KAI WAEHNER

2026

KAI WAEHNER LANDSCAPE 2026

PROCESS INTELLIGENCE LANDSCAPE

Vendors, Tradeoffs, and the Agentic AI Shift

Published: July 2026

Kai Waehner
ADVISORY FIELD CTO

kai-waehner.de

INSIDE THIS REPORT

Contents

Process intelligence combines process mining, process orchestration, and a decision gate into one enterprise architecture that shows how work actually runs, governs what happens next, and keeps automation and AI inside auditable boundaries. The process intelligence market is shifting faster than most enterprises realize.

Three forces are converging at once: the shift from periodic business process modeling to continuous AI-assisted operational intelligence, the rise of agentic AI that needs governed and data-fresh processes to function reliably, and a vendor landscape that is consolidating, crossing categories, and making architectural choices that will define enterprise lock-in for the next decade.

This landscape goes beyond vendor positioning. It maps the current landscape across platform openness and scope of process intelligence, explains why event-driven architecture and agentic AI are reshaping the requirements at every layer, covers the open standards that determine whether today's vendor decisions remain reversible tomorrow, and gives architects a practical framework for choosing tools based on their actual situation.

This is not a ranking. No vendor pays to appear here. The analysis is based on experience advising enterprises across the Global 2000 and organizations from scale-ups to digital natives, combined with ongoing research into product developments and enterprise adoption patterns. It is an independent practitioner perspective, not a formal research methodology.

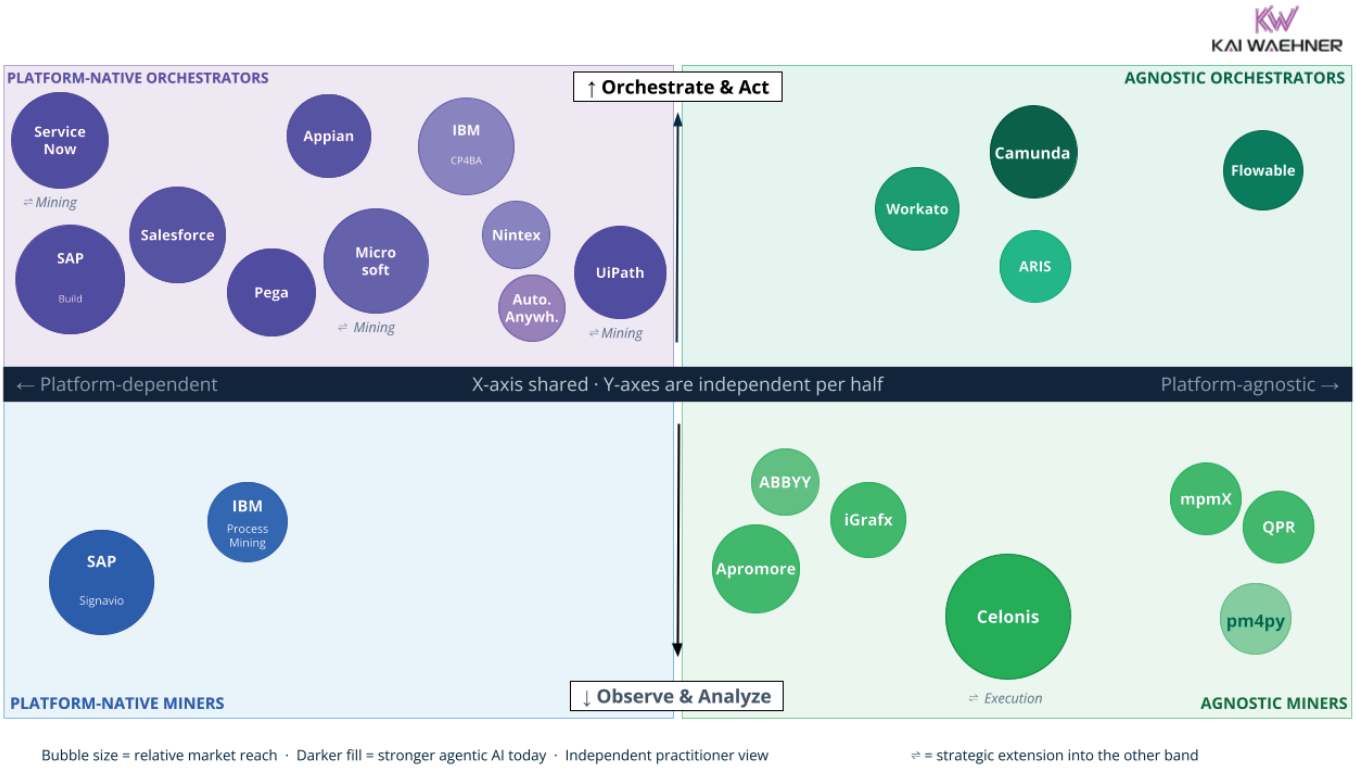

Figure 01. The Process Intelligence Landscape 2026. Bubble Size defines relative market reach while color depth stands for stronger agentic AI today.

01

CHAPTER 01

From BMP to Process Intelligence: Why the Label Changed

Process intelligence moves BPM beyond static models to continuous, real-time process optimization.

Business Process Management (BPM) has existed as a discipline for decades. The core idea was always sound: map how work should flow, automate the routine steps, give people visibility into what is happening. The execution fell short in predictable ways. Process models were built once and aged quickly and rarely reflected how work actually happened. Compliance teams relied on periodic audits. Process owners made decisions from dashboards populated by last night's data export.

The shift to process intelligence is not a marketing rename. It reflects a real architectural change. Process intelligence platforms continuously extract operational data from the systems where work actually runs. They identify where processes deviate from intended behavior. They surface automation opportunities based on real evidence. They increasingly act on those insights in real time. The label shifted from BPM to process intelligence because the discipline itself shifted: from designing processes to continuously observing, analyzing, and acting on them in real time.

The process mining market has shifted from specialist deployments to mainstream enterprise adoption, with growth accelerating as Agentic AI creates new demand for the operational context that process intelligence provides. **That growth is driven by three converging demands:** enterprises needing real visibility into operational reality before automating it, the push to identify automation opportunities with evidence rather than assumptions, and the rising urgency of agentic AI deployments that require governed, data-grounded process context to function reliably.

The vendor consolidation and platform maturity now visible across this market confirm that the pilot-stage experimentation phase is over.

A note on how the analysts scope this market, because it differs from how this landscape is organized. Gartner replaced its Process Mining Magic Quadrant with a Magic Quadrant for Process Intelligence Platforms in May 2026, scoped to process mining, task mining, modeling, and monitoring. Orchestration sits in a separate Gartner quadrant, the BOAT market for business orchestration and automation. Forrester splits it the same way, with one report for process intelligence and another for adaptive process orchestration. Two adjacent markets, two sets of vendors, two buying motions.

This landscape spans both on purpose. Analysts map how vendors go to market. Architects design the system those products plug into, and that system needs mining to see what runs, orchestration to govern what happens next, and a decision gate to keep automation and AI inside auditable boundaries. Buyers shop product by product. Architects design across the boundary. I make the full argument in a separate post on process intelligence as mining, orchestration, and the decision gate.

02

CHAPTER 02

What Process Intelligence Actually Covers

Enterprise process intelligence is built on three pillars: process mining, orchestration, and decision governance. This chapter will dive into all of them.

Three primary categories dominate enterprise investment and vendor competition.

Process Mining

Process mining and analytics form the observation layer. This covers discovering actual process behavior from event logs, system data, and user interaction traces. Conformance checking, root cause analysis, predictive analytics, simulation, and process digital twins sit on top.

The core input is the event log: a structured record of what happened, when, and in what system. Every enterprise already produces this data across SAP, Salesforce, ServiceNow, and other operational systems. **Process mining extracts it, reconstructs the actual process flow,** and compares it against the intended model. The gap between the two is where optimization opportunities live.

The most significant technical evolution in this layer is Object-Centric Process Mining (OCPM). Traditional process mining analyzes one business object in isolation: one order, one claim, one ticket. OCPM analyzes interrelated objects simultaneously: orders, items, shipments, invoices, and customers, and how they interact across the process. Wil van der Aalst, the originator of process mining as an academic discipline and Chief Scientist at Celonis, describes the shift as moving from a 2D X-ray of operations to a 3D MRI. Automotive manufacturers use OCPM for supply chain analysis. Airlines apply it to minimize flight delays by analyzing the interaction between aircraft, crews, gates, and maintenance simultaneously. **OCPM has become a standard evaluation criterion.** Its absence is already an eliminating factor for enterprises with supply chain, order-to-cash, or procure-to-pay processes in scope.

OCPM vendor support is uneven and matters for selection. The vendors with object-centric capability available today are Celonis, IBM Process Mining, mpmX, QPR, ServiceNow, and the open-source pm4py from Process Intelligence Solutions. Apromore supports it as well, now inside Salesforce. Microsoft Power Automate Process Mining has OCPM in public preview. SAP Signavio delivers it through Process Networks, which reached general availability in mid 2026. UiPath has it on the roadmap. iGrafx and ABBYY Timeline do not support it.

Process digital twins deserve a specific mention within this layer. Leading platforms now let enterprises simulate process changes in a virtual model before implementing them in production. The what-if layer is becoming a baseline expectation, not a premium feature. Celonis, ABBYY Timeline, and mpmX offer it as a core capability.

