INT | Empowering Visualization

# Liberate Your Upstream Exploration Data in the Cloud

A Guide to Architecture and Technology Best Practices

**February 2019**

aws partner network

**Standard** Technology Partner

# Contents

# Introduction

Over the past few years, the oil and gas industry has undergone a series of major—and likely permanent—changes, and exploration and production (E&P) companies are now operating in a radically different landscape. As oil prices continue to hover at low levels, companies must seek value over volume, improve operational efficiency, optimize production, and minimize non-productive time to compete successfully in the market.

To achieve this, companies must embrace digital transformation and make better use of their data. This means in many cases to make it more accessible and migrate it—exploration data in particular—properly to a data lake. They must fundamentally change the way they do business—a challenge for a historically risk-averse industry—and move quickly to take advantage of new opportunities, improve accuracy given that the stakes have become higher, and encourage cross-team and cross-department collaboration to ensure continued growth. Just migrating data isn't enough—data lakes are a way for companies to unlock their data, make it discoverable and searchable, and reduce the silos that currently plague the industry.

# Driving value by putting exploration data in the cloud

Traditionally, exploration data storage has been siloed, limited by user, department, and even location, with each user receiving a narrow, limited view of the overall picture. However, in today's hyper-competitive landscape, storing data and workflows on premises in proprietary, disjointed systems can have sweeping repercussions both in terms of analysis speed and overall operational cost.

For example, historically, when analysts needed access to data housed at another location, the process could often take a week or more to complete. Once they submitted a request, it had to be processed and assigned, the data found and checked, and then—anywhere from a week up to a month after the initial request—it could be sent to the analyst for review. Companies that still operate this way waste valuable time and lose opportunities to more technologically savvy, agile organizations that have found ways to centralize their data and workflows in the cloud.

The increasing volume of data produced by wells—now exceeding 10TB per day for a single well—can also create challenges in terms of costs for on-premises storage and compute resources. To avoid getting to the point at which there simply aren't enough servers for all the data, companies typically overprovision on-site servers, a needless expense. Moving exploration data to the cloud eliminates the risk of over- or underprovisioning, because storage and compute resources can be dynamically scaled to meet actual demand. There's no need to buy more servers or request 10,000 cores and keep them idle in slow times; scaling ensures that companies are only paying for the storage they use. In short, scalability in the cloud translates to agility in the business.

