COGNITE

The Industrial Data Architect Edition

Data Fusion Magazine

Discover the best in thought leadership and insights to help organizations in heavy asset industries, when it comes to DataOps, data contextualization, and data management.

In this issue

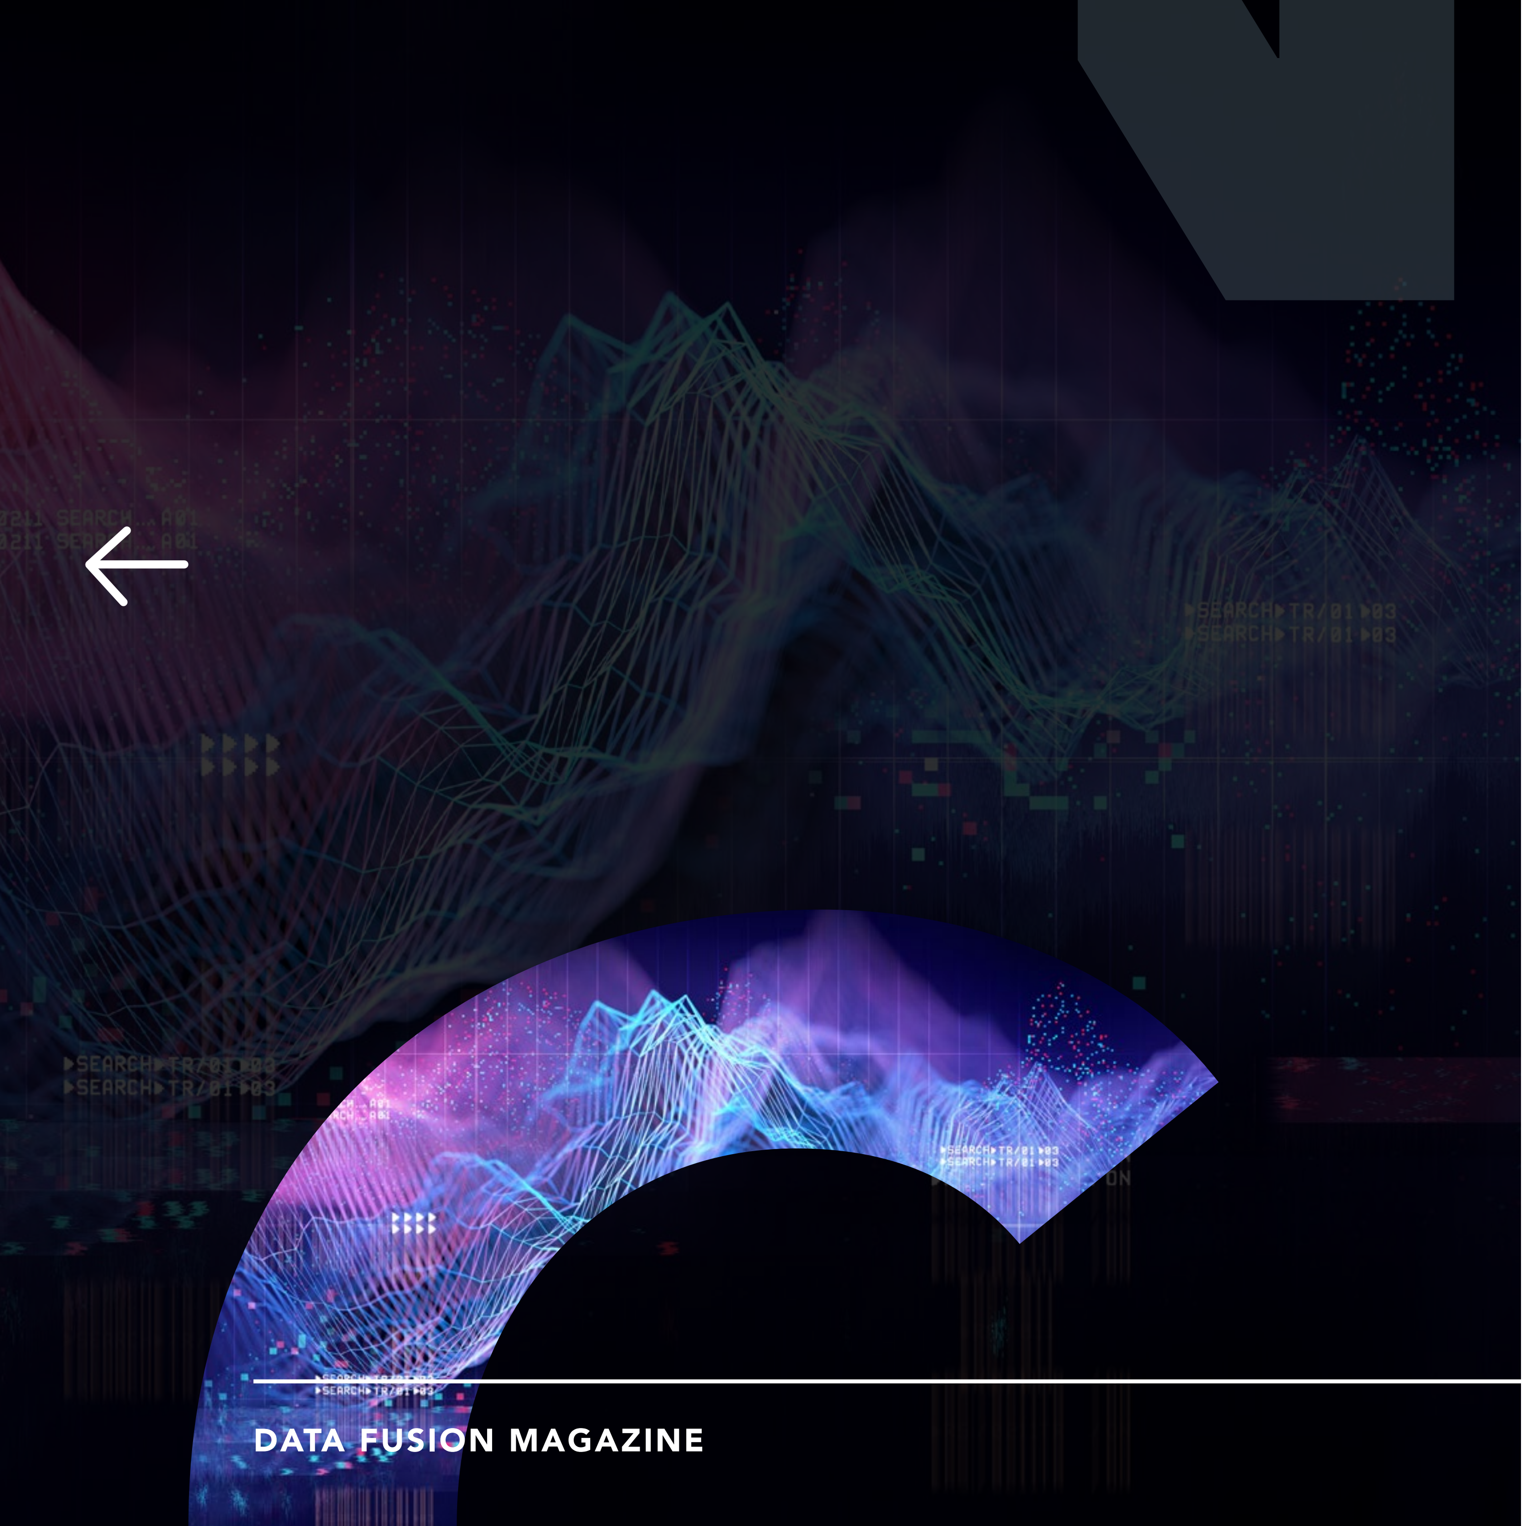

3 **Foreword**

5 **Reduce cost and drive data literacy with DataOps**

11 **DATA STORIES:** How Aker BP and Cognite secured correct allocation of oil and gas at Alvheim with minimal production loss

12 **The Data Liberation Paradox: Drowning in data, starving for context**

16 **DATA STORIES:** How Aarbakke and Cognite liberated and visualized data to optimize the use of cutting tools

17 **What is data fabric and does it complement my data warehouse?**

23 **The True Digital Industry 4.0**

26 **DATA STORIES:** How Cognite helped Hafslund E-CO optimize turbine start-up sequences

27 **Physics and AI hybrid delivers working AI for industry**

32 **Are you future-proof?**

FOREWORD

Data Fusion

This special edition focuses on DataOps and the industrial data architect.