Process Orchestration

Process orchestration and automation form the action layer. This covers executing workflows, enforcing business rules, coordinating humans, systems, bots, and AI agents, and producing the audit trails that compliance requires.

Where mining observes what is happening, orchestration changes what happens next. A workflow definition specifies the sequence of steps, the conditions under which decisions are made, the systems and people involved at each stage, and the boundaries within which automated actions can execute without human approval. For regulated industries, the orchestration layer is also where compliance is enforced: every step is logged, every human decision is recorded, and every deviation from the expected path is visible.

The technical foundation of enterprise process orchestration is BPMN 2.0 (Business Process Model and Notation), the open standard for modeling process flows as portable, auditable XML definitions. DMN (Decision Model and Notation) handles decision logic. CMMN (Case Management Model and Notation) handles unstructured case management. Together these three standards define the governance layer that separates enterprise orchestration from general-purpose workflow automation. A process built on these standards can be read by a compliance officer, modified by a developer, versioned in Git, and migrated between environments without rebuilding it inside a vendor's proprietary interface.

Orchestration platforms range from open standards-based engines like Camunda and Flowable, where all process logic is owned and operated by the organization, to platform-native orchestrators like Pega and Appian, where the process definitions live inside the vendor's proprietary environment. The architectural distinction between these approaches is covered in detail in section 3.

Decision Governance

A third dimension runs through both layers. Decision governance covers audit trails, conformance monitoring, and DMN-based decision modeling. It answers the question: did the process run as intended, and can the organization prove it? In operational terms this is the decision gate. Every automated action and every AI recommendation passes through it. The gate evaluates each one against business rules, confidence thresholds, and regulatory constraints. What passes proceeds. What does not is escalated to a human or rejected with a documented reason. Governance is the property. The decision gate is the mechanism that enforces it.

For process mining, governance means conformance checking: comparing actual process behavior against the intended model and flagging deviations in real time. For orchestration, it means producing immutable execution histories, enforcing approval gates, and providing the oversight mechanisms that regulators, auditors, and compliance teams require. For agentic AI deployments, it means defining in advance what an AI agent can decide autonomously and what must always escalate to a human.

Most mature mining and orchestration platforms provide decision governance as part of their core offering. It is a meaningful evaluation dimension for regulated industries and called out in vendor profiles where it is a differentiator. The EU AI Act has made governance a legal requirement for enterprises deploying AI in high-risk categories across EU markets, which is covered in section 7.

03

CHAPTER 03

How to Read this Landscape

The landscape organizes vendors across two dimensions, but it is not a pick-one-quadrant chart. It is two stacked decisions. The bottom half is the mining decision. The top half is the orchestration decision.

Most enterprises end up choosing at least one mining vendor and one or more orchestration vendors, because no single product covers both well across a multi-vendor environment. Read the two halves as separate maps, decide each one, then combine the results. Reading both correctly before evaluating individual vendors prevents the most common category errors in platform selection.

The axes and markers

The chart is split into two horizontal bands that share one X-axis. The X-axis measures platform dependency on the left and platform agnosticism on the right. The bottom band is mining, where vendors observe and analyze what is happening and why. The top band is orchestration, where vendors act and govern what happens next. The two bands are read independently. A vendor's horizontal position carries the same dependent-to-agnostic meaning in both bands, but agnosticism itself means something different in mining than in orchestration, which the next two sections make explicit. Together the chart identifies whether a vendor is locked to one ecosystem or portable across many, and whether it primarily observes or primarily acts.

A small number of vendors carry a bridge marker (two arrows), reserved for cases where the extension into the other band is a deliberate strategic direction, not just a secondary capability added through acquisition. Vendor position within each band reflects primary identity only. A vendor with a marker does not move between bands, and vendors whose secondary capabilities are real but not strategy-defining are covered in the profiles without a marker. A marker signals that the extension exists; the extension does not inherit the vendor's horizontal position or meet the other band's agnosticism criteria. Vendors with distinct products for each band, such as SAP and IBM, appear twice instead of carrying a marker.

Bubble size reflects relative market reach. Darker fill reflects stronger agentic AI capability in production today.

There is no single right quadrant, and there is no single right vendor. You are making two decisions with different agnosticism considerations, then combining them. Many organizations land on an agnostic miner and a platform-native orchestrator, or the reverse. The split is a normal and often correct outcome, not a contradiction. Every position represents tradeoffs that depend on existing commitments, team composition, regulatory environment, and strategic direction.

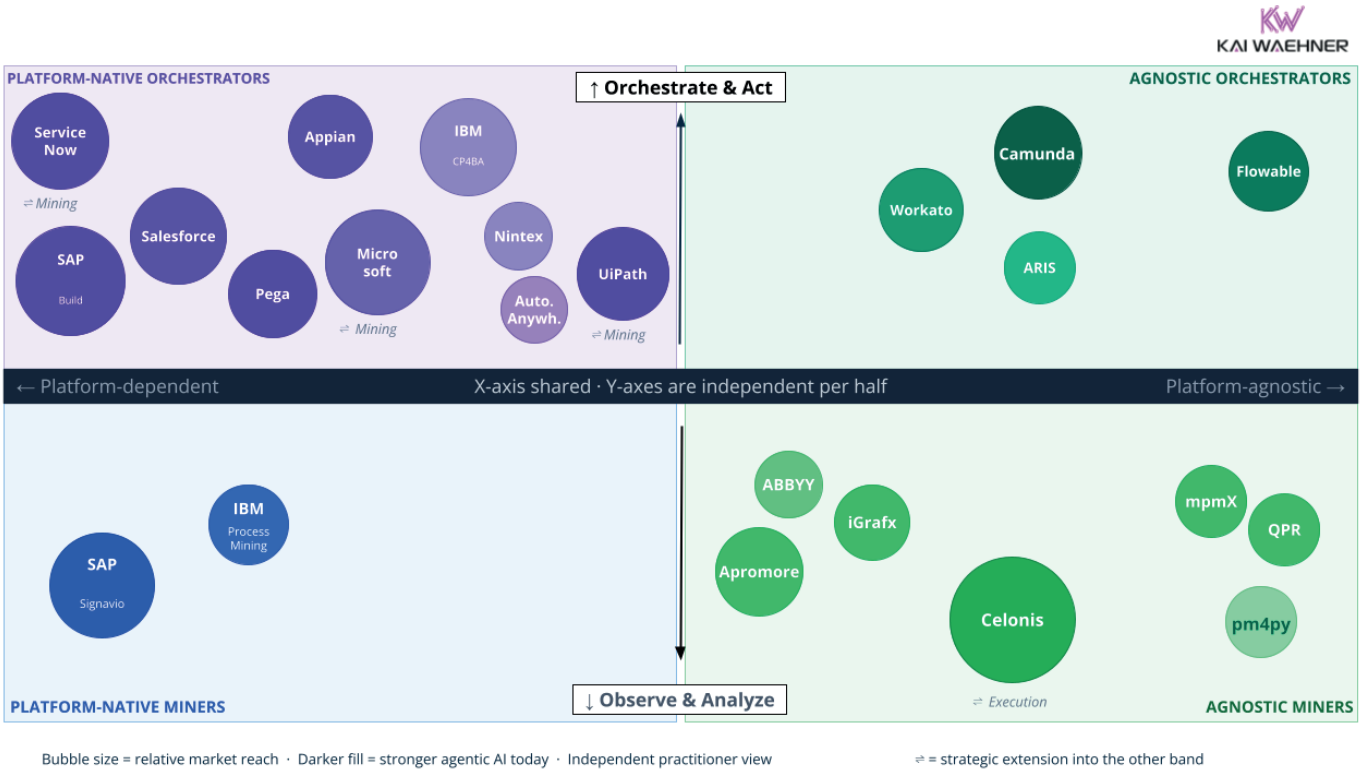

Figure 02. The Process Intelligence Landscape 2026. Bubble Size defines relative market reach while color depth stands for stronger agentic AI today.

Agnostic orchestration: open standards and code you own

In the orchestration zone, agnosticism is about technical and architectural independence. The question is not whether a platform connects to SAP, Salesforce, or ServiceNow. Most platforms can. The question is where the process logic lives, who owns the code, and whether the process definition is portable when the organization needs to change or extend it.

In platform-native orchestrators like Appian, Pega, or ServiceNow, the process logic is expressed in the vendor's proprietary environment. Appian uses its own low-code expression language. Pega uses its own proprietary rules engine. The business logic lives in the vendor's platform, not in code the organization owns and runs independently. The process definition is stored in a proprietary format. Adopting these platforms means committing to a development environment and an exit path that grows more complex with every workflow built. For organizations with long-term platform commitments, that is an acceptable and often efficient tradeoff.

In an agnostic platform like Camunda or Flowable, the process definition is a portable BPMN 2.0 XML file expressed in an open standard that any compatible engine can read, modify, and execute. The business logic is written in any language: Java, Python, Go, Node.js, or generated by GenAI coding agents, using standard APIs to connect to the orchestration engine. The worker code runs on any infrastructure and belongs entirely to the organization. Moving from one BPMN-compliant engine to another is possible with reasonable effort. Moving a Pega or Appian process definition to any other platform is not.