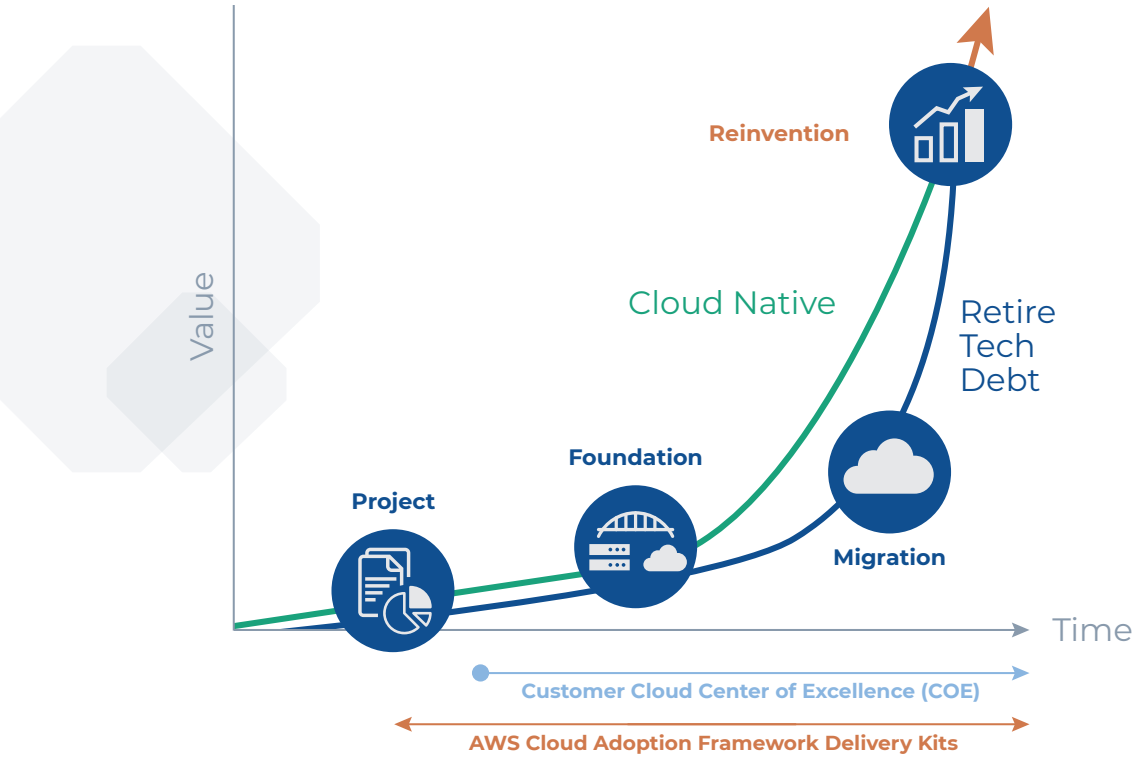


**Many E&P companies have already realized the substantial cost savings that are available. As part of its digital transformation, GE migrated more than 9,000 workflows, including 300 disparate ERP systems to AWS and achieved a 52% cost savings in total cost of ownership.**

## Faster, Better Access to Data

As E&P companies compete to find the best reservoirs, identify the best resources, or access a block before their peers do, they must have the most up-to-date data fast. When data is housed in disparate locations like data centers or on-premises servers, valuable time is wasted before key players can make informed business decisions. By storing data in the cloud, all of a company's exploration data can be accessed easily by any employee with the right permissions. This serves as critical advantage, increasing the speed of decision making and facilitating collaboration between teams who are not collocated.

## Data Discoverability

When data is stored on local disks or proprietary databases across different sites and users, E&P companies may not even know how much or what data they have. They may be able to give an estimate—"between 80 and 160 petabytes of data"—but that's still 80 petabytes of spread. Furthermore, on-premises storage opens the possibility for hundreds of different variations of the same file to be stored in multiple locations, with no insight into which file or version is most recent. When all of that data is moved to the cloud, employees at E&P companies can more easily find the data they're looking for or uncover new data that was previously housed in an inaccessible location. When that data is discovered, teams can ensure that only the most up-to-date data is stored and used—and unnecessary duplicates or variations removed or archived—to form a single source of truth.

## Cross-Department Cooperation

Another challenge of on premises storage is the lack of cross-team and cross-department collaboration. For example, teams in different locations may encounter similar problems, but without access to the same data, it can be difficult to collaborate and share strategies. Data lakes allow companies to start looking at relationships between data more efficiently; they can correlate seismic data with drilling data, seismic data with reservoir modeling, or production data with seismic data.

This effort typically results in workflows accelerating rapidly. Exploration cycles follow this trend, typically shrinking from six months to two months. And with more connected workflows, companies can identify opportunities to improve efficiency, save time, and reduce costs. To capture some of the value from accelerating workflows and enabling collaboration, companies like Woodside have begun a digital transformation and transition to the cloud to break down silos and share data across the organization.


Disparate Systems, Applications, and Databases Are Siloed, with Data Tied to Native Applications

Migrate Data to AWS Data Lake Using Snowball/Direct Connect and/or Amazon Kinesis (Streaming)

Data Lakes Like AWS Centralize Data to Allow Different Applications to Access Data in Real Time

**Companies have an opportunity to separate data from applications by migrating their assets to a data lake that can be consumed across multiple applications, services, and user populations. While this requires security and data governance considerations, this migration breaks down silos of information and better serves continuous delivery of applications that share the same data.**

# Moving exploration data to the cloud successfully

The benefits of moving exploration data to the cloud are well-established, but it can be challenging for E&P companies to accomplish this. Companies must keep in mind certain best practices to address some of the technical and cultural challenges associated with an undertaking of this size. Partnering with technology providers who have deep experience collaborating with E&P companies to make this transition—like Amazon Web Services (AWS) and INT, developer of the IVAAP™ data visualization platform—can help the process go more smoothly.

