verdantix

Industrial Transformation

Green Quadrant: Industrial Data Management Solutions 2025

By Joe Lamming
With Malavika Tohani

January 2025

This version of the report contains Verdantix's summary of Cognite's offerings to help prospective customers evaluate whether the vendor is a good fit for their requirements. It does not contain other vendor profiles.

Industrial Transformation

Green Quadrant: Industrial Data Management Solutions 2025

By Joe Lamming
With Malavika Tohani

January 2025

This report provides a detailed fact-based comparison of the 11 most prominent industrial data management (IDM) software providers in the market. Based on the proprietary Verdantix Green Quadrant methodology, our analysis comprised two-hour live product demonstrations with pre-set scenarios, desk research and vendor responses to a 111-point questionnaire covering 15 technical and nine market momentum categories. Verdantix also conducted interviews with 23 software users and reviewed data from our global survey of 304 industrial transformation decision-makers. Verdantix research finds that the IDM software market is driven by the increasing maturity of AI, the industrial Internet of Things (IIoT) and analytics-focused maintenance strategies, alongside continued demand for sustainability reporting. Vendors are differentiating themselves by offering robust out-of-the-box connectors, industrial DataOps capabilities facilitating unified namespaces (UNS), and even enterprise-scale knowledge graphs. Vendors also leverage partnerships with enterprise data platforms for greater scalability and AI-driven analytics. Among the firms analysed in this study, six providers – ABB, AVEVA, Cognite, Inductive Automation, Palantir and SymphonyAI – demonstrated leading IDM capabilities.

Table of contents

Summary for decision-makers 4
The state of the industrial data management (IDM) software market 5
Analytics for specific use cases will drive spend on data management solutions in 2025
Industrial DataOps requires a very particular set of skills
Growth in spend on IDM solutions stems from a renewed obsession with industrial AI readiness 12
While most facilities are still barely at Industry 3.0, the steps to 4.0 and beyond are easier than ever
LLM-driven unstructured data management is the next practical application for industrial AI
Challenges with AI model reasoning performance and reliability continue to limit their deployment outside carefully crafted IDM use cases
Green Quadrant for IDM software 2025 18
Green Quadrant methodology
Scope and methodology for the 2025 Green Quadrant IDM software study
Evaluated firms and selection criteria
Evaluation criteria
Cognite overview 28
Cognite provides a market-leading DataOps and AI platform for enterprise-scale, complex IDM projects

Table of figures

Figure 1. Industrial transformation leaders must choose between three distinct IDM approaches in 2025 6
Figure 2-1. Plant digital transformation spending increases are in part driven by the maturity of IIoT and AI solutions 8
Figure 2-2. Plant digital transformation spending increases are in part driven by the maturity of IIoT and AI solutions 8
Figure 3. Sensor monitoring was used for asset maintenance at the majority of industrial firms in 2024 9
Figure 4: Discrete manufacturing and oil and gas are the most digitally mature industries 11
Figure 5: Practitioners see automated prioritization, production optimization and data unification as the top high-value use cases for AI 12
Figure 6: Lack of awareness, trust and data quality issues hobble AI projects 13
Figure 7. The IDM vendor landscape 15
Figure 8. Technical capabilities criteria for IDM software 22
Figure 9. Momentum criteria for IDM software 24
Figure 10. Vendor category scores (capabilities) 25
Figure 11. Vendor category scores (momentum) 26
Figure 12. Green Quadrant for IDM software 2025 27

Organizations mentioned

4IR, ABB, Accel, Accenture, Airbus, Aker ASA, Aker BP, Aker Solutions, Alcon, Alibaba, Altair, Amazon Web Services (AWS), Anthropic, ANYbotics, Apache NiFi, Apache Software Foundation, APERIO, Apple, Artek Integrated Solutions, Asian Paints, AspenTech, Augury, Autodesk, AVEVA, B&L Information Systems, BASF, Bayer, Belden, Bentley Systems, Blue Yonder, Boston Dynamics, BP, Braincube, C3 AI, Canary Labs, Capgemini, Cargill, Catalent, Celanese, Cerence, Cirrus Link, Citibank, Coca-Cola, Cognite, Cognizant, Cosmo Oil Company, Crosser, Cursor, Cybus, Databricks, Dataiku, DataRobot, Dell, Deloitte, Drax, DXC Technology, Eclipse Mosquitto, Emerson, Enablon, ENGIE, Essex Furukawa, ETQ, Exposition Ventures, ExxonMobil, EY, Factry, Falkonry, FANUC, Ferroloy, FlowFuse, Fluid Framework, Fluke Reliability, Forbes, Fortune, FTI, General Mills, George Reed, Georgia-Pacific, GE Vernova, GitHub, Google, Gousto, Grafana, gRPC, Helios Solutions, Hemlock Semiconductor, Hexagon, HighByte, Hitachi Energy, HiveMQ, Honeywell, IFS, Inductive Automation, InfluxData, Infor EAM, Infosys, Intergraph, International Organization for Standardization (ISO), Intertape Polymer Group, Invensys, Itus Digital, j5, Jacobs, Johnson & Johnson, Johnson Controls, Jupyter, Keras, Klarna, KRONE, Kubeflow, Kubernetes, Lear Corporation, Leica Geosystems, Litmus Automation, MachineMetrics, MaestroHub, Maibornwolff, Maine Venture Fund, Meta, Microsoft, Miele, Mistral, Mitsubishi, Moderna, MPDV, MSC Software, MySQL, Nanoprecise, Nippon Gases, NIPPON SHOKUBAI, Node-RED, Northrop Grumman, NTT Data, NVIDIA, Oden Technologies, Omny, OpenAI, OpenText, Opto 22, Oracle, Palantir, PAS Global, PayPal, Perstorp, Pfizer, Pinnacle, Poclain, Porsche, PTC, PwC, Python, Quartic, Radix, Richemont, Rockwell Automation, ROI-EFESO, Salesforce, SANY, SAP, Saudi Aramco, Schneider Electric, SCHUNK, Seeq, Sepasoft, Shell, Siemens, Sift, Sight Machine, SKF Group, Skkynet, SLB, Snowflake, Standard Investments, SymphonyAI, Tata Consultancy Services (TCS), TCV, TeamViewer, Teradata, Tesla, Timeseer.AI, Transition Technologies PSC (TT PSC), TrendMiner, TSMC (Taiwan Semiconductor Manufacturing Company), TwinThread, Tyson Foods, UK National Health Service (NHS), UMH, UptimeAI, Vale Mining, Vandemoortele, Viewport.ai, Volkswagen, Vulcan Energy Resources, Wipro, Wood, Yokogawa.

Disclaimer

As an independent analyst firm, Verdantix does not endorse any vendor, product or service covered in our research publications, webinars and other materials. Verdantix does not advise technology users to select only those vendors with the highest ratings. Verdantix research publications consist of the opinions of the Verdantix research team based on its analysis of the market, survey data and review of vendor solutions. Verdantix disclaims all warranties, expressed or implied, with respect to this research, including any warranties of fitness for a particular purpose.

Summary for decision-makers

- Operations, maintenance, data science and technology decision-makers responsible for the procurement and management of IDM software should use this report to validate existing strategies or complete a vendor selection process. IDM software suppliers should use it to benchmark capabilities and roadmaps against competitors.
- The report leverages data from two-hour demonstrations, a 111-point questionnaire and desk research to form a view of the IDM software market. We also spoke to 23 IDM solution users to understand buyer demands and challenges.
- The IDM market is segmented into on-premises solutions for supporting basic analytics, hybrid edge-to-cloud architectures for industrial DataOps, and enterprise-scale DataOps ontologies for IT-OT-ET convergence. Growth in IDM spending is driven by the increasing maturity of AI analytics and machine connectivity, sustainability reporting requirements and the need for analytics-driven maintenance strategies to address worker shortages. Advanced IDM solutions are helping users build UNS for scalable data models, integrate unstructured data through document-understanding AI, and achieve enterprise-scale data unification via knowledge graphs and ontologies.
- Of the 11 vendors evaluated in this Green Quadrant, six emerged as leaders: ABB, AVEVA, Cognite, Inductive Automation, Palantir and SymphonyAI.

Figure 12

Green Quadrant for IDM software 2025

Source: Verdantix analysis

The state of the industrial data management (IDM) software market

Industrial data today comprise far more than just timestamped sensor measurements and event logs. Industrial control systems – manufacturing execution systems (MESs) and programmable logic controllers (PLCs) – historically operated in siloed systems, with human-machine interfaces (HMIs) enabling operators to manually adjust set points and start up and shut down equipment. These systems served as a bridge between operational processes and broader business objectives set by facilities managers. Over time, operational technology (OT) protocols matured. Simultaneous control and data acquisition (SCADA) systems evolved from basic control and data collection tools into platforms capable of orchestrating and automating complex, multi-step industrial processes.

Verdantix sees the industrial data management (IDM) software market today divided into three distinct segments (see **Figure 1**). The first focuses on connecting structured OT data – such as time series, events and control logs – to basic analytics. This is typically achieved through process historians and visualization tools, offering straightforward, on-premises, local-network data hub solutions with minimal cybersecurity overhead. The second segment represents a shift toward hybrid edge-to-cloud architectures, with data curated and contextualized via lightweight, purpose-built software and hardware. Solutions in this category align OT data with business operations (industrial DataOps) through unified namespaces (UNS) and hierarchical models, with ample scope for more sophisticated ML (machine learning)-driven analytics. The third segment targets global, enterprise-scale deployments, leveraging a DataOps ontology – rather than just a hierarchy – to unify OT, engineering technology (ET) and information technology (IT) data. Such complex undertakings are often driven by a need to make fundamentally distinct systems discoverable and analysable. Such an ontology – the industrial knowledge graph – delivers interconnected, domain-adapted data models that allow buyers to make decisions and build automations within an enterprise-wide data digital twin.

Buyers for all three are drastically different. The first are looking to simply provide visibility into a specific process, contextualized by trends over time, while the second are addressing operational excellence initiatives such as asset performance management (APM), production optimization and energy efficiency. Buyers for the third cohort want a tailored composition of solutions to deliver enterprise-wide digitization, implemented over months – and even years – to configure custom connectors, build scalable data models and train staff to maximize its utility.

Given the often blurred lines between the three IDM segments – and the relative novelty, complexity and IT-heavy origins of the third – this report provides individuals who are responsible for selecting, implementing and deriving value from data-driven initiatives with a detailed assessment of 11 prominent IDM solution providers and their product offerings. The customer questions answered by this report include:

- **What is the current state of the IDM software market?**
- **Which IDM software applications lead the market?**
- **Which IDM software applications will best match the requirements of my firm?**
- **How can I benchmark the capabilities of IDM software applications?**
- **What factors indicate that an IDM software vendor is a reliable partner for the future?**

To answer these questions, Verdantix assessed 11 suppliers using a 111-point questionnaire, two-hour live demonstrations and interviews with 23 existing IDM software customers. The resulting analysis is based on the proprietary Verdantix Green Quadrant methodology, which is designed to provide an evidence-based, objective assessment of vendors offering comparable products or services.