We'll explore how DataOps can reduce costs and drive data literacy, and we'll delve into the Data Liberation Paradox to learn how to avoid that trap.

Along the way, you'll see examples of heavy asset companies who are turning theory into action and getting great results, in sectors from oil and gas to manufacturing.

And Industry 4.0—we hear a lot about it. But what does it really mean in practice?

Discover all this and more in your issue of Data Fusion.

All the best,

John Markus Lervik
CEO, Cognite

Data points

"Data is abundant, but it's what organizations do with it that matters. Adding context to data improves usability and insight across business units.

But 85% of firms in this study face issues as a result of not providing contextually enriched data to data scientists, engineers, citizen developers, and citizen data scientists.

As a result, firms wade through avertible inefficiencies because it takes too much time to locate, access, use, or make sense of data."

Source: Contextualized Data And Digital Twins Amplify Digitization Value: Operationalize And Scale Models And Applications With Trusted Industrial Data. A Forrester Consulting Thought Leadership Paper Commissioned By Cognite, February 2020

WATCH THE WEBINAR

DOWNLOAD THE FULL REPORT

Reduce cost and drive data literacy with DataOps

Oil price fluctuations, capital market uncertainties, and the pandemic leave no part of enterprise operations unaffected. These unprecedented times compel us to examine our operational efficiency and resilience, including how we operationalize data to address the needs of our diverse data consumers.

What can digital and data management leaders do to reduce cost, drive data literacy, and improve data value capture in times of austerity and uncertainty?

What's wrong with the way we operationalize data today?

It is no secret that data management and analytics workflows[1] have always been complex, siloed, and costly to enterprises all over the world. Adding insult to injury:

- The rise of AI/ML with difficult to find data scientist is imposing its own set of - often very differentiated from conventional BI user - requirements on data modelling, data source availability, data integrity, and out-of-the-box contextual metadata on data;
- Data engineers working on industrial digitalization projects struggle with access to key source system data that is reminiscent of year 2010 in non-industrial verticals;
- Industrial companies are not only facing the same challenges as, for example, their retail peers, but are presented with a superset of challenges resulting from the IT/OT convergence and associated nonconventional-IT-only data velocity, variety, and volume; and
- Rushing to show digital execution, many have embraced the AI hype that has led to quickly demonstrable digital proof-of-concepts, yet is failing to yield truly operationalized - and even less scaled - concrete business OPEX value.

A typical data pipeline for analytics with associated workflow challenges. Source: Gartner