### Did you know?

Creation of data lakes may get easier due to the emergence of the **Open SDU** standard.

### What is it?

The goal of the Open Subsurface Data Universe (Open SDU) is to put all subsurface and well data (structured, unstructured, and real-time) at the center, stored in a single Data Platform, using the same standard data formats and providing a well-defined set of APIs to access the data ensures that all data can be located easily. And with multiple OSDU instances acting as one around the world, the latency distance limitations between 3D users and their applications / data can be virtually eliminated.

### Why is it important?

Currently, most E&P data is linked to applications and stored in silos. Metadata is not stored with the data, which means little or no search capabilities, making it not suitable for most data-driven applications. However, creating a common environment hosted in the cloud enables oil and gas companies to efficiently access massive amounts of data, reduce data silos, leverage cloud services, implement machine learning (ML), and lower the cost of collaboration between teams and organizations.

### Who are the key players?

Currently supported by the world's top O&G companies, including Shell, Chevron, ExxonMobil, Total, BP, ConocoPhillips, Devon Energy, Equinor, and more, the Open SDU is open to Operators, Software Suppliers, Service Companies, Academia, and others who want to help shape the development and make Open SDU an industry solution. INT and AWS are both members.

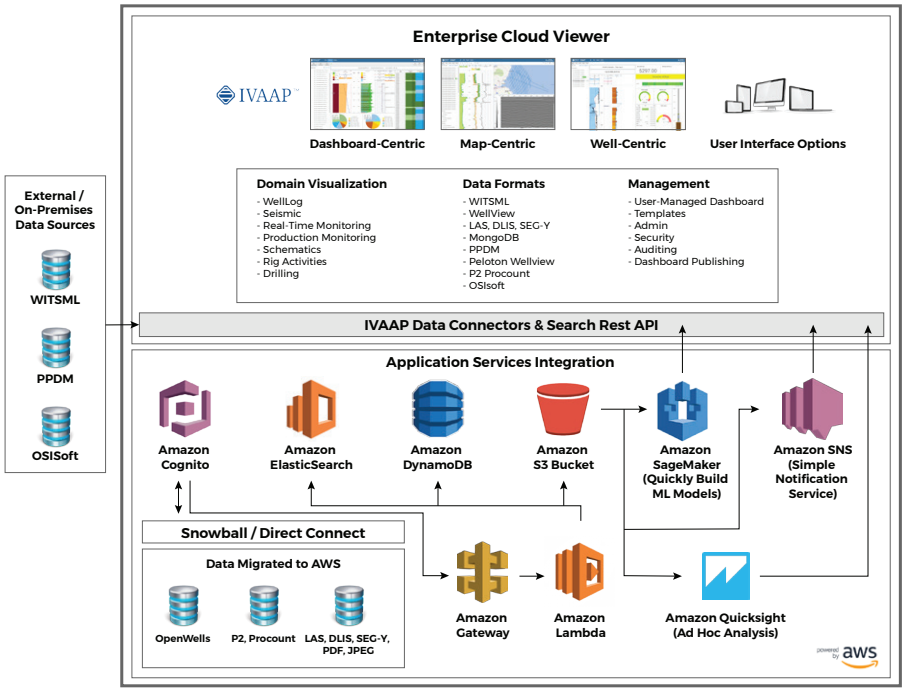


**INT is the industry standard, providing upstream data visualization libraries and enterprise platforms for 30 years. INT's IVAAP platform enables search and visualization of geophysical, petrophysical, and production data in the cloud, and it integrates seamlessly with AWS applications and services. This reference architecture illustrates the components of digital solutions requiring complex upstream data visualization.**

# Technical Challenges

## Security

E&P companies prioritize data security and minimizing risks at any cost. In addition to keeping data secure from outside the organization, these companies also want to ensure that only the right people within the organization have access to sensitive data and that data sovereignty is respected. However, migrating data without implementing the proper security controls can create new issues, such as confidential HR information showing up in unrelated searches.