Figure 1

Industrial transformation leaders must choose between three distinct IDM approaches in 2025

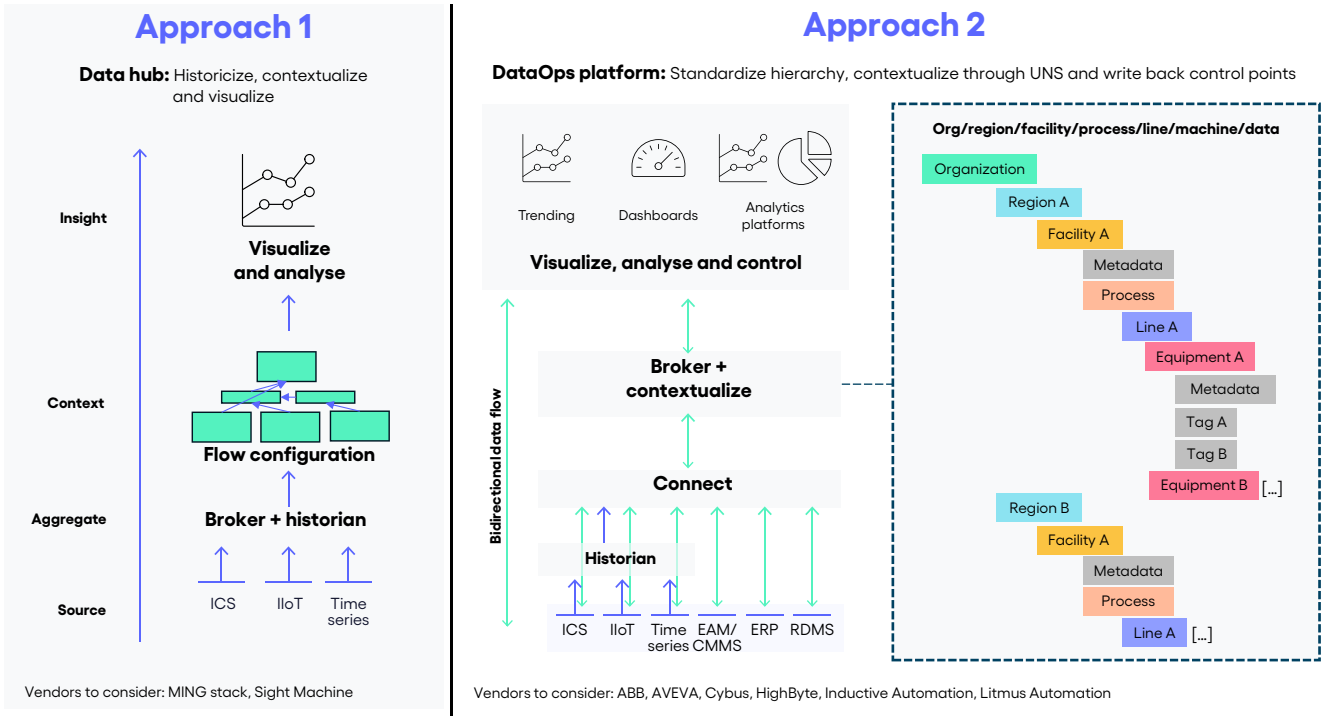

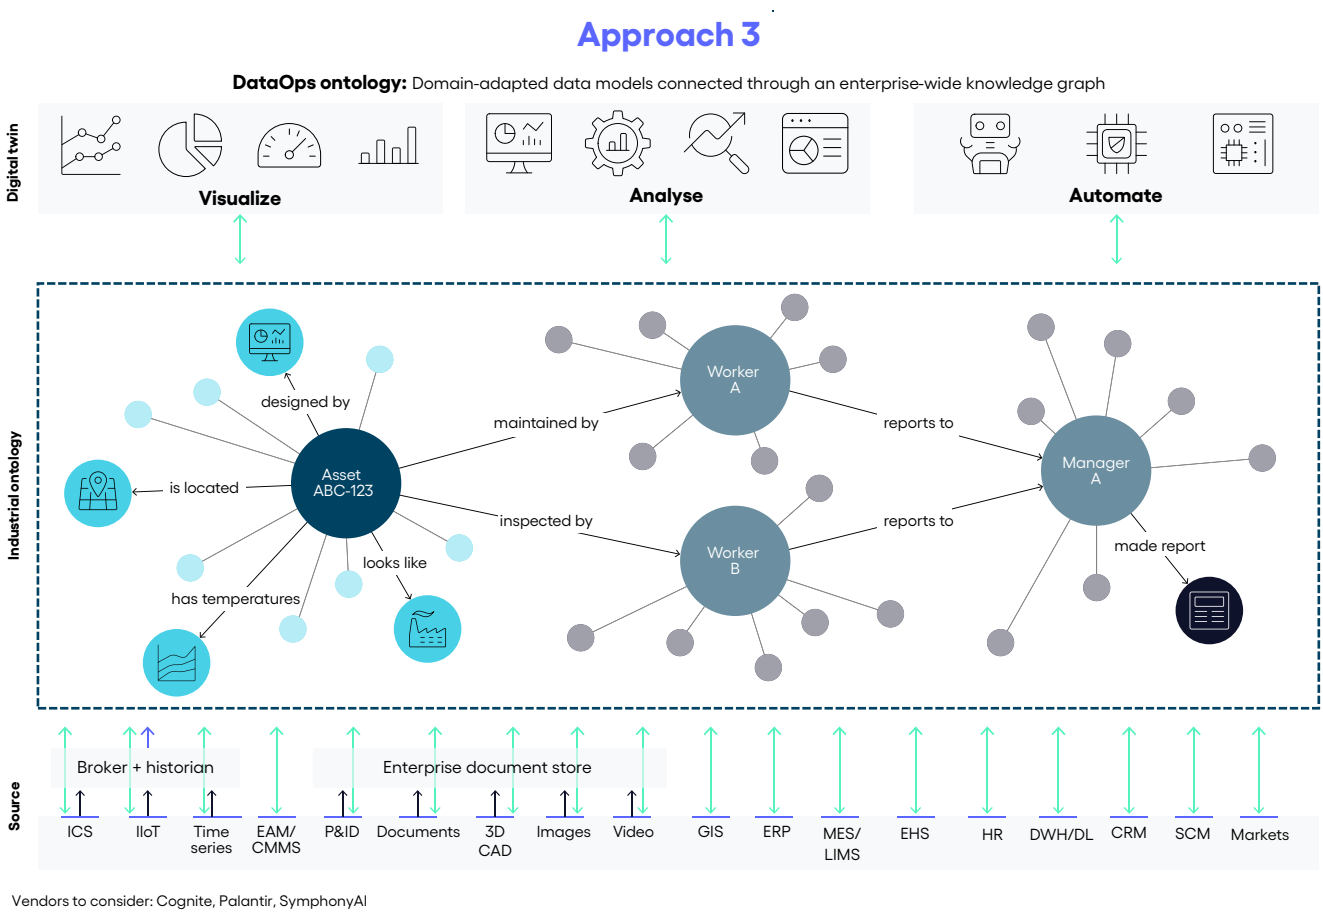

Source: Verdantix analysis

Analytics for specific use cases will drive spend on data management solutions in 2025

Today's operations, maintenance, engineering, process safety and industrial IT practitioners still largely rely on experience, insights from colleagues and basic dashboards, spreadsheets and even hour-by-hour boards to harmonize production. These approaches, while field-hardened and familiar, begin to fall short in addressing the scale – and emergent complexity – of modern industrial systems, as increasingly analytical management teams push for convergence with IT. Without such convergence, critical insights are siloed and decision-making reactive. Verdantix undertakes annual global corporate surveys with such industrial decision-makers to understand budgets, priorities and technology preferences. Our 2024 survey results suggest that growth in the IDM software market is driven by:

- **Increasing availability and maturity of AI, analytics and IIoT.**
  The industrial internet of things (IIoT) today owes much of its interoperability and divergence from the web-focused HTTP protocol to Arlen Nipper, who pioneered message queue telemetry transport (MQTT) from about 1999. This protocol, alongside OPC-UA, quickly saw uptake by disruptor IDM software vendors such as Inductive Automation from the mid-2000s, utilizing MQTT brokers' ability to reliably orchestrate numerous sensors publishing updates concurrently on bandwidth-limited networks. Recording to process historians, industrial firms, now time-series-rich, contributed to the growth in ML-driven analytics from the late 2010s. The Verdantix 2024 survey heard that 88% of respondents agree that the increasing maturity of such technology is driving plant digitization – with 46% planning single- or double-digit increases in spend (see **Figure 2-1** and **Figure 2-2**). Verdantix finds that the industrial AI analytics software market will grow from $1.7 billion in 2023 to $5.0 billion in 2028, at a CAGR of 23.9% (see Verdantix Market Size And Forecast: Industrial AI Analytics 2022-2028 (Global)).

- **Reconfiguring of O&M metrics to align with new sustainability reporting.**
  Despite regulatory uncertainty surrounding the return of the Trump administration in the US in 2025, sustainability reporting remains a critical focus for industrial firms, driven by global regulatory frameworks such as the EU Corporate Sustainability Reporting Directive (CSRD), which will impact over 3,000 US-based firms by 2028 (see Verdantix Market Size And Forecast: ESG Reporting Software Solutions 2021-2027 (Global)). The 2024 Verdantix industrial transformation survey shows that nearly two-fifths (36%) of 304 firms rank aligning operations and maintenance (O&M) metrics with sustainability frameworks as a high priority. Investor-grade ESG data are also increasingly essential as firms face rising demands from stakeholders and regulations such as the Uyghur Forced Labor Prevention Act (UFLPA) and the EU Corporate Sustainability Due Diligence Directive (CSDDD). With ESG software projected to grow from $0.9 billion in 2021 to $4.3 billion by 2027, firms are leveraging digital tools to meet reporting demands, streamline data-gathering from complex supply chains and minimize related risks.

- **Majority of firms choosing analytics-focused maintenance strategies.**
  The Verdantix industrial transformation survey heard 58% of respondents describe worker shortages due to aging and retiring workforces as a 'significant' or 'very significant' driver for the digital transformation of O&M. The same was true for 82% of those asked about the lack of availability of skilled workers. As a result, data-driven maintenance approaches, such as condition monitoring, predictive maintenance and reliability-centred maintenance (RCM), are increasingly chosen by industrial firms (see **Figure 3**). This is particularly pronounced in oil and gas, where 82% of respondents attest to using data gathered from asset-mounted sensors to deliver condition monitoring for at least critical assets. The IDM software market is expected to be worth $6.1 billion by 2029, growing at a CAGR of 19.9% – driven by increasing awareness of the return on investment (ROI) delivered by AI-focused analytics from industrial data (see Verdantix Market Size And Forecast: Industrial Data Management Software 2023-2029 (Global)).

Figure 2-1

Plant digital transformation spending increases are in part driven by the maturity of IIoT and AI solutions

How significant is the availability and increasing maturity of AI and IIoT in driving plant digitization?

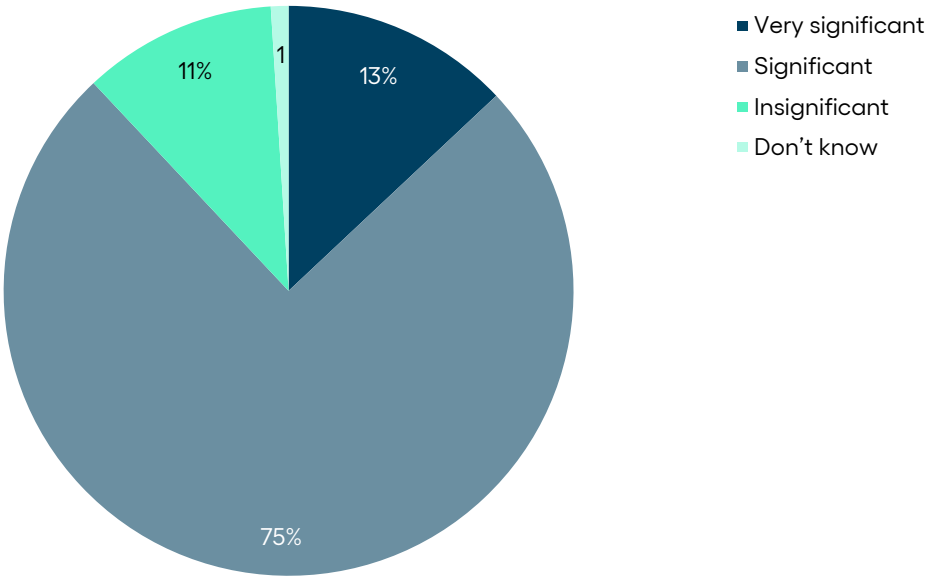

Note: Data labels are subject to rounding.
Source: Verdantix Industrial Transformation Global Corporate Survey 2024

N=304

Figure 2-2

Plant digital transformation spending increases are in part driven by the maturity of IIoT and AI solutions

How is your firm's spend on the following operational excellence initiatives expected to change in the next 12 months?