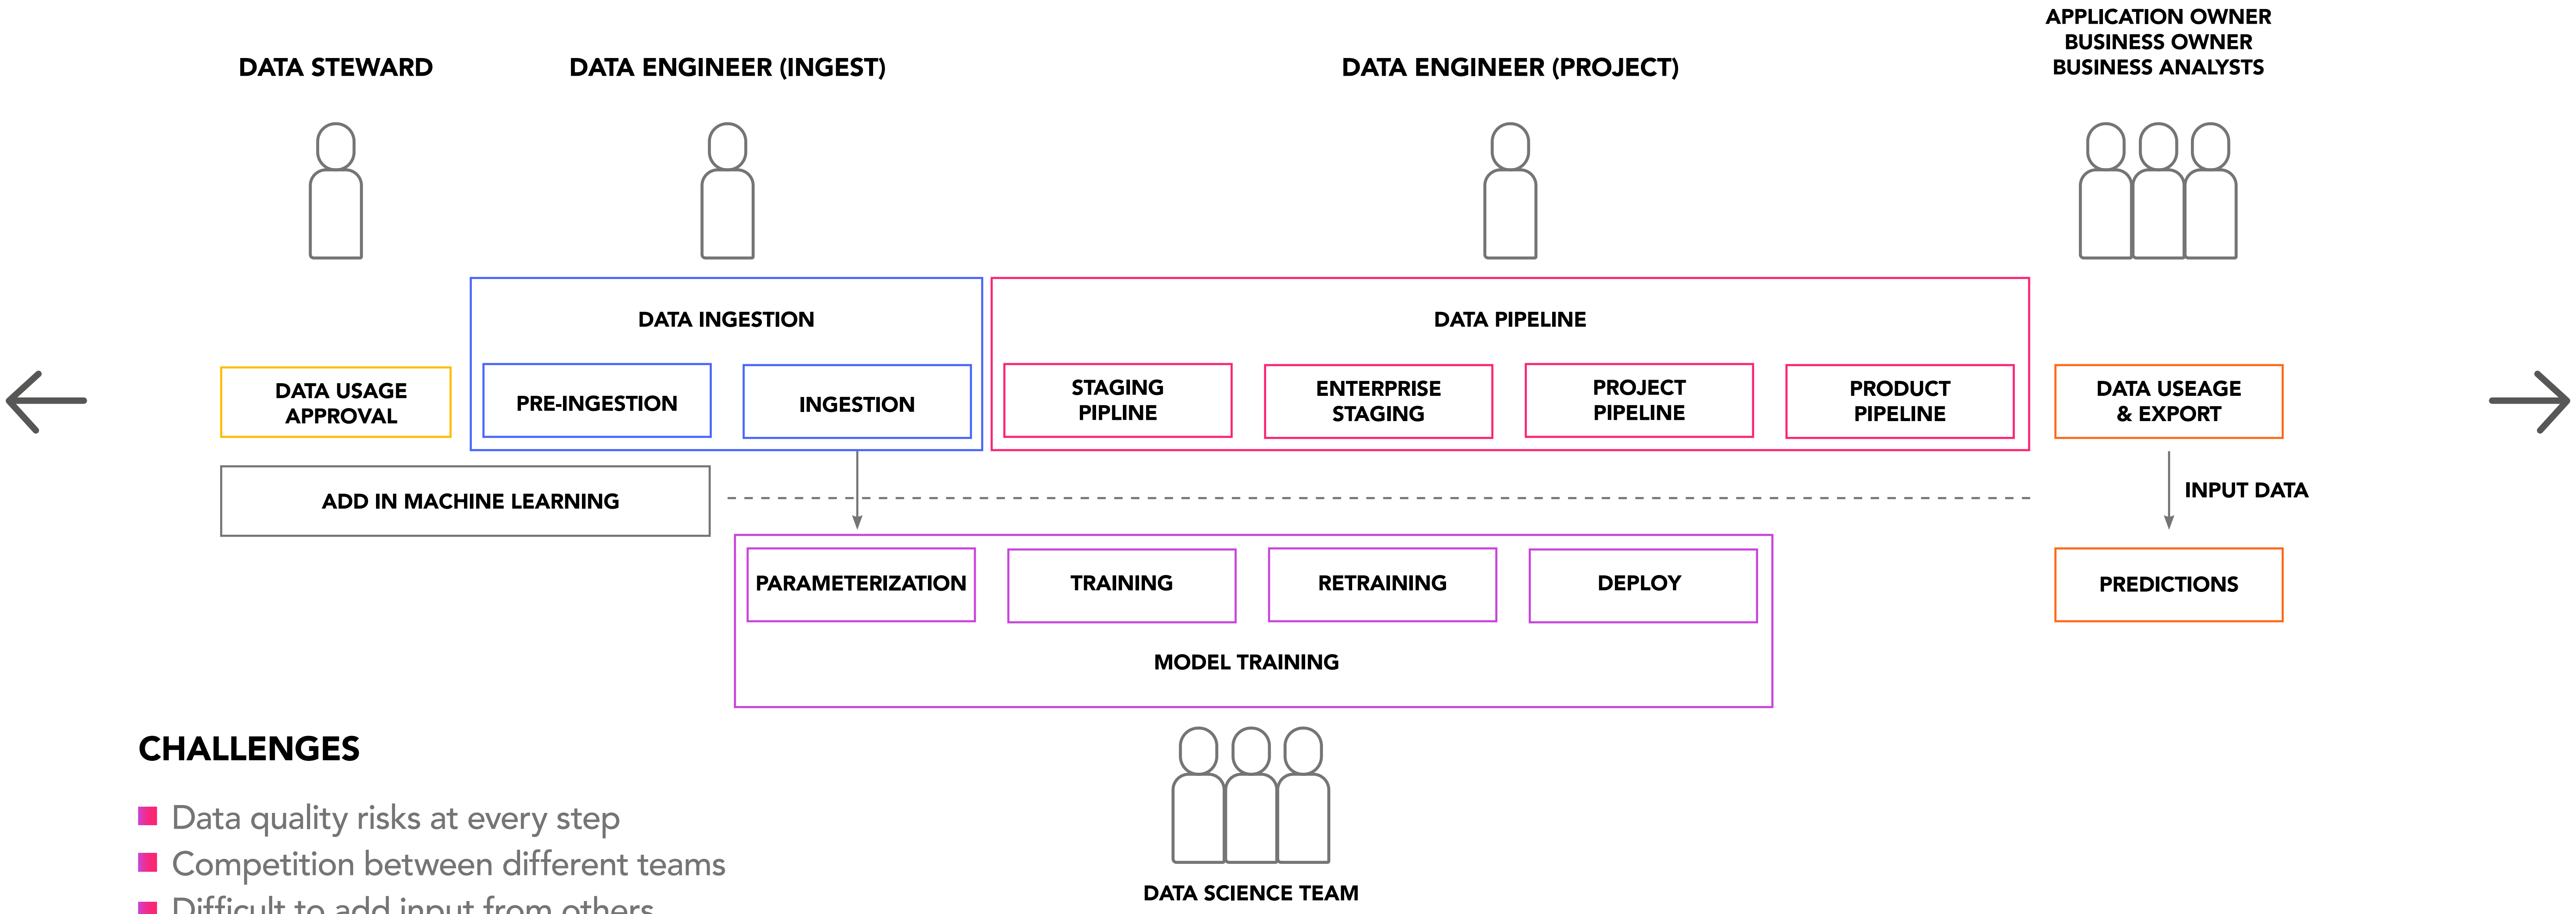

CHALLENGES

- Data quality risks at every step
- Competition between different teams
- Difficult to add input from others
- Lacks flexibility

The IT/OT divide is converging over time.

