The world is becoming increasingly competitive, and companies must do everything in their power to stay ahead of their peers. Efficiencies to achieve that competitive edge can come from pretty much anywhere. This could be done through a company reorganization, a shift in business strategy, streamlining processes around manufacturing and/or product distribution, for example. However, what cannot be ignored is that all aspects of our lives are now driven or empowered by digital technologies. Once all efficiencies around processes and people skills are achieved the only space left for improvement is digital. In many cases, improvements in processes and/or people depend on digital technologies to help us find and implement those improvements! Oftentimes the best approach for an organization to move up the digital maturity (and competitiveness) slope is by carefully advancing people, processes, and technology simultaneously.

What does it mean for an organization to be digitally transformed? There are a few traits that organizations will begin to exhibit as they achieve higher maturity in their digital transformation journey.

- Operate in a consistent manner: Companies (and teams within companies) no longer depend on the performance of individuals in order to achieve good results. Oftentimes we see many organizations depending on the "hero" of the team to come and save the day. The quality of the outcome no longer depends on what shift was on duty or who was on vacation that week.
- Results are always high quality, regardless of who performs the task: All employees are equally empowered through digital technologies, and those digital technologies have help embed the best knowledge and skills required to complete a given task. As new improvements are found, they are incorporated into the technology toolkit so that everyone benefits immediately.
- Organization is nimble, capable to switch direction in response to dynamic situations: Once the "tribal knowledge" is digital and there is fully transparent visibility of the performance of the entire value chain, companies feel much more comfortable making bigger changes much quicker in response to external influences (new competition, new customer trends, drastic changes in global economy, etc.) since all the digital tools needed to empower the change and stick to that change consistently are already in place. A digitally transformed organization has a much higher chance of survival during tough times than a more traditional company because of this ability to be flexible with a high degree of confidence.
- Does not have a dedicated IT department: This may be somewhat controversial, but if we think about it for a minute, we will see that it makes a lot of sense. Let's imagine for a minute how a typical accounting department works today. If people in accounting needed the IT department for every change needed in the configuration of their accounting package, their speed of delivery (and cost!) would not be acceptable for that company. It's far more efficient if someone within the accounting department is a "super-user" with the necessary privileges and knowledge to make changes whenever it makes sense. If all departments at all levels had their own concept of "superuser", and all the technology was hosted in the cloud (leaving it up to the cloud provider to take care of the nitty gritty aspects of the management) a company wouldn't need a dedicated IT department in order to function successfully. Please note that IT knowledge is still required, the difference would be that IT knowledge would be distributed among the organization rather than staying localized in a separate department (part of what oftentimes it is referred to as "people upskilling"). The first company that announces to the public that they are eliminating their IT department will be the first company to be digitally transformed.

Getting started, or even taken additional steps, on the digital transformation journey is never easy. One rule of thumb to apply here is to think of the company as if it was the product. What features would you improve? What is missing, what is extra? The answer will likely be a combination of people, processes, and technology changes that are needed, and it will be important for

all three to be addressed simultaneously in order to introduce the improvement that had been identified.

CLOSING

We started by highlighting the slow pace in which Oil & Gas has traditionally adopted technology, how that pace will be accelerating significantly to the point that even "slow" will feel to be very "fast", compared to before. A big contributor to that slow adoption pace has been the culture within the industry and the overall reluctance to try something new. Oftentimes that reluctance sparks due to that new stuff being unknown to people, or seemingly complex, or people don't understand very clearly how it works and therefore how it could fail (and therefore the default assumption is that it *will* fail).

We all have heard numerous stories of how advocates of technologies like flow computers and wireless sensors received a lot of pushback from the field back in the day because people wouldn't trust what was going on inside all that circuitry, and how people would slowly "convert" to these newer technologies as the trust level increased. For the technology advocates telling the story it was always clear that the new technology was beneficial for the organization, and in a way, they would mock people in the field that would show a lack of open mindset to adopt them.

The tables have turned!! Now the same technology advocates are being asked to adopt newer technologies that are that more complex and unknown, things like cloud, IoT, open source software, blockchain. We are beginning to see a similar level of hesitance and reluctance to adopt these new technologies for the very same reasons the previous generation of technologies were not adopted by the field. It is now our chance to break the vicious cycle, to step up to the plate as confident technology leaders, and demonstrate the level of open and forward-looking mindset that we asked our peers to have not too long ago, and help catapult Oil & Gas into an efficient, digitally transformed, technology driven industry vertical that it deserves to be.