SOLUTION SELECTION MATRIX™

2026 | Industrial AI Platforms Guidebook

LNS Research

Executive Summary

Few other technologies have disrupted the manufacturing sector with the speed and sheer scale that Industrial AI has had in the recent past. For an industry still running critical systems on Windows XP and decade-old PLC code, this paradigm shift has been a reckoning of sorts. Over the past few years, LNS Research has closely observed how manufacturing companies across the globe have responded to this reckoning with a strong sense of urgency, enterprise-level AI initiatives, and millions of dollars committed to AI-driven transformation.

At the same time, we have also observed parallel shifts in the technology vendor (and venture capital) landscape as the first industry-grade Chatbots and Copilots started coming out within six months of GPT 3.5. These initial breakthroughs quickly took the industrial world by storm, and before we knew it, almost every single industrial technology company became an Industrial AI company. We saw Copilots eventually get better at technical challenges like hallucination and traceability, and the momentum shift to AI agents that can take actions and are more powerful with reasoning and self-learning abilities. While we can debate the efficacy and relevance of such Chatbots, Copilots, and Agents in manufacturing, the fact remains that Industrial AI is here, and has already disrupted several technologies, particularly Advanced Industrial Analytics.

LNS Research has been a proponent of the notion that Industrial AI is neither synonymous with nor a replacement for Advanced Industrial Analytics. That said, we also believe the rapid disruption of generative and agentic AI capabilities has significantly altered the competitive landscape and warrants a repositioning of how we assess technology vendors in this category as part of the LNS Research Solution Selection Matrix (SSM) reports.

An investment in Industrial AI is not to be taken lightly, as most of these solutions, once installed, are not easily or quickly replaced, even if they are delivered via Software-as-a-Service (SaaS) or over cloud deployments. LNS Research's SSMs are intended to guide manufacturers in selecting technology categories and to help create a potential vendor shortlist. In addition to technical features, we also recommend that you understand the vendor's overall business strategy, their ability to continue to serve their target markets, and other external forces such as partnerships, mergers, and acquisitions. Ultimately, we hope you find this research useful for your Advanced Analytics and Industrial AI needs and use it to make well-informed decisions when selecting the right technology partner.

Industrial AI Definition

LNS Research defines Industrial AI as the combined application of machine learning and other Artificial Intelligence models, in conjunction with first principles, rules-based methods, and statistical, data-driven approaches, to solve industrial problems. The problems that Industrial AI applies to span both within the four walls of the factory, from asset monitoring, process optimization, material flow management, energy efficiency, workforce safety, to outside the factory across the entire supply network from procurement to customer delivery and field service.

As with Advanced Analytics, Industrial AI spans the spectrum from descriptive and diagnostic analytics to forward-looking predictive, prescriptive, and prognostic models that deliver data-driven insights across these use cases. However, what distinguishes Industrial AI from Advanced Analytics is its ability to provide Decision Intelligence capabilities that enable AI-first ways to monitor, synthesize, recommend, execute tasks, and adapt over time.

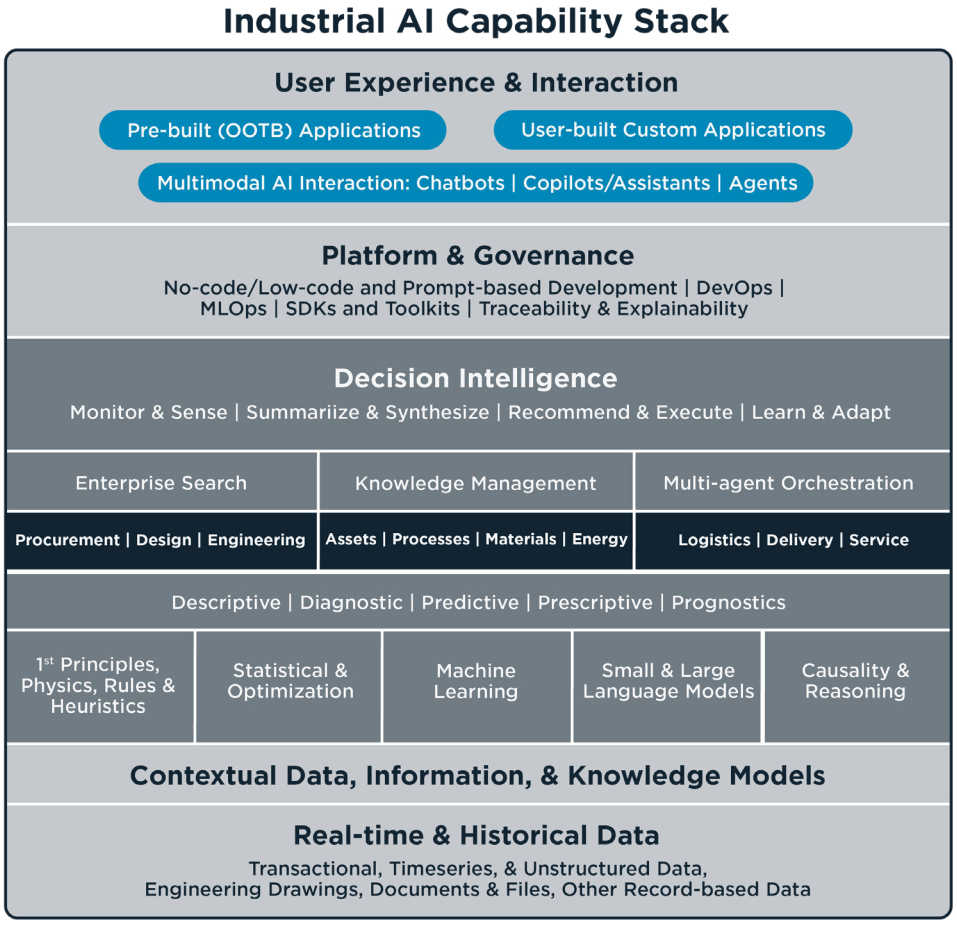

Figure 1: Industrial AI Capability Stack

The framework above provides a representation of the Industrial AI capabilities. As you can see, Industrial AI begins with foundational data capabilities that cover both real-time data and historical data, and builds contextual models of not just data, but also information and knowledge that have been further processed and understood. Built on this foundation of the contextual model is the Decision Intelligence layer, which includes the core capabilities of Industrial AI.

Industrial AI Definition (Cont.)

Decision Intelligence is an emerging umbrella term to describe agentic AI capabilities, though it currently lacks a clear definition. In the above Industrial AI capability stack, we define Decision Intelligence as the set of core analytical and computational capabilities of Industrial AI. It includes two complementary layers: the analytical layer that produces insights and the agentic layer that operationalizes the insights.

- The analytical layer encompasses a range of analytics algorithms, models, and techniques from first principles, statistical, data-driven Machine Learning (supervised, semi-supervised, and unsupervised) (ML), Deep Learning (DL), and Reinforcement Learning (RL) models, and emerging causal and reasoning models applied across the spectrum of descriptive – prognostic analytics.
- The agentic layer focuses on how said insights are consumed and acted upon through a continuum of chatbots, copilots, and agents that take on several roles from monitoring, sensing, summarizing, synthesizing, recommending, and acting – all done in a self-learning loop that makes sure it improves accuracy and reliability over time.