Cloud-based storage and data visualization platforms like IVAAP integrate a number of security features designed to protect and restrict access to sensitive data. In addition to basic security measures such as encryption in transit, user identities are tracked via a bearer token, which allows for detailed authorization checks. This can be integrated with security mechanisms through microservices, such as pluggable mechanism Node.js Passport Module System in user access control server, with little or no effort on the part of the company.

IVAAP, for example, features three access control options: Big Data, in which INT authenticates that the user has access to the data connector; Project, in which INT sets up access control groups and assigns those to projects, so users can only see certain data if they're in the right project group; and custom access control/ external access management. Depending on the needs of the team and the organization, these options mean that confidential information remains hidden from the wider user base and other data can be hidden from users who would find it irrelevant.

AWS provide a wide variety of best practices documents, encryption tools, and other guidance that Oil and Gas industry customers can leverage in delivering application-level security measures. In addition, AWS partners offer hundreds of tools and features to help customers to meet their security objectives, ranging from network security, configuration management, access control, and data encryption. AWS customers inherit all the best practices of AWS policies, architecture, and operational processes built to satisfy the requirements of our most security sensitive customers.

## Data Formats and API

At the most fundamental level, E&P companies need to be able to access, read, and display their aggregated geophysical and petrophysical data in the cloud. However, there are significant challenges to doing so successfully.

### Moving to the Cloud: More Reward than Risk?

One challenge of migrating to the cloud is the perceived risk of storing data online. While no solution is 100% risk-free, cloud solutions can be significantly more secure than on-premises storage, offering companies the additional benefit of tracking and real-time alerts and notifications in the event of a breach. On the other hand, on premises storage can be susceptible to attacks, but without the opportunity to take immediate action. For example, in Iran, a number of power plants without internet access were hacked by malicious software planted directly into the system, with no way to trace where the software came from.

By storing their data in the cloud, however, companies are able to take advantage of the expertise and experience of Amazon Web Services, which allows them to track how their data is used, who uses it, and who accesses it in real time. AWS customers can choose to encrypt their highly sensitive content as part of a standard security process, and AWS provides tools customers can use to encrypt their data at rest or in motion, or customers can choose from a number of supported third-party security solutions. Notifications of unauthorized party access are sent within minutes instead of days so remedial action can be taken immediately. And in cases like data sovereignty, if a company were to be required to turn over confidential data, AWS's double encryption capabilities would ensure that the data is unreadable by unauthorized users.

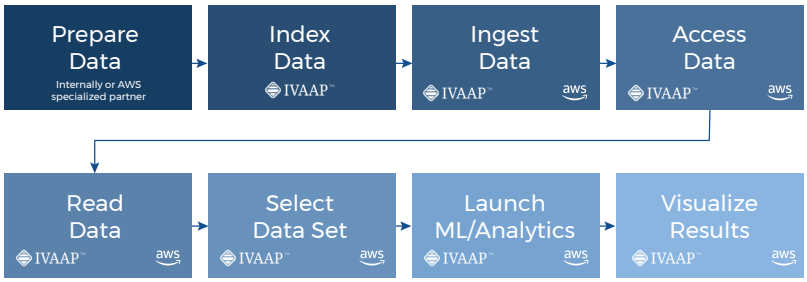


**The typical process for deploying Upstream Data Visualization on AWS Data Lake.**

### Accessing:

One of the challenges of working with exploration data is the number of file formats used in the industry, which can vary by data type, company, and location. Being able to open and access all of those files in different formats at the same time is often impossible. Backed by INT's deep domain knowledge and expertise, IVAAP reads all of the commonly used industry files formats and can evolve to add new formats quickly, thanks to its extensible microservices architecture.

