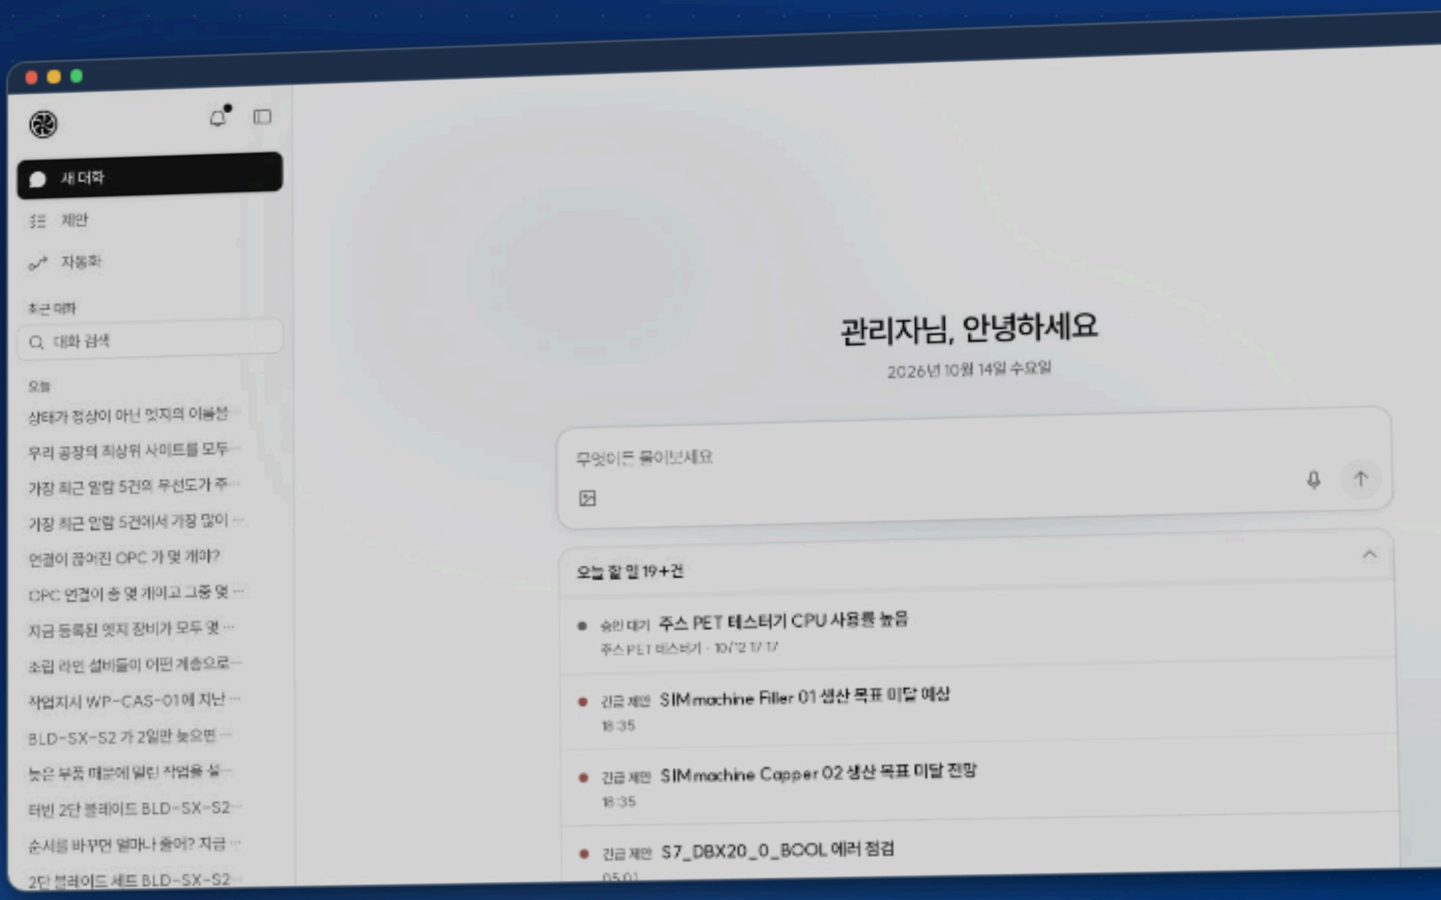
PlantPulse® **AI**

PLANT OPERATIONS COPILOT

The AI copilot for the plant floor · Product overview and proposal

An AI that knows your plant, right beside your crew.

It links scattered ERP · MES · PLM, reads documents, drawings and maintenance videos, and judges equipment condition to tell supervisors, plant managers and operators **«what needs doing today»** — with evidence. It runs on our own industrial model, **inside your plant**.



11 families
of customer systems linked

Drawings · video
evidence beyond documents

1 GPU
runs on-premises

October 2026

Plant operations AI copilot · On-premises / air-gapped

kopens®

CONTENTS

Contents

01	Why now — scattered data, buried know-how	3
02	What PlantPulse AI is — connect · read · judge	4
03	AURA — the industrial LLM	5
04	Connect — customer PO · thread ①–③	6–9
05	Read — docs · drawings · video ④–⑥	10–13
06	Judge — series · photos · actions ⑦–⑩	14–17
07	At work · accuracy ⑪–⑬	18–20
08	Architecture	21
09	Use cases — a day on the floor · roadmap	22–24
10	Global comparison · five things	25–26
11	Screen tour	27–28
12	Data sovereignty · security	29
13	Installation · hardware requirements	30
14	References · deployment · joint decisions	31–33

One-page summary

> The floor's questions do not stop at "what's wrong right now?" — **"Can I start this job today?"** · **"Is this the latest drawing revision?"** · **"How did we fix this last time?"** The answers are scattered across ERP · MES · PLM, drawing cabinets, training videos and sensors. PlantPulse AI brings them into one place.

Connect
11 families incl. ERP · MES · PLM · QM — read-only

Read
documents · drawings · maintenance videos as evidence

Judge
time-series engine · photo reading · action tracking

WHAT IT IS

A field decision app — it finds the plant's problems and **proposes actions with evidence, lets people approve and close them, answers when you ask, and hands repetitive checks to automations.** Three things to do — **Ask · Proposals · Automations**.

Not an MES or ERP — your systems stay the source of record; the only writes are approved drafts. **Not a BI dashboard** — the point is the cause and next action behind the number. **Not a chatbot** — chat is one way in; the digital thread, drawing reader and time-series engine decide. **Integrations are optional** — without them every screen is complete; with them only capabilities grow.