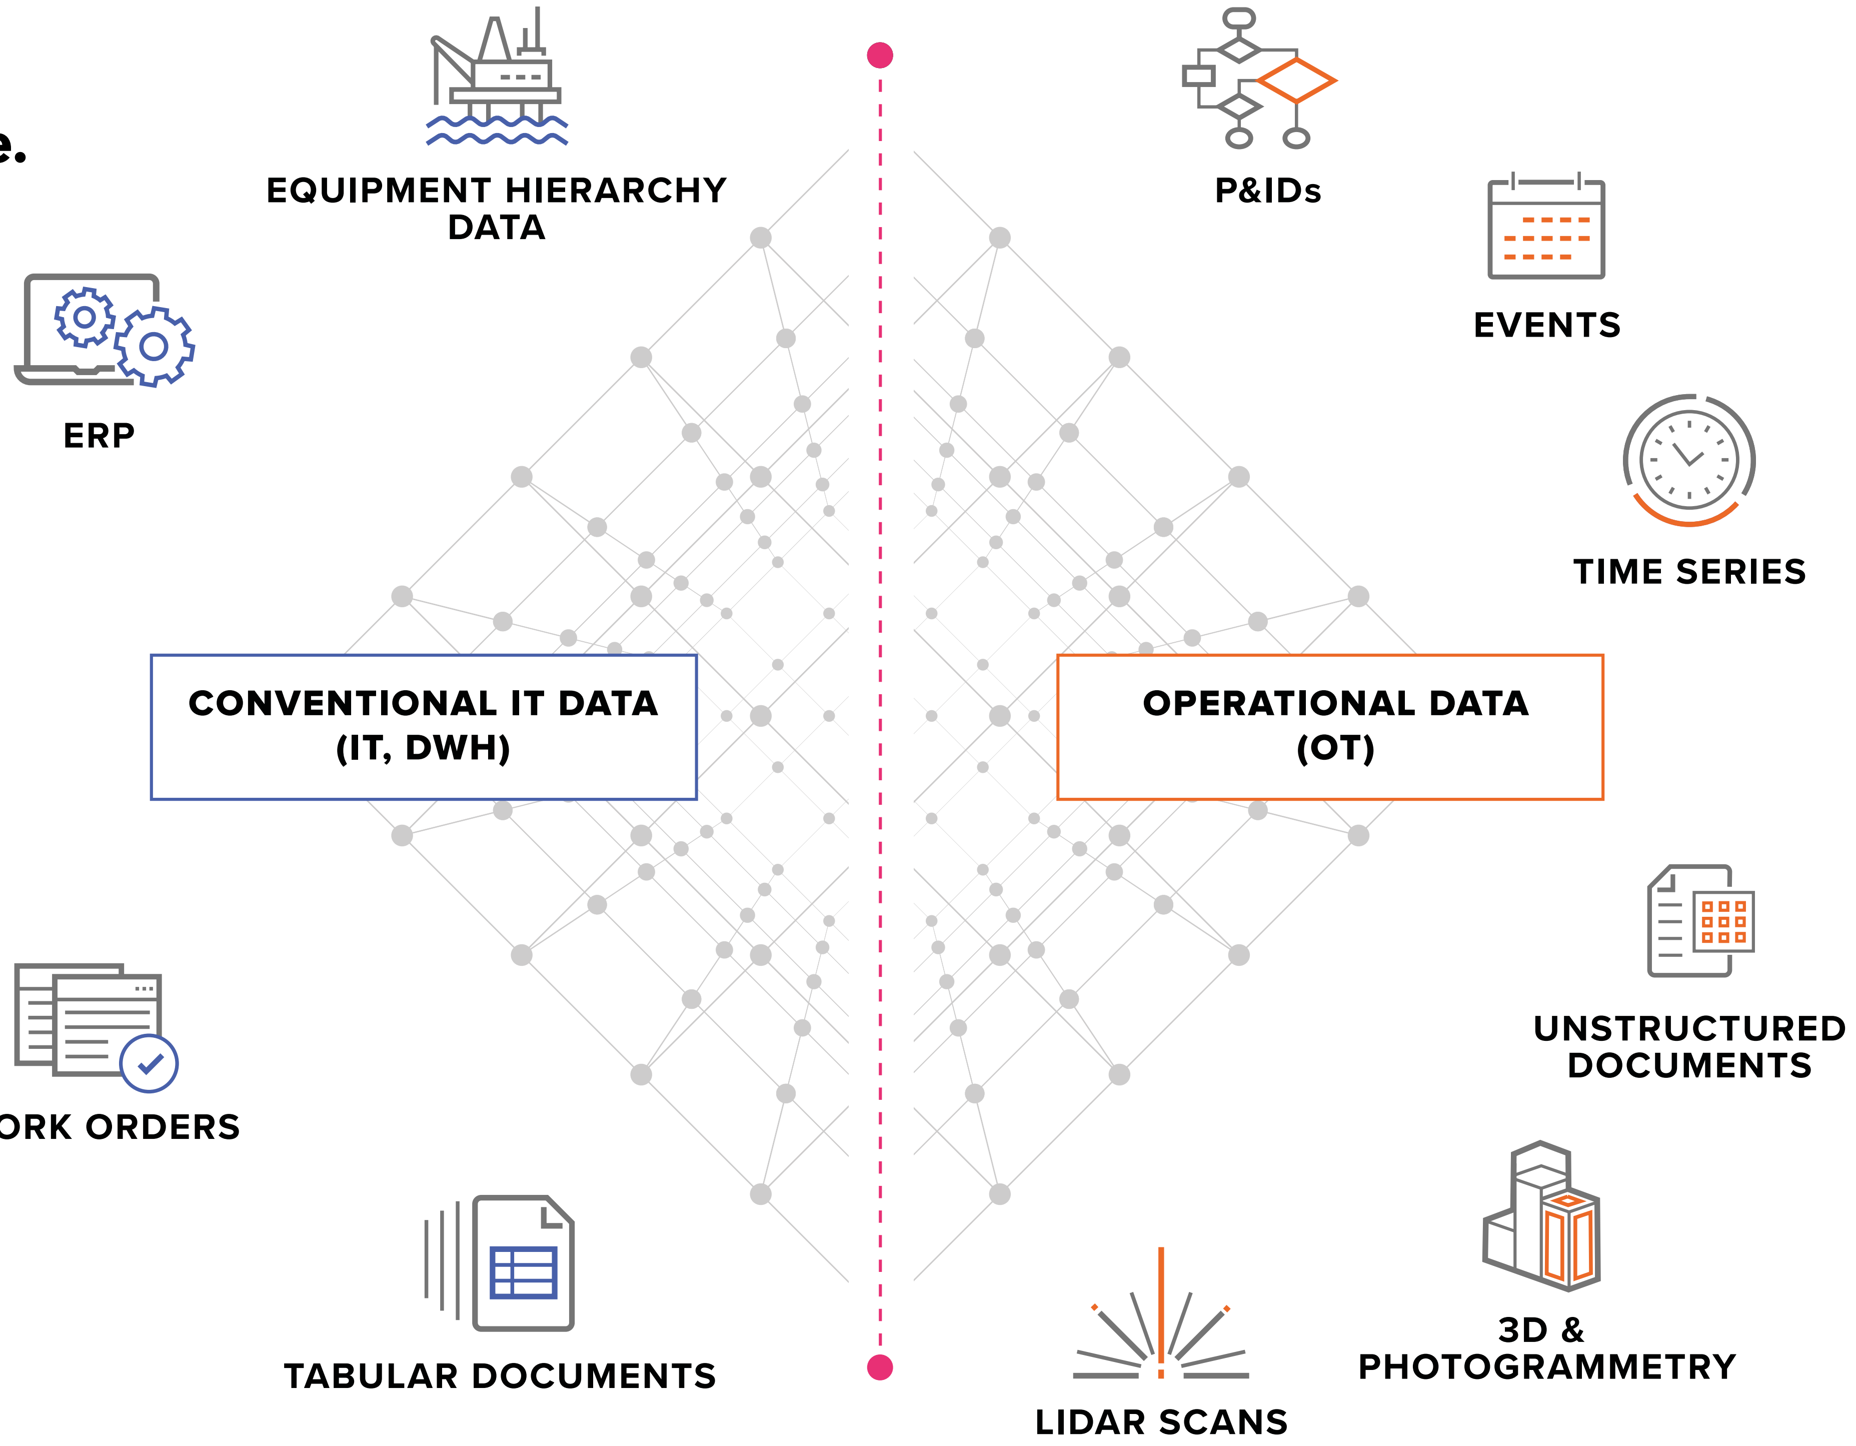

Your industry and your data culture are changing — you need DataOps

Similar to DevOps, DataOps has a compelling and topical value proposition. With DataOps, you reduce specialized roles in your data-to-value workflows and enable higher data consumer autonomy and empowerment, thus creating higher resilience and a more lean and cost-efficient core for digital transformation.

According to Gartner, 'the goal of DataOps is to create predictable delivery and change management of data, data models and related artifacts. DataOps uses technology to orchestrate and automate data delivery with the appropriate levels of security, quality and metadata to improve the use and value of data in a dynamic environment."

Not convinced DataOps is applicable to your organization and industry? Let's look at the recent market research data.

Forrester

Forrester (2020) found that 79% of digital leaders in industrial companies either "agree" or "completely agree" with the following two statments:

- "Data needs to become self-explanatory to data consumers without needing subject matter expert support"

- "A company is digitally mature when they can enable citizen data scientistt and citizen developers to do more with advanced data and analytics themselves"

Gartner

Gartner (2019) found that analytics and data mangement leaders' highest-ranked observation of change in environment is:

- "There is increased demand for new data types and the ability to model it independently"

Citizen data scientists - i.e., production and maintenance engineers proficient in Microsoft Excel - proactively upskilling themselves, are the fastest growing of organizations surveyed by **Gartner** (2019) listed citizen data scientists as their primary stakeholder group for developing and deploying AI initatives,

According to **Gartner** (2019), "the number of data and analytics experts un business is growing at three times the rate of experts in IT departments, which will force companies to rethink their organizational models and skill set" For those focusing on conventional BI and reporting needs being met as their first priority, **Gartner** (2019) found:

- "Only 12% of respondents claim their BI & reporting needs are being met. That leaves a staggering number of unsatisfied data consumers."

Similar to DevOps for professionally developed software, DataOps and low-code form the technological foundation for citizen data science- and citizen-developed

GARTNER (2019) FOUND ONLY

of respondents claim their BI & reporting needs are being met.

DataOps is key to increasing data literacy. Data literacy is key to securing value from data and digital.

Enabling all data consumers to have instant access to all data with context -- what we at Cognite call contextualized data -- is not easy. We know.

But imagine what your data consumers can do when they are all empowered to speak data, to independently and collectively access all relevant, contextualized data, and to safely develop the next generation of productivity-enhancing digital applications -- unleashing the transformative potential of Excel 2.0 for your business.

Share experiences and learn from peers

Leading change during times of disruption is not always straightforward. To guide digital leaders, Cognite is creating a research report to help digital leaders in industry with knowledge sharing in this time of collective upheaval. Participate now and receive a complementary copy of the report.

Be a part of our research report to help digital leaders in industry through this time of collective upheaval and receive your free copy.

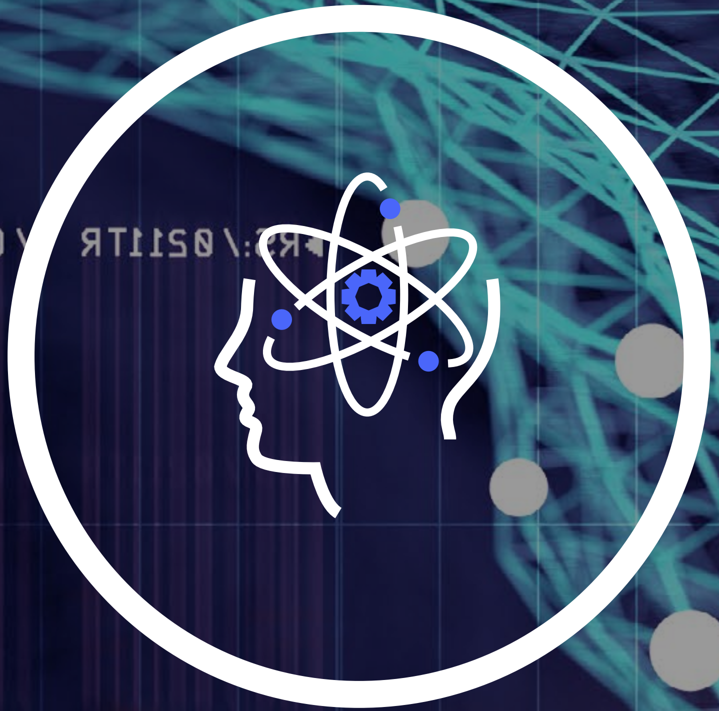

PARTICIPATE TODAY!

Data stories

How Aker BP and Cognite secured correct allocation of oil and gas at Alvheim with minimal production loss

$3.5 million a year.

That's how much Aker BP saved by optimizing a calibration process that had been unchanged for years, eliminating production losses without sacrificing accuracy.

Aker BP, Cognite, and Expert Analytics used Cognite Data Fusion and Asset Data Insight to develop a new algorithm for calibration of sensor arrays essential to the correct allocation of hydrocarbons at the Alvheim field.

FIND OUT HOW THEY MADE IT HAPPEN

The Data Liberation Paradox: Drowning in data, starving for context

Digital twins are in. Data lakes are out. Let's start with two alarming findings from Gartner:

- "80% of data lakes will not include effective metadata management capabilities, making them inefficient."

- "Data and analytics organizations that provide agile, curated internal and external datasets for a range of content authors will realize twice the business benefits of those that do not."

In other words, organizations that implement data catalogs to communicate data context to their internal and external data customers will realize 100% more business value than those who don't.

What happened to unlocking infinite data-driven potential by liberating data from siloed source systems and integrating it all into one single repository? Let's examine the anatomy of the paradox.

First, data lakes only store data in an untransformed raw form.

While raw data is theoretically available across realms of immediate, potential, and not-yet-identified interest, active metadata management is often a mere afterthought. It winds up as the technology project's flagship KPI but lacks enthusiasm or investment from the stakeholders along the way.

Raw data — absent of well-documented and well-communicated contextual meaning — is like a set of coordinates in the absence of a mapping service.

Those lucky few who intuitively understand the coordinates without a map may benefit. For all the rest, it's the map that provides the meaning. Without a map, coordinates alone are useless to the majority.

Second, data lakes lack contextualization.

While some applications benefit from raw data, most applications — especially low code application development — require data that has undergone some additional layer of contextual processing. This includes aggregated data, enriched data, and synthetic data resulting from machine learning processes.

Here is where the value of data contextualization becomes most pronounced. Aggregated, enriched, and synthetic data delivered as an active data catalog is far more useful to application developers.

