



# ARTA KYB

AS · Industrial Control & Cybernetics

---

## Advanced Process Control Platform

Architecture, Capabilities, and a Cooperation Framework for APC Use Cases

### Technical White Paper

Version 2.1 · June 2026

*Prepared for partnership discussions with Bouvet ASA*

## Contents

---

1. Executive Summary
2. The Advanced Process Control Opportunity
3. Platform Architecture
4. Intelligence — AI & Autonomous Agents
5. Technology Stack
6. Reference Use Case — Pump Cavitation Digital Twin
7. Deployment and Integration
8. A Cooperation Framework with Bouvet
9. Roadmap
10. About Arta Kyb

## 1. Executive Summary

---

Process industries — water treatment, oil and gas, chemicals, energy, and pulp and paper — operate plants that are increasingly instrumented yet still controlled by tuning loops in isolation. The gap between what the data could deliver and what classical regulatory control actually delivers is where margin, energy, and equipment life are lost.

Arta Kyb's Advanced Process Control (APC) Platform closes that gap. It is a modular, service-oriented platform that connects to plant instrumentation over industrial protocols, builds and maintains process models, runs model-predictive and adaptive control on top of existing regulatory layers, and continuously optimises operation against economic and safety objectives.

This white paper describes the platform's architecture and components, the technology stack behind it, a representative deployed use case — a pump-cavitation digital twin for an ozonation plant — and a concrete framework for cooperating with Bouvet to deliver APC use cases to industrial clients.

### In one sentence

*Arta Kyb turns historian data and live OPC tags into closed-loop predictive control and real-time optimisation, deployed as a containerised platform that sits beside — not inside — the existing control system.*

## 2. The Advanced Process Control Opportunity

---

Most plants run distributed control systems (DCS) or PLC/SCADA stacks with hundreds of PID loops. These loops are excellent at holding a single variable to a setpoint, but they are blind to interactions between variables, to constraints, and to the economics of operation. Three structural limitations recur:

- **Loop interaction.** Adjusting one loop disturbs others; operators detune to stay safe, leaving the plant running conservatively far from its optimum.
- **Constraint handling.** Classical control cannot push a variable toward a limit and hold it there safely. The most profitable operating point usually sits against a constraint.
- **No economic layer.** Setpoints are chosen by habit, not by continuously solving for the lowest-cost or highest-yield operating point given current prices, feed, and equipment health.

Advanced process control addresses all three by replacing or supervising groups of loops with a multivariable controller that respects constraints and is driven by an optimisation objective. The barrier has historically been cost and integration risk: traditional APC packages are heavyweight, vendor-locked, and expensive to commission. Arta Kyb's platform is built to remove that barrier.

## 3. Platform Architecture

---

The platform follows a layered, service-oriented design. A React client talks to a FastAPI gateway, which routes requests to six independent services. Each service owns a clear responsibility and can scale independently. Process data flows up from the plant through connectors; control and optimisation decisions flow back down to the plant in a closed loop.