What this document covers

What it takes off your hands

Pages 3–5. The floor's problem, the three-pillar product definition, our own model.

How it manages that

Pages 6–21. A digital thread centred on the customer PO, document · drawing · video RAG, time-series · photo judgement, Ask · Proposals · Automations, structure.

How it is actually used

Pages 22–28. A day on the floor as use cases, the roadmap, five differences, and screens from a running system.

What deployment looks like

Pages 29–33. Data sovereignty, install requirements, references, schedule.

01 · Background

Why now — scattered data, buried know-how

Collecting the data has largely succeeded. Yet the floor still **shuttles between systems and filing cabinets looking for answers.**

PROBLEM 1

Systems that don't talk

Work orders live in MES, materials in ERP, design changes in PLM, nonconformances in QM. To know “can this operation start today?”, a supervisor has to cross-check four or five screens.

PROBLEM 2

Know-how buried in drawings and videos

Tolerances sit on drawings; repair tips sit in maintenance footage and training videos. None of it is searchable, so “asking someone who knows” is the only way.

PROBLEM 3

Working from an old drawing revision

Even after a design change is approved, the drawing posted at the workstation may be an old revision. The difference shows up as a defect after machining is done.

PROBLEM 4

Many alarms, and data that cannot leave

A threshold alarm only says “something was exceeded”. Yet sending process data to a commercial AI gets stopped at security review.

So what is needed

> An AI that **connects and reads** scattered systems **without touching them**, answers with documents, drawings and videos **as evidence**, and **never leaves the plant**.

PlantPulse AI was designed on exactly those terms. Our own industrial language model runs on-premises and answers from the platform's real data and ontology, the records in your systems, and documents, drawings and videos. Neither prompts nor process data ever leave the site.

02 · Introduction

What PlantPulse AI is — connect · read · judge

An AI copilot for plant operations. It does three things, and delivers the results through **one field decision app — Ask · Proposals · Automations**.

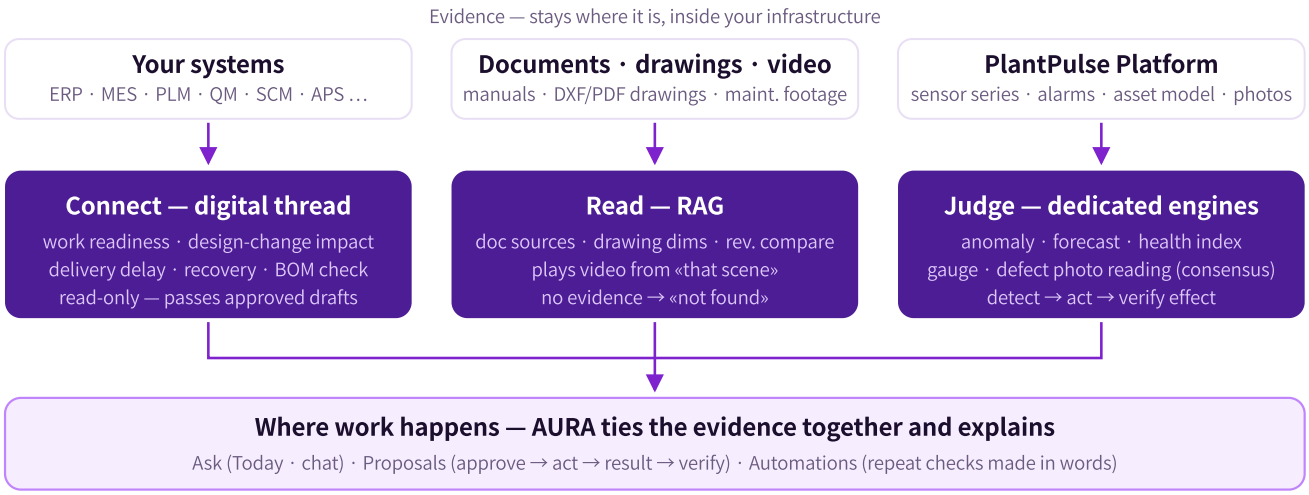


Figure 1. Sources stay where they are and are only read. Numbers and verdicts come from dedicated engines; the AI explains them.

CONNECT

Can it start today?

For each operation it gives Go / No-Go and the blocking reason across eight items: predecessor · material · tooling · staff · permit · quality · equipment · drawing.

READ

Where is it written?

“Where is Ø62 H7?” is answered with drawing · sheet · zone; “What do I check before press-fitting?” with that moment in the video.

JUDGE

What is wrong right now?

Sensors are checked continuously every 10 minutes; anomalies are grouped into incidents and carried through to an action plan and a recurrence check after the action.



Figure 2. Home (Ask) — below the greeting and the input box, **«Today»** expands to list up to five items, most urgent first (it is folded to one line by default). The screen is monochrome; colour is used only for status dots.

> **It runs on the platform.** The AI does not collect its own data. It works directly from the asset model, time series, alarms and ontology that **PlantPulse Platform** has standardised, and from the records in your systems — so “adopting AI” never turns into a separate data-building project. That said, the better organised the data, the more accurate the answers.

03 · Model

AURA — an industrial LLM that runs inside the plant

The brain of PlantPulse AI is **AURA, our own industry-specific language model**. Rather than renting a general-purpose model, we run a model trained on industrial documents and floor vocabulary **inside the customer's infrastructure**.

What is different

Item	Content
Industrial training data	Fine-tuned on occupational-safety SOPs, quality standards and domestic and international safety procedures
Floor vocabulary	A Korean-native model that knows terms such as degradation, interlock and trip
Long context	Ultra-long context that takes a whole equipment manual in one go and answers from it
Multimodal	Reads charts and equipment photos — meter reading, defect verdicts, classification, text recognition. It reads several times and answers **only when the readings agree**
On-premises	The language model, embedding and reranker share a single GPU in-house — NVFP4 quantised
Own model names	Served as **aura-35b** (language · vision), **aura-embedding-4b** (search embedding) and **aura-reranker-0.6b** (reranking)
Cost structure	No token billing — a **fixed-cost structure**; cost does not rise with use

WHY OUR OWN MODEL ①

Data never leaves

Questions contain equipment names, tags and process conditions — trade secrets in themselves. With the model in-house, not even the prompt goes outside.