Strong API and SDK support, designed for use by external data customers, further amplifies the value of this processed data.

It's also something raw data in a large, unified container fails to address.

Data lakes that only store data in an untransformed, raw form offer little relative value. These vast amounts of expensively extracted and stored data are rendered unusable to anyone outside the data lake project team itself. (And too often remain somewhat useless to that team, as well.)

Read also: Reaching a Critical MaaS

The solution is to complement existing data lake practices with data contextualization.

Data contextualization goes beyond conventional data catalog features (see table below) by providing relationship mining services using a combination of machine learning, rules-based decision-making, and subject matter expert empowerment.

DATA CATALOG FEATURES

- Allow users and data curators to assign different levels of trust to both raw and processed data sets
- Allow approved users to rate integrated data sets for accuracy
- Offers transparency to actual data usage and query usage
- Single source of truth on which data sets are related and the nature of that relationship
- Track data lineage

UNCONTEXTUALIZED DATA

Many mid-sized organizations, operating mostly with IT data, may benefit from a simpler data catalog solution. However, large industrial asset operators dealing with the synthesis of OT and IT data—not least the on-going proliferation of IoT data together with very complicated brownfield data realities—call for an enterprise-grade data contextualization solution.

ML

28 732 TIME SERIES

132 132 ASSETS

The most powerful and practical application of data contextualization is the creation of a foundational digital twin. This digital representation of an organization's full, live, connected operational reality is a huge benefit to IT organizations struggling to showcase their data platform capabilities, including data contextualization, to a less data-savvy audience.

RULES ENGINE

55 RULES INDENTIFIED

Contextualized data generates immediate business value and significant time-savings in many industrial performance optimization applications, as well as across advanced analytics workstreams. (We have found that data science teams are one of the greatest beneficiaries of data contextualization.

SME W/TOOLING

99.9% MATCHED

Digital twins are proving to be a fantastic tool for communicating the potential value and opportunities presented by liberated, enriched, and contextualized data.

The distance between a data lake and a digital twin may not be that far after all. Contextualization, offered as a service for data already aggregated in one place, offers an instant upgrade.

Data stories

How Aarbakke and Cognite liberated and visualized data to optimize the use of cutting tools

Cognite and Aarbakke used Cognite Data Fusion to break down the walls between separate systems that tracked machine operation data and work orders, helping the manufacturing company analyze and optimize cutting tool usage.

-60%
reduction in tool assemblies

+10%
increase in cutting tool efficiency

22,000+
unique cutting tools overviewed

AARBAKKE

SEE HOW THE DATA LIBERATION MADE A DIFFERENCE

What is data fabric and does it complement my data warehouse?

Connecting DWH to advanced analytics at scale requires a data fabric, not just data availability.

For industrial companies, the path to ultimate value from data liberation requires three crucial steps. Many organizations have already achieved step one: liberating data from siloed source systems and aggregating it in a traditional data warehouse (DWH).

Three ladders of data liberation

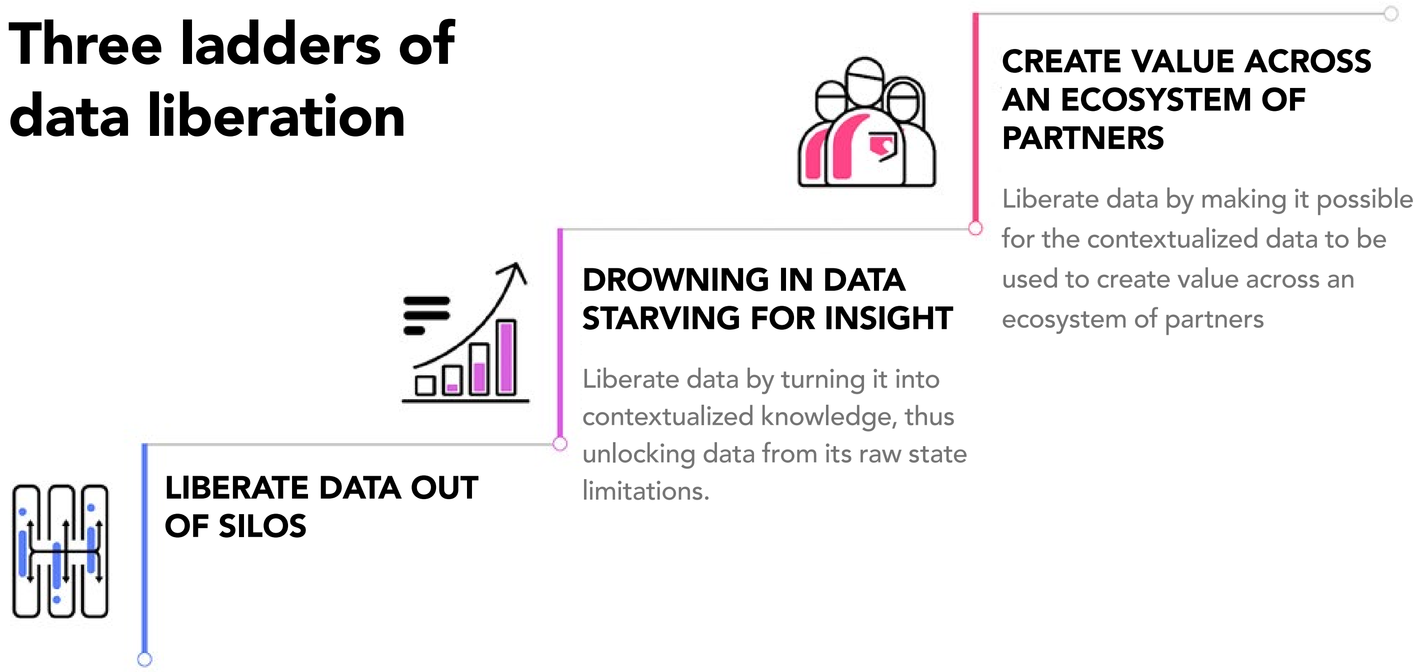

The bad news: Step two is far more difficult to achieve. The good news: The rewards for successfully taking it are correspondingly higher.

In today's mature DWH market, progressive data-driven organizations are actively utilizing data fabric solutions as a complement to existing DWH strategies.

With data fabric, organizations can liberate their data once again. Lifting it from the pool of aggregation and turning it into contextualized knowledge to deliver on their ambitions for advanced analytics.

What is data fabric? And what sets it apart from data warehousing?

The two main pillars of data fabric are Context and Discovery. These define data fabric and make it both distinctly different from and complementary to existing DWH.

1. Data context is the sum of meaningful use case supportive relationships within and across different data types and data artifacts. It is the result of data relationship mining and curation in a so-called contextualization pipeline. The process of adding context to data is often referred to as data contextualization or data fusion.

Contextualization pipeline

Prior to contextualization, data is often integrated from a multitude of source systems and co-located in a common data repository, similar to traditional DWH. Alternatively, data integration is virtualized through data federation, avoiding the need for data duplication and transfer. More recently, a hybrid approach has become common, especially for latency-sensitive IoT data applications, where data aggregation and data synthesis must be performed close to the data.

	DATA INTEGRATION	DATA FEDERATION	HYBRID
DATA RESIDENCY	Data duplication	No data duplication	Both
DATA LINEAGE	Complicated but manageable	Resides at source	Both
DATA QUALITY	Complicated but manageable	Resides at source	Both
ABILITY TO PERFORM DATA OPERATIONS	Rich support for aggregration & synthetic	No close-to-data operations	Both
DATA RESIDENCY	Low latency included	High latency only	Both
DATA LINEAGE	Long	Up to 50% less	Medium
DATA QUALITY	Large	Up to 40% less	Medium

What is the industrial status quo?

In oil and gas, digitalization efforts have long been limited to pilot projects, proofs of concept and case studies, with no large-scale operationalized projects.

This is mainly due to outdated IT infrastructures that rely on legacy systems and only enable point-to-point integrations for application providers. These one-off solutions–sometimes including limited digital twins–can actually complicate digitalization goals because the resulting projects are as siloed as the original data, making them impossible to scale and, therefore, costly to the point of wastefulness.

Complementing existing DWH solutions with data fabric has dramatically reduced costs, while simultaneously enabling scalability, speed of development, and data openness throughout our many complex customer organizations.

2. Data discovery is about making data effortlessly available to the right user in the right format. This has always been the goal of data and information architects.

Discovery in B2C technology is instantaneous, autonomous, and continuously self-learning. In other words, it's far ahead of enterprise and IoT data discovery. But that's where we're going: shifting from active search to passive discovery based on personalized relevance.