The next level of Industrial AI capabilities is the platform and governance layer, including a no-code/low-code development environment, DevOps/MLOps capabilities, and other development and integration capabilities, in addition to generative AI capabilities such as prompt-based development, Retrieval Augmented Generation, etc. Finally, the user experience layer must include a combination of pre-built apps, custom apps, and a continuum of chatbots, copilots, and agents that can be interacted with not just text, but also across images, audio, and video.

Industrial AI Definition (Cont.)

Due to the wide range of capabilities across data plumbing and infrastructure, storage and modeling, Advanced Analytics, and, of course, the generative and agentic capabilities, there have been several paths for Industrial AI vendors to take, resulting in a highly diverse and fragmented competitive landscape. These companies don't always compete head-to-head or at the same level in the organization. As each of them has approached Industrial AI from a different perspective, we have had to break down the segment into six distinct sub-categories:

- Advanced Analytics
- Asset Optimization
- Process Optimization
- Enterprise Application Suite
- Industrial AI Platforms
- Machine Vision

In an ideal world, each of these sub-categories would have its own Solution Selection Matrix to enable end users to make apples-to-apples comparisons. However, we are far from that. We live in a world – the industrial world – where there is always an overlap and a nuance and a technicality that makes apples-to-apples comparisons very difficult. As a result, we will be assessing Industrial AI vendors separately and grouping them where there is significant overlap in product features, scope across the value chain, and approach to architecture.

The first Industrial AI Solution Selection Matrix Guidebook focuses exclusively on two categories of Industrial AI: Data Platforms and Advanced Analytics, where the companies are too similar to separate, yet too different to combine. Let's take a closer look at these sub-categories and how we plan to cover them in this guide.

A) Industrial AI: Advanced Analytics

The first category of Industrial AI SSM – published in December 2025 - includes software applications that offer Decision Intelligence capabilities listed in Figure 1, primarily through pre-built applications. These vendors provide data collection, analytics, and visualization functionalities as part of their application(s). However, they do not provide a licensable platform for end users to build custom applications. These companies also differentiate in their scope, which focuses on not just one or two but several industrial use case categories across assets, processes, material flow, energy efficiency, workforce safety and training, etc.

List of vendors assessed:

- Augury
- Canvass AI
- Falkonry
- Oden Technologies
- Proficy (Predictive Analytics)
- Seeq
- TrendMiner

B) Industrial AI: Data Platforms

The second category includes solutions that have Decision Intelligence capabilities mentioned above, plus a full-stack platform that brings together multiple data sources across IT and OT through pre-built connectors, APIs, webhooks, etc. These vendors also provide robust data connectivity, quality, conditioning, and contextualization, and data models to store this processed and contextualized data in a cloud infrastructure. Ideally, these Industrial Data Platforms should collect, store, and model structured, time-series, and unstructured data (excluding historians and time-series databases).

End users will then be able to perform advanced analytics through a set of pre-built, out-of-the-box applications on the platform. Some vendors in this category provide a no-code/low-code environment for building custom applications. Still, their primary focus is on adding value through their data platform and pre-built applications. While companies in this category all offer both a data platform and applications built on top of it, they differ significantly in how they prioritize these two pieces. Some emphasize the data platform as their core offering, with applications playing a supporting role. Others lead with their applications and treat the data platform as a secondary consideration.

To reflect this variation in our assessment and enable fair comparison with the other two categories, the data platform will be assessed on the Advanced Analytics capabilities plus Industrial Data Platform capabilities. However, the overall product score will reflect their performance across both these capabilities combined.

List of qualified vendors:

- Braincube
- C3.ai
- Cognite
- Palantir
- Quartic.ai
- Sight Machine
- SymphonyAI
- TwinThread

Coverage of Industrial AI: Advanced Analytics and AI Platform Providers:

The first guidebook will cover the Advanced Analytics companies, evaluated primarily on their Decision Intelligence capabilities, along with a subset of their data connectivity, conditioning, and contextualization features. The second guide, to be published shortly after, will cover the Data Platform companies, assessed on both Decision Intelligence and a broader set of DataOps capabilities. Finally, because companies in these categories compete with one another, we will publish a combined SSM that evaluates them side by side.

Adjacent Categories and Honorary Mentions

As mentioned earlier, Industrial AI is one of the most diverse technology categories, with a broad array of providers offering competing solutions. While this assessment focuses on two types of vendors, many companies from adjacent categories also offer overlapping capabilities.

Emerging startups in such categories leverage data-driven and AI-based approaches, such as supervised/unsupervised machine learning, neural networks, and reinforcement learning, to solve problems traditionally addressed using first principles, physics-based models, and rules-based logic. Some of these companies are included in this guide, while others are not, based on their primary focus areas, as explained below:

Industrial AI Applications Suite

This category includes the companies that provide a suite of Industrial AI solutions at the highest level. This group includes several types of companies, including large-scale established enterprise software companies and automation & control providers who have put together a suite of Industrial AI capabilities through organic growth and or acquisition. These companies extend their scope to not just IT and OT, but also the rest of the value chain's data together. It's important to note that these vendors also provide a set of out-of-the-box applications on their platforms to perform advanced analytics, along with a no-code/low-code environment for the end user to build custom applications and other agentic capabilities.

List of qualified vendors:

- ABB Genix
- AVEVA CONNECT
- Dassault 3DXperience
- Emerson/AspenTech
- Honeywell (Forge for Industrial)
- IFS.ai
- Oracle (Fusion)
- QAD
- Rockwell Automation
- SAP (Digital Manufacturing Cloud)
- Siemens (Xcelerator)

Industrial AI: Asset Optimization

This sub-category of Industrial AI includes emerging AI-native startups that use proprietary industrial-grade sensors, cloud-native platforms, and services to provide predictive and prescriptive analytics for asset monitoring. While most of these vendors will be assessed in a separate guide, Augury, with its strong presence across both asset and process analytics, will be included in the Advanced Analytics guide as well.

List of qualified vendors:

- AssetWatch
- Augury
- Dynamox
- Fluke Reliability
- Infinite Uptime
- KCF Technologies
- Nanoprecise
- Petasense
- Samotics
- Tractian
- Viking Analytics

Industrial AI: Closed-Loop Process Control

This category includes companies that deploy a combination of first principles, physics, data-driven, and AI approaches to solve closed-loop process control, a critical step in the ultimate goal of Autonomous Operations. These companies include both emerging startups that are AI-first and legacy APC providers that have upgraded their legacy platforms.

List of qualified vendors:

- ABB (APC Software and Systems)
- AspenTech (DMC3™)
- AVEVA (APC)
- Honeywell (APC)
- Imubit
- Intelecy
- Kelvin
- Rockwell Automation (FT Pavilion8 MPC)
- Sorba.ai

Honorary Mentions #1: Agentic Operations and Knowledge-Driven Industrial AI Startups

Finally, there is a growing category of Industrial technology companies that are building AI-first Industrial AI solutions. These vendors are focused on decision intelligence from the ground up—using approaches like knowledge graphs, causal reasoning, and autonomous agents to move beyond dashboards and deliver measurable outcomes. They represent the next generation of Industrial AI providers and directly address many of the limitations that held back earlier analytics-focused solutions. That said, most of these companies are still in the early stages of growth and are not yet mature enough to be featured in this guide. We will continue to track their progress closely and expect to include some of them in future editions.