WHY OUR OWN MODEL ②

Cost is predictable

Usage-billed models cost more the more they are used after roll-out. Letting the floor use it freely calls for a fixed-cost structure.

External models can be used too. Where security policy allows, a commercial API can be connected. The choice rests with your policy; the product supports both paths.

04 · Connect

All work traces back to the customer PO

In a make-to-order plant, the floor's questions come down to one — **«if this job is late, which customer order is at risk?»** The digital thread takes the **customer PO (sales order)** as the root and links project · work order · operation · equipment in one line, binding supplier **material POs** to those work orders. Your systems are **only read**.

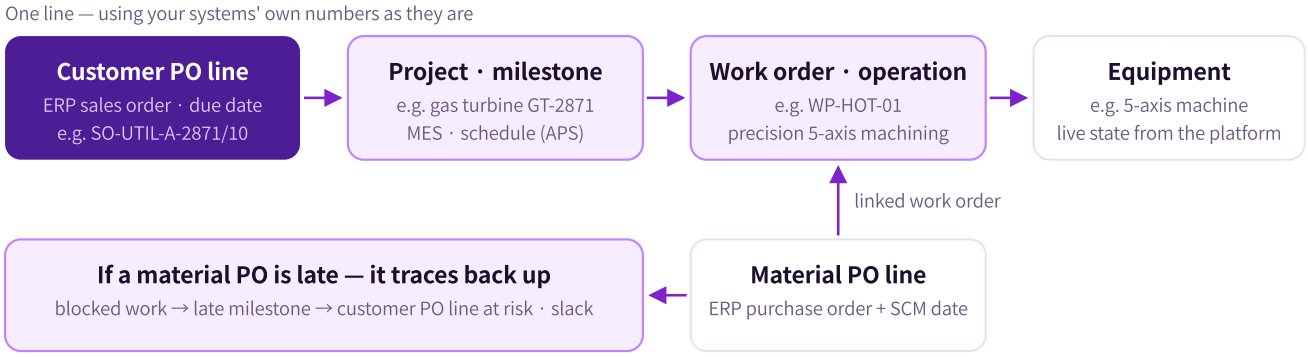


Figure 3. Links rooted in the customer PO. The master records stay in your ERP · MES · SCM and the platform; the digital thread links only numbers · status · timestamps · sources.

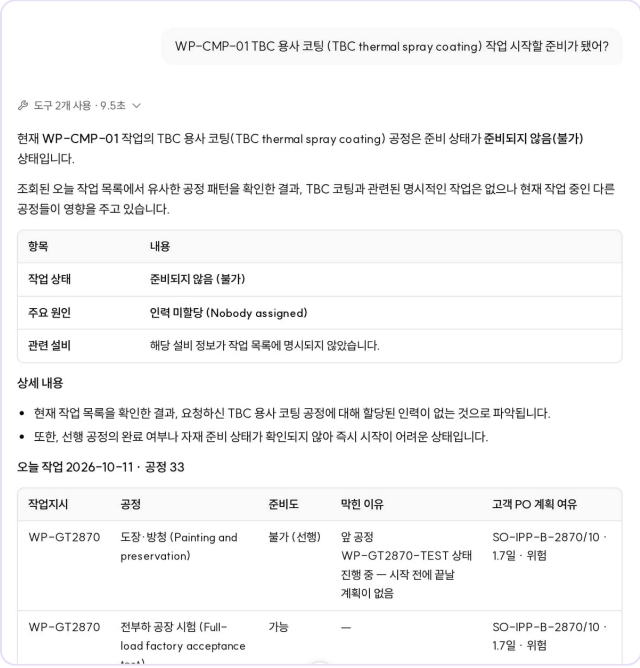


Figure 4. In Ask, “Is this job ready to start?” — the verdict and blocking reason, followed by the «Today's work» table: per operation, readiness · blocking reason · **customer PO days of slack** (at risk · slack) · drawings to check.

ERP is the master of POs

Purchase order lines (SAP MM · Oracle Purchasing etc.) and sales order lines are read through the ERP connector; with SCM, the supplier-confirmed delivery dates are merged in. Where the two disagree, it is flagged as a conflict.

Days of slack are calculated

It calculates, in days, how a delay passes through jobs · milestones to **leave how many days before which customer PO line is due**. Demo example: material PO-G4001 late → job WP-HOT-01 → GT-2871 shipment milestone → customer PO SO-UTIL-A-2871/10 — **0.7 days of slack after the delay, at risk**. The «planned slack» in the answer is the upper bound if everything finishes as planned.

You ask it in Ask

“What's on today's work? Start with what's blocked and customer PO slack” — the answer carries the calculated table and cards for the drawings to check. Numbers in the table come only from the calculation, not the model. There is no separate dedicated screen.

04 · Connect

① Can this job start today?

Before a shift starts, a supervisor has one question — **"Can the planned operations start, and if not, why?"** The digital thread **only reads** your MES · ERP · PLM · QM · SCM · maintenance (CMMS) · safety (EHS) · tooling systems and answers with **Go / No-Go and the blocking reason** for each operation.

Eight items decide it

Item	To pass	Read from
Predecessor	The preceding operation is done, or will be done before the planned start	MES
Material	The component part numbers are in stock, or the pending order is due before the start	SCM (or MES)
Tooling	The required tools · gauges are available and their calibration is valid at the start time	Tool management
Staff	Someone with the required qualification is assigned, and the qualification is valid	ERP (or MES)
Permit	The required work permits · isolations are valid within the planned window	EHS
Quality	No open nonconformance blocks this job · operation	QM (or MES)
Equipment	No maintenance window or breakdown overlaps the planned window	CMMS
Drawing	No pending design change blocks the job, and **the drawing at the workstation is the revision approved in PLM**	PLM · drawing index

The verdict follows the worst item

Blocked · at risk · unknown · ready (+ not applicable). **Blocked items come first** in the answer; each row gives the blocking item and a one-line reason, along with **which record in which system** the verdict came from.

«Unknown» is not «ready»

If a source is missing or unreadable, the item stays «unknown» and says which source could not be read. If connector data is older than its freshness limit, **no verdict is given at all** — so stale data never produces a «ready».

The drawing item — so no one machines from an old revision

It checks whether a design change (ECO) is pending on the part numbers the operation uses, what its stock disposition and effectivity are, and whether the new revision is usable before the operation starts. If none of the drawings for that part number posted at the workstation is **the approved revision, the verdict is «blocked — an old drawing revision is at the workstation»**. If drawings are not yet indexed, the revision check is skipped and the verdict uses the design change alone — a missing drawing is never called «wrong».