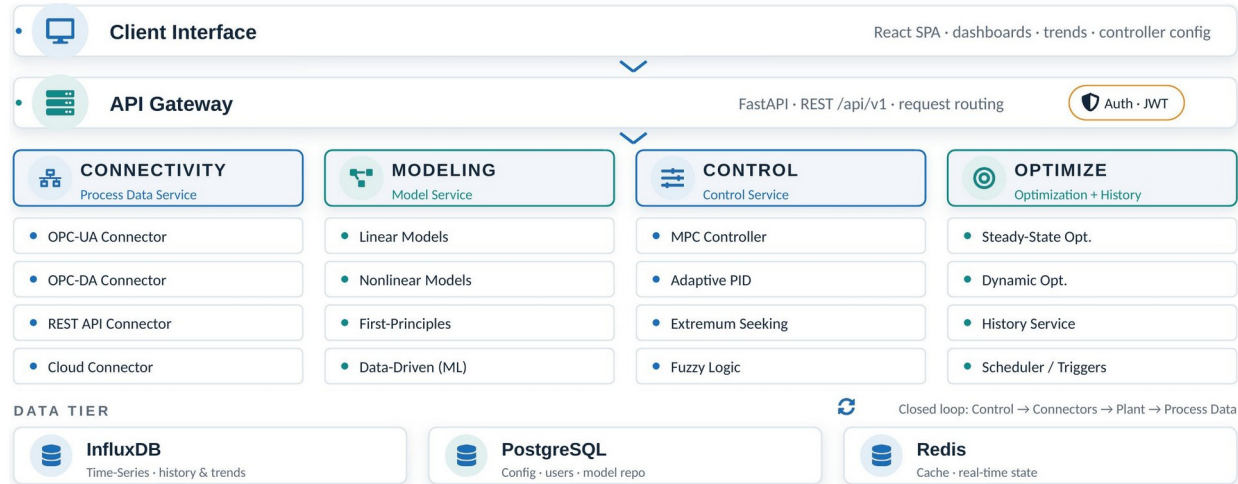


Figure 1 — Layered architecture of the Arta Kyb APC Platform.

### 3.1 Presentation and API layers

**Client Interface (React).** A single-page application providing operator dashboards, live trends, model inspection, and controller configuration.

**API Gateway (FastAPI).** A versioned REST surface at /api/v1 with CORS handling and request routing. Every protected route is guarded by JWT authentication, so identity and access are enforced once at the boundary.

### 3.2 Core services

Six services make up the application core. The first, Authentication, issues and validates JSON Web Tokens. The remaining five carry the process-control workload:

Service	Responsibility	Key sub-components
Process Data	Acquire and normalise live and historical signals	OPC-UA, OPC-DA, REST, Cloud connectors
Model	Build, store, and serve process models	Linear, nonlinear, first-principles, data-driven
Control	Compute manipulated-variable moves	MPC, adaptive PID, extremum seeking, fuzzy logic
Optimization	Find the best operating point	Steady-state and dynamic optimisation
History	Persist and query time-series	Trend retrieval, backfill, audit

### 3.3 Connectivity layer

The Process Data Service abstracts the plant behind a uniform interface. OPC-UA and OPC-DA connectors integrate with modern and legacy control systems respectively; a REST connector ingests data from line-of-business and IIoT systems; and a cloud connector bridges to historians and analytics in Azure, AWS, or private clouds. Reconnect behaviour is configurable, so a dropped link self-heals rather than stalling the controller.

### 3.4 Modeling layer

Control quality is bounded by model quality. The Model Service supports a spectrum so the right tool fits each problem:

- **Linear models.** Step-response and state-space models for well-behaved, near-linear units — fast to identify and robust.
- **Nonlinear models.** For units with strong gain or dynamic nonlinearity across the operating range.
- **First-principles models.** Mass, energy, and momentum balances where physics is known — they extrapolate where data is sparse.
- **Data-driven models.** Machine-learning models identified from historian data, including hybrid grey-box forms that combine physics with learned residuals.

### 3.5 Control layer

The Control Service supervises the regulatory layer rather than replacing it, which keeps commissioning safe and reversible:

- **Model Predictive Control (MPC).** The multivariable workhorse — predicts future behaviour over a horizon and computes constraint-respecting moves that drive the plant to setpoints and limits.
- **Adaptive PID.** Self-tuning single-loop control for loops not yet under MPC, or as a robust fallback.
- **Extremum seeking.** Model-free optimisation that hunts for the best setpoint on slow processes without an explicit objective model.
- **Fuzzy logic.** Rule-based control encoding operator expertise where a crisp model is impractical.

### 3.6 Optimization and history

**Optimization Service.** Steady-state optimisation chooses the economically best targets for the controllers to hold; dynamic optimisation plans the best trajectory through transitions such as grade changes or start-ups.

**History Service.** Backed by a time-series database, it serves trends, supports model re-identification, and provides the audit trail needed in regulated environments.

### 3.7 Closed-loop operation

These layers form a continuous loop. The Process Data Service reads the plant; the Model Service predicts; the Optimization Service sets economic targets; the Control Service computes moves; and those moves are written back through the connectors to the plant — whose response is read on the next cycle. Default cycle times are configurable (control on the order of a minute, model and optimisation refresh hourly), and every step is logged.