Vendors that fall in this category include:

- Arch Systems
- Faclon
- IntuigenceAI
- Parabole
- Phaidra
- UptimeAI
- XMPro

Honorary Mentions #2: Composable Operations Platforms

There is also a category of industrial technology companies offering half-stack or full-stack platforms focused primarily on operations and execution use cases, rather than analytics or decision intelligence. These platforms should not be interpreted as traditional MES/MOM systems; instead, they offer many core MES capabilities while adding elements such as Industrial DataOps, lightweight data modeling, and application-building tools.

Some also incorporate functionality typically found in quality management systems and connected frontline worker (CFW) solutions, including digital work instructions and recipe management. Importantly, these platforms are designed with a composable or modular architecture, reflecting lessons learned from earlier generations of monolithic systems. Because of this hybrid positioning, they are not explicitly included in this guide. However, we will continue to monitor their evolution for potential inclusion in future editions.

Vendors that fall in this category include:

- BEET
- MachineMetrics
- Factbird
- Fuuz
- Tulip

Industrial AI Platforms Definition

This category represents a class of full-stack, industrial-grade software platforms designed to enable manufacturers and industrial operators to deploy AI at scale across a wide range of industrial use cases. Unlike Advanced Industrial Analytics solutions, Industrial AI Platforms are able to provide integrated capabilities across three core layers: Industrial DataOps, Data Platforms, and Advanced Industrial Analytics.