Ask it just like that in Ask. "Tell me which operations can't start this shift, why, and what to clear first." The same check can run as an **automation** every day before the shift and reach the supervisor as a proposal card (p. 18). The digital thread is an **optional integration** — without it Ask · Proposals · Automations run as they are; with it, this answer is added. Nothing is **written** to your systems.

04 · Connect

② Days lost to design changes · quality · delays

The floor's question soon moves from “what's wrong?” to **“so what happens to the delivery date?”** The digital thread links equipment, operations, parts, design changes and quality records into **one time-aware plant graph** and calculates, **in days**, which operations and milestones a delay reaches.

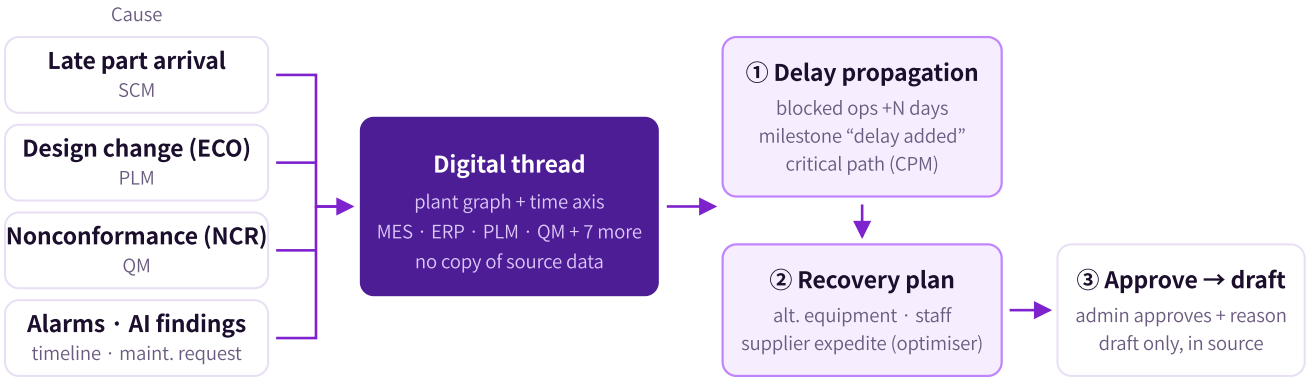


Figure 5. Wherever the cause comes from, its impact is traced on the same graph. The numbers come from the calculation and the optimiser; the AI **explains** them.

What it shows

- **Equipment timeline** — platform alarms and events plus the AI's findings, actions and outcomes on one line per asset. “What happened when, and what did we do?” — ask it in Ask.
- **Impact of design changes and quality issues** — for each ECO or NCR: which part numbers it blocks, which operations it pushes back and by how much, and **how many days this one item added** to each milestone.
- **Supply recovery** — judges whether expediting a late part from the supplier **can or cannot recover** the schedule.
- **No made-up numbers** — when inputs are missing it states the scope of the calculation and what was missing, and with no recovery levers it reports “impact only” rather than an empty “0-day recovery”.

Connects your existing systems

MES, SCM, scheduling (APS), ERP, PLM, QM, maintenance (CMMS), safety (EHS), tooling, warehouse (WMS) and lab (LIMS) — **11 families** — connect through **open MCP specifications**. A conforming connector plugs in without product changes and is enabled **only after passing a conformance test**. Several systems of the same kind per site are fine. Connections are managed on the admin console's **integration map**.

The AI never writes to your systems directly

No equipment control, no changes to committed MES records; PLM, QM and ERP are **read-only**. A material expedite (plant manager or admin) or maintenance request (supervisor, maintenance, plant manager or admin) reaches the source system as **a single draft** only after it is **approved with a reason**; every attempt is logged and one approval never creates two drafts.

The timeline works out of the box. Operation and milestone delays and recovery plans need MES and SCM (ERP for staff); design-change and quality impact need PLM and QM.

04 · Connect

③ See what is bound to what, on one screen

Your systems stay as they are. Which system connects through which tools, and what it is bound to in the plant graph, is checked and managed on one screen — the admin console's **integration map**.