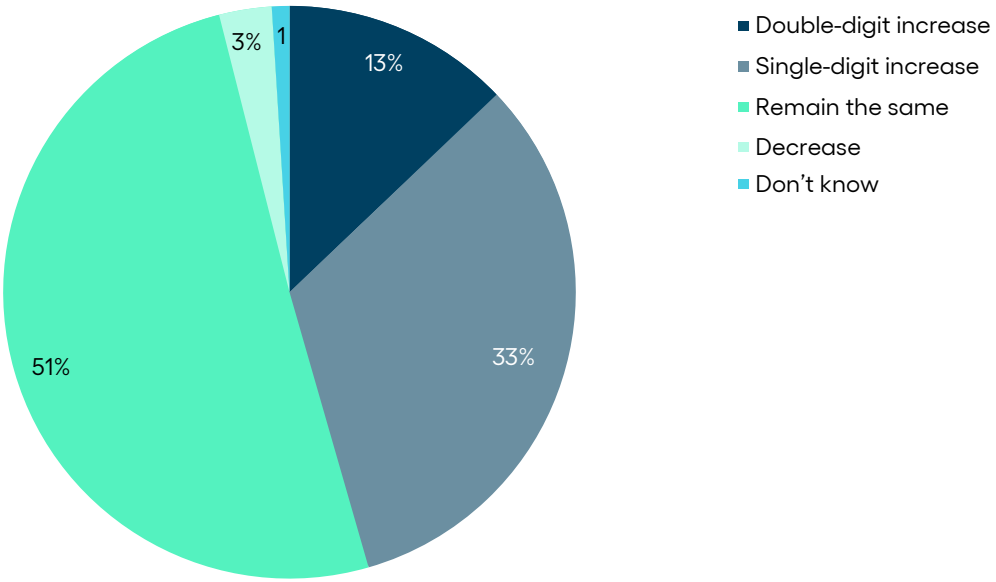

Note: Data labels are subject to rounding.
Source: Verdantix Industrial Transformation Global Corporate Survey 2024

N=304

Figure 3

Sensor monitoring was used for asset maintenance at the majority of industrial firms in 2024

Which of the following maintenance approaches do you currently follow at your firm?

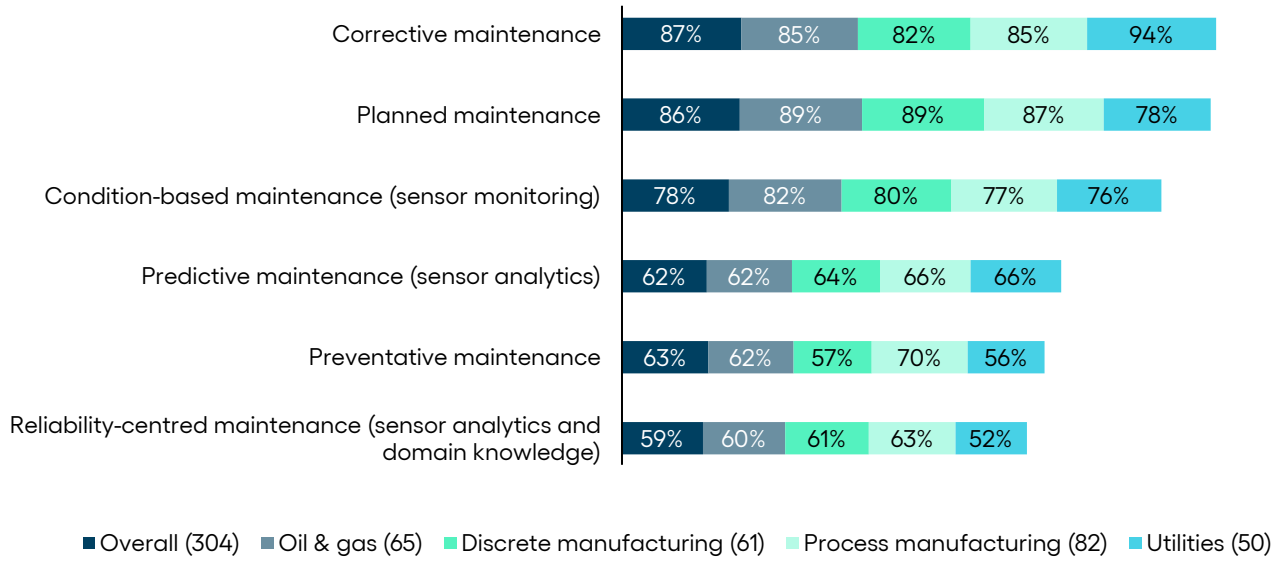

Note: Data labels are subject to rounding.

Source: Verdantix Industrial Transformation Global Corporate Survey 2024

N=304

Industrial DataOps requires a very particular set of skills

Digital transformation can no longer be described as an 'aspiration' for today's industrial facilities. The majority are already planning double-digit spending increases on operational excellence initiatives, requiring deeper visibility into both operations and maintenance. Others see the need for accurate, real-time data from across their operations as vital to meeting regulatory and reporting requirements. As such, the IDM software market is expected to see rapid growth – driven in part by a perception that analytics and AI are a compelling investment requiring a robust data foundation. Based on our analysis of IDM software market trends and vendor strategies over the previous two years, Verdantix finds that:

- **Bridging OT and IT requires more than just APIs.**
  Low-latency processing is fundamental to industrial automation solutions. Adjusting control points, raising alarms and responding to start-up and shutdown commands are safety-critical functionality aspects that cannot rely on potentially interruptible connections to the global internet. Industrial control systems naturally evolve more cautiously than the 'move fast and break things' approach pioneered by web technology firms from the early 2000s. As such, OT data sources such as asset PLCs and SCADA systems are almost exclusively deployed on secure, local, wired networks employing specific protocols such as MQTT, Modbus and OPC-UA. Firms such as Cybus, HighByte and Litmus Automation offer edge-deployed software solutions to act as highly scalable, reliable data brokers for these sources, ingesting OT data and building payloads of related data. This allows users to send these payloads securely to cloud services such as Amazon Web Services (AWS), Google Cloud and Microsoft Azure for multi-facility analytics.

- **Enterprise data platforms are partnering with incumbent IDM providers.**
  The 1980s witnessed the emergence of data warehousing, led by firms such as IBM and Teradata. Improvements in reliability and high performance at vast scales continued into the 2000s, with projects such as Hadoop released by the Apache Software Foundation. Improvements in ML, especially for image and natural language processing (NLP), meant that data lake analytics platforms – encompassing more than just structured data – began with vendors such as Databricks and Snowflake. Incumbent IDM software providers, evolving in parallel within their aforementioned industrial environment constraints, came under pressure from users demanding the scalability, cost efficiency and MLOps features afforded by cloud analytics. AVEVA announced a partnership with Databricks in April 2024, aiming to combine its own time series and industrial analytics capabilities with Databricks's large-scale trend analysis tools, to improve fuel efficiency for customers such as Drax.

- **Unified namespaces are making industrial data more self-explaining.**
  Point-to-point communication systems are less scalable than those utilizing a central point of trust – a central data broker. However, such centralization of data orchestration requires standardization. In the late 1990s, ISA-95 evolved as a standard for communication between OT systems such as the MES, and IT systems such as enterprise resource planning (ERP) systems. UNS architecture built upon ISA-95 in response to the proliferation of IIoT, enforcing not just consistency in how data sources are referenced, but encouraging the use of OT-focused, lightweight data transfer protocols such as MQTT. Tags for data sources encode the data hierarchy within the name itself, and require that all tags follow the same naming convention. The data hub architecture eliminates the layer-by-layer data transfer regime often followed by ISA-95 and enables buyers to build and quickly iterate on ML-driven analytics projects covering multiple assets, processes and even facilities. As such, today's IDM solution providers, such as Cybus, HighByte, Inductive Automation and Litmus, are racing to offer ever-more-complete UNS enablers as part of their portfolios, making it easier for buyers to adopt – technically and culturally – consistent data modelling.

- **Long-running time series data sets are largely unique to asset-heavy industries.**
  The polling of sensors – for temperature, acceleration (vibration) and electromagnetic field strength – at fixed time intervals is common in lab environments. Similarly, web data collected by consumer- and enterprise-facing technology firms are largely event-driven and metadata-heavy. Industrial facilities, however, utilize time series data stretching back decades as part of automation systems, condition monitoring and even predictive maintenance. Industrial automation and software provider Siemens boasts of having tens of billions of time series data points available to it, while SymphonyAI released its Industrial LLM in November 2023, trained on 20 years of time series sensor, events and other asset data. Similarly, AVEVA's PI System takes in more than two billion continuously polled time series data streams installed at nearly 30,000 industrial sites. The ubiquity of sensors and prevalence of process historians within asset-heavy industries – alongside the ease of compression of time series data to manage storage costs – mean that such data remain a dominant source to be managed by IDM solutions.

- **White-labelled IDM solutions are pervasive in the industrial software market.**
  Similarly to how enterprise-focused cloud analytics software vendors choose to partner with industrial software vendors rather than fully bridge the IT-OT divide, hyperscalers opt for partnerships with OT-focused edge vendors. Google worked with Litmus Automation to develop Google Manufacturing Connect, to tap into its library of more than 250 machine protocols, translate these into cloud-ready payloads and run hybrid cloud visualization and low-latency ML inferencing for bidirectional control. Similarly, Rockwell Automation worked with Cognite to develop FactoryTalk DataMosaix, built on the industrial knowledge graph solution Cognite Data Fusion (CDF). Such software distribution strategies help smaller vendors by increasing reach, while offering general-purpose hyperscalers the ability to focus on more fundamental challenges, such as multi-region cloud resilience and cost efficiency.

- **High operational complexity industries have mostly bridged IT, OT and ET.**
  ET data – comprising technical documentation, 3D models and even point clouds – have long been served by software vendors such as Autodesk, Bentley Systems and Hexagon, catering to engineering, procurement and construction (EPC) customers. Engineering document management systems offer version control, approval workflows and even mark-up capabilities for collaboration (see Verdantix Best Practices: Six Top Strategies For Seamless Construction-To-Operations Data Transfer). Handover of such data from EPC to owner-operator, however, depends on the engineering software stack compatibility of both parties – often defaulting to interoperable formats such as PDFs and spreadsheets. In the 2024 Verdantix industrial transformation survey, half of the firms in industries with greater variety within day-to-day operations, such as oil and gas and discrete manufacturing, attested to managing OT, ET and IT data with robust governance to support analytics. A further 8% stated that their OT, ET and IT data for plant operations were fully connected, delivering automation and optimization using AI-driven analytics (see **Figure 4**).

Figure 4

Discrete manufacturing and oil and gas are the most digitally mature industries

Which statement best describes your firm's current data management strategy for intelligent plant operations?

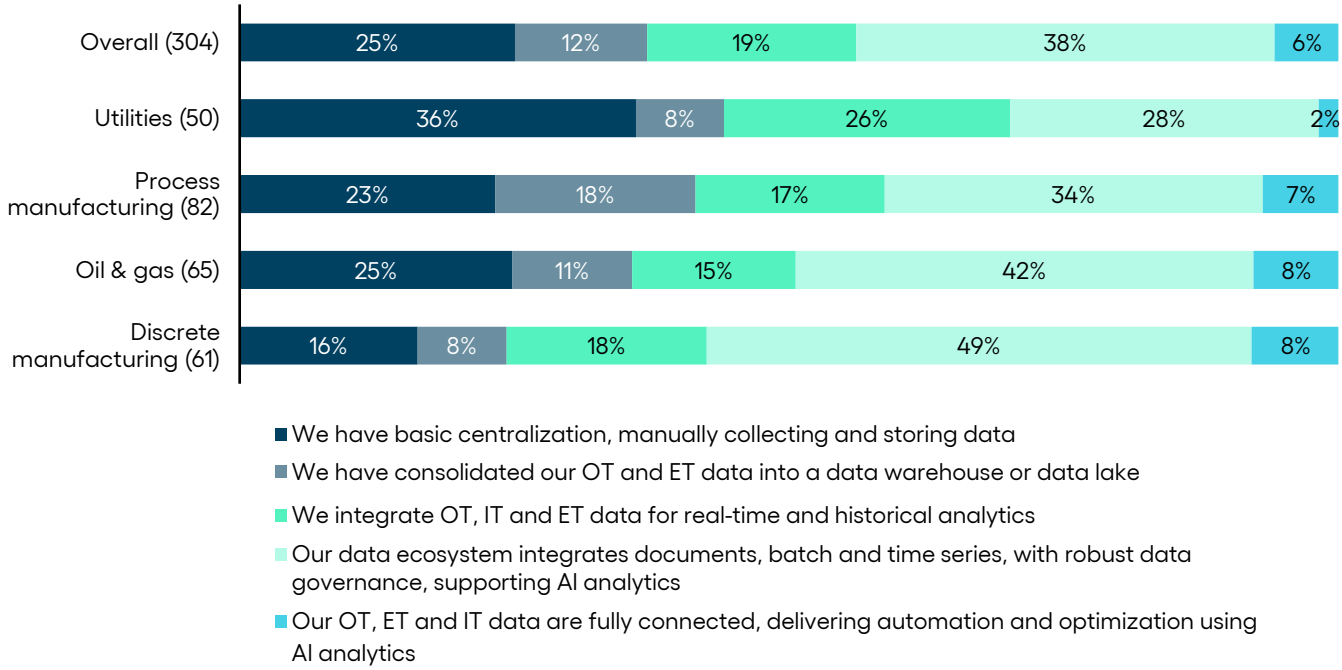

Note: Data labels are subject to rounding.
Source: Verdantix Industrial Transformation Global Corporate Survey 2024