### Reading:

Additionally, seismic files are typically enormous, making it challenging to effectively process them before they can be parallelized and displayed. IVAAP reads through the various formats of data and extracts the correct information, which can then be indexed.

### Aggregating:

Though data may come from multiple sources, companies still want to be able to collect that data together and form a holistic view. IVAAP aggregates and streams multiple data sources simultaneously in real time and displays the results in one view.

### Integrating:

IVAAP leverages microservices technology to allow users to extend the system, augment functionality, and integrate with services and application stacks. This technology also enables IVAAP to access different data types and formats with a uniform access layer from a viewer's point of view.

## Intelligent Searching

If companies want to store their exploration data in the cloud, they need to be able to quickly find the data they need. However, with such a diverse range of file formats and data types, companies need a platform and process that can intelligently ingest, index, and apply metadata to ensure that users can successfully search through them in the cloud.

For platforms like IVAAP, all of the native file formats are stored in a data lake. INT then indexes the seismic datasets, selectively reading certain sections of particularly large datasets. By using combined filters and queries, Amazon ElasticSearch is able to comb through those files—even massive seismic data or well log files—to find the dataset in question. As previously mentioned, access control groups and security measures ensure that sensitive or confidential information is hidden from view.

## Deployment

On the deployment side, some platforms will soon be in a state where they can accommodate continuous deployment—with no downtime or risk of human error—for bug fixes, platform improvements, or new features. They can also be validated in pre-production environment, which reduces debugging time and the risk of a failed deployment. AWS natively provides the ability to perform blue-green deployments with A/B testing enabled, where a small subset of users can be exposed to a new feature, thus minimizing the risk of system-wide failures. And if something goes wrong, AWS provides the tools to rapidly rollback to the stable environment. This is an exciting development for E&P companies looking to increase efficiency, and they'll need to look specifically for tools that utilize a microservices-based architecture through infrastructure designed to work with AWS.

## Challenges of Upstream Data Visualization

- Indexing of specialized upstream data formats
- Ingestion and reading of E&P data types (LAS, DLIS, SEG-Y, WITSML)
- Visualization of very large data volumes
- Security and entitlements
- Search capabilities (geospatial, metadata, per type)
- Support of an integration with existing upstream workflows
- Data usage for ML and advanced analytics