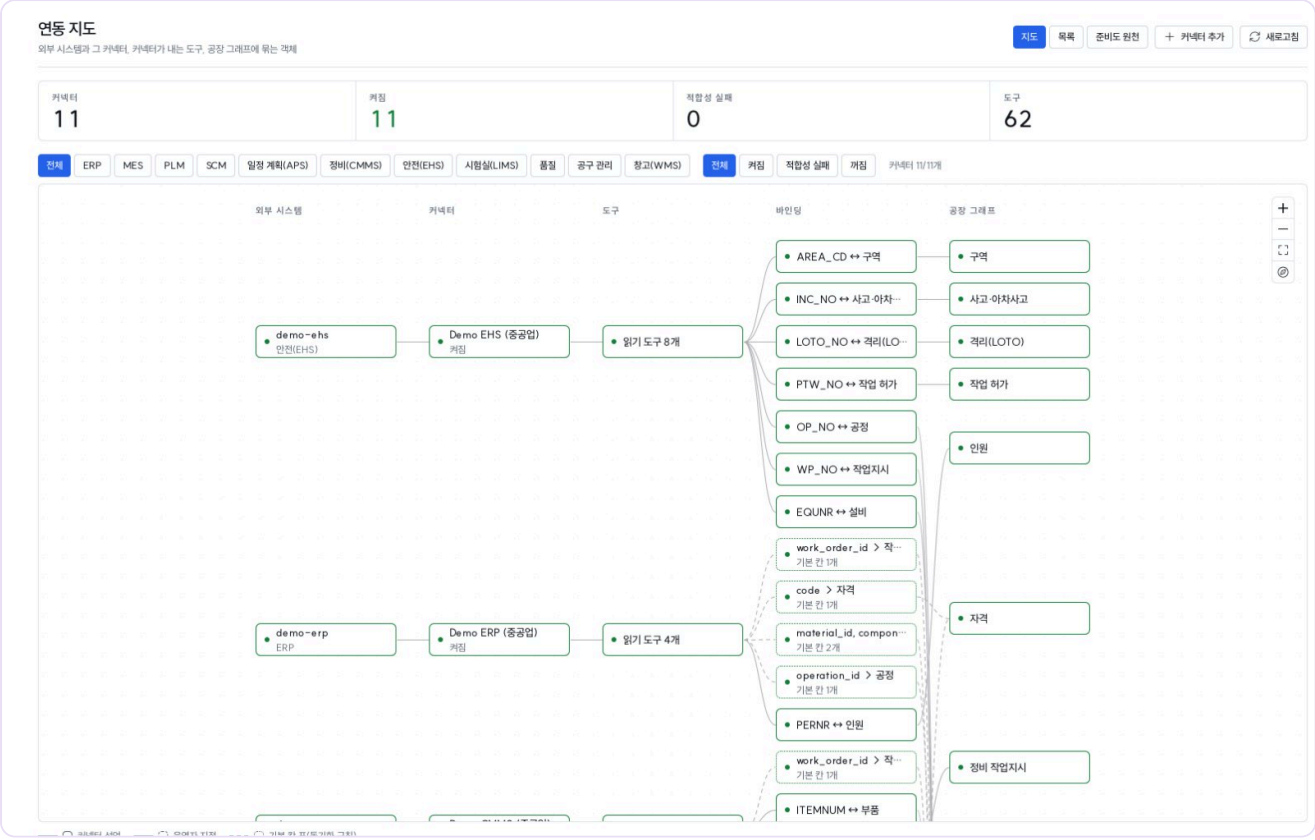


Figure 6. Integration map (admin console › Digital thread) — left to right: **external system → connector → tools → binding → plant graph**. Here 11 demo families (ERP, MES, SCM, PLM, QM, CMMS, EHS, APS, tooling, warehouse, lab) connect through 62 tools and bind to people, equipment, work orders, parts, isolations, work permits and more; the buttons on top filter by family. Each row shows which code in your system (e.g. `EQ_CD`) is which object in the graph (equipment).

How a system is connected

Step	What happens
Register	Your IT or SI registers a connector built to the open MCP specification. Credentials are stored encrypted
Conformance test	Checks that it answers to spec — **it can be enabled only after passing**
Binding	Links each system code to an object in the graph. Operators can override the spec defaults on screen
Graph projection	Loads the platform's equipment and tag hierarchy and the relations the AI has learned into the graph **every 10 minutes** — references only, no copy of source data

> **Connects without product changes.** Any vendor's MES or ERP connects the same way as long as its connector follows the spec. With several systems of the same kind on one site, each is connected and one is made primary — recovery queries all of them, and an approved draft goes to the primary.

05 · Read

④ Answers from documents, with sources

On the floor, «the AI said so» is not evidence. PlantPulse AI searches equipment manuals and work standards and shows **which document** the answer came from. The same search engine also takes drawings (⑤) and videos (⑥) — there is one window to ask in.

Document search (RAG)

- **Vector search combined with the knowledge graph** — it finds documents **related to this machine**, not merely documents with similar words.
- **Source shown** — «which document of which asset» is displayed with the answer.
- **Kept current automatically** — new, changed and deleted documents on the platform are picked up every 10 minutes. Nobody has to maintain it.
- **Many formats** — PDF, Word, PowerPoint, Excel, Markdown and text are indexed automatically. One upload window takes documents, drawings, video, audio and subtitles together; disabled formats are blocked with a reason before sending.
- **Equipment attachments too** — videos, audio, subtitles and drawings attached to equipment documents on the platform are also picked up by the 10-minute sync and indexed (when video and drawing ingestion is on).
- **No evidence, no answer** — when nothing matching is found in the indexed documents, it returns «not found» instead of an answer.
- **Not exposed** — the search engine and its stores open no ports to the outside and are reached only through a keyed search API.

Preview how a file was read

Upload a file to the admin console's **document parse check** screen and it immediately shows the text and confidence read with **the same parsing options** as the search index — the text shown is exactly what the AI would have indexed. When an answer disagrees with a document, this tells you whether the cause is in the «reading»; if the text comes out empty, as with scans, compare it against an OCR profile. Drawings are checked in the same console's **drawing parse test** (⑤).

THE DIFFERENCE IT MAKES

An answer that ends at «bearing temperature is high» and one that says **«bearing temperature rising — check the lubrication procedure in this machine's maintenance manual first (source: machine ○○ · maintenance manual)»** are worth entirely different things on the floor. The second can be **acted on immediately**. Open “How this answer was made” on any answer to see which tools the AI called, with what conditions, and what came back.

05 · Read

⑤ Reads drawings, down to sheet and zone (optional)

The drawing is the master of design information. PlantPulse AI reads the **title block, dimensions, tolerances, geometric tolerances, surface roughness, parts list and notes** from DXF and vector-PDF drawings **together with their location**, and answers down to **which sheet and zone** a value sits in.

① Reads as structure

Title block (drawing no., part no., material, revision, scale), dimensions with tolerances · fits · threads, geometric tolerances, Ra, parts list, notes and revision history come out item by item, with a **zone** from the border grid.

② Finds dimensions exactly

Values match to ±0.01 mm, and fits **tell H7 (hole) from h7 (shaft)**. Drawing and part numbers must match in full — it finds the same dimension, not similar text.

③ A drawing card on answers

Under the chat answer: thumbnail · "drawing no. rev X" · "sheet 1 B3"; click to open it zoomed, the evidence **outlined in red**.

How accurately it reads — complex machined-part test drawings

Test drawing	Title block	Dimensions	GD&T	Ra	Parts list	Notes	Rev. history
Hydraulic manifold HMB-2040 rev C · D, PDF and DXF each	12/12	27/27	3/3	3/3	6/6	5/5	—
Gearbox housing GBH-7300 rev B PDF	10/10	58/58	8/8	6/6	16/16	6/6	2/2
GBH-7300 rev B DXF	10/10	39/39	6/6	4/4	16/16	5/5	2/2
GBH-7300 rev C PDF / DXF	10/10	58/58 · 38/38	8/8 · 6/6	6/6 · 4/4	16/16	7/7 · 6/6	3/3

For dimensions, both the number found (recall) and the number read correctly (precision) equal the table values. Zero parse warnings; 0.04–0.54 s per file (GB10). **Measured on synthetic test drawings and answer keys made in-house** — not results on real customer drawings. We measure again on your drawings during deployment validation.