This is the same architectural principle that made Kafka the standard for event-driven integration. The parallel is not about messaging. It is about where the logic lives. Kafka producers and consumers are implemented in any language, with any tool, on any infrastructure. The broker decouples systems from each other. The business logic lives entirely in the client code, owned and operated by the organization. Camunda works the same way: the orchestration engine coordinates the flow, but the business logic lives in worker code that the organization writes, owns, and runs independently of the vendor.

The open standards that make this durability possible are BPMN 2.0 for process flow definitions, DMN for decision logic, and CMMN for case management. These standards and their implications for vendor selection are covered in full in section 5.

For agentic AI deployments, MCP and A2A are the open standards for agent tool connectivity and agent-to-agent coordination. Both are now governed by the Linux Foundation. Vendors adopting them create integration surfaces that work across the entire AI ecosystem. Vendors that build proprietary agent protocols risk the same lock-in critique that has historically limited closed BPM platforms.

That risk is not theoretical. Every major platform vendor is now repositioning as agentic AI infrastructure, applying the same proprietary logic to the agent layer that previously applied to process definitions. The platform-native orchestrators in this landscape are each making this move. Salesforce has repositioned its entire platform, including data, workflows, and business logic, as an API-first infrastructure for AI agents. ServiceNow has AI Agent Fabric. IBM has watsonx orchestration. SAP has the BTP agentic layer. Every platform vendor is making the same claim: run your agents on our infrastructure, within our guardrails, on our data. The governance offered within each of these platforms is real. The portability is not. An agentic process built entirely inside one vendor's platform is not portable. The lock-in is identical to what closed BPM platforms created a decade ago, just applied to a new layer of the stack. Organizations building process intelligence for a multi-vendor future need an agnostic layer that sits above any single platform's agentic infrastructure.

Agnostic mining: connect to any system, owned by none

In the mining zone, agnosticism has a different meaning. Here the question is data source independence and commercial independence: not the portability of code, but the freedom to connect to any enterprise system without being owned by any of them.

An agnostic miner connects to SAP, Salesforce, and ServiceNow simultaneously because it is licensed and governed independently from all three. It uses its own connectors for each system without requiring any single vendor's integration infrastructure. SAP Signavio can technically connect to non-SAP data sources, but doing so requires routing through SAP's own middleware: SAP BTP, SAP Datasphere, or SAP Data Intelligence.

Choosing SAP Signavio for cross-system mining means choosing SAP as the integration intermediary for all data sources, not just SAP systems. For organizations running complex multi-vendor environments, the constraint is real.

A larger shift sits underneath this. The systems holding the most critical business data are tightening how external tools and AI agents reach it. SAP's 2026 API policy restricts third-party AI agents from its APIs and routes agentic traffic toward its own endorsed stack. Salesforce and ServiceNow apply quieter versions through Slack API limits and inbound rate throttling. The lesson for mining is direct. A platform that depends on a vendor's own middleware to reach that vendor's data inherits that vendor's access policy. One that consumes from an independent, owned data layer does not. The durable answer is the pattern enterprises used to offload mainframes. Extract once into a layer you own, keep it current with CDC or event streaming, and let mining tools and agents consume from there. This is covered in depth in a separate post on data ownership in the age of agentic AI.

Commercial independence matters as much as technical independence. A platform-native miner is always a line item within a larger platform contract. Its pricing, roadmap, and support terms are bundled with the broader vendor relationship. An agnostic miner is licensed and governed independently. Pricing, roadmap, and support do not depend on any other platform decision the organization makes.

Data format independence is the third dimension. Some miners use proprietary query languages regardless of their agnosticism on data sources. Celonis uses PQL. SAP Signavio uses SIGNAL. XES export support is the relevant portability check for classic event logs, and OCEL 2.0 export is the equivalent check for object-centric data. Together they determine whether historical process data can move if the vendor changes.

Some agnostic miners take this further by analyzing data where it already lives, without requiring movement to a third-party platform. mpmX integrates directly into existing BI stacks like Qlik and Databricks. QPR delivers process mining as a Snowflake Native App, keeping analysis inside the organization's own Snowflake environment. For organizations in regulated industries that cannot move operational data to a third-party cloud, this data-in-place approach is a real differentiator.

04

CHAPTER 04

Reading the Zones

Each zone in the landscape reflects a different architectural commitment. The profiles in this chapter describe what that commitment means in practice for vendor selection, deployment, and long-term flexibility.

Platform-native orchestrators

Organizations with deep long-term commitment to one platform often find this an acceptable constraint and the fastest path to results. Organizations that need process intelligence to span SAP, Salesforce, ServiceNow, and other platforms simultaneously face a structural limitation that compounds with every workflow built.

Appian

Appian integrates BPA, RPA, low-code, AI agents, and a data fabric into one platform. Agent Studio is generally available, enabling business users to define agent goals in natural language while agents inherit process context, data access, and guardrails from the underlying platform. Process mining capabilities are included through its Lana Labs acquisition, complementing the orchestration and automation core. Strong deployment history in financial services, public sector, and healthcare, where its low-code approach and pre-built process content reduce time to production. Process logic is expressed in Appian's proprietary environment, which means the platform commitment deepens with every application built.

Automation Anywhere

Automation Anywhere has evolved its RPA foundation toward full agentic process automation through the Process Reasoning Engine (PRE), which adds AI-driven reasoning, context awareness, and adaptive workflow execution that recovers from failures without manual intervention. Process Composer provides BPMN 2.0-based workflow design for orchestrating AI agents, bots, APIs, and humans in a single flow. The Mozart Orchestrator extends this into multi-agent coordination across internal and external agent ecosystems. MCP and A2A interoperability are supported. Strong enterprise connectivity including Salesforce, SAP, and mainframe systems. The orchestration and BPMN capabilities are recent additions. Organizations with complex, long-running process governance requirements should evaluate maturity depth alongside the broader automation capabilities.

IBM (CP4BA)

IBM combines process orchestration, case management, and process intelligence across two primary products. CP4BA provides enterprise-grade process orchestration and case management with strong compliance and governance capabilities. watsonx Orchestrate is the agentic AI layer, enabling organizations to build, deploy, and coordinate AI agents across SAP, Salesforce, ServiceNow, Workday, and other enterprise applications. IBM supports MCP and A2A protocols for agent interoperability. Process intelligence, built on the myInvenio acquisition and enhanced with watsonx AI, provides OCPM support and prescriptive analytics. The full stack requires significant implementation expertise and is typically sold as part of a broader IBM infrastructure commitment rather than as a standalone product. Best suited for enterprises already deeply inside IBM's infrastructure who need enterprise-grade governance alongside agentic AI.

Microsoft Power Platform

Microsoft's primary process intelligence contribution is the Power Platform: a low-code development, automation, and agent orchestration layer deeply embedded in the Microsoft 365 and Dynamics 365 ecosystem. Copilot Studio is the agent orchestration surface, supporting multi-agent coordination, A2A protocol for cross-vendor agent interoperability, and MCP server connectivity. Power Automate has evolved toward goal-based process orchestration: users define business objectives and AI agents determine the execution sequence across systems. Power Automate Process Mining adds process intelligence within the stack, with OCPM in public preview. No on-premises option across the platform.

Nintex

Nintex serves mid-market organizations and enterprise departments with a full-cycle automation platform covering process management, workflow, RPA, document automation, and agentic orchestration. Nintex Orchestration enables phase-based workflow design that coordinates AI agents, human decision points, and business systems within a governed framework. Agent Designer allows teams to embed AI agents directly into orchestration flows. The K2 product line provides an on-premises option for regulated deployment requirements. Strong ease of use and broad connectivity make it a practical choice for organizations that need broad automation coverage without enterprise-scale complexity.

Pegasystems

Pega combines BPM, RPA, case management, low-code development, and AI agents in one integrated platform. Its architectural approach to agentic AI is distinctive: rather than relying on runtime prompts, Pega anchors agents in pre-approved workflows built in Blueprint, its AI-driven design environment. This makes agent behavior more predictable and auditable, which resonates particularly in regulated industries. The Agentic Process Fabric orchestrates agents across Pega and third-party systems. Pega supports MCP and A2A protocols for agent interoperability. Process mining and task mining capabilities are included in the platform alongside its core orchestration and case management strengths. Audit trail and compliance capabilities are strong throughout. TCO is high and specialist skills are required.

Salesforce

Agentforce delivers agentic AI deeply embedded in Salesforce CRM data and workflows. The Atlas Reasoning Engine handles multi-step reasoning and planning. The Einstein Trust Layer provides data governance, audit trails, and privacy controls for agent actions. Flow Orchestration handles deterministic workflow automation inside the ecosystem, while Agentforce handles more dynamic AI-driven processes. Salesforce supports both MCP and A2A protocols, enabling Agentforce agents to connect to external tools and coordinate with agents from other vendors. Process mining now comes from Apromore, which Salesforce acquired in November 2025 and is integrating into Agentforce. Cross-system orchestration outside the Salesforce ecosystem is handled through MuleSoft, which requires separate licensing.

SAP Build

SAP Build is SAP's low-code platform for building applications, automating processes, and orchestrating workflows on the Business Technology Platform. Joule Studio, now generally available within SAP Build, enables developers to build and deploy custom AI agents with support for MCP and A2A protocols.