N=304

Growth in spend on IDM solutions stems from a renewed obsession with industrial AI readiness

Industrial and adjacent enterprise data connectivity is a key capability delivered by pre-built connectors. Robust software around such connectors offers more than just application programming interface (API) layers. It allows buyers to make data interoperable, assemble that data with metadata and hand-picked related data, and pipe the end-result towards valuable dashboards, analytics and automation. Analytics, and its more recent variant – AI analytics – are seen by Verdantix as a busy market with huge variation in approaches to harnessing time series data and beyond (see Verdantix Market Overview: Industrial AI Analytics Solutions). More advanced AI systems, such as those using Transformer-based large language models (LLMs), have spent a full two years in the spotlight. As such, Verdantix survey data find that 58% of practitioners rank the automated prioritization of operational issues as their top high-value use case for AI. Conversely, expectations for other use cases, such as GenAI copilots for extracting operational data from existing data sets, are more muted (see **Figure 5**). As for the most significant barriers to industrial AI projects, 72% of respondents cite a lack of awareness of AI analytics solution vendors. Similarly, 68% lack trust in AI-derived insights – with 75% seeing poor or incomplete data quality as a 'significant' or 'very significant' barrier (see **Figure 6**).

Figure 5

Practitioners see automated prioritization, production optimization and data unification as the top high-value use cases for AI

Rank the top three high-value use cases of AI for your firm from the following.

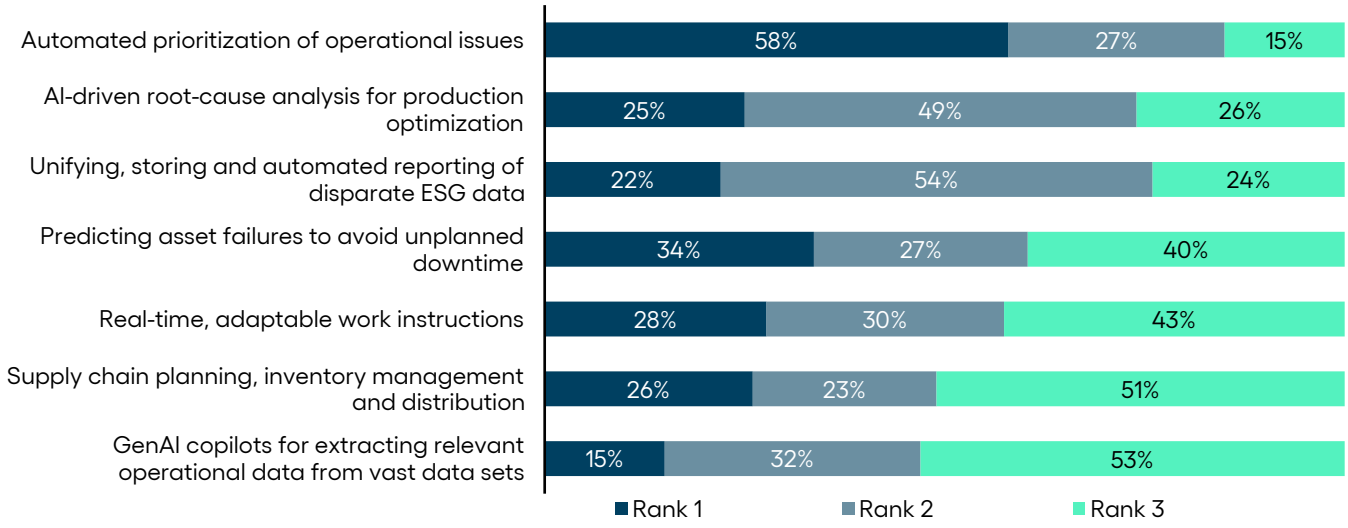

Note: Data labels are subject to rounding.
Source: Verdantix Industrial Transformation Global Corporate Survey 2024

N=304

Figure 6

Lack of awareness, trust and data quality issues hobble AI projects

How significant are the following barriers in impacting the success of your industrial AI analytics projects?

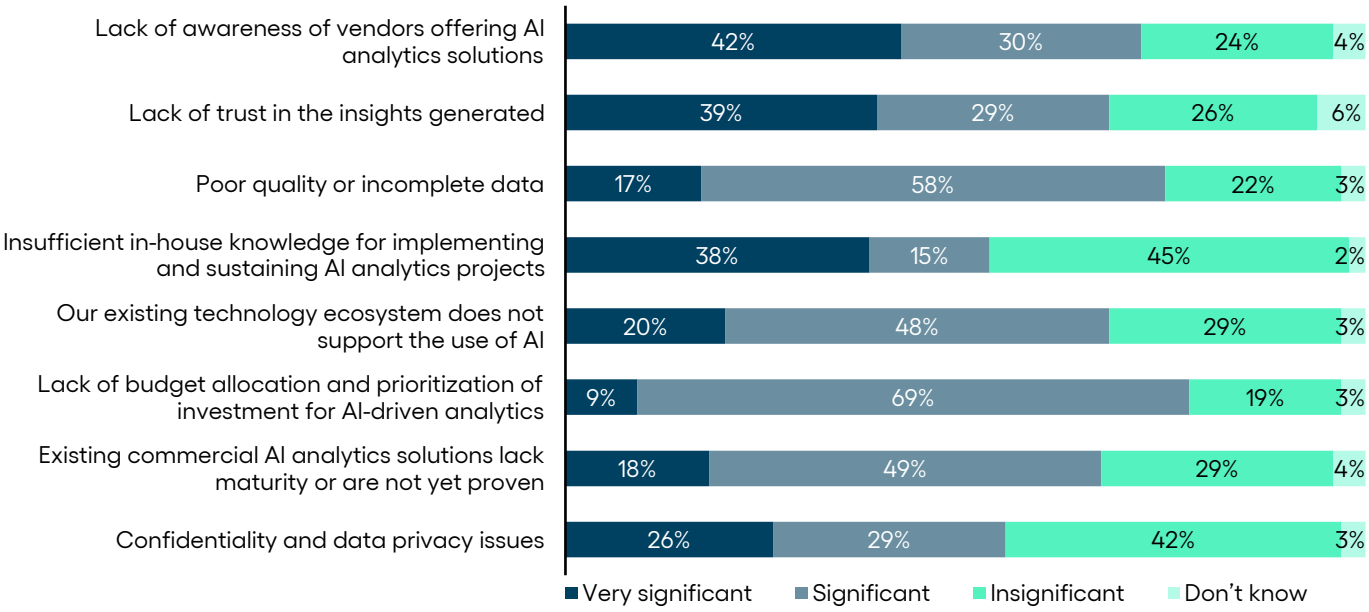

Note: Data labels are subject to rounding.

Source: Verdantix Industrial Transformation Global Corporate Survey 2024

N=304

While most facilities are still barely at Industry 3.0, the steps to 4.0 and beyond are easier than ever

Industrial and enterprise software firms long ago jumped on the AI bandwagon, attempting to capitalize on investor and buyer interest. However, when it comes to the implementation of effective solutions, the reality is that in 2024 nearly 90% of industrial firms admit that some assets still receive only corrective and planned maintenance (see **Figure 2**). Technical, cost and cultural barriers mean that the current state of digitization is insufficient for most buyers to deploy enterprise-scale knowledge graphs or automate back-office operations with AI agents. Nevertheless, Verdantix continues to witness an explosion of open-source, established, commercial, low-cost and powerful start-up offerings. This makes the leap to full IT-OT connectivity – laying the foundation for AI – a feasible goal for buyers in 2025, as:

- **Modern IDM solutions facilitate both metadata and human factors.**
  Metadata – that is, data about other data – are the most common form of context in IT, OT and ET data systems. Creation date, data type, owner and date modified fields persist across most file formats, providing the context needed for human and machine users to interpret and use data to make decisions. Oden Technologies allows users to explore time series data plots and add metadata in the form of annotations – complete with pre-defined categories and free text notes – associated within its low-code Explore interface to the related asset within its data model. Similarly, the success of any IDM strategy relies on user uptake. Forty-six per cent of respondents to the Verdantix 2024 industrial transformation survey see improving collaboration between process safety and engineering and operations teams as a 'high' or 'top' priority in the year ahead. Ease-of-use, training and the presence of domain experts – especially at the data ingest and data analysis ends of the IDM stack – are critical to successful data-driven projects.

- **SCADA and MES solutions increasingly deliver off-the-shelf industrial DataOps.**
  Since the early 2000s, traditional SCADA solutions have been evolving beyond their traditional roles of monitoring and control. Today, they aggregate data from PLCs, but also – via MQTT brokers or OPC-UA servers – IIoT devices and even enterprise systems such as ERP and MES. Inductive Automation's Ignition, launched in 2003, has a significant customer base, especially in North America, thanks to its server-based unlimited tag and connector licensing model, offering deeply configurable data flow between edge devices, HMIs, historians and even cloud services. By integrating real-time data collection with user interface (UI)-configurable contextualization and even bring-your-own model analytics, today's manufacturing-focused IDM solutions offer solid DataOps for individual facilities.

- **Open-source architectures such as the MING stack are cheap and effective for pilot projects.**
  A simple first step in operationalizing IDM involves connecting a handful of time series data sources, such as temperature sensors and alarm tags, to a historian, and visualizing the data through a trending tool (see **Figure 1**). Today, a wide variety of solutions are available to customers to buy and implement. Connectivity is enabled through open-source MQTT brokers such as Eclipse Mosquitto or HiveMQ's Community Edition. For the historian, commercial solutions such as AspenTech IP.21, AVEVA PI System, Canary Labs and Factry offer technical support, while open-source alternatives such as InfluxDB present more developer-focused, customizable data historicization. Node-RED, launched in 2013 as an open-source project from IBM's Emerging Technology Services group, offers flow-based, low-code connectivity for small, local-network MQTT devices. For visualization, another open-source project, Grafana, was launched in 2014. This combination – MQTT, InfluxDB, Node-RED and Grafana – is immensely popular in the IIoT community, although issues relating to cyber security, concurrent performance and a lack of data model enforcement limit scalability and enterprise robustness. Regardless, the MING stack presents a compelling first step for firms looking to begin their IDM journey and provide simple trending and dashboards.

- **Commercial OT-focused industrial DataOps solutions provide easy GUIs.**
  For buyers managing data from a mix of legacy equipment, and even facilities acquired through corporate mergers, taking the time to implement a UNS provides scalable structure to data. Tools such as Cybus Factory Data Hub, HighByte Namespaces, Inductive Automation's Ignition and Litmus UNS help enforce consistent data models. Once established, the onboarding of new equipment requires systems integrators to simply follow the naming convention. However, to use this newly standardized data model, users must create data pipelines that deliver industrial DataOps – data quality management, contextualization and discoverability. Crosser's Flow Studio offers a robust, scalable alternative to Node-RED for users to connect disparate data sources in a low-code interface. Meanwhile, HighByte's Pipelines and Litmus Edge Flow's UI provide low-code configuration of data ingest from specific sources, validation, transformation and delivery to the cloud analytics provider of choice.

- **Industrial DataOps ontologies are increasingly available to bridge IT-OT-ET.**
  At the highest level of complexity, whereby organizational data span not just time series, events and alarms, but unstructured sources such as documents, and semi-structured ERP, customer relationship management (CRM) and even markets data, the confines of the UNS begin to break down. The world is not always hierarchical. Raw data flow to multiple consumers, for radically different reasons. As such, data models attempting to create a digital twin of complex systems tend towards an ontology – the philosophical study of 'being'. IDM solutions adopting DataOps ontologies are often described as knowledge graphs, using graph databases and semantic modelling to represent data and their relationships. Commercial providers such as Cognite, Palantir and SymphonyAI offer such modelling of complex relationships – semantic, kinetic and dynamic – and even mass transfer – between industrial and enterprise entities (see **Figure 1**). For most organizations, a knowledge graph is overkill. Similarly, wild diversity in data models results in vendor lock-in, while data quality issues can cripple large-scale utilization in analytics and even AI agent tasks.