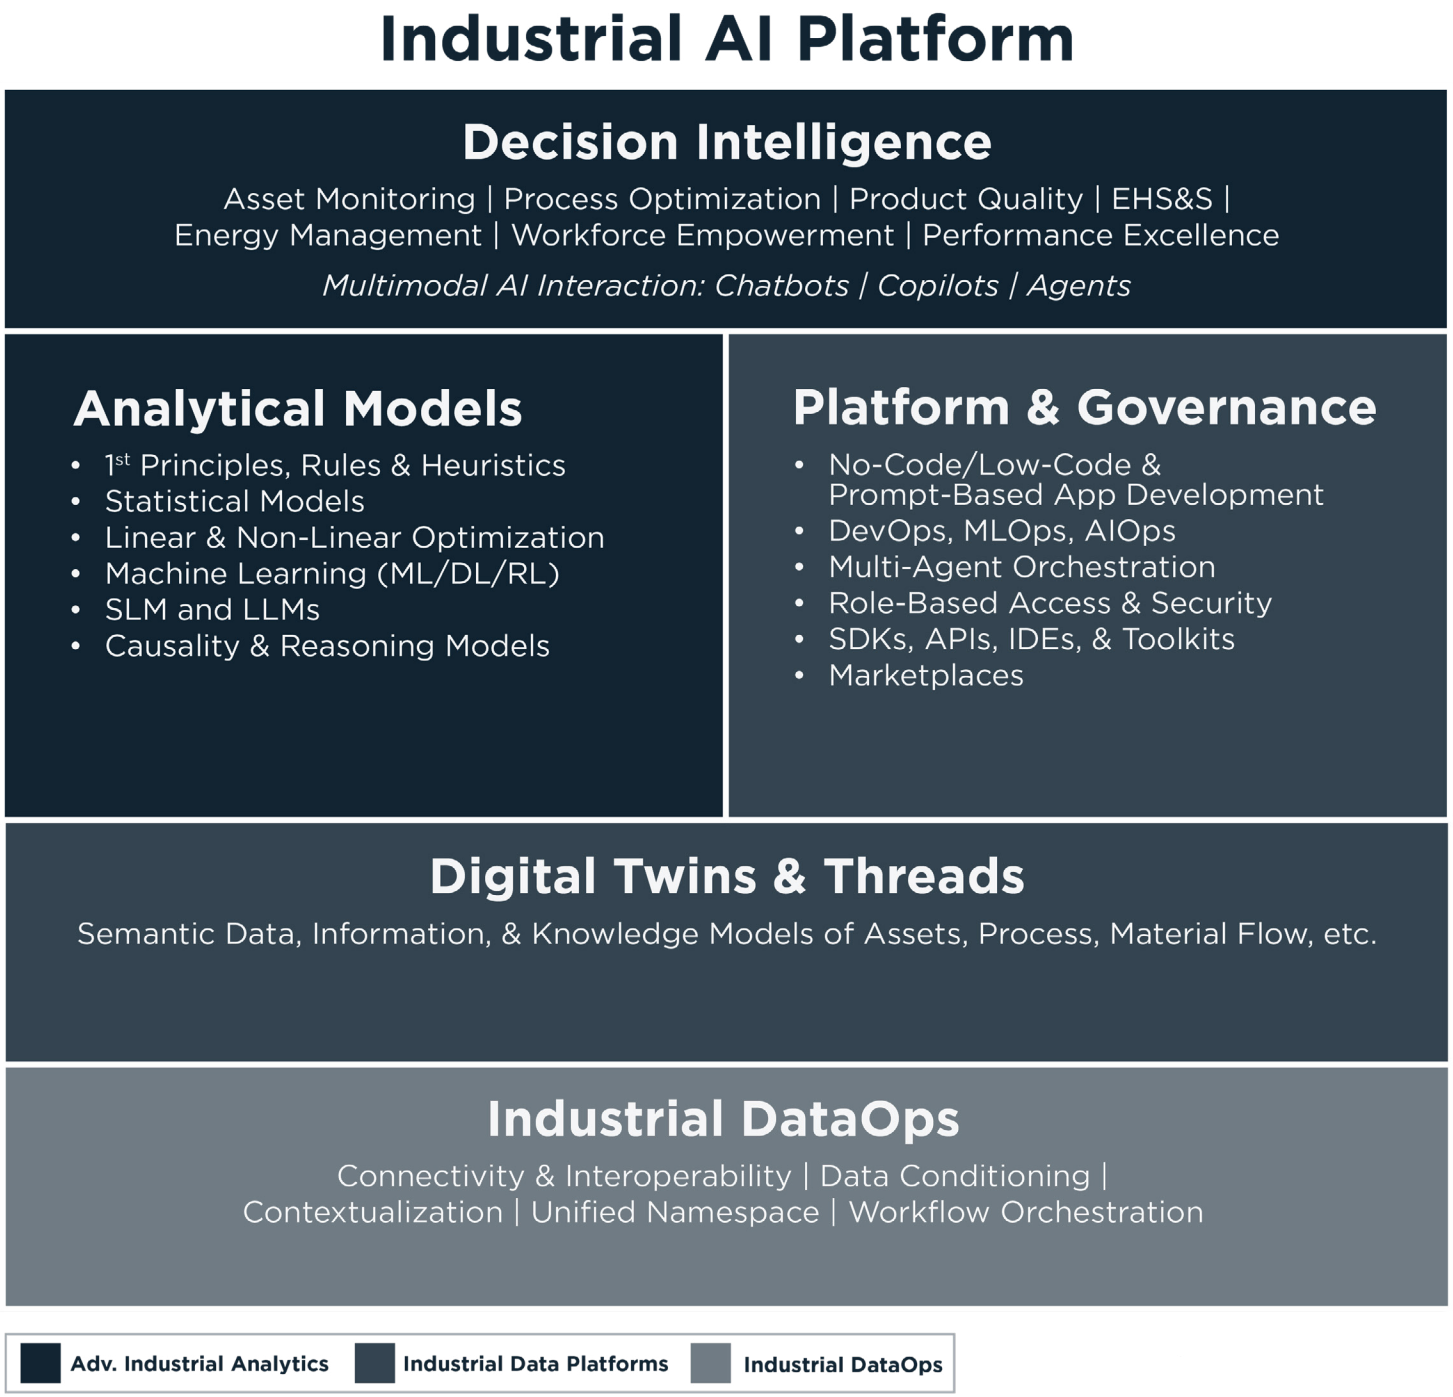

Figure 2: Industrial AI Platforms Definition

Industrial AI Platforms Definition (Cont.)

Together, these three layers support the end-to-end lifecycle of Industrial AI—from industrial connectivity and contextual data modeling, through governance and model development, all the way to productionized analytics applications embedded into operations. Let's take a look at each:

Industrial DataOps:

Not surprisingly, the first part of Industrial AI begins with data- specifically Industrial DataOps. LNS Research recently defined Industrial DataOps as an emerging discipline that unlocks the value of industrial data through connectivity, interoperability, contextualization, and modeling of data flow across industrial environments. In practice, it is the set of capabilities required to ensure that real-time and historical data from OT and IT systems is consistently collected, conditioned, structured, and made usable to support everything from analytics and AI models to applications and operational decision-making.

Built on the principles of DevOps, Agile, and Lean, Industrial DataOps is both a methodology and an enabling capability set that streamlines industrial data collection, transformation, contextualization, and orchestration across manufacturing environments. It bridges the persistent IT/OT gap by standardizing how data is managed and deployed across multiple use cases, sites, and initiatives—reducing duplication, improving governance, and accelerating time-to-value.

A growing ecosystem of providers is emerging with Industrial DataOps as a standalone offering. These companies will be assessed in a separate guide; however, the providers assessed in this guide deliver Industrial DataOps as part of a broader full-stack industrial AI platform offering, integrated into their end-to-end capabilities. While the Industrial DataOps category is not as mature as Industrial AI, these vendors are continuing to win their fair share of the market, and will be covered as part of an upcoming Solution Selection Guide (without a 2X2 matrix), including the following vendors:

- Cybus
- Flow Software
- Highbyte
- HiveMQ
- Kepware
- Litmus
- Rhize
- Timeseer
- United Manufacturing Hub

Industrial AI Platforms Definition (Cont.)

Industrial Data Platform:

The Industrial Data Platform layer consists of four tightly integrated components. It begins with Digital Twins and Digital Threads, which provide semantic models across increasing levels of abstraction—from raw data to information and operational knowledge. These models range from foundational asset hierarchy structures based on standards such as ISA-95 and ISA-88 for batch processes, to more advanced information models and knowledge graphs that capture complex, non-hierarchical relationships such as material flow, genealogy, and cross-process dependencies.

Built on top of these models is the Decision Intelligence Engine, which supports a broad spectrum of analytical and AI techniques. On the deterministic side, this includes first-principles models based on physics, thermodynamics, and chemistry; rules and heuristics; and operations research methods such as linear and nonlinear optimization. On the probabilistic side, the engine supports machine learning approaches, including supervised, semi-supervised, and unsupervised learning, deep learning, and reinforcement learning, as well as generative AI foundation models—both general-purpose and industry-specific LLMs and SLMs—along with causal and reasoning models.

The third component is Platform Enablement and Governance, which provides the infrastructure required to build, deploy, and manage industrial AI at scale. This includes no-code/low-code, and prompt-based development environments for creating custom applications; DevOps, MLOps, and AIOps capabilities to manage models and applications across multiple environments; and enterprise-grade security, integration, and marketplace capabilities.

Finally, the platform supports a rich interaction layer, spanning a continuum of chatbots, copilots, and autonomous or semi-autonomous agents. Importantly, this interaction layer combines both deterministic systems—where predictable, rules-based behavior is required—and probabilistic AI assistants, enabling safe, explainable, and context-aware human–AI collaboration in industrial settings.

Industrial AI Platforms Definition (Cont.)

Advanced Industrial Analytics:

While there is no shortage of full-stack platforms with Industrial DataOps and Digital Twin capabilities, a key differentiator of Industrial AI Platforms is the application layer. To qualify as an Industrial AI Platform, providers must deliver Advanced Industrial Analytics capabilities built on top of the DataOps and Data Platform layers, enabling outcomes for several use cases across the industrial value chain.

These analytics capabilities are delivered through software applications that embed the Decision Intelligence functions described in Figure 1, most commonly through pre-built, out-of-the-box applications, along with capabilities for users to build their own applications without the need for custom coding.

The Industrial AI platform providers also differentiate themselves through the breadth of their application scope. Rather than focusing on one or two specialized use cases, Industrial AI Platforms support multiple categories of industrial value creation—spanning assets and reliability, process optimization, material flow and production execution, energy efficiency, and workforce safety and training, among others.

Vendor Inclusion and Exclusion Criteria

LNS Research's definition of Industrial AI Platforms incorporates several key technical inclusion criteria for this Solution Selection Matrix:

Significant presence in industrial operations:

- Providers must demonstrate a meaningful footprint in industrial operations, with clear adoption in manufacturing, process industries, energy, utilities, mining, or other industrial asset-intensive domains.
- Platforms must support industrial workflows and operational decision-making, not solely general enterprise analytics.

Company Maturity & Commercial Presence:

- Mid-to-high seven figures in annual recurring revenue (ARR) or above.
- A commercially mature organization, including dedicated teams across product management and engineering, sales and account management, marketing and go-to-market, customer success/implementation/support
- Proven customer deployments in industrial operations, with multiple production customers (not only pilots), referenceable customers upon request, and evidence of repeatability across sites/use cases (where applicable)

Full-stack platform with robust features across three key layers:

- **Advanced Industrial Analytics:** Prebuilt and/or configurable applications for high-impact industrial use cases
- **Industrial Data Platform:** Semantic Digital twin/thread models, decision intelligence engines, governance, and other platform capabilities.
- **Industrial DataOps:** Connectivity, interoperability, conditioning, contextualization, and edge-to-cloud orchestration

Vendor Inclusion and Exclusion Criteria (Cont.)