The key distinction from most other orchestration vendors in this landscape is architectural intent. SAP Build is designed primarily as an extension layer for SAP's own application suite rather than as a standalone cross-platform orchestrator. It is most valuable when the processes being automated predominantly live within the SAP landscape. Connecting to non-SAP systems requires SAP Integration Suite. For organizations deeply committed to SAP, it is a natural and efficient choice. For organizations coordinating processes across SAP, Salesforce, ServiceNow, and other platforms simultaneously, an agnostic orchestrator is the more practical starting point.

ServiceNow

ServiceNow is a workflow orchestration and IT service management platform that has evolved into an enterprise control plane for agentic AI. The Now Platform coordinates workflows across ITSM, HRSD, CSM, and other business domains, with AI agents operating within deterministic workflow boundaries.

The Context Engine, built on ServiceNow's Service Graph and Knowledge Graph, provides AI agents with real-time cross-enterprise context for informed decision-making. The agentic AI suite covers agent development through AI Agent Studio, multi-agent coordination through AI Agent Orchestrator, and enterprise-wide governance through AI Control Tower, including oversight of agents deployed on Microsoft and third-party platforms.

Process mining, added through the UltimateSuite acquisition, provides workflow analytics within the Now Platform. Object-centric analysis is supported through multidimensional process maps.

UiPath

UiPath started as an RPA leader and built process and task mining specifically to feed that automation engine: mining identifies what to automate, RPA and AI agents execute it, Maestro orchestrates the end-to-end flow. That discovery-to-execution loop is tighter in UiPath than in most other vendors in this landscape. Maestro adds BPMN and DMN-based process orchestration with an event-sourced execution engine, durable execution, and human-in-the-loop governance. The AI Trust Layer governs agentic behavior across the platform. UiPath supports MCP for agent interoperability. The February 2026 acquisition of WorkFusion added domain-trained KYC and AML agents, and Maestro Case, launched in June 2026, extends the orchestration suite into dynamic case management for long-running, judgment-heavy work such as claims, disputes, and financial crime investigations. OCPM support remains on the roadmap.

Agnostic orchestrators

In this zone, agnosticism means the platform is not owned by, or structurally biased toward, any of the application ecosystems it orchestrates. In the strongest cases, process definitions are expressed in open standards: versioned in Git, reviewed in pull requests, tested in CI/CD pipelines, and migrated between environments using standard tooling, with business logic written in standard code the organization owns. Not every vendor in this zone meets that bar, which is exactly what the profiles below distinguish.

This zone contains tools at different governance depths. Camunda and Flowable are full BPMN orchestration platforms with compliance-grade audit trails and human-in-the-loop governance. ARIS spans process modeling, governance, and execution support across the full BPM lifecycle. Workato is an integration-first platform with workflow coordination capabilities and lighter process governance. These orientations serve different use cases and are not interchangeable.

ARIS

ARIS originated at IDS Scheer, a pioneering BPM company acquired by Software AG, and has been a fixture in enterprise process governance for decades before becoming a standalone business under Silver Lake ownership. It spans process modeling, mining, simulation, GRC (governance, risk, and compliance), and execution support within one suite. Silver Lake's multi-year investment commitment is confirmed, though enterprise architects should factor the PE ownership structure into long-term vendor stability assessments.

Whether the independence accelerates or constrains product innovation is still an open question. The separation is too recent to judge. Strong in regulated industries and governance-heavy environments. A high learning curve is consistent user feedback. Execution capabilities are less mature than its modeling and governance strengths, making it a stronger choice for process governance and analysis than for operational orchestration at scale.

Camunda

Camunda is the reference case for platform-agnostic process orchestration. The Zeebe engine is cloud-native, horizontally scalable, and event-sourced by design. BPMN 2.0 and DMN are open standards throughout. The process definition is a portable XML artifact that any tool or AI coding agent can read, modify, and redeploy without touching a proprietary GUI. Workers are written in any language: Java, Python, Go, or Node.js, using standard REST and gRPC APIs. Camunda ships production-ready MCP and A2A connectors, enabling AI agents to interact with Camunda processes as tools and coordinate across agent systems while maintaining BPMN-based governance and full audit trails. Agents can be embedded directly within BPMN process models, blending deterministic control flow with AI-driven decision steps.

Camunda provides operational control over process definition versioning that most orchestration tools do not offer. When a process definition changes, running instances can continue on their original version or be migrated without interruption. For long-lived processes like insurance claims or regulatory workflows, this is a production requirement, not a convenience feature. Camunda also ships Optimize as a lightweight process analytics component, complementing dedicated mining tools rather than replacing them.

Camunda includes a visual BPMN modeler for process design and review, accessible to non-developers. The implementation of business logic still requires developer-written worker code. Organizations that need business users to build and configure automation end to end without developer involvement should evaluate the platform-native orchestrators zone first.

Flowable

Flowable is a BPMN and CMMN-native orchestration platform with deep roots in case management for knowledge workers. Its core engines are Apache 2.0 licensed open source, making it a natural choice for organizations that want to embed process orchestration directly into custom Java applications or products and ship them to customers without engine licensing fees.

The platform supports both embedded deployment inside existing applications and standalone enterprise use. The enterprise platform adds governed agentic AI: the latest release includes multi-agent orchestration via MCP and A2A, human-in-the-loop controls, and full audit trails for AI-driven decisions. This positions Flowable well for regulated industries like banking, insurance, and healthcare where AI governance and compliance auditability are non-negotiable. A mature codebase with a governance-first design philosophy and a strong fit for compliance-heavy environments.

Workato

Workato is an iPaaS platform that has evolved into a full agentic orchestration layer. Its recipe-based, low-code builder allows business users and non-developers to connect systems, build workflows, and deploy AI agents without writing code. Enterprise MCP is a core capability: existing workflows, integrations, and APIs are transformed into governed, agent-ready skills that any MCP-compatible AI agent can invoke.

Pre-built Genies cover key business functions including sales, HR, IT, and customer experience. Agent Studio enables custom agent development without code. Governance and audit capabilities are enterprise-grade for integration and agentic workflows.

The platform does not offer BPMN-based process governance or the long-running process execution model that compliance-heavy regulated industries require. Organizations with those requirements should evaluate dedicated BPMN orchestration platforms alongside Workato.

Agnostic miners

In this zone, vendors connect to any enterprise system regardless of which vendor built it, licensed and governed independently from all of them. They act as a neutral analytical layer across the entire enterprise application landscape without structural bias toward one vendor's data or processes. The vendors in this zone vary significantly in analytical depth, OCPM support, and data deployment options. The profiles below cover those distinctions.

ABBYY Timeline

ABBYY Timeline combines process mining with ABBYY's core strength in intelligent document processing (IDP), making it distinctive among agnostic miners. Where most process mining tools analyze structured event logs from enterprise systems, ABBYY Timeline can enrich process context with insights from document content, combining system-level event data with document and task-level analysis in a single platform. This is particularly valuable for processes where document handling is a significant part of the workflow, such as invoice processing, KYC, and contract management.

The platform creates explicit process digital twins and supports task mining alongside system-level mining. No OCPM support. ABBYY is actively investing in agentic AI capabilities through its Process AI platform, with governance and compliance as explicit design priorities.

Apromore

Apromore covers the full process mining lifecycle in a single no-code environment: task mining, OCPM, process modeling, digital twin simulation, root cause analysis, and conformance monitoring. The GenAI copilot supports process redesign by generating improvement recommendations directly from mining findings. Apromore originated from research at the University of Melbourne, giving it deep roots in process mining theory.

Salesforce completed its acquisition of Apromore in November 2025 and is integrating it into Agentforce. The technology still connects across multiple systems, so it stays in the agnostic miners zone for this edition. What changed is commercial independence. Pricing, roadmap, and support now run through Salesforce rather than a standalone vendor. Its no-code lifecycle coverage and OCPM support remain real technical strengths. Treat Apromore as a capable cross-system miner whose direction is now set by a single platform owner, and watch whether its neutrality narrows as the Agentforce integration deepens. On that trajectory it becomes a candidate to move into the platform-native miners zone in the next edition.

Celonis

Celonis holds the largest market share in process mining and is the platform that defined what enterprise process intelligence means in practice. Its OCPM-compliant Process Intelligence Graph is a unified cross-system data model that continuously ingests real-time process events from SAP, Salesforce, ServiceNow, and any other operational system, reconstructing how work actually flows across the enterprise rather than how it was designed to flow. Identifying and closing that gap is the core value proposition. OCPM support, process digital twin simulation, AI-generated recommendations, and conformance monitoring are all built on top of this graph. The cross-system view this produces is the primary architectural reason organizations with multi-vendor environments choose an agnostic miner over a platform-native one.

Celonis uses its own query language (PQL), and its object-centric mining runs on this proprietary graph, not on OCEL 2.0. OCEL export, not the OCPM label, is the portability check. In standard deployments, event log data from source systems is extracted and replicated into Celonis's cloud platform to build and query the Process Intelligence Graph. This is a data sovereignty consideration for regulated industries. Celonis Data Core enables integration with existing data lakes like Snowflake without requiring full data replication, reducing this constraint for organizations that already have a centralized data environment.

Celonis has moved beyond mining into execution and decision intelligence. It acquired Symbio for process design and modeling, and the Orchestration Engine from Emporix for the workflow execution layer that closes the loop from insight to action. In May 2026 it launched the Celonis Context Model, an evolution of the Process Intelligence Graph that it positions as the context layer between enterprise data and AI agents. Alongside it, Celonis announced the acquisition of Ikigai Labs to add simulation, planning, and forecasting through large graphical models built for structured enterprise data. As part of the agreement, Celonis gains exclusive rights to MIT-owned patents, MIT becomes a Celonis shareholder, and Ikigai co-founder Devavrat Shah, a professor of AI at MIT, joins Celonis as Chief Scientist of Enterprise AI.