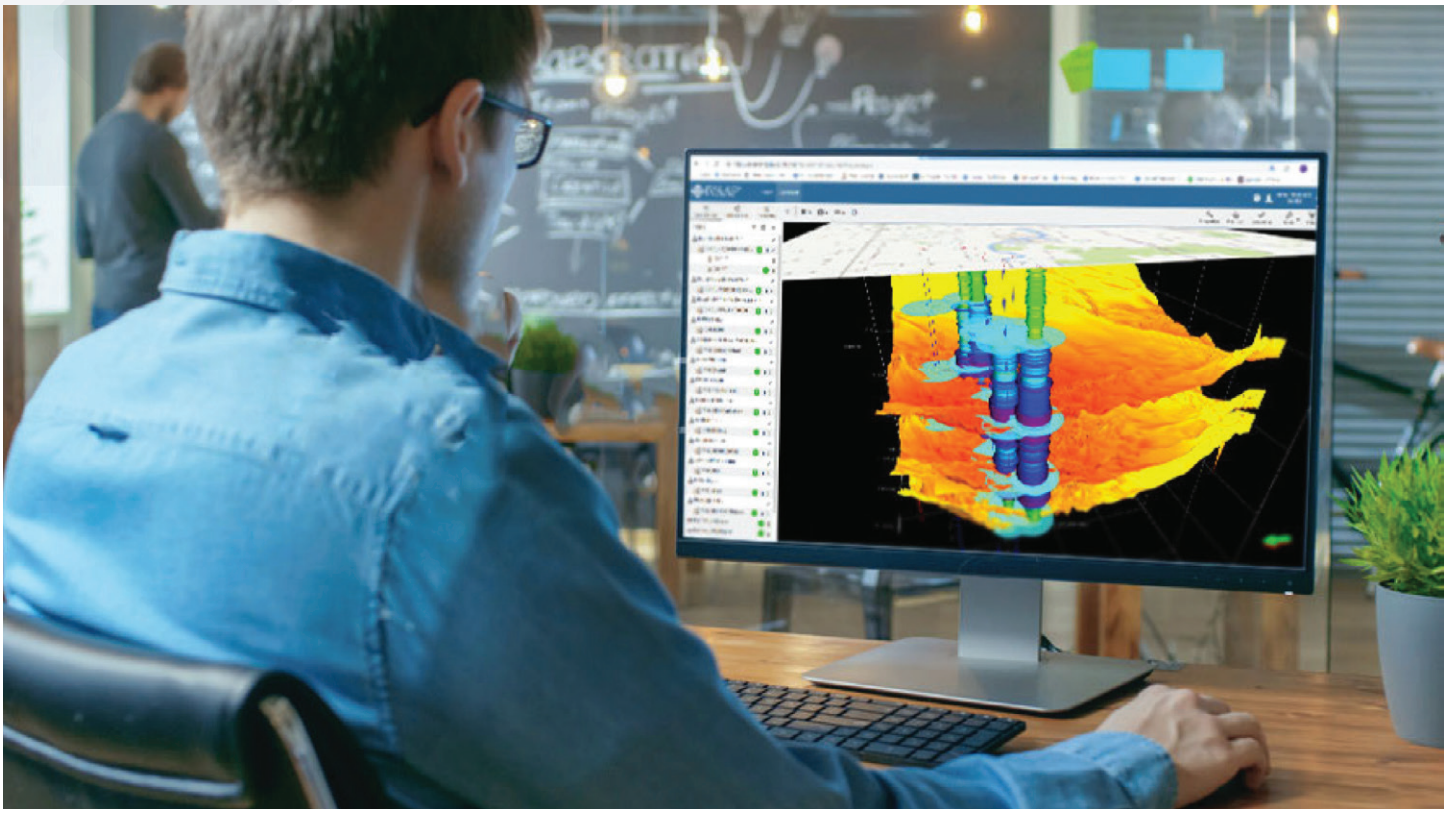

# Cultural Challenges

### User Adoption

One of the biggest cultural hurdles to moving exploration data to the cloud and using a new platform is user adoption. Culture can be hard to change, and people quickly grow used to their current way of working and become resistant to change, even though the old way may be inefficient. In this regard, education, outreach, and highlighting solutions to specific pain points and concerns can be critical. Both AWS and INT offer training and workshops to help E&P companies and their employees successfully use new platforms. INT also constantly improves the IVAAP user interface to be more intuitive for users and to fit more seamlessly in various workflows.

# Cloud Platform Challenges

### Data Lake Ingestion

Relying on data lakes enables companies to ingest and store vast quantities of data—both structured and unstructured. AWS provides a number of migration tools to move resources and data from on-premises to the cloud. Existing large amounts of data can be migrated using AWS services like Snowball and SnowMobile, each capable of ingesting hundreds of terabytes to exabytes of data respectively. And core services such as Amazon Kinesis, AWS IoT Greengrass, and Amazon Storage Gateway enable real-time ingestion.

However, getting the data lake architecture correct can be difficult to achieve. IVAAP and AWS use Amazon S3—built to store and retrieve any amount of data from anywhere—for object storage. With virtually unlimited scalability, companies can easily increase their storage to house petabytes of exploration data at a lower cost than on-premises storage. And with intelligent tiering, customers can archive unused data to Amazon S3 Deep Archive.