Modular, Flexible, and Scalable Industrial-Grade Architecture:

- Platforms must be built on a modular and scalable architecture supporting composability across data, models, apps, and interactions.
- Must support multi-site deployment patterns and reusability (templates, accelerators, reusable models).
- Platforms must support deployment across:
  - Public cloud, private cloud, and hybrid configurations
  - On-premise environments (or strong validated support for on-prem/hybrid deployments
  - Edge deployments for low-latency use cases, resilience, buffering, and offline modes
- Auditability, lineage, traceability, version control
- Lifecycle management for DevOps, MLOps, and AIOps
- Operability and Management across multiple plants/sites and environments

Breadth of Industrial Use Cases & User Personas:

- **Must-address use cases:** Asset Monitoring, Process Optimization, Product Quality, EHS&S, Energy Monitoring, Performance Excellence, etc.
- **Must-serve user personas:** Maintenance, Reliability, Industrial, Process Engineers, Plant-level subject matter experts, frontline operators, Data Scientists, Plant Management, Operations leadership

Vendor Inclusion and Exclusion Criteria (Cont.)

Furthermore, to narrow the list of vendors to a reasonable and relevant field of suppliers for evaluation, LNS Research is explicitly excluding the following categories of analytics-related tools and applications. These offerings may be critical enablers within the broader ecosystem, but they do not meet the inclusion criteria for a full-stack Industrial AI Platform

- General-purpose cloud services and analytics building blocks that are not delivered as an integrated industrial AI platform, and generic analytics toolkits and statistical packages, spreadsheets, BI, and reporting tools
- Databases, data warehouses/lakes/lakehouses, and other database infrastructure tools (including hyperscaler components)
- Industrial IoT platforms focused mainly on connectivity/device management without advanced analytics and decision intelligence
- IT/OT systems of record or execution (e.g., ERP, MES, SCADA, historians, CMMS/EAM, PLM, QMS, LIMS) that include limited analytics but are not intelligence-first platforms
- Execution-centric platforms that may span DataOps + data platform capabilities, but whose primary focus is systems of execution/record, such as MES, CFW, EQMS, EHS, ERP, PLM, etc., and not decision intelligence that wraps and extends around said systems
- Standalone no-code/low-code, workflow, and app development tools
- Point analytics solutions focused on a single or limited set of industrial use cases

The 3P Evaluation Model

LNS Research uses the "3P Evaluation Model" covering vendors' Product, Potential, and Presence to evaluate their ability to serve the Industrial AI: Advanced Analytics market.

Product

LNS Research **evaluates the products** against our Industrial AI Capabilities Stack and Advanced Analytics definition. In each **Vendor Profile**, we use Harvey Balls to score the degree of functionality or feature coverage and then calculate an overall product score from 1 to 6. Finally, these individual scores are weighted and averaged to produce an overall Product score, as seen in the scoring matrices in section VIII.

Product Definitions and Scale

6. Market leading capabilities across the spectrum of key functionality as defined by LNS Research[1]; proven success meeting all requirements in markets specifically targeted by vendor[2].
5. Robust spectrum of capabilities for applicable target markets. Few shortcomings that are recognized.
4. Broad though not complete spectrum of capabilities applicable to target markets.
3. Limited but sufficient capabilities applicable to a subset of target markets; lagging in product development and functionality.
2. Some gaps in functionality required to be addressed to fully meet target markets.
1. Newly launched minimal viable product and/or significant gaps in functionality required to be addressed to fully meet target market requirements.

---

[1]Assessments are made against a specific LNS Research-defined "Reference Architecture." For example, IIoT vendors would be judged against the LNS Research IX Reference Architecture.

[2]Target markets are the geographies and application areas specifically pursued by vendor. For example, vendors targeting only English-speaking markets will NOT be penalized for lack of two-byte character support.

Industrial AI: Advanced Analytics & Data Platform Functionalities

Capability | Functionality	Comments | Insights
Advanced Industrial Analytics	
Analytics levels supported (Descriptive – Prognostic)	Descriptive – Predictive analytics are prevalent. Braincube, Cognite, and TwinThread stand out with proven prescriptive and prognostics capabilities at scale
Use cases supported	Most cover asset, process, and energy use cases, with predictive/ prescriptive maintenance as the starting point. Quartic scores well in batch use cases, while C3.ai and Palantir perform well in use cases across the supply chain
Deterministic models: Applied math, first principles, statistics, optimization, etc.	Most platforms incorporate first principles and other deterministic models
Probabilistic models: Machine learning, deep learning, reinforcement learning, Large and Small Language Models, Vision AI, causal/reasoning, and other AI models	Supervised and unsupervised ML models are prevalent across most platforms. SymphonyAI has the potential to significantly differentiate with its industrial-grade foundational LLM model
Analytics layer (focused, systematic, global)	Most are focused on systematic primarily, with some extending to global, and focused use cases. C3.ai has selected case studies across all three
Self-learning models (closing loop after events)	Manual and semi-autonomous closed-loop predominantly
Monitoring and/or control (open- & closed-loop)	Monitoring for the most part. Braincube, Quartic, and TwinThread have proven to be deployed in open and closed-loop control
Industrial Data Platforms	
Deployment models supported	All platforms support edge-to-cloud deployments with partnerships with major hyperscalers. Palantir stands out with its FDE model that accelerates deployment notably
Multiple data types: Structured, semi-structured (time-series), unstructured	Time series data is the primary focus, with a secondary focus on structured data (transactional and other relational data) and unstructured data. Notable mentions: •Cognite extracts rich context from P&ID •SymphonyAI provides context from a wide range of unstructured data
Digital twins (asset, product, process twins)	All platforms provide Digital Twin models. Notably, •Braincube's Product Clones model products attributes to analyze operational and financial performance •Sight Machine takes a unique approach to model schemas of parts, cycles, defects, and downtime. •TwinThread's product and process twins are robust in depth of features.
Modular, flexible, scalable architecture	Most platforms are built to be scaled across multiple sites. Palantir's Apollo platform enables rapid deployment at scale, and SymphonyAI is known for its composable architecture

No-code/low-code development environment	Non-developers and data scientists are the primary target persona for all platforms, and empower these SMEs to build apps using a no-code/low-code environment. Notably, SymphonyAI prioritizes custom app building using NC/LC and generative AI tools over OOTB apps
Integrated workflow/orchestration engine	Cognite and TwinThread lead in building and templatizing workflows of data and insight flows
Extensible with SDKs, APIs, and custom code (Python, R, Java, C/C++, etc.)	C3 and Palantir are most malleable for customization and extending with external services
Industrial DataOps	
Connectivity and communications protocols supported	All platforms provide extensive connectivity through OPC-UA, MQTT, Modbus, and REST APIS. Quartic stands out with its unique approach to DataOps using Kafka for real-time streaming data
Data quality, cleansing, conditioning	All platforms prioritize data quality and provide conditioning capabilities. Sight Machine's unique schema approach stands out here
Contextualization	Edge-to-cloud contextualization capabilities are provided by all. Cognite and Sight Machine stand out with their features here

**The above table only includes features and functionalities of vendors scored in the Industrial AI Platform category, namely Braincube, C3.ai, Cognite, Palantir, Quartic, Sight Machine, SymphonyAI, and TwinThread. Other SSM guidebooks, such as the Industrial AI: Advanced Analytics Guide, will feature insights for corresponding vendors.*

Potential

LNS Research assesses the vendor's potential for growth in product and presence dimensions. Potential may be impacted by scale, focus, financial resources, market positioning strategies, the management team (especially for smaller companies), merger and acquisition plans, partnering strategy, or any number of other factors.