Figure 7
The IDM vendor landscape

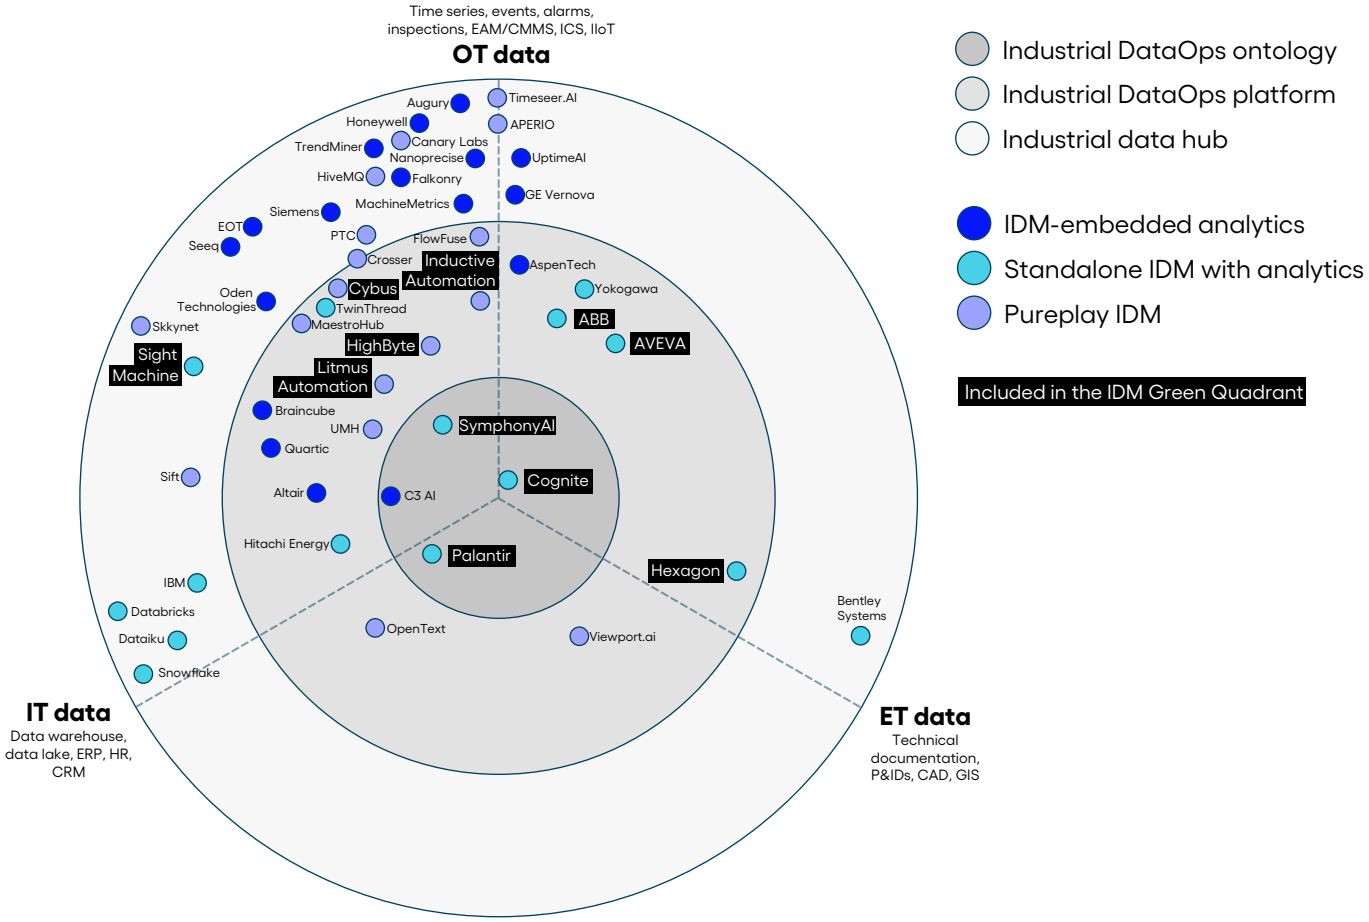

Note: This market map is not exhaustive. Verdantix sees IDM focus more on IT-OT convergence, hence there are fewer vendors in ET.
Highlighted vendors met the inclusion criteria for the Green Quadrant analysis.
Source: Verdantix analysis

LLM-driven unstructured data management is the next practical application for industrial AI

IIoT is enabled at the simplest level with a MING stack, implemented more robustly with DataOps platforms and applied at a multinational enterprise scale by DataOps ontologies. The vendor landscape here is diverse – and is especially crowded at the data hub segment (see **Figure 7**).

Despite compelling developments from open-source projects and commercial providers, in industrial data modelling, the challenge of acquiring data from unstructured sources remains. Enter IDM in the world of large AI models. OpenAI singlehandedly brought the LLM to the mainstream at the end of 2022. Two years on, technology multinationals, start-ups and their investors, analysts and regulators have not slowed the development of larger models. Investment continues in new architectures, novel learning and inference strategies, trained on more multimodal data, with more sophisticated scaffolding. In fact, capital expenditure at the four largest US technology firms was expected to grow to $209 billion in 2024, according to a forecast from Citibank – 42% greater than in 2023. For reference, the UK's National Health Service (NHS) in England was expected to spend around £190 billion ($239 billion) in 2024. Amid such frenzy, however, this spending in computational resources – such as graphical processing units (GPUs) from firms such as

NVIDIA – has delivered startling breakthroughs. Today's commodity models such as Anthropic's Claude 3.5 Sonnet and OpenAI's GPT-4o are adept at writing boilerplate code, and are already improving the productivity of software developers through applications such as Cursor and GitHub Copilot. Beyond this, Verdantix sees several other trends in IDM solutions utilizing frontier AI technologies, with:

- **Document-understanding AI models accelerating the ingest of unstructured data.**
  PDFs and spreadsheets represent the bulk of high-value industrial data. The portable document format (PDF) and spreadsheet still serve today as the primary formats for both unstructured and semi-structured information. Spreadsheets offer structure through rows and columns, while PDFs dominate as a literal ISO standard for sharing unstructured documents – from technical manuals to compliance disclosures. Through their ubiquity as a human-readable format, these modes of information transfer create challenges for IDM software. A spreadsheet might embed unstructured elements such as charts, macros and annotations, while PDFs prioritize visual consistency over mechanistic interpretability. Commercial document readers offered by hyperscalers AWS and Microsoft Azure, alongside open-source projects such as IBM's Docling, offer out-of-the-box functionality; however, domain-tuned solutions from Cognite, Palantir and SymphonyAI allow specific industrial data to be extracted more consistently. Such firms offer piping & instrumentation diagram (P&ID) readers not just to extract tags, but to recognize symbols and even perform rudimentary path tracing, thus helping them hydrate their industrial DataOps ontologies.

- **Retrieval-augmented generation (RAG) thriving on DataOps ontologies.**
  The popularity of RAG rose after instruction-tuned LLMs demonstrated strong in-context learning abilities (see Verdantix Market Insight: Ten Applications Of Large Language Models For Industry). These enable generative models to use new information provided in prompts to deliver more accurate answers, reducing reliance on intelligence gained through training alone. However, RAG is the toothpick in the Swiss army knife of AI-driven IDM – barely more than cosmetic, compared with the precisely engineered tools at its core. Naive approaches use vector similarity search, while commercial RAG solutions rely on sophisticated retrieval systems – matching a query and context provided by a user with relevant data from an indexed corpus. Vendors offering DataOps platforms and ontologies to users – and knowledge management systems built on top of them – provide far more sophisticated retrieval systems than vector search alone. Knowledge graphs offered by Cognite and SymphonyAI, alongside Palantir's Foundry Ontology, provide not just data matching the user query, but semantic, kinetic and dynamic context for the retrieved data. If executed upon with strong guardrails for source attribution, such systems can deliver reliably accurate answers to users in natural language.

- **Domain-specific visual data understanding accelerating rapidly in 2024.**
  Multimodal AI models – capable of understanding inputs across not just text and time series, but also images, audio and even frame-by-frame video – were first glimpsed in 2021 through OpenAI's contrastive language-image pre-training (CLIP). This utilized a modification to the existing Transformer architecture to recognize visual features – while in early 2023 Salesforce Research open-sourced BLIP-2, a CLIP-based model combined with an LLM to allow conversational interaction with images. Multimodal AI models accelerated with OpenAI's GPT-4 with vision (GPT-4V) in September 2023, followed by its more lightweight and more capable GPT-4o (omni) in May 2024. Today, multimodal AI models are commoditizing fast, with strong open-source offerings such as Mistral's Pixtral-12B released in November 2024. As such, software developers and ambitious industrial firms are building compelling use cases. Witness, for example, chemicals manufacturer Celanese showcasing its CELIA chatbot solution, built on Cognite's CDF. This custom solution gives operators access to deeply contextualized structured data within CDF's industrial knowledge graph, alongside the ability to ask questions to frontier LLMs, grounded with domain knowledge, regarding images of corrosion and even specific types of material fracture. Computer vision is also offered through SymphonyAI's Vision AI, helping manufacturers autonomously detect defects.

- **Industrial agents representing far greater utility than chatbots alone.**
  Autonomous systems capable of receiving user instructions and navigating the possible answer space, collecting data, assessing intermediate answers and using available tools are the next frontier for AI systems. Palantir Foundry integrates closely with the provider's Artificial Intelligence Platform (AIP) to provide users with the ability to build automated agents for data retrieval, reranking and intermediate interpretation. Meanwhile, Cognite offers a similar agent orchestration solution through its Atlas AI, alongside an industrial-focused benchmarking report to help buyers select the most appropriate LLM for their use case. Industrial agents will also benefit from faster modes of interaction, such as through voice – with AI labs OpenAI and Google competing intensely in December 2024 on price, capabilities and latency.

Challenges with AI model reasoning performance and reliability continue to limit their deployment outside carefully crafted IDM use cases

Practical applications of frontier AI technologies such as LLMs largely comprise data extraction, enrichment and transformation to use-case-optimized data models (see Verdantix Market Insight: Practical High-Value Industrial Use Cases For Generative AI In 2024). As chatbots, RAG helps ground responses in up-to-date, proprietary data, while multimodal AI models offer users the ability to mix text, visual and audio data into their interactions with AI systems. At the leading edge, autonomous AI agents are offered by IDM software vendors as part of robust software development kits (SDKs) and even low-code interfaces, to build multi-step, multi-reasoning LLM data management pipelines. Limitations, however, persist across not just commodity LLM offerings, but even the largest, most expensive and most capable models. Verdantix finds that:

- **Reasoning performance of frontier AI models still depends on memorization.**
  LLMs learn to generate coherent text by making small adjustments to model weights while pre-training on trillions of tokens (words, word fragments, symbols and visual patterns). However, this process encourages the imitation of training data rather than understanding and improvement. Solving true reasoning – and especially visual reasoning tasks – is, today, still not a capability of frontier AI models. François Chollet, creator of the deep learning library Keras, proposed a visual reasoning challenge in 2019 that cannot be solved through memorization alone. An untrained human will score 80%, but even today's most powerful autoregressive models, such as Anthropic's October 2024 Claude 3.5 Sonnet, barely pass 20%. Reinforcement learning LLM systems such as OpenAI o1 and o3 are, however, gaining traction as an avenue for improvement through 'test-time compute' – where strong models are allowed to 'think' through possible solutions and converge on an answer previously considered improbable using conventional autoregressive generation. OpenAI's o3, previewed in December 2024, aced the challenge by scoring nearly 90% – albeit after task-specific tuning and at a retail cost of around $3,500 per question. Open-source projects, such as the Alibaba MarcoPolo team's Marco-o1, are part of a vibrant community exploring ways to unhobble existing models' limitations. Such approaches train models to memorize robust, generalizable reasoning steps to converge on the correct answer, rather than just trying to memorize the answer itself.