Recently, the exponential increase in data volume, velocity, and business value, coupled with the meteoric rise of low code and citizen data science programs, is making data discovery more important than ever before.

	TRADITIONAL DATA	CLOUD DWH	DATA FABRIC
DATA RESIDENCY	On-premise	Stored in cloud	Integrated and federated in cloud
GOVERNANCE	Strong	Strong	Strong
ELASTICITY	Limited	High	High
SETUP TIME	Long	Short	Short
SUPPORT FOR IOT	No	Yes	Yes
SUPPORT FOR LOW COCE	No	Limited	Very strong
SUPPORT FOR DATA SHARING	Limited	Yes, discovery difficult	Yes, easy discovery
SUPPORT FOR 3RD PARTY APPLICATION DEVELOPMENT	Limited	Yes, Limited	Yes, strong
TIME TO VALUE	Long	Short	Short
COST	Fixed	SaaS	SaaS

In the context of enterprise data management, enabling the right data to be easily discoverable relies on much the same recipe: the right metadata, labeling, linkages to other data, and data cataloging to make it readable by both machines and humans.

Outdated manual metadata management is gradually being replaced by active, machine learning-supported metadata practices, used to discover and infer additional metadata from relationships and clustering.

This is why progressive organizations are seeking out data fabric solutions to complement their DWH strategies. Data fabric adds critical context and discovery to existing DWH data assets.

Complementing DWH with data fabric is the only way to push through all three steps to true data liberation.

Download our eGuide to learn how to extract value from contextualized data in the oil and gas industry.

DOWNLOAD NOW

Data points

Tools are crucial – 98% expect benefits from using an industrial data and AI platform to consolidate and contextualize data from disparate IT and OT systems.
85% of industrial enterprises face issues as a result of not providing contextual data to data consumers

As digitalization is increasing, data contextualization is becoming crucial to make informed business decisions.

Source: Contextualized Data And Digital Twins Amplify Digitization Value: Operationalize And Scale Models And Applications With Trusted Industrial Data. A Forrester Consulting Thought Leadership Paper Commissioned By Cognite, February 2020

WATCH THE WEBINAR

DOWNLOAD THE FULL REPORT

The True Digital Industry 4.0

Platform fatigue and the rise of the ecosystem

Those working in enterprise architecture have witnessed how the walls of the decades-old, centrally-governed, and waterfall-process-loving enterprise have come down. The replacement: a more agile, more adaptive, and more innovative organism... the ecosystem. A network of ecosystems, in fact.

And as the ecosystem model usurps traditional value chains, on-premise IT architectures are being replaced by cloud software ecosystems. This new model is comprised of loosely integrated yet tightly coupled IaaS, PaaS, and SaaS. Innovation architecture is moving from an internal IT organization with a digital CoE sidekick to an ecosystem of partners and their ecosystems of partners.

Across all of these ecosystems, the life blood is data. High-quality, contextually relevant, easily available and easily shareable.

And yet, in the new race to become digitally-native, cloud-first, and open, many are falling for an old trap: Restricting themselves to the narrow and outdated definition of a software platform, thereby constraining their own potential for innovation.

Industrial software vendors are, of course, guilty of platform mania by promoting their-logo-alone Industry 4.0 future evangelism. But they are not alone. Many industrial enterprises are equally complicit.

Rightfully fearing vendor platform lock-in, they have gone about building their own in-house data platforms as the centerpieces of digitalization projects.

Such single-supplier, end-to-end digital visions are missing the very crux of what makes today's cloud era transformative.

In the industrial world, building a complex physical asset like an oil platform requires tens of thousands of parts, machines, and systems. These are delivered by a wide variety of vendors. Each of those vendors selected for their product expertise and cost-effectiveness (in maintenance, etc.). Often, those vendors are in direct competition. In other words, the delivery of a complete oil platform is the result of an ecosystem.

The owner of the oil platform obtains the best result by being able to pick and choose during the design and build phase, then piecing everything together.

Digital Industry 4.0 should look just like this. Only by embracing the truly disruptive nature of today's cloud-operating architectures can we realize the fourth industrial age. Together.

It's time for industrial asset operators and others to move beyond endless proof of concepts and relentless omnipotent platform vendor pitches.

In order to deliver intelligent industrial optimization at scale, fitting unique operating environments, the solution architecture must be a hybrid of many best-in-class services and microservices. These will be sourced from a multitude of partners, across a multitude of IaaS, PaaS, and SaaS offerings.

Tomorrow's technology stack will run from ingestion to storage to contextualization to the hundreds (or thousands) of machine learning models and niche industry applications that will power true industry 4.0.

No matter how good they are, no one vendor alone will build—much less maintain—that stack.

It is also not a question of Build vs. Buy. There are at least 50 shades of grey between the two extremes. And business-first IT organizations are rightfully active in pursuing partnerships with open, modular software partners to accelerate their time to market. This is the only way to make continuous innovation and maintenance realistic also over the longer term.

So, the next time you—as an industrial operator—sit across one more omnipotent Industry 4.0 "platform" provider, do yourself a favour. Suggest to that vendor they stop pretending to be the platform with the ecosystem off the shelf.

Then ask how they see themselves working within the broader ecosystem that is true Industry 4.0.

Download our eGuide to learn more about innovation and collaboration in the oil and gas industry.

DOWNLOAD NOW

Data stories

How Cognite helped Hafslund E-CO optimize turbine start-up sequences

Hafslund E-CO and Cognite optimized the start-up sequence of the turbines at one of the company's hydroelectric plants. Using Cognite Data Fusion, Asset Data Insight, and CDF's connectivity delivered:

Hourly savings from avoiding downtime

Potential issues flagged before they cause a problem

Greater ability to stop and start production efficiently

DISCOVER HOW THEY CHANGED THE GAME

Physics and AI hybrid delivers working AI for industry

Forget hybrid cloud, hybrid analytics is the new black.

Until recently, purely data-driven artificial intelligence (AI) — machine learning most notably — has been looked upon as the most attractive technology for enabling new data across industries, including digital twins deployed by heavy asset industries such as oil and gas (O&G). More established, though much less hyped, physics-based modeling has rarely enjoyed the spotlight in recent years.

COGNITE DIFFERENTIATES FROM PURE AI COMPANIES WITH A HYBRID DATA SCIENCE MODEL UNIQUE INDUSTRIAL REALITY

Because of AI's inherent 'black box' nature, however, pure AI-based approaches are failing to gain full acceptance with field operations whose culture is rooted in engineering sciences with zero risk tolerance for critical systems.

In addition, mounting empirical evidence from hundreds of proof-of-concepts involving promising AI startups by O&G industry leaders is debunking the omnipotence of AI to solve production optimization and predictive maintenance use cases as boldly as claimed.

This more informed reality of AI in industry is driving the future of hybrid machine learning, a blend of physics and AI analytics that combines the 'glass box' interpretability and robust mathematical foundation of physics-based modeling with the scalability and pattern recognition capabilities of AI.

Both physics-based models and machine learning (the most common form of AI applications) can be used to make future predictions — so which one to use for what, and when is a hybrid the best solution? The answer depends on the problem you are trying to solve, with the problem classes falling mainly in two categories.