How it is used

Who	They ask	What the AI returns
Machining · QC	"Bore tolerance of housing GT-2871-104?"	"Ø62 H7 — DWG-GT-2871-104 rev C, zone B3", that spot highlighted
Finding it	"Where is the 120±0.1 dimension?"	Drawing · sheet · zone with exactly that value
Assembly	"What is balloon 7 and how many go in?"	Parts-list part no., name, quantity, material
Purchasing · materials	"Material and latest revision of this part?"	Title-block material and revision, plus a PLM-based revision badge

05 · Read

⑤ Compares revisions, filters out old drawings (optional)

Most drawing mistakes come from the «revision». PlantPulse AI points out **what changed between old and new revisions**, and checks PLM through the digital thread to tell you **whether the drawing in front of you is the latest**.

Revision compare — changes in colour

Two revisions of one drawing number sit side by side, zooming and scrolling together, with **changed areas** painted green (added) · red (removed) · amber (modified). Below is a change list for title block (rev · ECO · date · approval), dimensions (before/after and zone), parts list, notes and GD&T; **what could not be compared is stated separately**. If the revisions cannot be overlaid, it invents no areas and shows only the list.

Test pair	Dimensions	Parts list	Notes	GD&T	Dim. changes inside changed areas	False pos.
HMB-2040 C→D (PDF · DXF)	4/4	1/1	1/1	—	4/4	0
GBH-7300 B→C PDF	7/7	6/6	2/2	1/1	8/8	0
GBH-7300 B→C DXF	5/5	6/6	2/2	1/1	6/6	0

Title-block field changes and added revision-history entries are also detected; 0.17–0.36 s per pair. Same synthetic test set as the previous page.

Revision badge — chat · drawing register

✓ latest · ⚠ old revision (latest X) · ⏳ change in progress · ? revision unconfirmed. From the revision approved in PLM and changes in progress; unconfirmed is never called «latest». The admin console's **drawing register** groups drawings by number, one row per revision, filterable to «old revisions only».

Drawing · E-BOM · M-BOM three-way check

The digital thread compares the drawing's parts list, the design BOM in PLM and the manufacturing BOM to find **missing items · quantity differences**, and flags a **revision mismatch** if the job follows an old revision. The «drawing» item of work readiness (p. 7) uses this.



Figure 7. Admin console › Drawing register — hydraulic manifold HMB-2040, revision C → D. Title block, four dimensions, parts list and note changes are listed with their zones, and the changed spots are marked on the new revision (staging, synthetic test drawing).

Switching on · scope. Enabled at installation (off by default). No DWG — **export to DXF or PDF** to upload. Text from scanned drawings (image PDFs without a text layer) and P&ID (piping and instrumentation) come next. Originals and index stay on the customer's servers; the digital thread holds only numbers and references — no outside calls.

05 · Read

⑥ It plays maintenance · training videos from «that scene» (optional)

Much of the floor's know-how lives in **video** rather than documents — equipment training, repair recordings, operating tips, vendor demos. Most of it piles up where no one can find it again. PlantPulse AI indexes the **speech, subtitles and on-screen text** of videos and recordings with timestamps, answers procedure questions with the matching scene as evidence and **plays it from that moment**.

① It listens and transcribes

Speech becomes timestamped sentences (Korean and English detected automatically). Embedded subtitles and separately uploaded subtitle files (srt · vtt) are read too.

② It reads on-screen text

Burned-in captions and values shown only on screen are read. Numbers are used **only when two frames agree**. (separate switch)

③ Playback at that moment

A video source card sits under the answer; click it to play **that segment** in the right-hand panel while the conversation stays in place. Only segments the answer cites get a card.

How it is used

Who	They ask	What the AI returns
Maintenance	"What do I check before press-fitting a new bearing?"	The scene in the maintenance video where the check is explained — played from that moment
Safety training	"What is the LOTO procedure, step by step?"	The matching segment for each step of the training video
Tool handling	"How do I fix the reaction arm of the hydraulic torque wrench?"	The segment of the demo video showing that action
On-screen values	"What bolt torque was shown in that video?"	A value **only written on screen**, never spoken, with its scene
Asset context	"Press #3 anomaly — how was it handled before?"	Videos bound to that asset come in as evidence alongside documents and history
Onboarding	"First mould change — show me the whole flow"	A step-by-step list of scenes from the training video
Vendor material	"How do I reset the alarm in the vendor demo?"	English vendor videos are searchable through subtitles and speech

Ask by voice

Hold the microphone in the input box and speak (up to 60 s); the transcript lands in the input field — check it, then send. It uses the same transcription engine as video, and recordings never leave the plant.

How well it finds · switching on

In-house eval set v2 (83 questions × 3 runs): right video 1st 90%, top-5 97%; median scene-time error 7.6 s; «I don't know» on 95% of not-in-material questions. Enabled at installation (off by default). Transcription starts 8 pm–6 am plant time by default, keeping the GPU free by day (window adjustable). Picking the exact sentence of a **Korean video from an English question** is still weak.

06 · Judge

⑦ It knows what a tag *means*

A plant has tens of thousands of tags, and for most of them the name alone says nothing. PlantPulse AI **profiles each tag on real data**, infers its meaning and stores it.

Two stages of understanding

1. **Statistical profiling (no AI)** — tags are classified from real data: continuous measurement, state flag, counter, setpoint or dead tag, with units and the measured distribution attached. **Almost free**, and refreshed periodically.
2. **Semantic inference (once per machine)** — on top of that, the AI is called once per machine to infer each tag's role and "what it means when it is abnormal", which is stored and refreshed periodically.

It does not make things up

Every inference carries a **confidence score**. Low-confidence interpretations are not used in diagnosis and are shown as low-confidence on screen — a safeguard against **plausible wrong answers**.

It sees relationships too

Correlation and lag analysis between tags captures relationships such as "when A rises, B rises five minutes later". So a diagnosis does not stop at **"temperature anomaly"** but goes on to **"temperature rise followed five minutes later by rising vibration — a classic wear pattern"**.

WHY THIS IS THE CORE

The main reason AI fails on the shop floor is **missing context**. Answering from tag names alone produces plausible but wrong statements. This layer closes that gap, and the confidence display **does not hide whatever gap remains**.

06 · Judge

⑧ Time series belong to a dedicated engine

Language models handle sentences well but time series poorly. So anomaly detection and forecasting are done by a **dedicated insight server**, and the AI interprets its results.