- **Response latency is significant and unpredictable in LLM-driven IDM systems.**
  Deploying an LLM within private infrastructure requires significant computational resources. Massively parallel matrix multiplication hardware, such as graphics processing units (GPUs) designed by NVIDIA and produced by chip fabricators such as TSMC, continue to drive hundreds of billions in capital expenditure (CAPEX) – largely in the service of LLM training and inferencing. While lightweight models such as Llama-3.2-8B can produce coherent text at conversational speed running on a modest spec of Apple's latest MacBook Pro, frontier LLMs require far more horsepower. In enterprise-scale IDM solutions, LLM inferencing is often handled by hyperscalers such as Azure OpenAI Service, who take advantage of vastly higher GPU parallel throughput afforded by serving multiple inferencing tasks concurrently. Open-source projects such as vLLM also allow IDM solution providers to serve open-weight LLMs to customers within their private tenants. The unpredictability of AI agent systems retrieving data, exploring potential solutions and taking additional actions means that back-end processes may not respond for minutes. In all, this results in a significant user-experience and expectations-management challenge.

- **Consistent data modelling is difficult to enforce with automated graph hydration.**
  Deploying a UNS architecture for OT-focused data systems helps with scalability. Such an approach requires industrial workers, practitioners, managers and systems integrators to carefully follow a clearly defined hierarchy across facilities, lines and equipment, down to individual sensors. The same is true for automated systems. LLM-driven knowledge graph hydration requires identification of tags, checks for similar tags already within the ontology, and entity resolution to ensure that duplicate data sources are eliminated. This means providing not just human reviewers, but the LLM itself, with clear examples of how to systematically apply each data model to different domains. In practice, for large organizations, this takes months – hence the current lack of strong case studies for such a process (see Verdantix Generative AI With Knowledge Graphs: A Giant Leap For Industrial Data Management).
- **Retrieval-augmented generation is still more art than science.**
  Effective RAG requires far more than just providing models with a large quantity of relevant data. In practice, naively retrieved, isolated paragraphs and a small amount of metadata will rarely provide enough context for an LLM – trained on general web data – to reliably discard irrelevant results, distil correlations and generate a helpful response. Some approaches involve the user of a 'reranker' – a small LLM designed to compare initially retrieved results directly with the user query and metadata – to provide more accurate relevance scores. However, today's most powerful models still hallucinate, or are simply not intelligent enough, and even reranking models can sometime mistakenly assign high relevance to spurious results. Effective solutions require domain expertise across both AI model orchestration and application domains.

Green Quadrant for IDM software 2025

Buyers of IDM software from both heavy and light industries seek comprehensive, configurable and scalable solutions that build strong foundations for asset information, asset health, failure prediction, reliability analysis, maintenance optimization and integrity management, alongside adjacent functionality such as risk management, asset lifecycle management and environmental performance management. For the purposes of this report, Verdantix defines IDM software as:

> *"Software facilitating a professional approach to managing data in asset-heavy industries, improving data quality and facilitating collaboration with subject-matter experts (SMEs) and data scientists – all while constructing data pipelines for tracking and verifying data origins from the moment of acquisition through to usage and eventual deletion."*

This definition does not include software services related to implementation, integration, training and consulting, or software or applications with a focus on a single or a select few impact areas. The assessment encompasses both applications deployed on-premises and those that are single- or multi-tenant cloud-hosted.

Green Quadrant methodology

The Verdantix Green Quadrant methodology provides buyers of specific products or services with a structured assessment of comparable offerings at a certain point in time. The methodology supports purchase decisions by identifying potential vendors, structuring relevant purchase criteria through discussions with buyers and providing an evidence-based assessment of the products or services in the market. To ensure objectivity of the study results, the research process is guided by:

- **Transparent inclusion.**
  We aim to analyse all providers that qualify for inclusion in the research. For those providers that offer insufficient information or are unwilling to cooperate fully on the questionnaire and product demonstration, we include them in the report based on public information, where this provides an accurate analysis of their market positioning.
- **Analysis from the market perspective.**
  We integrated findings from our latest global corporate industrial transformation survey of 304 decision-makers, many of whom have bought or plan to buy software products such as those analysed in this Green Quadrant. The data-driven survey findings inform how we define the relevant software categories, sub-categories and weightings that propel the Green Quadrant graphical output.
- **Reliance on professional integrity.**
  As it is not feasible to check all data and claims made by vendors, we emphasize the need for professional integrity. Assertions made by software providers are put in the public domain via this Verdantix report and can be checked by competitors and existing customers. Verdantix also retains previous iterations of vendors' Green Quadrant questionnaire responses and makes comparisons and scoring adjustments as needed, to ensure accuracy.
- **Scores based on evidence, briefings and customer interviews.**
  To assess software vendors' expertise, resources, business results and strategies, we gather evidence from public sources and conduct interviews with multiple spokespeople and industry experts. When providers claim to be "best in class", we challenge them to present supporting evidence.
- **Scores based on relative comparisons.**
  We construct measurement scales ranging from 'worst in class' to 'best in class' performance at a certain point in time. A provider's position in the market can change over time, depending on how its offering and success evolves relative to its competitors. As a result, a vendor's Quadrant positioning may not necessarily improve – even if it adds new applications, makes a strategic acquisition or receives investment – as the assessment is relative to what other vendors are offering or have been doing since the previous Green Quadrant study. The Green Quadrant analysis is typically repeated every one-and-a-half to two years.

Scope and methodology for the 2025 Green Quadrant IDM software study

Verdantix studies reflect the current state of customer requirements and product capabilities. As such, we have developed assessment criteria to ensure alignment with the present state of the market. In this 2025 iteration of the Green Quadrant IDM software study, Verdantix:

- **Developed IDM scenarios from capability assessments.**
  For this study, we established a set of the most important and relevant capability areas in which customers expect vendor functionality. Using the 2023 Smart Innovators and 2024 Buyer's Guide studies on IDM (see Verdantix Smart Innovators: Industrial Data Management Solutions and Verdantix Buyer's Guide: Industrial Data Management Solutions (2024)), alongside feedback from vendors and customers, we developed a framework of 15 technical capability areas covering data acquisition, modelling, governance, contextualization, interoperability and utilization.

- **Weighted the questionnaire categories to reflect market priorities.**
  The Verdantix Green Quadrant evaluates the latest customer technology preferences to ensure that the weightings of all high-level criteria reflect global buyers' current priorities across all IDM software components. Following extensive interviews with 304 senior operational excellence decision-makers, we applied adjusted weightings for each high-level capability criterion to mimic its relative priority for improvement and to reflect IDM software spending plans for 2025 amongst customers.

- **Included coverage of customer success and adoption.**
  Customer success strategies are often overlooked in assessment criteria for buyers. To account for these, Verdantix included questions around total customer count, renewal rates and strategy. Furthermore, we undertook 23 customer interviews with users of vendor solutions highlighted in this Green Quadrant.

Evaluated firms and selection criteria

For this Green Quadrant, Verdantix defined vendor inclusion criteria to ensure that the analysis only compared firms providing identifiable IDM offerings. The 11 IDM software providers included in this study were selected because they have:

- **Ability to sell IDM as a standalone software solution.**
  Vendors must demonstrate the ability to sell IDM as a standalone software solution or as part of an integrated suite where the data management component can operate independently within an existing third-party analytics and/or automation tech stack.

- **Strong functionality to offer all three of the core IDM capabilities assessed.**
  We identified vendors offering comprehensive IDM solutions capable of supporting the three core IDM capabilities: (1) direct and/or indirect industrial data acquisition from control systems, sensors or historians or enterprise systems such as document stores, geographic information systems (GIS), enterprise asset management (EAM) systems, computerized maintenance management systems (CMMS) and ERP systems; (2) industrial data quality management and/or lineage visibility and governance; and (3) industrial data contextualization through hierarchies or ontologies. Vendors must rely on their own developed or directly licensed technologies to deliver these capabilities. To ensure comparability, we excluded vendors that focus on only one or a subset of these areas, depend heavily on third-party tools, or lack the ability to provide robust, end-to-end IDM functionality.

- **At least 40 employees, $3 million IDM software revenues and five named customers.**
  This Verdantix Green Quadrant study is intended to assess the most prominent vendors offering IDM software solutions. The vendors included in this study have IDM-designated employee numbers ranging from 40 to more than 3,000 and annual IDM software revenues ranging from $3 million to $2 billion. All vendors disclosed at least five customers who adopted and deployed their software for IDM use cases in 2024.

- **Resources and strategic incentive to continue developing their IDM suite.**
  We focused the study on vendors with the human, financial and technological resources to meet the needs of IDM customers for the foreseeable future. This reflects the desire of customers to invest time and resources in IDM solutions, with the expectation that vendors will not discontinue them.

Based on the inclusion criteria above, this report looks in depth at the IDM software offerings available from 11 vendors: ABB, AVEVA, Cognite, Cybus, Hexagon, HighByte, Inductive Automation, Litmus Automation, Palantir, Sight Machine and SymphonyAI. With the exception of Inductive Automation, Palantir and Sight Machine, which were invited to take part but did not actively participate or did not respond, all vendors actively participated in the research through responses to a 111-point questionnaire, by allowing customer interviews and by taking part in a two-hour product demonstration.

Evaluation criteria

Verdantix defined the evaluation criteria for the Green Quadrant for IDM software using a combination of interviews with corporate practice managers and software executives, desk research, discussions with multiple customers and staff expertise. Our analysis was also informed by responses to the Verdantix global corporate industrial transformation surveys. In full, this year's Green Quadrant analysis compares offerings from 11 software vendors using a 111-point questionnaire covering 15 categories of technical capabilities and nine categories of market momentum. In our analysis:

- **Capabilities measure the breadth and depth of functionality.**
  The capabilities dimension, plotted on the vertical axis of the Green Quadrant graphic, is a measure of the breadth and depth of each software provider's functionality. To assess this, we evaluated data for 15 technical capabilities. The technical capabilities are: direct industrial data acquisition; indirect industrial data acquisition; data modelling; data storage and single source of truth; data quality management and lineage traceability; cyber security; unstructured data utilization; data discoverability; data contextualization; data visualization and analytics; deployment options; platform APIs; data interoperability; platform development environment; and user interfaces (see **Figure 8**).
- **Momentum measures strategic success factors.**
  The momentum dimension, plotted on the horizontal axis of the Green Quadrant graphic, measures each software vendor on a range of strategic success factors. The criteria that make up the momentum score are grouped into nine high-level categories: brand preference; vision and strategy; product strategy; innovation process; funding; acquisitions and partnerships; organizational resources and growth; revenue growth; and customers (see **Figure 9**).

The evidence provided by all the software vendors was assessed using a quantitative model that started with the sub-criteria scores. Each sub-criterion was individually weighted to generate the overall score for each capability area. For example, data modelling is one of the high-level criteria evaluated in the capabilities section, but is composed of five sub-criteria covering time series, relational databases, document stores, graph databases and domain-specific data models. These are individually weighted to determine the overall data modelling score.

All sub-criteria were scored between the values of zero ('no capability') and three ('best in class'). Subsequently, each high-level criterion was allocated a percentage weighting that determined its contribution to the overall score for the specific capability. Weightings were based on customer survey data regarding the IDM software functionality that is most widely used, along with analyst perceptions of the broader IDM software landscape. The combination of high-level criteria scores in the capabilities and momentum sections generated the Green Quadrant rankings (see **Figure 10** and **Figure 11**) and graphic (see **Figure 12**).