AgentC connects process context and recommended actions to AI platforms including Microsoft Copilot Studio, Amazon Bedrock, and Salesforce Agentforce. Celonis exposes its graph as an MCP server for agents to query directly. The direction is clear. Celonis wants to be the context layer for enterprise AI, not only the market-leading mining tool.

iGrafx

iGrafx is one of the few vendors in this landscape that still leads with BPM as its primary positioning. It covers the full traditional BPM scope: process modeling, mining, simulation, and GRC (governance, risk, and compliance) in one platform. The Pia assistant enables natural language process discovery, allowing business users to explore and analyze processes without requiring technical expertise. No OCPM support. iGrafx is actively extending toward agentic AI through a partnership with NTT DATA INTRAMART, aimed at unifying process intelligence with enterprise execution and autonomous decision support. Strong presence in regulated industries where governance depth and compliance auditability are primary evaluation criteria. Unlike ARIS, whose center of gravity is governance and execution support, iGrafx is used primarily for process analysis, which is why it sits in the mining band.

mpmX (MEHRWERK)

mpmX follows a data-in-place philosophy. Rather than replicating data to a separate mining environment, it analyzes process data where it already lives inside existing platforms including Qlik, Snowflake, and Databricks. This directly addresses data sovereignty concerns for organizations in regulated industries that cannot move operational data to a third-party cloud. OCPM is supported alongside process digital twins and AI-driven recommendations. A production MCP server enables AI agents to query process intelligence context directly. Strong footprint in the DACH region with active expansion across broader Europe and North America. MEHRWERK was acquired by Fortino Capital, a European technology investor, with a mandate to accelerate product development and international expansion.

pm4py (Process Intelligence Solutions)

pm4py is the leading open-source process mining library, written in Python and originally developed at the Fraunhofer Institute for Applied Information Technology. It is now managed by Process Intelligence Solutions, a Fraunhofer spin-off, under an AGPL license with a separate commercial license for closed-source use. It comes out of the academic process mining community associated with Wil van der Aalst, and was an early, complete implementation of object-centric process mining and OCEL 2.0. It is a library, not a packaged platform. There is no business-user GUI, no managed connectors to SAP or Salesforce, and no built-in governance suite. For data science teams that want to embed process mining into custom Python applications, run OCPM on their own infrastructure, or avoid commercial licensing for analytical work, it is the natural open-source option. For business-led deployments that need extractors, dashboards, and support, a packaged platform fits better.

QPR Software

QPR ProcessAnalyzer was the first process mining solution to run natively inside the Snowflake AI Data Cloud, keeping analysis where the data already lives without requiring any data movement or separate integration setup. This Snowflake-native architecture delivers enterprise-grade process mining with strong data sovereignty guarantees and a lower total cost of ownership than most alternatives in this zone. QPR now supports object-centric process mining and ships an MCP interface that lets AI agents query process context directly. GenAI capabilities include synthetic event log prediction for scenarios where historical data is incomplete. Geographic presence is strongest in the Nordics with active international expansion underway.

Platform-native miners

The right time to buy mining from your platform vendor is when the processes to be analyzed are predominantly native to that platform, procurement is already consolidated there, and cross-system mining is not a priority. Platform-native miners have deep pre-built extractors, pre-configured data models, and ready-to-use accelerators for their own ecosystem. That depth is the real advantage. The limits appear when analysis needs to span multiple vendor ecosystems simultaneously or when OCPM is required.

Each vendor in this zone built its mining capability through acquisition rather than organic development. That history explains both the depth of native ecosystem integration and the variation in analytical maturity across the zone.

IBM Process Mining

IBM Process Mining is built on the myInvenio acquisition and deeply integrated with watsonx AI. OCPM support combined with watsonx-powered prescriptive analytics distinguishes it analytically: the platform identifies process issues and recommends specific corrective actions with quantified expected benefits.

A generative AI copilot powered by watsonx enables natural language process discovery and optimization, reducing the technical expertise required to extract insights from complex process data. The platform connects directly to watsonx Orchestrate, enabling process mining insights to feed AI agent deployments and automation initiatives within the IBM stack.

Delivered via Red Hat OpenShift, it supports hybrid cloud and on-premises deployment, making it a viable option for regulated industries with strict data residency requirements. Implementation and configuration typically require IBM expertise and are best supported through IBM's partner network. The strongest fit is for enterprises with existing IBM infrastructure commitments where OCPM depth, prescriptive AI, and on-premises deployment options are all evaluation requirements.

SAP Signavio Process Intelligence

SAP Signavio Process Intelligence covers process mining with 184 industry benchmarks and deep integration with SAP systems. It uses SiGNAL (SAP Signavio Analytics Language), its proprietary query language, for process analysis. CDC-based continuous data replication via SAP Datasphere enables real-time extraction of process-relevant data directly into the process mining layer.

SAP Signavio can connect to non-SAP data sources, but doing so requires routing through SAP's own integration middleware, adding complexity and deepening SAP ecosystem dependency. Its object-centric process mining capability, Process Networks, reached general availability in mid 2026 as part of a next-generation Process Intelligence release, with what-if process simulation following in open beta. SAP Signavio also introduces agent mining capabilities for governing and monitoring AI agents operating across business processes. SAP offers both Process Intelligence (the full suite) and Process Insights (a simplified offering). Clarifying which product is in scope is an important early step in any SAP Signavio evaluation.

Tools that operate at adjacent architectural layers like pipeline orchestration or durable execution infrastructure are covered in section 8 of this article.

05

CHAPTER 05

Open Standards on Both Sides of the Landscape

Open standards reduce vendor lock-in across the process intelligence stack, in orchestration, mining, and agentic AI. Prioritizing vendors committed to openness keeps your options flexible and makes best-of-breed combinations possible.

Open standards create portability and reduce vendor lock-in at every layer of the process intelligence stack. They matter in orchestration, in mining, and in agentic AI. Choosing vendors that commit to open standards across all three reduces the cost of changing tools and increases the ability to combine best-of-breed solutions.

Orchestration standards

BPMN 2.0 defines the process flow. DMN defines the decision logic. CMMN defines case management. All three are OMG standards with stable, widely supported specifications. A process definition built on these standards is a portable XML artifact: versionable in Git, reviewable in pull requests, deployable via CLI, and readable by AI coding agents. A process definition built in a vendor's proprietary low-code environment is none of those things.

This matters increasingly as AI coding agents become part of how enterprise process automation is built. These are software systems that analyze codebases, write code, deploy and test, observe failures, and iterate toward a working result. When an AI coding agent hits a closed platform with no programmatic foundation below the GUI, it cannot read the process definition, cannot modify it, and cannot verify what it deployed. The conversational layer sits in front of a wall, not on top of a foundation.

Visual-first platforms still serve organizations where business users, not developers, own and maintain processes. That use case remains valid. The shift is in what becomes possible at the engineering layer when a real programmable foundation exists underneath the visual surface.

Mining standards

XES (eXtensible Event Stream) is the IEEE standard for process mining event logs. It defines how event data is structured and shared between systems, making it the portability standard for mining data. An event log exported in XES format from one system can be imported into any compatible mining platform. XES export support is worth checking explicitly in vendor evaluations: it determines whether historical process data is accessible if the mining platform changes.

OCEL 2.0 (Object-Centric Event Log) is the emerging open standard specifically for OCPM event logs, developed by van der Aalst's research group. It defines how multi-object event data is structured and exchanged between platforms. One distinction matters for selection. Object-centric process mining is a capability, and most vendors deliver it on their own internal data model. OCEL 2.0 is the standard that makes that data portable. So OCEL export is the real lock-in check, not the OCPM label. Commercial adoption is still early. The open-source pm4py has full native OCEL 2.0 support, while most commercial platforms run object-centric mining on proprietary models today. The standard still defines the interoperability direction, and vendors committing to OCEL 2.0 are signaling readiness for OCPM at scale.

Agentic AI standards

MCP and A2A are now the open standards for agent tool connectivity and agent-to-agent coordination. Both are now governed by the Linux Foundation. MCP is one of the three founding projects of the Linux Foundation's Agentic AI Foundation, formed in December 2025 alongside goose from Block and AGENTS.md from OpenAI. A2A was contributed to the Linux Foundation by Google in 2025 and is maintained there as its own project. MCP has reached near-universal adoption across AI providers. A2A reached its v1.0 production release in March 2026 with signed agent cards, and is now deployed in production at more than 150 organizations.

Vendors adopting MCP and A2A create integration surfaces that work across the entire AI ecosystem, enabling agents from different vendors to collaborate without custom integration code. Vendors that build proprietary agent protocols risk the same lock-in critique that has historically limited closed BPM platforms.

06

CHAPTER 06

Event-Driven Architecture: Scalability, Data Freshness, and True Decoupling

A process intelligence platform is only as useful as the data it operates on, and only as reliable as the infrastructure that delivers it. Scalability is no longer a future concern. It is a present architectural requirement.