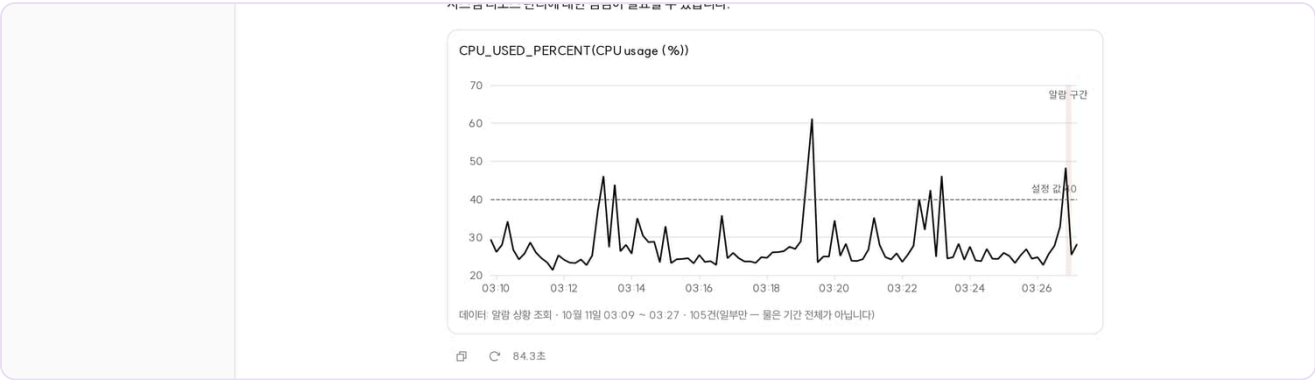


Figure 8. In Ask, “This asset raised an alarm — why?” — the time series the engine queried is drawn inside the answer. The dashed line is the set value, the shaded band the alarm window, and the line below says which query the data came from.

What it does

Function	Content
Anomaly detection	Sensors, OEE, reliability and energy checked every 10 minutes; alarms and work orders analysed on request
Forecasting	Forecasts for the main domains; **confidence intervals** when there is enough calibration data, otherwise the reason they were not produced
Equipment health index	Penalties for anomalies, alarms and idling give a 0–100 score and A–D grade — “insufficient data” instead of a grade when coverage is low
Attention ranking	Ranks “assets to watch first” per site from open findings, health score and trend
Production analytics	OEE loss breakdown (six big losses · stop causes) · SPC charts · work-order completion forecast (P10–P90)
Feature store	Per-tag and per-machine features generated automatically — used for similar-equipment search and as AI evidence
Alarm diagnosis	Narrows the root cause from five angles: frequency, duration, sensor correlation, trend forecast, statistics

Why they are separate

- **Accuracy** — a time-series model reads time series better.
- **Cost** — the language model is not called for every judgement, so continuous monitoring is affordable.
- **Explanation** — the AI puts words on numeric evidence, so it can say **why**.
- **Honesty** — when data is short it does not invent intervals or grades; it says what it could not produce and why.

06 · Judge

⑨ Floor photos are read «only when the answers agree»

Analogue gauges, equipment nameplates and defects such as leaks and corrosion are still checked by eye. Attach a photo in chat and the vision engine reads it — but because in a plant **a wrong number is worse than no answer**, it reads the same photo several times and answers **only when the readings agree**.

What it does

Reading	What it covers	Consensus rule
Gauge reading	Reads gauges such as pressure and temperature gauges, and digital displays, with units	Read five times, **four or more** agree
Defect verdict	Judges whether there is a leak · corrosion · crack · damage · foreign matter · overheating	Two of three agree
Text reading	Organises the text on nameplates · labels, Korean included — for checks before ordering parts	Two of three agree

Disagreement means «don't know»

A sharp photo gives the same value all five times. A blurry or noisy one gives five different values. PlantPulse AI **returns that disagreement as «unreadable» and asks for a sharper photo** — blocking by structure the fabrication that simply telling a language model "say you don't know when you don't know" cannot stop.

Every verdict leaves a record

Whether consensus was reached, and how many readings agreed, is recorded for each verdict, so admins can filter to «rejected verdicts only». The agreement rule can be made stricter in settings. Photos are discarded after processing and not stored in the vision engine. Free-form descriptions ("describe this photo") cannot be checked for consensus, so they are kept apart as **answers that give no numbers or verdicts**.

How it is used

When	You ask	You get
Logging analogue gauges	Attach a gauge photo: "What does this pressure gauge read?"	The reading and unit — if blurry, it **rejects** and asks for a retake
Before ordering parts	Attach a nameplate photo: "Summarise this nameplate"	Maker · model · rating text organised — rejected if readings disagree
Inspection rounds	Attach a pipe photo: "Any leak?"	Defect type and verdict — «cannot judge» if readings disagree

06 · Judge

⑩ Unbroken from detection to action

If it ends at finding the anomaly, a person still has to tidy up afterwards. PlantPulse AI ties detection, alarm, root-cause diagnosis, remedy and post-repair verification into **one flow**.

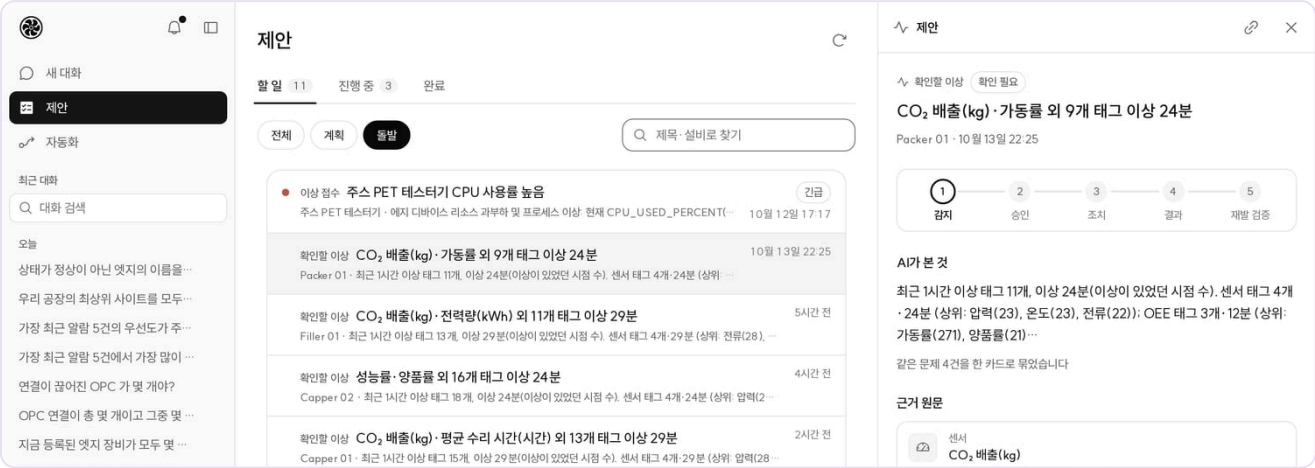


Figure 9. Proposals › To do — planned (proposals) and unplanned (anomalies to check · incidents) items in one list, most urgent first. Click a row and the right side opens what the AI saw, the evidence and **one next button**.