Figure 8

Technical capabilities criteria for IDM software

Capabilities	Questions
Direct industrial data acquisition (9%)	Provide details of how your IDM solutions acquire and integrate data from asset programmable logic controllers (PLCs) and industrial control systems (ICS). If you have an official partnership with equipment providers or original equipment manufacturers (OEMs), please describe the partnership, the value it adds to your offering and any customer success stories. Provide details of how your IDM solutions acquire and integrate data from industrial Internet of Things (IIoT) sensors. Provide details of how your IDM solutions acquire and integrate data from production microphones and cameras. Provide details of any other direct industrial data sources from which your IDM solutions can ingest data.
Indirect industrial data acquisition (10%)	Provide details of how your IDM solutions acquire and integrate data from process historians, computerized maintenance management systems (CMMS), enterprise resource planning (ERP) and enterprise asset management (EAM) systems. Provide details of how your IDM solutions acquire and integrate data from document management systems. Provide details of how your IDM solutions acquire and integrate data from engineering design, geographic information system (GIS) and computer-aided design (CAD) software. Provide details of how your IDM solutions integrate with data warehouses/data lakes/lakehouses. Are there any other industrial or enterprise data sources from which your IDM solutions integrate data?
Data modelling (7%)	Do your IDM solutions model time series data or other sequential data? Do your IDM solutions model relational data (interlinked tables, relational database management systems, DataFrames)? Do your IDM solutions model data in a document store (e.g. unstructured and semi-structured data)? Do your IDM solutions model graph data (e.g. labelled property graphs)? Provide details of asset/process-specific data models where the overall data modelling strategies above (time series, relational, document store and graph databases) are modified and combined to produce a hybrid data model for specific industrial assets, systems and processes.
Data storage and single source of truth (5%)	Do your IDM solutions facilitate the creation of a catalogue of references to a consistent, synchronized version of industrial data? Provide details of how your IDM solution delivers a centralized directory where data points from multiple sources are assigned to specific identifiers or topics, flattening the data architecture and allowing for straightforward data retrieval without the complexity of disparate naming conventions or isolated data silos. Does your IDM solution offer the ability to configure rules for data storage prioritization, cache frequently used data, archive unused files and schedule deletion? What options are available to your IDM customers for back-up, the selection of storage media types and storage locations?
Data quality management and lineage traceability (5%)	Provide details of how your IDM solutions offer pre-built tools for monitoring and alerting for common sensor failures. Provide details of how your IDM solutions enable customization for industry- and application-specific sensor failures. Provide details of how your IDM solutions aggregate and visualize industrial data from the moment of acquisition, transfer, transformation, utilization and eventual archiving or deletion. Provide details of pre-built tools delivered by your IDM solutions that apply effective data cleaning through interpolation, cross-referencing or smoothing, to improve signal-to-noise ratio for analytical use cases. Provide details of other industrial data quality management capabilities offered by your IDM solutions.
Cyber security (3%)	Explain the use of encryption for securing data in transit and at rest, ensuring data privacy and integrity. Provide details of your existing cybersecurity certifications (e.g. CCPA, GDPR, ISO 27001). Provide details of any notable real-world cyber-related near-misses your customers experienced that were defended against by your IDM solutions. Provide details of how your IDM solutions enable fine-grained access management, such as role-based, location-based or other techniques.
Unstructured data utilization (6%)	Describe your IDM solutions' ability to extract metadata, interpreted data and other useful data from unstructured sources such as documents, images, video, audio, and piping and instrumentation diagrams (P&IDs). Describe your IDM solutions' ability to extract free text from unstructured data sources such as documents, images, video, audio and P&IDs. Describe your IDM solutions' ability to extract tabular data from unstructured data sources such as documents, images and diagrams. Describe your IDM solutions' ability to generate useful descriptions and metadata using document understanding and diagram understanding systems. Describe your IDM solutions' ability to generate labelled property graphs or knowledge graphs from structured and/or unstructured data, and to perform entity resolution, source attribution and consistent discoverability.
Data discoverability (10%)	Describe how your IDM solutions deliver user interface (UI)-level search engine functionality for asset tag names, tag content and tag metadata to serve fast, relevant results while accounting for minor typographical errors. Describe how your IDM solutions deliver UI-level search engine functionality for asset tag names, tag content and tag metadata to serve fast, relevant results from typed or spoken queries in natural language. Describe how your IDM solutions deliver UI-level search engine functionality for images, videos, charts and diagrams used in industrial operations and maintenance to serve fast, relevant results from typed or spoken queries in natural language. Describe how your IDM solutions deliver UI-level search engine functionality for information stored in relational databases to serve fast, relevant tables and/or specific answers from typed or spoken queries in natural language. Describe how your IDM solutions deliver UI-level search engine functionality for information stored in graph databases to serve fast, relevant network visualizations and/or specific answers from typed or spoken queries in natural language.

Figure 8 (continued) ↓

Figure 8 (continued)

Data contextualization (9%)	Describe how your IDM solutions help users, as part of the graphical user interface (GUI) or otherwise, visualize data points representing industrial events, incidents and other specific moments in relation to contemporary events, to boost situational awareness. Also explain how your IDM solutions visualize industrial assets, workers and processes and their spatial, temporal and systemic relationships with each other. Provide details of how your IDM solutions assemble, as part of the GUI or otherwise, disparate pieces of data and package them together to provide a self-explaining payload of information that accurately describes a useful state of industrial operations and/or asset management. Provide details of how your IDM solutions facilitate, as part of the GUI or otherwise, the annotation of industrial data and documentation, such as P&IDs, charts and other visualizations, such that adjacent teams and decision-makers can quickly share and persistently record analysis. Explain how your software, through low-code user experience (UX), facilitates cross-functional collaboration within a common data ingestion, transformation, contextualization and visualization environment between the frontline workforce, data scientists and senior decision-makers. Discuss how your IDM solutions abstract away IT-focused difficulties involved with data discovery through the automation of best practices via robust, low-code GUIs.
Data visualization and analytics (5%)	Provide details of the pre-built, application- and industry-specific data visualization dashboards offered by your IDM solutions to expedite understanding of operational data. Provide details of how your IDM solutions provide pre-built trend analysis for anomalies, time series forecasting or other analytical methods. Provide details of how your IDM solutions assist with input data feature selection, model architecture selection, training and deployment monitoring.
Deployment options (5%)	Describe the minimum and recommended hardware and software requirements for on-premise deployment of your IDM solutions. Also describe supported operating systems, installation and configuration processes and any dependencies or prerequisites – or note whether pre-built hardware is provided to customers and installed by your engineers. Also discuss procedures for updates, maintenance and scaling capacity. Provide details on whether or not you offer edge deployment of aspects of your IDM solutions, optimized for low latency and close proximity to industrial assets and processes, and how these differ from on-premise deployment capabilities. List the cloud service providers supported by your IDM solutions and explain the cloud deployment models supported (e.g. SaaS, PaaS, IaaS, IaC) that are available to customers and systems integration partners. Describe any partnerships or certifications with these cloud providers. Describe the support for hybrid and multi-cloud environments, enabling data integration and management across on-premise and cloud-based systems. Explain the mechanisms in place for data synchronization and consistency across on-premise and cloud environments and the tools available for managing data replication, conflict resolution and latency.
Platform APIs (8%)	Provide details of how your IDM solutions provide application programming interfaces (APIs) for sending data to third-party IDM and industrial DataOps providers. Provide details of how your IDM solutions provide APIs to third-party industrial analytics, AI trending and other data science software and service providers. Describe how your IDM solutions utilize RESTful APIs for integration with both first-party and third-party IT, operational technology (OT) and engineering technology (ET) systems, such as support for standard HTTP methods (GET, POST, PUT, DELETE) and JSON or XML for data interchange. Also describe your support for GraphQL APIs and gRPC for high-performance, complex industrial deployments. Provide details of how your IDM solutions implement OAuth 2.0 for secure, token-based data transfer for third-party solutions and how keys and tokens are managed within your solutions. Discuss the availability of comprehensive API documentation, such as endpoint descriptions, request/response examples and usage best practices, alongside developer portals and software development kits (SDKs).
Data interoperability (6%)	Explain how your IDM solutions support extract, transform, load (ETL) and extract, load, transform (ELT) processes for data integration – and how this works with the data quality management discussed earlier in this questionnaire. Also describe the tools and features available for schema mapping for data integration between different systems, both first- and third-party, with data modelling strategies across time series, relational, document-oriented or graph. Describe how your IDM solutions use event-driven architecture, MQTT and OPC-UA to enable real-time data exchange and interoperability across first- and third-party industrial systems. Discuss the use of webhooks, message queues and publish/subscribe (pub/sub) models for real-time and asynchronous data exchange. Detail your support for MQTT, including quality of service (QoS) levels and OPC-UA and their roles in secure and reliable data communication. Explain how your IDM solutions support the development of a unified namespace (UNS) architecture. Explain how these provide a consistent, synchronized view of industrial data.
Platform development environment (5%)	What development tools can customers use to develop/customize analytics models to support self-service analytics? This includes integrations with tools such as Python/Jupyter and R for custom model development and 'no-code/low-code' applications.
User interfaces (7%)	What is the usability/user-friendliness of the enterprise app interface (including mobile functionality)? Do you offer any accessibility functionality? How many languages are offered out of the box? Can users easily switch to other units of measurement? In what ways do you engage customers with regard to obtaining user feedback/strengthening user experience?

Note: Figures in brackets represent the weighting given to each criterion in the flexible multi-criteria model that generates the Green Quadrant graphical analysis.
Source: Verdantix analysis

Figure 9

Momentum criteria for IDM software

Capabilities	Questions
Brand preference (10%)	Based on Verdantix analysis
Vision & strategy (10%)	What is your firm's vision for the target customers/addressable market opportunity you seek to target over the next 5 years (industries, geographies, revenue-size customers, functional budget-holders, functional solutions)? How do you intend to achieve your vision? What is your firm's strategy to meet the needs of your target customers over the next 5 years?
Product strategy (15%)	What is your firm's IDM and DataOps product strategy for the next 2-5 years? What does your product roadmap look like? What are your plans with regard to offering on-premise IDM and DataOps solutions over the next 5 years? Discuss whether or not you will discontinue on-premise options and whether cloud deployments will move towards private tenant or multi-tenant.
Innovation process (10%)	What is your firm's innovation framework and strategy? Comment on whether your firm has global innovation hubs or runs hackathons and other initiatives. What is your firm's strategy when it comes to R&D investment allocations to support long-term viability and maintain competitive advantage(s)? Please feel free to detail % of revenue invested in R&D and/or % of R&D invested in new product development or other metrics to demonstrate re-investment in new product development.
Funding (10%)	Have you received any funding over the last 3 years, e.g. public, capital markets?
Acquisitions & partnerships (5%)	How many acquisitions have you made over the past 2 years? Please provide details of the acquisitions. How many formalized partnerships have you made over the past 2 years? Please provide details of the partnerships and the value these generate for your clients. (This includes partnerships with service providers, software resellers and cloud partnerships.)
Organizational resources & growth (15%)	Total number of employees? How many employees are dedicated to IDM? In how many/which countries does your firm have an office?
Revenue growth (15%)	What were your firm's annual revenues for your most recent reporting year? What were your firm's revenues from IDM software for your most recent reporting year? By how much did your firm's total revenues grow in the most recent reporting year compared with the year prior?
Customers (10%)	How many of your customers/accounts (not individual users) are using your firm's IDM and DataOps solutions? How many customers/logos have you added over the past 12 months? Who are the key buyer personas/stakeholders you target? How does your go-to-market (GTM) messaging vary with different stakeholders?