The batch and consistency problem

The core problem is not any single integration pattern. APIs, batch exports, and CDC all have legitimate roles in enterprise architecture. The problem is what happens when process intelligence platforms consume from these sources directly, without a shared event streaming layer in between. A process mining platform ingesting from a nightly SAP batch export reflects yesterday's operations. A direct API pull from Salesforce reflects data from a minute ago. A snapshot from ServiceNow reflects whatever state existed three hours ago when the last extraction ran. The result is inconsistent data across systems, where the same business object (an order, a customer, a claim) appears in different states depending on which system was queried and when. Point-to-point connections compound this by creating direct dependencies between every source and every consumer. Adding a new source system means building new connections to every downstream consumer. Removing or changing one can break processes in unpredictable ways.

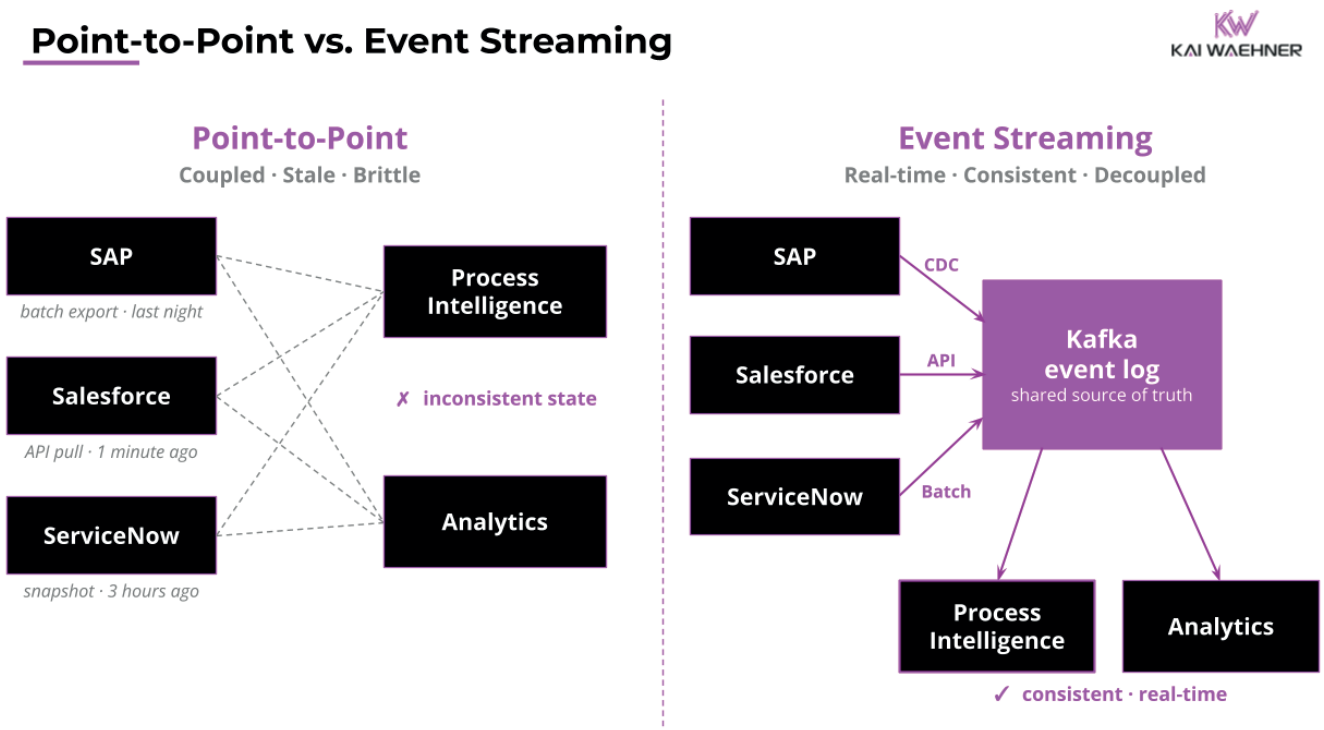

Figure 03. Point-to-Point vs. Event Streaming

How event streaming solves it

Event streaming solves this at the architectural level. A shared, durable event log replaces the web of point-to-point connections. All systems publish to the log as changes happen, through whichever integration pattern fits the source: real-time CDC for transactional systems, API-based ingestion for cloud services, or batch loads where real-time extraction is not yet available. Downstream consumers read from the same log and work on a consistent, current view of operational reality. Stream processing engines handle stateful transformations, aggregations, and enrichment at scale, ensuring data arrives at the process intelligence layer in the right shape and at the right freshness.

True decoupling is the other critical benefit. When operational systems publish to an event log, the process intelligence layer consumes events from that log rather than maintaining direct integrations with each source system. Adding a new source means publishing to the log. Nothing downstream breaks. The process intelligence layer sees one consistent stream of events regardless of how many source systems feed it or which integration pattern each one uses.

Apache Kafka is the de facto standard for event streaming at enterprise scale. Other options exist: Apache Pulsar, AWS Kinesis, and Azure Event Hubs serve similar architectural roles. Kafka has the broadest ecosystem, the deepest enterprise adoption, and the most mature tooling for CDC and connector infrastructure. Apache Flink is the leading stream processing engine for stateful transformations on top of Kafka.

For agentic AI deployments, data freshness is not a performance preference. It is a functional requirement. An agent reasoning from stale data will make confident decisions based on a version of reality that no longer exists.

Why execution architecture matters

Traditional message brokers were designed for reliable message delivery in contained environments with predictable transaction volumes. Apache Kafka rearchitected what a message broker could do: persistent, replayable, horizontally scalable, built for the throughput and distribution patterns of modern distributed systems. The same generational shift that moved enterprise integration from server-bound runtimes to event-sourced streaming infrastructure is now playing out in process orchestration. The operational consequences are identical: better scalability, fewer failure modes, and a fundamentally lower cost of running automation at enterprise scale.

The full argument for event streaming as the integration foundation, including the vendor landscape across all three communication paradigms, is covered in the Data Integration Landscape 2026.

Most legacy BPM platforms persist state to a relational database on every workflow step and handle concurrency through locking mechanisms. That architecture does not scale to the volumes, speeds, and concurrency that modern enterprise automation requires. The vendors that have addressed this built natively event-sourced workflow engines from the ground up: append-only log, event-sourced state, horizontal scalability across partitions. Most platform-native orchestration tools have taken a different path, adding Kafka connectors and eventing interfaces on top of database-backed execution models designed for a different scale.

The same pattern appears on the mining side. A small number of mining platforms have built real-time connectors that consume events continuously rather than ingesting scheduled snapshots. Others address data movement constraints through native integration with existing data platforms, though the freshness of analysis in those cases still depends on how frequently source data is loaded. That gap between natively event-driven architectures and connector-extended legacy models matters increasingly as automation volumes grow.

07

CHAPTER 07

Agentic AI and the Process Layer: What Agents Actually Need

Agentic AI is only as reliable as the process layer governing it. Without defined boundaries, live context, and audit trails, even the most capable agents will fail in production.

Why the process layer is non-optional

Agentic AI systems take actions, execute workflows, interact with operational systems, and make decisions with downstream consequences. That is what makes them powerful. It is also what makes the process intelligence layer non-optional for any serious deployment.
An AI agent without a process layer is ungoverned. It may perform well in development environments. In production, under real load and real compliance requirements, it will fail in ways that are difficult to predict and expensive to fix. The pattern repeats in production deployments: ungoverned agents make confident decisions on stale data, exceed their intended scope, and produce outcomes that are difficult to audit or reverse.

What agents need from the process layer

Operational boundaries define what the agent can decide autonomously, what requires human approval before action, and what must always escalate. These are process-level guardrails enforced before any action executes, not model-level configurations. Live context means agents work on current, consistent operational reality, not on data that was already outdated when they received it. Audit trails mean every agent action, model decision, and human-in-the-loop checkpoint must be recorded in a way that satisfies compliance requirements after the fact.

How agents connect to data and systems

MCP is the de facto standard for connecting AI agents to tools and data sources, including knowledge bases, search indexes, and process intelligence platforms. A2A handles agent-to-agent coordination, enabling specialist agents from different vendors to collaborate on complex multi-step tasks without custom integration code. Kafka handles the operational integration backbone: core business data such as orders, claims, and financial transactions requires the reliability, scalability, and governance that a data streaming platform provides. Kafka, MCP, and A2A each have a distinct role in a well-designed agentic architecture and are not interchangeable. A detailed comparison is covered in a dedicated post: MCP vs REST/HTTP API vs Kafka: The Architect's Guide to Agentic AI Integration.

Mining before automation

Process mining plays a specific and underappreciated role here. Before automating a high-volume process with AI agents, an enterprise should understand how that process actually runs, where it deviates from its intended model, and where exceptions occur that will break automated handling. Mining provides that evidence. The combination of mining and orchestration is what makes agentic automation safe to deploy in regulated environments.

One distinction matters that is easy to overlook. Introducing agentic AI is not the same as embedding AI into an existing process. An existing process optimized for human execution is rarely the right foundation for an agentic workflow. The step before automation is redesign: using process mining evidence to understand what the process actually does, then rethinking it from the ground up for the capabilities and constraints of autonomous agents. Organizations that skip this step automate their inefficiencies at machine speed.