In the Advanced Industrial Analytics solution market, hundreds of vendors seek to carve out a niche. Therefore, we have only included those companies we believe have an opportunity for growth and long-term viability in this report. In fact, we have specifically included certain smaller startups/pure plays because we believe their focus and technology make them competitive, and they have a good chance of surviving, even thriving, in the market.

Potential Definitions and Scale

6. Likely overall market leader (across many industries, geographies, and application areas); currently outpacing all competitors.
5. Among small set of likely overall market leaders evidenced by current leadership in target markets and proven record of innovation.
4. A likely leader in some targeted markets with growth potential to move up; could rise to leading position in specific markets.
3. Likely a significant player in target markets with defensible barriers to competition and growth prospects.
2. Likely a niche player in target markets with some known risks to future growth in product and presence.
1. Early-stage company with wide dispersion of potential long-term performance and/or a niche player with significant risks to future growth.

Presence

LNS Research develops a composite score assessing a vendor's market penetration vis-à-vis geographies, industries, and customer sizes served. Score factors both "capability to serve" and "proven success."

- **Capability to serve:** Focused (experienced in the specific domain) sales and service resources, along with product enhancement to serve specific target markets.
- **Proven success:** Market success (satisfied installed base) in the markets served.
- LNS Research asks all Industrial AI solution vendors to provide insight into the competitive landscape. References to clients' "proven success" are used as inputs to the evaluation.

Presence Definitions and Scale

6. Market leading sales, service[3] and successful customers globally, in all relevant industries, and with companies of all sizes.

5. Robust sales and service capabilities and successful customers across all major geographies (North America, Western Europe, Middle East[4] and industrialized Asia-Pacific), a range of industries and company sizes.

4. Broad-based sales, service, and successful customers across most major geographies (North America, Western Europe, and selected Middle East and Asia-Pacific countries), specific industries and/or companies of a specific size (SMB, for example).

3. Regional sales, service, and customer success in select major western geographies; select industries and sizes.

2. Sales, service, and customer success in limited market(s), industry, and company sizes.

1. Still in early stages of launching new product and/or company with a very limited group of customers.

---

[3]Presence is measured by a vendor's capabilities to serve AND proven implementations. Vendors' capabilities to serve are measured by resources able to support a specific technology, not overall company capability.

[4]Middle East is a "major" geography for many process industries but not for discrete.

Frameworks for Vendor Assessments

LNS Research scores technology vendors across two frameworks: the Solution Selection Matrix and a new drill-down tabular framework, the Market Fit Grid. Let's take a close look at each in the following section.

Solution Selection Matrix: Industrial AI – Advanced Analytics

The Solution Selection Matrix follows the 3P evaluation outlined in our Evaluation Criteria blog post, How LNS Research evaluates vendors in technology domains, and then plots the vendors in a 2X2 matrix assessing their product, potential, and presence.

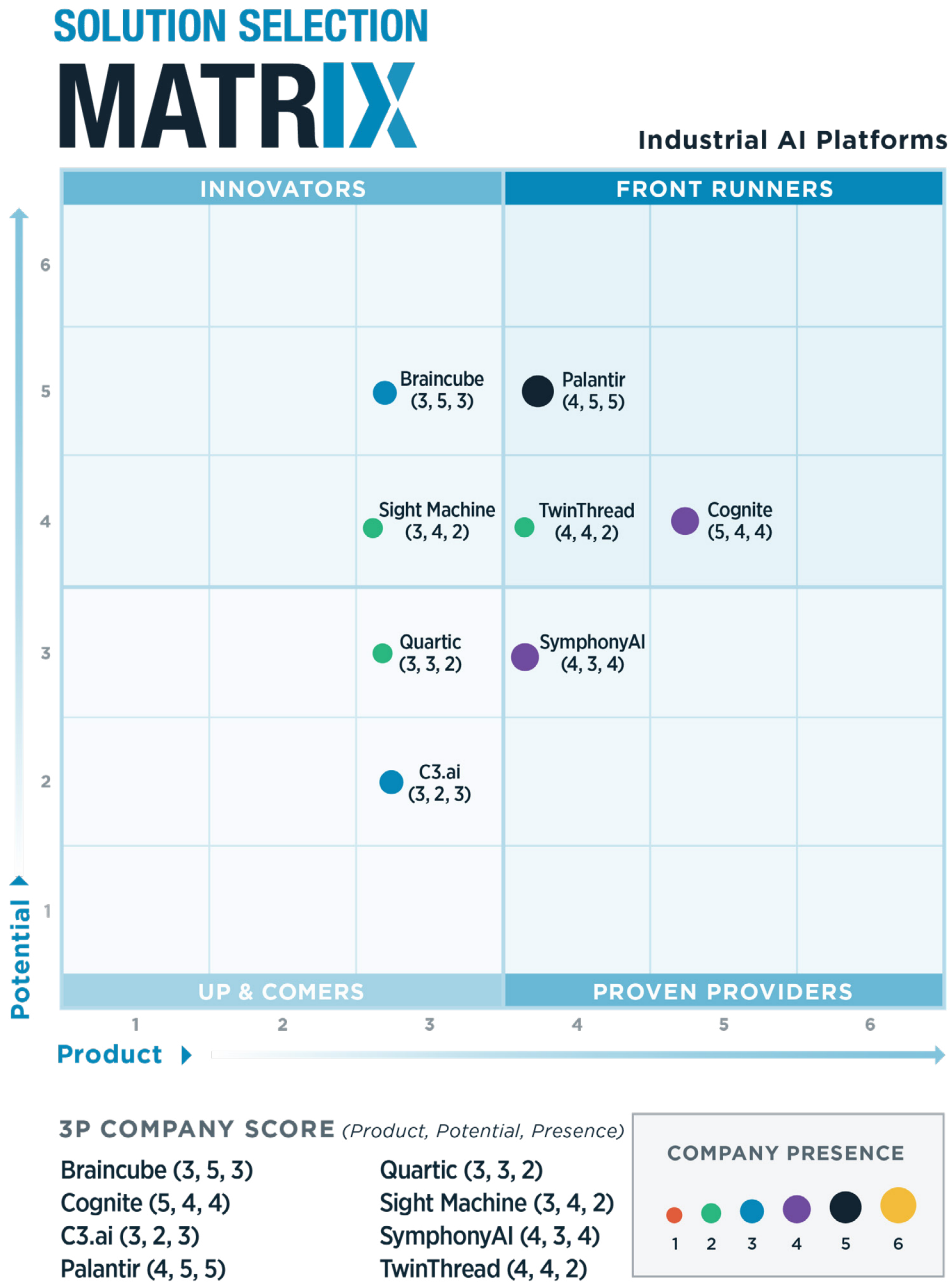

Figure 3: Industrial AI Platforms SSM

As you will also see, many vendors will be placed in the Up & Comers category. It is important to interpret this placement properly. It does not mean that the vendor's Product is inferior, but it may mean that it lacks some key features or functions for targeted markets, addresses only certain use cases, and/or lacks referenceable customers or has limited deployments.