Note: Figures in brackets represent the weighting given to each criterion in the flexible multi-criteria model that generates the Green Quadrant graphical analysis.
Source: Verdantix analysis

Figure 10

Vendor category scores (capabilities)

	ABB	AVEVA	Cognite	Cybus	Hexagon	HighByte	Inductive Automation	Litmus Automation	Palantir	Sight Machine	SymphonyAI
Direct industrial data acquisition	2.5	2.1	2.4	1.8	0.8	1.8	2.6	2.6	1.4	0.5	2.1
Indirect industrial data acquisition	2.0	2.4	3.0	0.8	2.7	0.8	0.7	0.8	2.7	0.6	2.3
Data modelling	2.2	2.4	3.0	1.4	1.2	1.7	1.6	1.7	2.3	1.5	2.8
Data storage and single source of truth	2.0	2.0	2.4	1.4	1.7	1.5	3.0	1.5	3.0	0.4	2.3
Data quality management and lineage traceability	2.3	1.3	1.8	0.2	2.0	1.5	0.5	1.7	2.0	1.0	2.1
Cyber security	2.0	2.1	2.3	2.3	2.2	2.1	2.0	1.3	3.0	1.6	2.3
Unstructured data utilization	1.7	1.1	2.7	0.0	2.1	0.0	0.0	0.3	2.7	0.0	1.8
Data discoverability	1.7	1.3	2.5	0.0	0.3	0.0	0.3	0.3	2.9	0.5	2.1
Data contextualization	1.3	2.2	2.8	0.9	2.3	0.6	1.5	1.1	2.7	1.3	2.0
Data visualization and analytics	2.8	2.4	2.2	1.0	2.0	0.0	1.8	1.8	2.6	1.6	2.0
Deployment options	2.3	2.2	1.7	2.1	1.5	2.1	1.5	2.4	2.7	1.4	2.6
Platform APIs	1.3	2.0	3.0	2.5	1.5	2.3	2.3	2.1	2.5	1.2	2.1
Data interoperability	1.7	1.7	2.8	2.5	0.7	3.0	2.0	2.5	2.0	1.0	2.7
Platform development environment	3.0	2.0	3.0	0.0	1.0	0.0	2.0	1.0	3.0	0.0	3.0
User interfaces	1.5	2.0	2.4	1.4	1.9	1.5	1.9	1.7	2.2	1.6	2.4

Note: Verdantix research teams determine all scores at either sub-criteria level (for capabilities) or criteria level (for momentum), using the scoring framework below. These assessed scores are then weighted and compiled into derived scores at criteria or capability/momentum level. The scoring framework is shown on Figure 11.
Source: Verdantix analysis

Figure 11

Vendor category scores (momentum)

	ABB	AVEVA	Cognite	Cybus	Hexagon	HighByte	Inductive Automation	Litmus Automation	Palantir	Sight Machine	SymphonyAI
Brand preference	2.6	2.0	2.4	1.0	2.0	1.4	2.0	1.4	1.8	0.6	1.6
Vision & strategy	2.5	2.5	3.0	1.5	1.0	3.0	1.5	2.5	2.5	0.5	3.0
Product strategy	1.7	2.0	3.0	1.0	2.0	2.4	2.0	2.3	2.3	0.0	2.0
Innovation process	2.5	0.5	2.5	1.5	1.0	2.0	1.5	2.0	1.5	0.0	2.5
Funding	2.0	2.0	1.0	1.0	1.0	1.0	0.0	1.0	3.0	1.0	2.0
Acquisitions & partnerships	1.8	2.4	2.4	0.8	1.6	1.6	1.6	2.4	2.4	1.6	1.6
Organizational resources & growth	1.7	1.2	1.8	1.3	1.3	1.6	1.8	1.8	1.9	0.7	2.3
Revenue growth	1.3	1.5	1.9	2.4	0.9	2.4	2.0	2.4	1.9	0.9	2.3
Customers	1.2	1.8	2.6	1.2	1.0	2.2	2.4	2.8	2.2	0.4	2.0

Scoring framework	
Evidence of market-leading functionality or positioning	3
Evidence of strong, above-par functionality or positioning	2
Evidence of on-par functionality or positioning	1
Lack of evidence, or evidence of sub-par or a lack of functionality or positioning	0

Verdantix research teams determine all scores at either sub-criteria level (for capabilities) or criteria level (for momentum), using the scoring framework above. These assessed scores are then weighted and compiled into derived scores at criteria or capability/momentum level.

Source: Verdantix analysis

Figure 12
Green Quadrant for IDM software 2025

Capabilities

This dimension measures each service provider on the breadth and depth of its IDM solutions across 15 capability areas, as outlined in **Figure 8**.

Momentum

This dimension measures each service provider on nine strategic success factors, as outlined in **Figure 9**.

Note: A white plot indicates a non-participating vendor.
Source: Verdantix analysis

Cognite overview

Information

Cognite, founded in 2016 by John Lervik (Chief Strategy Officer), Geir Engdahl (Chief Product Officer) and Stein Danielsen (Chief Solutions Officer), is headquartered in Oslo, Norway. Lervik, formerly CEO of FAST Search & Transfer (acquired by Microsoft for $1.2 billion in 2008), Danielsen (also from FAST) and Engdahl (Google) built the software vendor now led by CEO Girish Rishi, ex-CEO of SCM firm Blue Yonder. Beyond Oslo, Cognite has offices in Austin, Houston, Phoenix, Tokyo, Riyadh and Abu Dhabi. In 2022 it partnered with SLB for oil and gas solutions and with Rockwell Automation to enter manufacturing. It formed alliances with Accenture and Pinnacle in 2022 and 2021, and Radix in July 2024. Ownership is divided between Aker ASA (50.5%), Accel (13%), TCV (10%) and Saudi Aramco (7.4%).

Vendor info

Firm name	Cognite
Headquarters	Oslo, Norway
Employees	600
Revenues	$100m to $500m
No. of offices	7
Example customers	Aker BP, Celanese, Cosmo Oil Company

Customer regional presence

Asia	◑
Oceania	○
Europe	◕
Middle East and Africa	◑
Latin America and the Caribbean	◔
North America	◕

% Customer base

○ 0% ◔ <10% ◑ 10%-25% ◕ 25%-50% ● above 50%

Cognite's top three industry penetration

1. Oil and gas
2. Manufacturing
3. Utilities

Technical capability scores

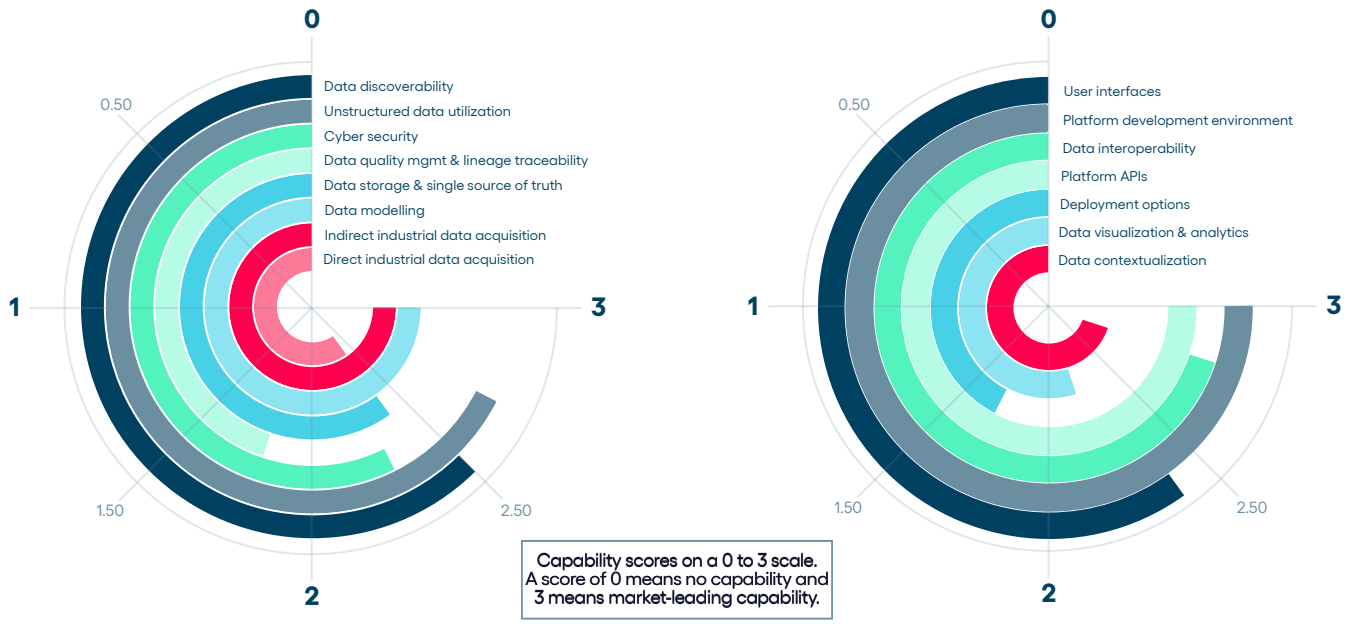

Capability scores on a 0 to 3 scale. A score of 0 means no capability and 3 means market-leading capability.

Cognite provides a market-leading DataOps and AI platform for enterprise-scale, complex IDM projects

Cognite's core DataOps product, Cognite Data Fusion (CDF), is a platform built on a cloud- and scale-optimized industrial knowledge graph, fed by a rich assortment of data connectors and integrations. In June 2023 Cognite announced new GenAI features to augment its DataOps platform, with subsequent launches such as Canvas and Atlas AI delivering promising glimpses of the use of large language model (LLM)-driven chatbots and AI agents for industry. The Verdantix Green Quadrant analysis finds that Cognite provides:

- **Deeply configurable unstructured and structured industrial data orchestration.**
  Cognite's foundational industrial DataOps platform, CDF, presents firms with a meticulously documented set of standardized building blocks for representing industrial data. Its core data model, CogniteCore, represents data with two main components: core features and core concepts. The former are sets of properties – names, descriptions and sources – which can be applied to core concepts (specific data entities such as time series, assets or files). This core data model is extensible, theoretically enabling it to represent an entire industrial facility within a data digital twin – as a user-traversable industrial knowledge graph. For this, the vendor scored 3.0/3.0 for data modelling, with a similarly perfect score for indirect industrial data acquisition, including for unstructured data. It was also one of the top scorers (2.7/3.0) for unstructured data utilization, owing to its ability to extract tags names and locations from piping and instrumentation diagrams (P&IDs) and to read tables and free text in technical documentation, all within its low-code CDF software as a service (SaaS) interface. For developers, CDF offers perfect-scoring platform application programming interfaces (APIs) and model development, with a still best-in-class 2.4/3.0 for its mix of developer-focused and low-code user interfaces (UIs).

- **Fast-paced product development that occasionally frustrates maturity-focused customers.**
  Like most vendors working with industrial data, Cognite increased its focus on AI and GenAI development after November 2022, subsequently launching Cognite AI, Canvas and Atlas AI. Similarly, its InRobot autonomous robotics solution now integrates with quadrupeds from ANYbotics and Boston Dynamics to acquire data directly from the field. Such capabilities, while powerful, present customers and systems integration partners – including Cognite's services division – with additional technical overhead and systems to orchestrate, alongside a steeper learning curve. While customer interviews revealed near-universal satisfaction with the reliability of Cognite's in-production IDM solutions, the vendor should maintain its balance between innovation and upkeep of its existing products. Although its UI is market-leading in this study, Cognite has some room for improvement in its new low-code Atlas AI and Canvas products, especially as its GenAI offering opens up its user base not just to developers, but to non-programmer process engineers and executives.

- **World-class offering for complex asset-heavy industrial data and AI projects.**
  Cognite's field-hardened and highly productized industrial DataOps platform, CDF, presents a smorgasbord of pre-built capabilities for buyers, deployed SaaS and even private SaaS, to quickly implement and get returns. Cognite Embedded allows original equipment manufacturers (OEMs) and third-party software vendors to build with CDF using pre-built templates and infrastructure provided by Cognite. In addition, it delivers DataOps at a truly global scale, thanks to its long-time partnership with engineering firm Aker Solutions and oilfield services vendor SLB, alongside more recent alliances with operational technology (OT) cyber security firm Omny, engineering firm Radix and manufacturing-focused Rockwell Automation. Cognite worked with chemicals manufacturer NIPPON SHOKUBAI to gather nearly 20 years of operations and maintenance (O&M) data and unstructured inspections and engineering data from disparate systems, unifying them with a customized low-code interface powered by CDF. For field workers, Cognite's InField solution enables direct data acquisition from O&M processes, while desktop users can select groups of tags within a P&ID, jump to plots of time series data associated with them, and offload to a third-party analytics provider – as well as build plant human-machine interfaces (HMIs) within Canvas. The entire NIPPON SHOKUBAI project, covering 100 plants within a 900,000m^2 facility, took a year to implement. It is expected to eliminate 9,000 hours of workload per year, while promoting far greater ability to make data-driven decisions.

verdantix

Independent insight and analysis

Our research is a trusted source for some of the largest and most innovative businesses in the world. With over a decade of reports, data and analysis, our subscribers have access to depths of insight that cannot be found elsewhere.

Whether you are implementing a leading-edge technology strategy, or developing the products and value propositions of the future, our analysis will help you futureproof your thinking.

Our expertise

EHS & Quality

ESG & Sustainability

Industrial Transformation

Net Zero & Energy Transition

Real Estate & Built Environment

Risk Management

Contact

Verdantix Ltd, Woolyard,
52-56 Bermondsey Street,
London SE1 3UD, United Kingdom

contact@verdantix.com
@Verdantix

Opportunities at Verdantix

Since 2008, Verdantix has been delivering high-quality research and advice to its clients. If you're interested in joining a world-class team with an unwavering focus on success, apply to join us today. We are delighted to be hiring across all teams and have a variety of opportunities in both London and Boston