### Applications

Oil companies have thousands of applications in use across all of their organizations. They are frequently challenged by the need to migrate only some of them to the cloud while still maintaining access to others. Migrating exploration data to the cloud and deploying a platform is a multi-year journey, so it's important to find a platform that can integrate with existing applications and workflows while that process is ongoing. AWS has a deep understanding of the complexity associated with keeping existing applications accessible and functional while moving to the cloud.

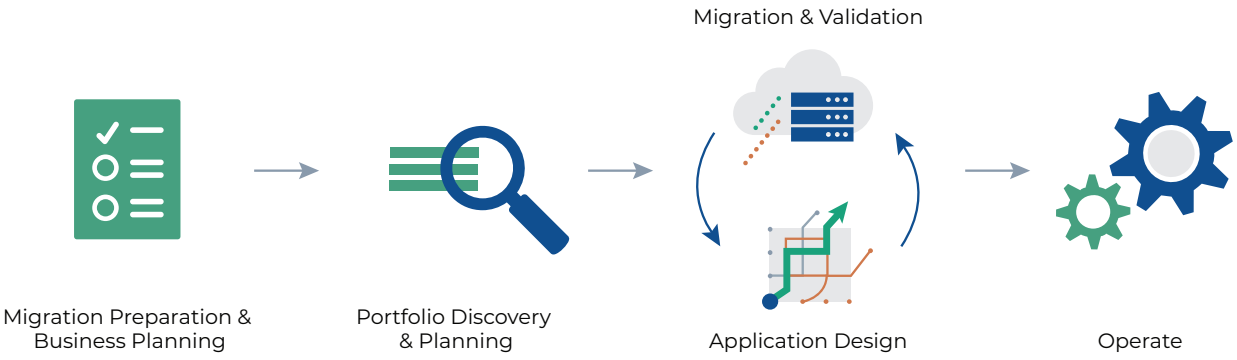

# Conclusion

Moving exploration data to the cloud can be a game-changer for companies looking to become faster and more efficient than their competitors in this new energy environment. However, the simple act of moving data to the cloud is not enough to ensure success. E&P companies must go beyond just storing data in a central location in the cloud and instead take a holistic approach to data visualization. Partnerships like the one between INT and AWS are particularly valuable because of their deep understanding of how the other functions, as well as an extensive prior working relationship. These factors enable INT to implement solutions designed to work effectively and seamlessly with AWS.

Designed specifically to address the challenges that E&P companies face with their existing workflows, IVAAP is a cloud-based data visualization platform built to empower product owners, developers, and architects to accelerate the development of subsurface digital solutions for oil and gas. Rather than develop a new application from scratch, companies can customize and implement IVAAP in their organization in a fraction of the time it would take to build in-house. Deployable in any environment, IVAAP can scale to meet the needs of tens to thousands of users, and it can be customized to ensure that the right people are accessing the right data.

IVAAP's well log visualization and analysis framework makes it easy to input data, add algorithms, and launch and visualize the results. Robust out-of-the-box features include log displays, well section, lithology editing, curve filtering, formation tops, well correlation, and much more. All of this is available within an intuitive, easy-to-use interface that makes it easy to find and draw connections.

To successfully make the move, E&P companies will need to assess the architecture of a platform. Is it an efficient data structure? Does it enable support? Does it allow for seismic processing and adequate data searching? How does it perform under a heavy load? How is it organized, and who has access to what? Answering questions like these can narrow down the list of tools to only the most promising. Once that tool is selected, companies can begin moving their exploration data in earnest, ensuring their swift progress toward greater accuracy, cooperation, efficiency, and speed of decision making.

INT | Empowering Visualization

# About INT

Interactive Network Technologies, Inc. (INT), provides data visualization software using the latest technologies such as HTML5 and JavaScript. For 30 years, INT has created cloud-enabled and mobile-responsive solutions and platforms to empower the leading oil & gas and service companies to visualize, monitor, and analyze their complex data.

For more information or to schedule a demo of our software, contact us at **+1 713-975-7434** or **info@int.com.**

aws