Many companies, especially startups, have a small Presence, operating only in a few markets with a select number of customers, and the majority with revenues typically less than $10 million annually. Additionally, their future is relatively uncertain early on in their journey, so their potential may be rated lower.

These scores, of course, are expected to change as all companies mature. We will explore this in more depth in the individual vendor profiles. In addition, we will discuss each vendor's relative strengths and weaknesses specific to the markets and industry verticals they intend to serve.

Solution Selection Matrix: Industrial AI Platforms + Advanced Analytics

As mentioned earlier, Industrial AI is a broad and diverse category, with vendors taking very different approaches to capabilities, architecture, and delivery models. Within this landscape, Advanced Analytics vendors and Industrial AI Platforms represent two distinct subcategories, yet they also share significant overlap in the value they deliver.

For this reason, we evaluate these two groups separately in separate Solution Selection Matrices to enable a more apples-to-apples comparison and to clearly highlight the differences between a provider that delivers value primarily through self-service tools and pre-built applications versus one that offers a full-stack, extensible platform. At the same time, it can be useful to view both sub-categories together, since they often compete for the same budgets, particularly in top-down, enterprise-led Industrial AI initiatives.

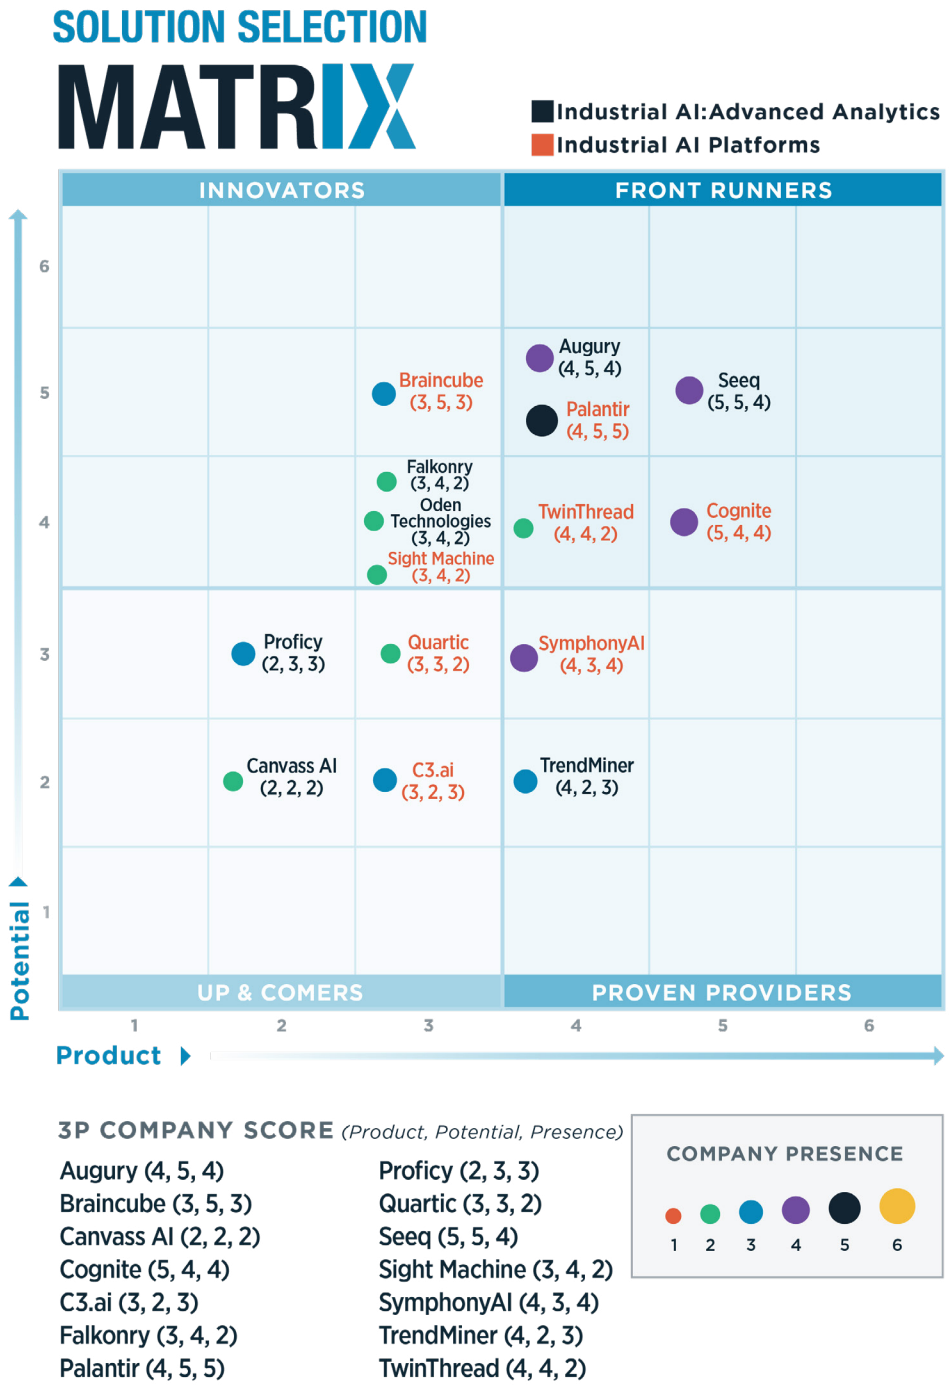

Figure 4: Industrial AI Platforms + Advanced Analytics SSM

The figure above presents a combined view of both Advanced Analytics vendors and Industrial AI Platforms, with each sub-category clearly distinguished. However, this combined view should be interpreted carefully: the product scores are not based on the same scoring scale across both groups. Advanced Analytics vendors are not evaluated on full data platform capabilities, so it would be misleading to compare a product score for an Advanced Analytics vendor directly against the same for a full-stack platform, since the scores reflect different evaluation criteria. Even so, end users should still consider both groups together when building a shortlist, because they overlap in key use cases and may compete head-to-head depending on the scope of the initiative.

Market Fit Grid

While the Solution Selection Matrix offers a three-dimensional framework to evaluate vendors across product, potential, and presence, it still needs additional context. Factors such as industry presence, key capabilities, service-intensiveness, and other attributes are required to make a well-rounded decision.

LNS Research's new Market Fit Grid framework is introduced to address this. This framework provides additional perspective on how well the solution fits its target markets across several aspects. This tabular framework provides binary (yes or no) insights across five dimensions—industry focus, core strengths, product types, solutions offered, and target user personas. Cells are highlighted where a vendor demonstrates significant focus or capability, offering a clearer picture of their market alignment.

In the below example, a sample Advanced Analytics vendor is illustrated as primarily targeting the Discrete industry, with core strengths in process monitoring and control. The vendor delivers its offering through a software-based model but does not include an open platform layer. To clarify, even if this vendor has one or two customers outside the discrete space or can be configured to support asset monitoring use cases, these aspects will not be highlighted in the grid because they do not represent the vendor's primary focus or strengths.