The EU AI Act and agentic AI governance

The EU AI Act brings a firm regulatory timeline to this architecture. The Digital Omnibus agreed in mid 2026 deferred the obligations for high-risk AI systems, including automated credit scoring, candidate screening, critical infrastructure, and healthcare, to December 2027 for stand-alone systems and August 2028 for AI embedded in regulated products. The transparency obligations were not deferred and apply from August 2026. The deferral is a delay, not a dismantling: the requirements for transparency, auditability, and human oversight remain intact, and the conformity, documentation, and oversight work they demand cannot be built in the final months before the deadline.

Process orchestration platforms with full execution histories, human-in-the-loop checkpoints, and conformance monitoring capabilities are compliance infrastructure for any enterprise with EU operations deploying agentic AI in high-risk categories.

This is not three separate technology decisions. It is one converged architecture. Event-driven integration provides live, consistent data. Process intelligence provides governed boundaries and observability. Trusted agentic AI executes within those boundaries. The Trinity is covered in depth in a dedicated post: [The Trinity of Modern Data Architecture: Process Intelligence, Event-Driven Integration, and Trusted Agentic AI.]

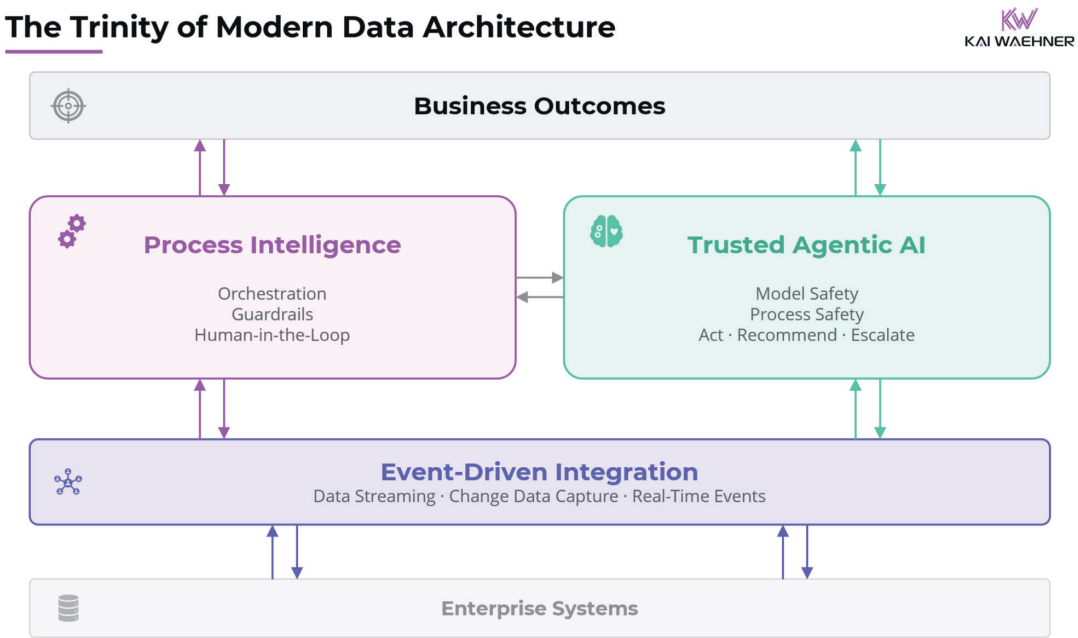

Figure 04. The Trinity of Modern Data Architecture

08

CHAPTER 08

The Orchestration Stack: Where Process Intelligence Fits

Process intelligence platforms do not operate in isolation. They sit within a broader orchestration stack that architects need to understand when designing the full architecture.

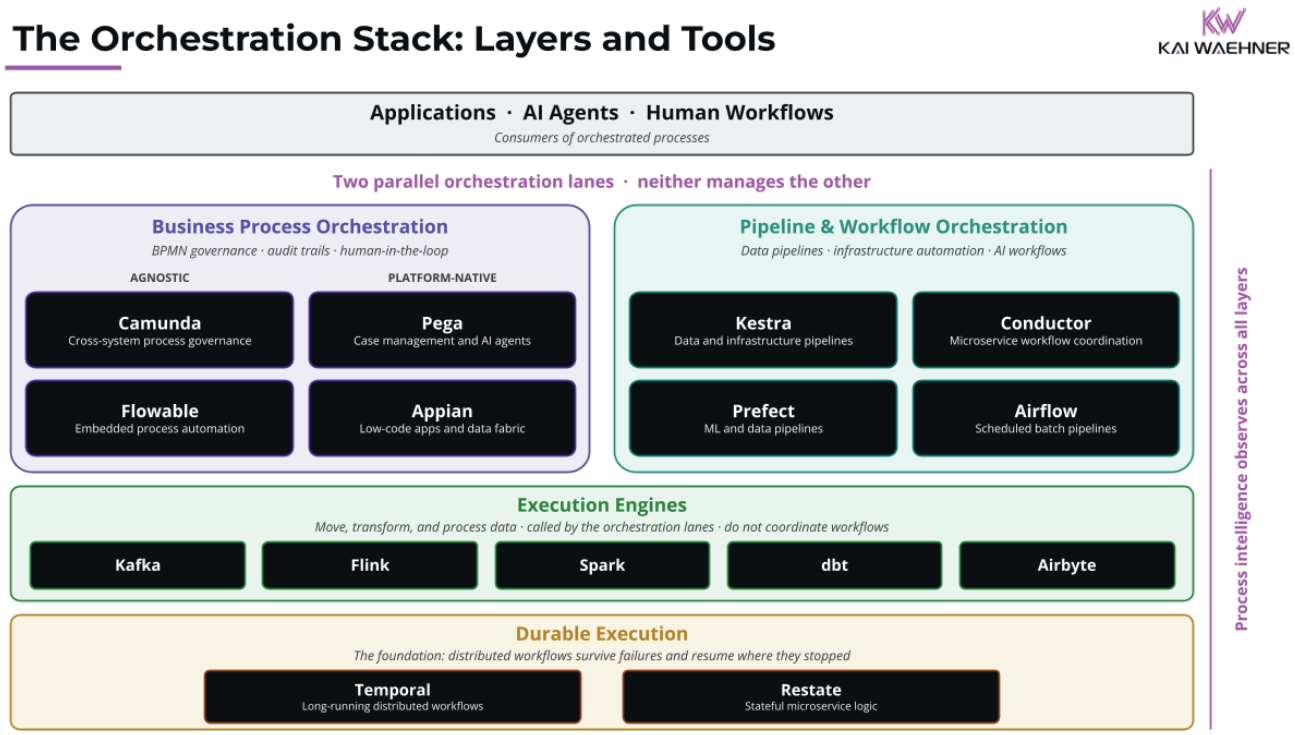

Figure 05. The Orchestration Stack: Layers and Tools

The stack has four layers.

At the top sit applications, AI agents, and human workflows. These are the consumers of orchestrated processes that deliver business value.

The next layer contains two parallel orchestration lanes at the same abstraction level. Neither manages the other. This is a common source of confusion: pipeline orchestration tools like Kestra are sometimes evaluated as alternatives to business process orchestration platforms like Camunda. They are not alternatives. They serve different primary users at the same architectural level. Business process orchestration governs business processes with BPMN standards, compliance audit trails, and human-in-the-loop checkpoints. Within this lane, agnostic orchestrators like Camunda and Flowable serve cross-platform deployments, while platform-native orchestrators like Pega and Appian are strongest inside their own ecosystems. Pipeline and workflow orchestration coordinates data pipelines, infrastructure automation, and AI workflows for engineering teams. Kestra, Conductor, Prefect, and Airflow are the primary examples in this lane.

The execution engines sit one level down: Kafka, Flink, Spark, dbt, and Airbyte. They move, transform, and process data. They do not coordinate workflows. They execute specific technical operations that the orchestration lanes call.

At the foundation sits durable execution infrastructure. Tools like Temporal and Restate ensure that distributed workflows survive failures, resuming exactly where they stopped when a service recovers. Most enterprises run this layer underneath an orchestration platform without evaluating it directly.

Process intelligence platforms feed from this stack and observe what is happening across the execution layer. Platforms like Celonis that have added execution capabilities increasingly act on those observations directly. Others surface improvement opportunities for human decision-makers or connected orchestration platforms.

The tools in the orchestration stack adjacent to process intelligence, how the layers relate, when to use each tool, and where durable execution infrastructure fits will be covered in full in a dedicated companion post on the Enterprise Workflow Orchestration Landscape 2026.

09

CHAPTER 09

Process Intelligence: Key Trends

Several forces are reshaping the process intelligence market simultaneously. Understanding them before starting a vendor evaluation prevents decisions that look right today but create constraints within twelve months.

Market consolidation is accelerating. Mining vendors are acquiring execution capabilities. Platform vendors are pulling independent tools into their ecosystems. Private equity interest in smaller independent mining specialists remains high. Enterprise architects should factor vendor stability and independence into selection decisions. Smaller vendors may be acquired within an existing contract period, changing the support, roadmap, and pricing calculus.

Every major platform vendor is repositioning as the agentic operating system for the enterprise. The organizational implication is the same as it was with BPM platforms a decade ago: the more process logic and agent behavior lives inside a single vendor's proprietary environment, the more expensive the exit becomes.

OCPM is crossing from early adoption to mainstream. Vendors without OCPM are already losing evaluations where supply chain, order-to-cash, or procure-to-pay processes are in scope, because those processes inherently involve multiple interacting objects that case-centric mining cannot adequately analyze. Checking OCPM support at the start of an evaluation keeps vendors that cannot meet it off the shortlist.