## 4. Intelligence — AI & Autonomous Agents

---

Classical control relies on fixed laws and hand-tuned models. The Arta Kyb platform adds a layer of intelligence on top: models that learn from data, and software agents that perceive, reason, and act within bounds set by the operator. This is what lets the platform improve over time and shoulder routine supervisory work — without ceding authority.

### 4.1 Hybrid, learning models

Pure first-principles models extrapolate well but are costly to build and never quite match the plant; pure data-driven models fit the data but fail outside it. The platform combines them. A first-principles core captures the known physics — mass and energy balances — while a machine-learning component absorbs the unmodelled residual from historian data. The result is a hybrid “grey-box” model that is both accurate

and trustworthy, and that is re-identified as new data arrives so it tracks the plant as equipment ages and operating regimes shift.

Every prediction carries an uncertainty estimate. Controllers and optimisers use that uncertainty to stay cautious where the model is unsure and to push harder where it is confident — which is exactly the behaviour an operator expects from a competent engineer.

### 4.2 Autonomous agents, supervised

On top of the models sits an agentic layer. An APC agent continuously perceives the plant through live data and the digital twin; reasons about what to do by calling the platform’s own tools — models, the optimiser, and history; and proposes concrete actions such as setpoint moves or alarms. Crucially, the agent does not act unilaterally: proposed moves pass through an operator-approval gate, and the agent then observes the result and adapts. Autonomy is bounded and reversible by construction.

The same pattern scales from a single loop to a plant: specialist agents can monitor, diagnose, and optimise individual units while a supervisory agent coordinates them. Because every agent action is logged and gated, the approach fits the audit and safety expectations of regulated industrial environments — intelligence is added without giving up accountability.

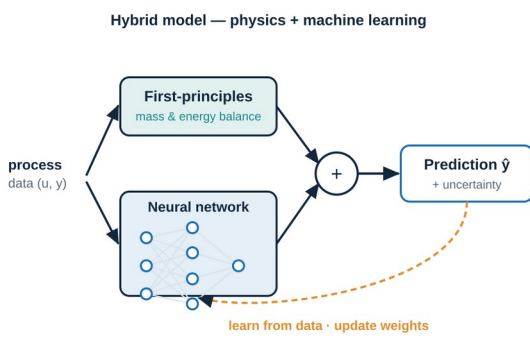


Figure 2 — Hybrid model: physics fused with a learned residual.

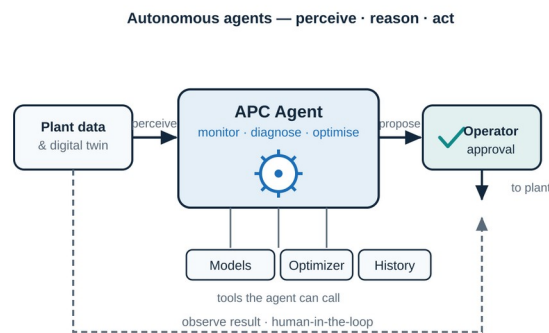


Figure 3 — Agentic control: perceive, reason, act — operator in the loop.

## 5. Technology Stack

The platform is containerised end to end, so a full stack runs identically on a laptop, an on-premise server, or a cloud host. This is central to low-risk deployment: the same image that is validated in test runs in production.

Layer	Technology	Role
Frontend	React	Operator dashboards, trends, configuration UI
Backend	FastAPI (Python)	REST API, service orchestration, control logic
Relational store	PostgreSQL	Users, configuration, model repository
Time-series store	InfluxDB	Process history and trends
Cache / state	Redis	Real-time controller state, fast lookups
Admin	PgAdmin	Database administration

Layer	Technology	Role
Orchestration	Docker Compose	Reproducible multi-service deployment
Migrations	Alembic	Versioned database schema evolution
Security	JWT + hashed credentials	Authentication and access control

The codebase is organised into database models, validation schemas, and CRUD APIs for processes, models, and controllers, with a development toolchain (build scripts, test suite, and notebook-based component testing) that keeps iteration fast and changes verifiable.