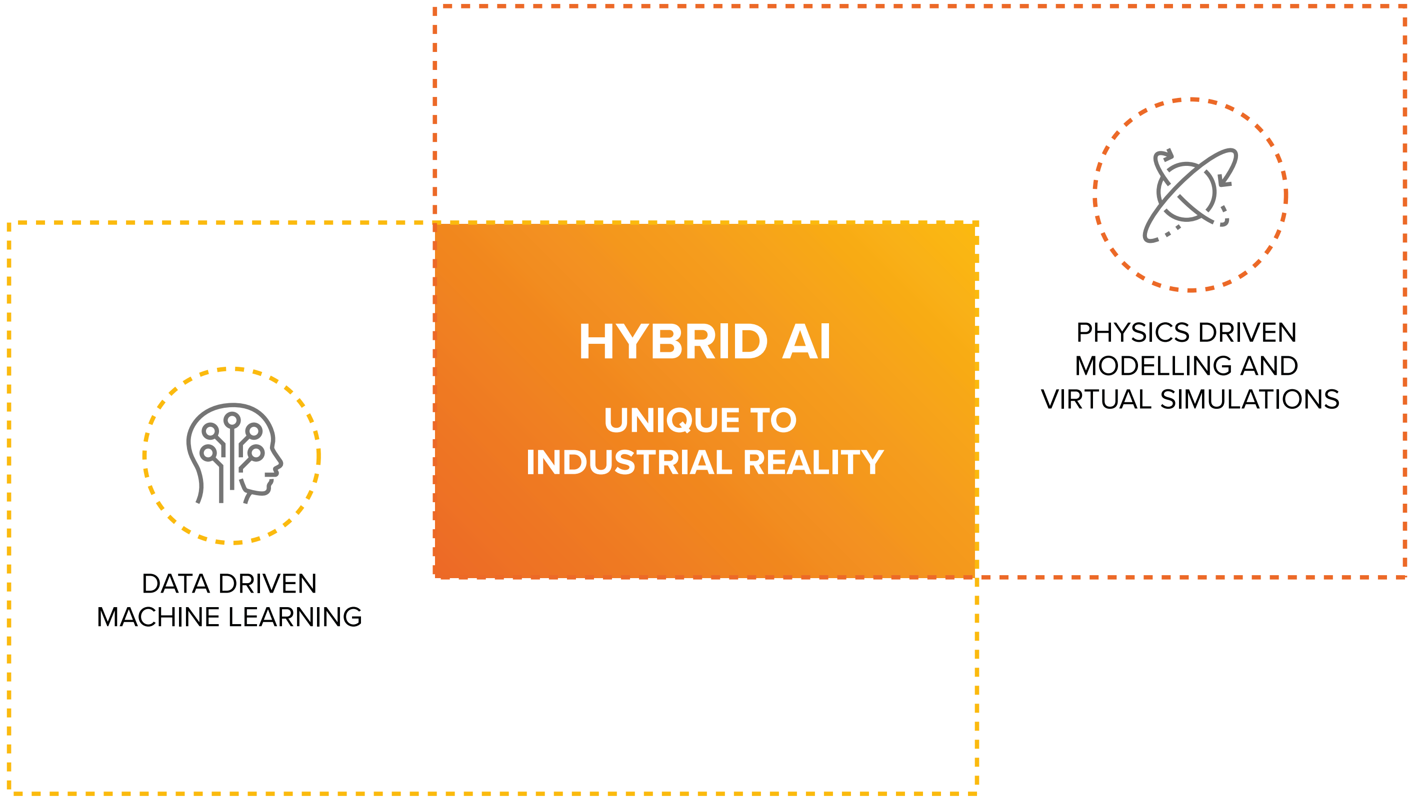

- Systems with lots of experimental data about historical behavior, but no theoretical knowledge framework

- Systems with good mathematical theory framework in place (commonly matched with equally robust empirical behavior data).

One advantage of a physics simulator is that it can predict with a certain confidence even when NO historical data exist. Which means it works from 'first oil' (and it works during the design phase). Historical data is used to increase accuracy and estimate uncertainty

For systems in the first category, a physics-based model is not possible as we are not able to formulate a robust mathematical model to describe the system. Assuming enough contextualized training data is available, a machine learning model should be able to learn any underlying pattern between the system and its outcomes, and ultimately also make predictions.

Two caveats remain, however: the questionable confidence level in resulting predictions (i.e., the precision and recall challenge), and the often absent teaching sample of true failures in critical systems, as traditional scheduled equipment maintenance is designed to prevent such costly failures above all else.

For systems in the second category, a physics-based model can offer a good solution. Physics-based modeling is tried, tested, and validated for even the most critical of simulations. Its most notable limitation is the computational cost of persisting physics-based models in runtime environments with live data, especially across computationally heavy IoT use cases. It is here where hybrid analytics machine learning is offering an attractive solution.

Hybrid models are best suited for complex industrial process problems where a mathematical theory framework exists that can be used to teach a machine learning model that is then used on real-time data for predictions.

The result is a high confidence tailored hybrid model combining strong domain knowledge (physics) with machine learning for cost efficiency and scalability. Especially in the proliferating space of digital twins, hybrid analytics is showing great potential.