Industrial AI Platform				
Industries	**Discrete**	**Batch/Hybrid**	**Process**	**Infrastructure**
Core Strengths	**Asset Monitoring**	**Process Monitoring**	Process Control	**Data Model**
Products	Hardware	**Software**	**Services**	
Solutions	**Applications**	**Data Platform**	**DataOps**	
User Personas	**Maintenance/ Reliability Engineers**	**Process Engineers**	**Frontline Operators/ Supervisors**	**Data Scientists/ Analyst**

Figure 5: Sample Market Fit Grid

Market Fit Grid (Cont.)

Each vendor profile in this guide will include a Market Fit Grid that highlights where the vendor demonstrates meaningful focus and proven capability. As noted earlier, these highlighted callouts are not meant to be simple "checkbox" indicators, but represent true differentiators supported by real customer outcomes and validated case studies.

Additionally, the table below consolidates all these features across all covered vendors in one place. Because most vendors deliver a "software + services" model and span Advanced Industrial Analytics, Industrial Data Platforms, and Industrial DataOps capabilities, the "Products" and "Solutions" sections are not emphasized in this consolidated view.

Industrial AI Platforms: Market Fit Grid

	Industries	Core Strengths	User Personas
Braincube	Batch/Hybrid Process	Asset Monitoring Process Monitoring Data Model	Maintenance Engineer Process Engineer Data Specialist
C3.ai	Process Infrastructure	Asset Monitoring	Maintenance Engineer Data Specialist
Cognite	Process Infrastructure	Asset Monitoring Process Monitoring Data Model	Maintenance Engineer Process Engineer Frontline Ops Data Specialist
Palantir	Discrete Batch/Hybrid Process Infrastructure	Asset Monitoring Process Monitoring Data Model	Maintenance Engineer Process Engineer Data Specialist
Quartic.ai	Batch/Hybrid	Process Monitoring Process Control	Process Engineer Data Specialist
Sight Machine	Discrete Batch/Hybrid	Process Monitoring Data Model	Process Engineer Frontline Ops
SymphonyAI	Batch/Hybrid Process Infrastructure	Asset Monitoring Process Monitoring Data Model	Maintenance Engineer Process Engineer Frontline Ops Data Specialist
TwinThread	Batch/Hybrid Process Infrastructure	Asset Monitoring Process Monitoring Process Control Data Model	Maintenance Engineer Process Engineer Frontline Ops Data Specialist

Figure 6: Market Fit Grid for Industrial AI Platforms Overview

Even here, it is important to note that the callouts reflect what is most noteworthy, and not everything a vendor's product is technically capable of. For example, a vendor may offer a data model, but it may not be listed as a core strength if the vendor's unique selling proposition is stronger elsewhere. The grid is designed to highlight what the vendor is best at and most recognized for, rather than providing an exhaustive list of capabilities.

Summary and Recommendations

Software solution selection is a complex undertaking that demands multi-level, multi-regional, cross-functional, and inter-departmental collaboration. There are many pitfalls and challenges throughout the selection process. Manufacturers should apply a proven methodology to make the right choice quickly and confidently to eliminate worries around alignment, time, cost, and risk in solution selection.

Here are some final recommendations for manufacturers looking into Industrial AI solutions:

1. **Don't do AI for AI's Sake:** An Industrial AI and Advanced Analytics initiative, like any other strategic initiative, should be aligned with the company's overall business objectives. Understand your C-suite's objectives and build your Analytics strategy to deliver on those objectives.
2. **Balance Data and Analytics Needs:** At this point, it is well-known that good data is a fundamental prerequisite to good (and accurate) analytics. This includes establishing connectivity, conditioning and contextualization, hiring the right personnel (data engineers, solution architects, data stewards, etc.), and building governance processes. However, putting your analytics initiative on hold for months and years to fix data issues is not a practical solution; data issues must be continuously solved along with your analytics initiatives.
3. **Take a User Persona-Based Approach:** The most common way to apply Industrial AI is by beginning with the use cases- identifying business problems needing analytics. However, LNS Research's findings have shown that focusing on use cases is not the best approach to pivot from vision to value. Instead, take a user-persona-based approach, where you identify analytics needs for each type of user, such as engineers, business users, data scientists, executives, etc.
4. **Predictive Analytics is Not the Beginning or the End.** Most companies in the early stages of Industrial Transformation (IX) maturity have much to gain from having data available and deploying predictive analytics. For companies a little ahead in their Industrial Transformation (IX) journey, there's a lot more to gain beyond predictive to prescriptive and prognostic capabilities and build a closed-loop self-learning model.

Summary and Recommendations (Cont.)

5. **Create the Best Combination:** A typical medium-to-large size company will need two to three types of software/platforms. Let your business priorities, operational architecture, IT and OT organization structure, and manufacturing network's technical debt decide what combination of industrial applications, data platforms, and application platforms you will need for your Industrial AI needs.

6. **Beware of the Shiny Object Syndrome:** Time to value, ability to provide insights at scale to the targeted user personas, and technical capabilities often trump the precision of the model or the elegance of the solution.

7. **Understand Your Vendors' Core Focus and Financial Strength:** Many, but certainly not all, vendors are growing in today's manufacturing environment. When challenges arise, vendors retreat to their strengths and core product offerings. Make sure that your most pressing needs align with your vendors' core strengths and understand their current financial position and strategy to avoid unpleasant surprises down the line.

8. **Resist the Temptation to DIY:** Ask yourself, "What is our core competency? Are we in the software product development and support business?" Beware of sets of analytics components that are pitched as easy to integrate but, in reality, have initial costs that are the tip of the TCO iceberg, degree of difficulty, and time to value. Apply the Industrial Transformation Reference Architecture (section II) to decide if you need.

9. **Don't be Deterred:** Most importantly, do not let the challenges around this solution selection process prevent you from embarking on or accelerating your Industrial Transformation (IX) journey. Step change improvement is possible, and LNS Research is here to help you along that path.

Methodology References

This **Industrial AI: Advanced Analytics Solution Selection Matrix** research was conducted based on the following published materials:

1. What is Industrial AI: The Hype, The Facts, and The Path Forward
2. What is Advanced Industrial Analytics (AIA)?
3. Five Ways Industrial AI is Shaking Up Manufacturing (and Who's Doing It)
4. Recipes for Success: A Culinary Guide for Advanced Industrial Analytics
5. From the Industrial Internet of Things (IIoT) Platform to Industrial Transformation Reference Architecture
6. Changing Dynamics of the IIoT Market
7. How LNS Research Evaluates Vendors in Technology Domains

Vendor scores are summaries from their respective Vendor Profiles that will be available in the LNS Research membership library. Visit www.lnsresearch.com to learn more.

SOLUTION SELECTION MATRIX™

Industrial AI Platforms | 2026 Guidebook

Authors:

Matthew Littlefield
President & Research Leader
matthew.littlefield@lns-global.com

Vivek Murugesan
Research Analyst
vivek.murugesan@lns-global.com