## 6. Reference Use Case — Pump Cavitation Digital Twin

To make the platform concrete, Arta Kyb built a digital twin of pump cavitation for an ozonation stage in a water-treatment plant, driven by the plant’s own P&ID. Cavitation silently erodes impellers, raises energy use, and ends in unplanned downtime; it is rarely instrumented directly.

The twin frames the problem in standard APC terms — controlled, manipulated, and disturbance variables — and continuously estimates each pump’s net positive suction head (NPSH) margin from available measurements:

- **Controlled variables (CV).** NPSH margin and the cavitation state of each pump (healthy, incipient, developed, cavitating).
- **Manipulated variables (MV).** Pump speed via the variable-frequency drive, and valve positions that set suction conditions.
- **Disturbance variables (DV).** Feed flow, suction pressure, temperature, and ozone and reagent dosing across treatment trains.

Operators see each pump on the live P&ID with a status ring and NPSH-margin badge, and can inspect suction-head curves and trends per pump. The same modelling and control machinery that powers this twin — connectors, models, and the closed loop — generalises directly to compressor surge, heat-exchanger fouling, distillation control, and energy optimisation.

### Why this matters for a partnership

The digital twin demonstrates the full chain end to end on a real flow diagram: live data in, physics-plus-data model, operator-facing visualisation, and a control path out. It is a ready-made anchor for a first joint proof of value.

## 7. Deployment and Integration

Deployment is designed to be incremental and reversible, which is the only way industrial clients adopt new control technology.

1. **Connect (read-only).** Stand up the stack alongside the existing control system and stream live and historical data through the connectors. No write access, no risk.
2. **Model and shadow.** Identify models and run controllers in advisory mode — the platform shows the moves it would make while operators stay in control.

3. **Close the loop.** Once trust is established, enable write-back on a bounded set of variables, supervising the existing PID layer.
4. **Optimise.** Turn on the economic optimisation layer to continuously retarget the controllers against cost, yield, energy, or emissions objectives.

Because the whole stack ships as containers, each stage is validated in a test environment that is byte-for-byte identical to production, and any stage can be rolled back instantly.

## 8. A Cooperation Framework with Bouvet

Arta Kyb brings deep control-engineering and platform technology; Bouvet brings client relationships, delivery scale, and complementary software and data engineering. Together they can package APC as a repeatable offering rather than a bespoke project each time.

### 8.1 Complementary roles

Arta Kyb	Bouvet	Joint
APC platform & IP	Client access & domain context	Use-case selection & scoping
Control & modelling expertise	Project delivery & scale	Proof-of-value execution
Platform operation & support	Data engineering & integration	Productised APC offering

### 8.2 Proposed first steps

- **Select a lighthouse use case.** Choose one well-instrumented unit — the pump-cavitation case is a natural candidate — for a time-boxed proof of value.
- **Run a read-only pilot.** Connect to historian and live data, build the model, and demonstrate advisory control with a quantified benefit estimate.
- **Define the commercial model.** Agree how platform, delivery, and support are packaged and priced for clients.
- **Build a reusable playbook.** Capture the pilot as a repeatable pattern — connectors, model templates, and dashboards — to shorten the next engagement.

## 9. Roadmap

- **Near term.** Complete the remaining service APIs, finish full-stack integration testing of the containerised deployment, and harden the connector suite.
- **Mid term.** Expand the model library (hybrid grey-box identification), add controller templates for common unit operations, and extend the digital-twin gallery beyond pump cavitation.
- **Longer term.** Plant-wide optimisation across coupled units, and a marketplace of pre-built APC use cases that partners can deploy with minimal configuration.

## 10. About Arta Kyb

Arta Kyb AS is an industrial control and cybernetics company focused on bringing advanced process control within reach of plants that have historically been priced out of it. The company combines research-level

control engineering with modern, containerised software so that APC can be deployed quickly, safely, and reversibly.

**Contact.** For partnership discussions or a proof-of-value pilot, reach out to the Arta Kyb team. This document is intended as a basis for cooperation conversations with Bouvet and may be shared under the terms agreed between the parties.