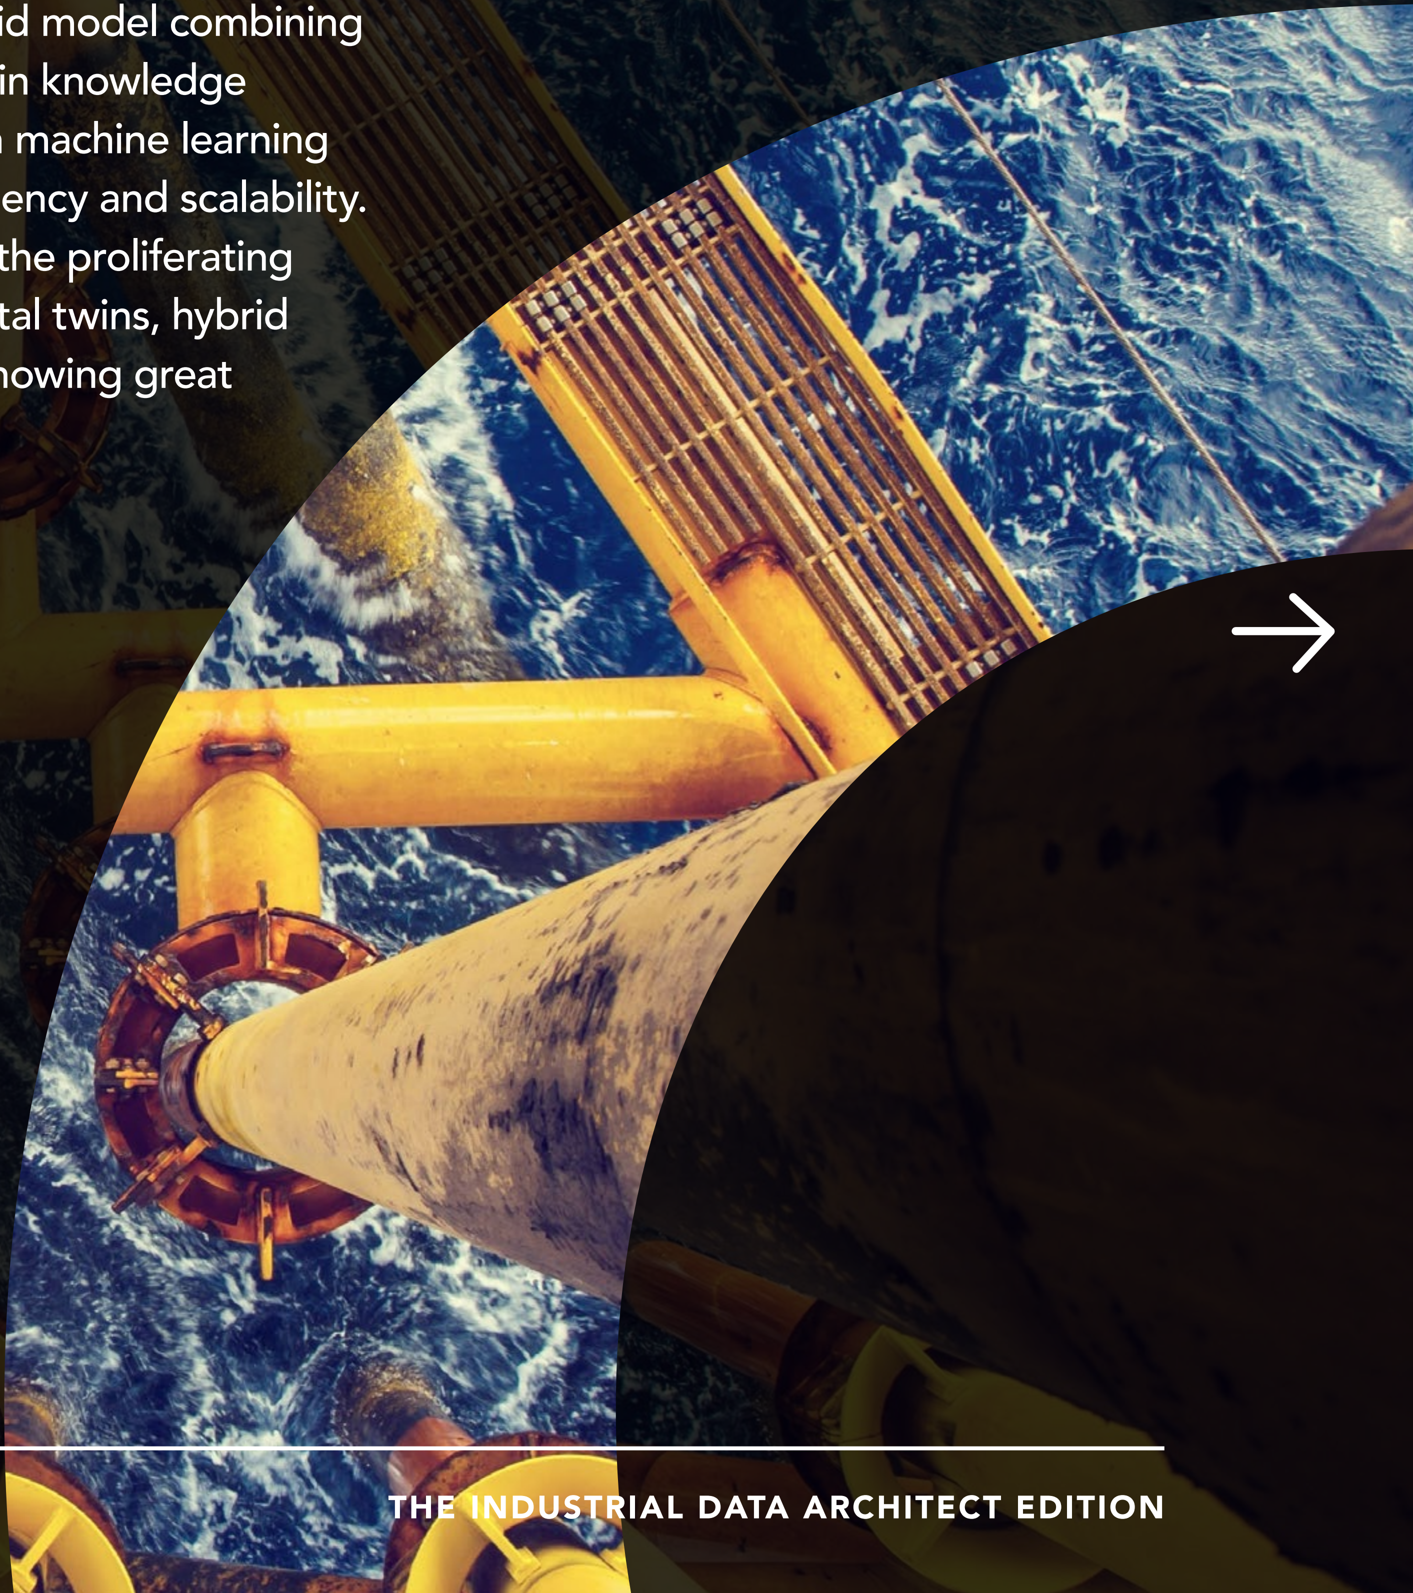

PHYSICS-BASED MODELING

- Tested, tried, proven across industries, and even the most critical applications.
- The uncertainty in the models has been extensively studied and can be taken into account during design and operation.
- Requires a good mathematical theory framework describing the system.
- Can be calibrated and validated with limited experimental data sets.
- Can predict outside the range of the existing data.
- Requires a complete set of boundary conditions for the mathematical equations in addition to information like geometry and other fluid/material properties. This is not always available.
- Commercial simulators are often very expensive.
- Computationally expensive to run at scale in live IoT environments.
- Requires extensive subject matter expertise.
- Difficult to scale across fleets of assets (BC, geometry and fluid/material properties must be changed, but the mathematical models scale. It may not scale with respect to time needed to set up for new asset, but it often scales well with respect to accuracy for different assets).
- Can predict future events.

DATA-DRIVEN AI

- Requires a large contextualized teaching data set, including critical mass of failure events.
- Does not require information like geometry, fluid/material properties and can work on a much smaller set of sensors (but influences accuracy...).
- After initial model training, cost effective to run with real-time streaming data.
- Requires data science expertise.
- Scales very well across fleets of assets (Although not always true: A trained model can only be transferred for very similar assets. A new asset may have different sets of sensors, or on different locations, requiring retraining. A new asset may also be dominated by different physical phenomena requiring different sensors to get reliable predictions -- the physics may not scale).
- Can predict future events (assuming the future event is inside the training set).
- Not interpretable on prediction logic; 'black box'.

HYBRID ANALYTICS

- Combines physics-based modeling and machine learning. Highly suitable for industrial systems analysts across many scenarios.
- Requires both subject matter and data science expertise.
- Offers semi-interpretable prediction logic.
- Cost efficient compared to pure physics-modeling in production and at scale across fleets of assets.
- Can be applied before any historical data exist (first oil).

Are you future-proof?

Find out where your digital strategy is leading, and how to use data to operate with greater precision and efficiency.

Answer a few short questions to gain valuable insights into your data, digital, and operational initiatives, highlighting opportunities to improve your performance.

We'll gauge your capabilities in areas such as:

- Whether you can liberate and consolidate all data in one location.
- Whether you can use data to transform the way you work with partners and suppliers.
- Whether you can use edge analytics to run linear and nonlinear performance analyses on systems.

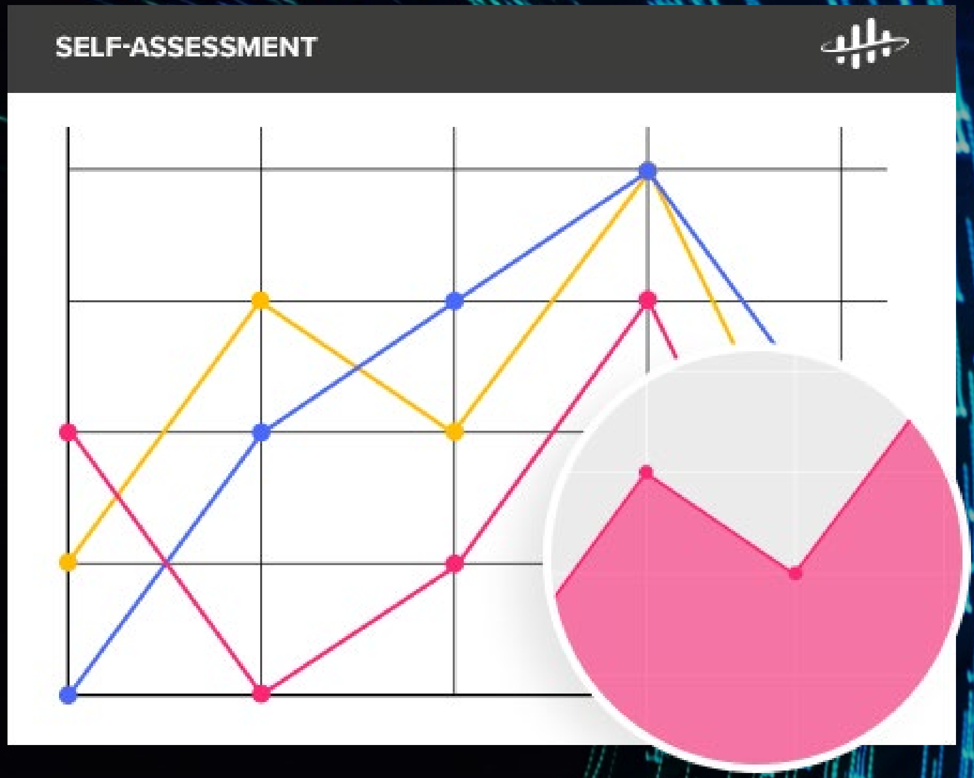

TAKE THE SELF-ASSESSMENT

COGNITE

Cognite is a global industrial AI Software-as-a-Service (SaaS) company supporting the full-scale digital transformation of heavy-asset industries around the world. Our key product, Cognite Data Fusion (CDF), empowers companies with contextualized OT/IT data to drive industrial applications that increase safety, sustainability, and efficiency, and drive revenue.

Learn more at cognite.com

Norway
Fjordpiren,
Laberget 22,
4020 Stavanger

Oksenøyveien 10
1366 Lysaker
Norway
Tel: +47 95 09 21 00
contact@cognite.com

US
Domain Place,
3200 Gracie Kiltz Ln, Ste 200
Austin, TX 78758
Main +1 (512) 593-7120
Toll Free +1 (888) 417-6484
contact@cognite.com

Japan
Shin-Marunouchi Building 10F,
1-5-1 Marunouchi Chiyoda-ku
100-6510
Tokyo, Japan
info-j@cognite.com