Sustainability is an emerging process intelligence driver. Enterprises use process mining to identify process steps that drive unnecessary energy consumption, material waste, and carbon output, turning operational improvement into ESG evidence. This is becoming a board-level topic in Europe and increasingly in North America.

Process digital twins are becoming a standard expectation. The ability to simulate a process change before implementing it is no longer a premium feature. It is a baseline evaluation criterion alongside conformance checking and predictive analytics.

Open standards are moving from technical preference to procurement filter. Vendors that cannot demonstrate MCP, A2A, XES, and OCEL 2.0 support will find themselves excluded from evaluations where interoperability is a requirement.

The boundary between process intelligence and the broader orchestration stack is narrowing. Mining vendors are acquiring orchestration capabilities. Orchestration vendors are adding analytics layers and partnering with specialist miners. Developer-first pipeline tools are gaining business-facing interfaces. Where those trajectories meet is where the next competitive disruption in this market will occur.

10

CHAPTER 10

Decision Framework: 6 Questions to Ask

Six questions guide platform selection, applied separately to mining and orchestration, with the most important being whether your processes live in one platform or span many.

Six evaluation questions determine which zones of the landscape fit the organization's situation. Mining and orchestration are separate decisions with different agnosticism considerations, so apply the questions to each decision independently: the orchestration zones for the orchestration decision, the mining zones for the mining decision. Then combine the results into a platform selection. Many organizations will find that different zones survive for mining and orchestration. A split result is the expected outcome and the correct architectural starting point, not a contradiction.

The question that does the most work is the second one: whether processes live inside one platform or span many is what separates the platform-native zones from the agnostic ones. The remaining questions narrow the shortlist within the zones that survive. Use the landscape diagram above to identify specific vendors.

Q1. Who builds, maintains, and needs to verify the processes?

These are three related but distinct considerations. Who implements the process logic determines whether a pro-code or low-code platform fits. Who maintains it over time determines whether the platform needs to be accessible to non-developers. Who needs to read and verify it determines whether an open standard like BPMN is required. BPMN-based platforms serve all three simultaneously: developers implement workers in any language, process architects model the flow in a visual diagram, and compliance teams read and audit that same diagram without touching code. Camunda and Flowable are the clearest examples of this combination. Visual-first platforms like Appian and Pega prioritize the non-developer builder but store process logic in proprietary formats that cannot be versioned in Git, audited outside the vendor's interface, or migrated to another platform. The selection should match the full set of stakeholders who will interact with the process, not just the team that builds it.

Q2. Are processes inside one platform or do they span multiple systems?

If most processes live inside SAP, Salesforce, or ServiceNow and those are long-term commitments, buying process intelligence from the same vendor is legitimate and often the fastest path. The pre-built extractors, native data models, and ecosystem depth justify the platform dependency. The same logic applies to mining: an organization committed to SAP long-term may find SAP Signavio the most efficient starting point for mining, accepting the ecosystem dependency as a known tradeoff. If processes span multiple platforms, the agnostic zone is the right starting point. Platform-native tools will require additional middleware and integration work for every system outside their own ecosystem, and that cost compounds with every new source added.

Q3. Do you need mining, orchestration, or both?

The right starting point depends on what the organization already knows. Mining comes first when existing processes are opaque: when the organization cannot confidently describe how work actually flows across systems, where bottlenecks occur, or why previous automation attempts underdelivered. Mining provides the evidence before committing to automation. Orchestration comes first when the process is already well understood or does not yet exist in a running system. This covers greenfield automation, known process redesign, and organizations coming from a BPM background with existing process definitions.

Q4. Is agentic AI a near-term priority?

If agents are planned for production within twelve months, the event-driven integration layer and process governance capabilities become immediate selection criteria, not future considerations. Agents need live data, operational boundaries, and audit trails from day one. Stale data produces confident wrong decisions. Ungoverned agents exceed their intended scope. Verify MCP and A2A support explicitly in any vendor evaluated for agentic AI use cases. These protocols are becoming the connectivity standard for agent tool access and agent-to-agent coordination across vendor ecosystems. The differentiator between the zones is not protocol support, which most major vendors now offer, but portability: agents built inside a platform-native environment inherit that platform's lock-in, while agnostic platforms keep the agent layer portable across vendors.

Q5. Do you have data residency or on-premises requirements?

Several vendors offer only SaaS deployment. Organizations with data residency regulations, air-gapped environments, or strict sovereignty requirements need to verify deployment options before shortlisting. For orchestration, Camunda, Flowable, IBM, and Nintex through its K2 line offer on-premises deployment. For mining, SaaS-only platforms require data to leave the organization's environment entirely. mpmX analyzes data where it already lives inside existing BI stacks like Qlik and Databricks without requiring any data movement. QPR delivers process mining as a Snowflake Native App, keeping analysis inside the organization's own Snowflake environment. Identifying this constraint early prevents investing evaluation effort in vendors that cannot meet it.

Q6. Do your processes involve multiple interacting business objects?

Some business processes involve a single object moving through a sequence of steps: one order, one claim, one invoice. Many of the most valuable enterprise processes involve multiple interacting objects simultaneously: an order with multiple line items, each with its own shipment, invoice, and delivery status. Traditional process mining analyzes one object at a time and produces an incomplete picture of these multi-object processes. Object-Centric Process Mining (OCPM) analyzes all interacting objects together and how they affect each other across the process. If the processes in scope include supply chain, order-to-cash, or procure-to-pay scenarios, OCPM support is a hard requirement. iGrafx and ABBYY Timeline do not support OCPM today. Celonis, IBM Process Mining, mpmX, QPR, ServiceNow, Apromore, and the open-source pm4py do support it. SAP Signavio has Process Networks generally available since mid 2026, and Microsoft Power Automate has it in preview.

A NOTE FROM THE AUTHOR

The Architecture is the Strategy. Process intelligence, event-driven integration, and trusted agentic AI are three layers of one converged architecture. Decisions made in each layer constrain what is possible in the others.

A process orchestration platform that cannot consume live events will automate yesterday's decisions with confidence. A process mining platform that ingests data from batch exports will discover problems after they have propagated through the business. An AI agent without a governing process layer will be fast, capable, and ungoverned. A capable but ungoverned agent is a different problem from the ones it was deployed to solve.

The vendor landscape reflects this convergence in progress, with every category moving toward the others simultaneously. No vendor has fully converged all three layers into one coherent architecture. The ones that come closest will define where this market goes next.

For enterprises, the implication is practical. Process intelligence platform selection, data integration architecture, and agentic AI strategy are not sequential decisions. Each constrains the others. Organizations that treat vendor selection as the end of the decision, rather than the beginning of an architecture commitment, will discover the cost of fragmentation in the first production incident.

Processes must be observed, governed, and acted on in real time. Everything else follows from that.

ABOUT THE AUTHOR

Kai Waehner

ADVISORY FIELD CTO

Kai Waehner is an Advisory Field CTO who has spent over 20 years helping enterprises make their most consequential data, process, and AI architecture decisions. He works with Fortune 500 and Global 2000 companies across Europe, North America, the Middle East, Asia, and Australia, and follows a deliberately vendor-neutral approach in his advisory work that prioritizes the right architectural choice over the easiest sell.

His effectiveness as an advisor comes from range. He moves fluidly between a strategic conversation with a CIO and a deep architecture review with an engineering team, and across more than a dozen industries, from financial services and manufacturing to telecom, retail, and the public sector. That range is grounded in 100+ speaking engagements, from technical conferences like AWS re:Invent and QCon to CIO and CTO executive summits.

Kai is known for his independent technology landscapes for data integration, process intelligence, and trusted agentic AI. He also writes the blog at kai-waehner.de, covering industry use cases, technical best practices, and emerging topics for enterprise architects, CTOs, CDOs, and data engineers. Kai is available for advisory engagements, workshops, and keynotes worldwide, as well as media collaborations.

ABOUT THIS LANDSCAPE

This landscape is an independent analyst perspective, not a quantitative ranking. Vendor selection, tier placement, and market footprint reflect the author's assessment based on public information, vendor announcements, and two decades of enterprise architecture work with the platforms covered. No vendor paid for inclusion or placement, and no vendor reviewed or approved its section before publication. Revenue and customer figures are drawn from public vendor statements where available. The author advises technology vendors and enterprises through Kai Waehner GmbH and previously served as Global Field CTO at Confluent; this landscape was produced independently of any vendor engagement.

WEBSITE	**kai-waehner.de**
LINKEDIN	**linkedin.com/in/kaiwaehner**
NEWSLETTER	**Subscribe for real-time data integration insights**

KAI WAEHNER

2026

KEEP READING · STAY INFORMED

Data must be clean, current, and connected. Everything else follows from that.

BLOG & NEWSLETTER

Regular updates from the industry

Subscribe for the latest thinking on enterprise architecture, data integration, process intelligence, and trusted agentic AI.

kai-waehner.de

FREE BOOK

The Ultimate Data Streaming Guide

A practical resource covering streaming use cases, architectures, and real-world industry case studies.

Download Now

CONNECT

LinkedIn & X

Follow along for ongoing analysis of the enterprise integration landscape.

LinkedIn | X (Twitter)

MORE READING

The Trinity of Modern Data Architecture

How data integration, process intelligence, and agentic AI fit together.

Read more