Grouped into incidents, closed on one page

- Continuous automatic detection raises an **anomaly to check**; once a person confirms it, it becomes an **incident** — high severity is reported automatically. False alarms are **dismissed** with a reason, which feeds the accuracy calculation.
- Duplicates with the same cause are **grouped into one row** — «◯◯ and N more assets». You handle one cause, not each asset separately.
- The AI diagnoses the cause and writes an **action plan** with a draft maintenance request. Evidence comes attached: sensor trend · drawing · video segment · past cases.
- Five steps — **Detected · Reported → Approve → Act → Result → Verify** — show on one detail page. After the result is recorded it watches whether the same problem returns, and if it does, sets it back to «Recurred».

It actually reaches the owner

Delivered immediately in-app, by web push and by email — and to Slack once a webhook is registered. The point is that nothing is missed even when nobody is looking at the screen.

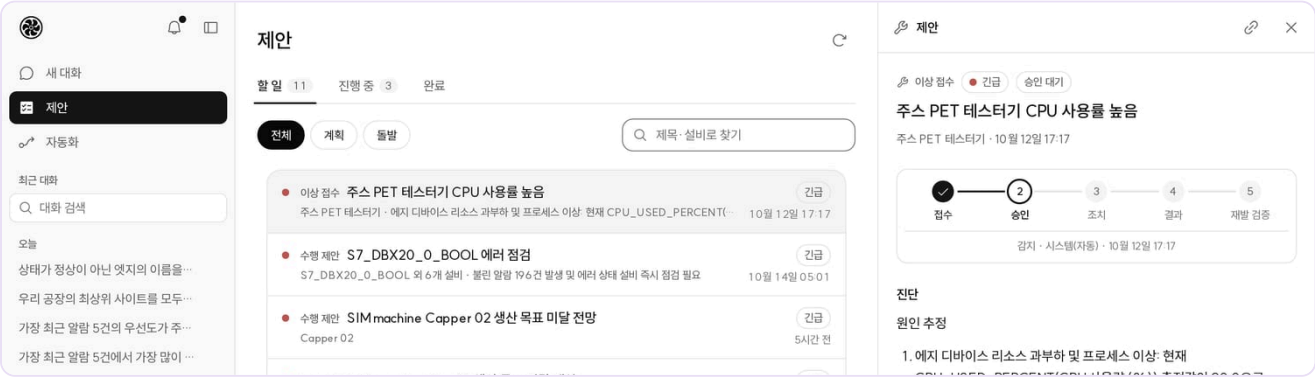


Figure 10. Incident detail — the five-step indicator, the result record and the AI diagnosis on one page. The former «Resolve» menu no longer exists separately; it is the later stage of Proposals.

07 · Where work happens

⑪ Three things to do — Ask · Proposals · Automations

Floor people open only three screens. **Ask when you don't know, review proposals and act, and hand repetitive work to automations.** Built for one-handed use on a phone; work roles decide who may approve what.

Screen	What it does
Ask	The home screen. Ask by text, voice or photo and get an answer with evidence. Below the input box, «**Today**» lists up to five items, most urgent first
Proposals	**To do · In progress · Done**. Planned (proposals) and unplanned (anomalies to check · incidents) items share one list, and one detail page closes them from approve → act → result → verify (p. 17)
Automations	One sentence — "Every morning at 7, check unit 3 vibration and tell maintenance if it differs" — becomes a draft, switched on after a **Preview run**. Sharing · runs · reports · versions and restore

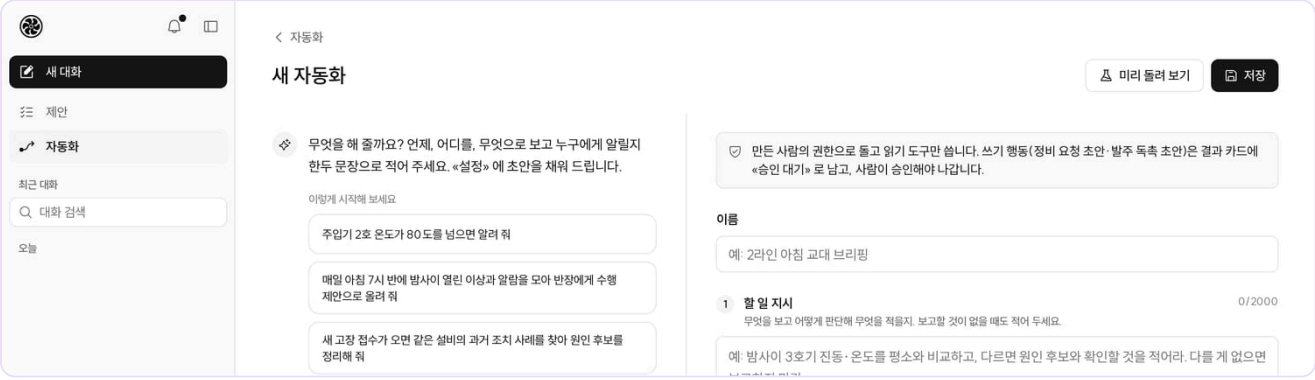


Figure 10-1. Create automation — describe it in words on the left and the fields on the right (Instructions · Tools · When to run · Where results go) fill in.

Work roles and approvals

Work role	Sees in Proposals	Approves material-expedite drafts	Approves maintenance-request drafts
Operator	Planned · incidents	—	—
Maintenance	+ anomalies to check	—	✓
Supervisor	+ anomalies to check	—	✓
Plant manager · admin	Everything	✓	✓

Five default roles — operator · supervisor · maintenance · plant manager · admin; quality · design · purchasing and others add bundles. Everyone starts in **Ask**; drafts are approved, with a reason, in the proposal detail. A material expedite bound for a supplier goes out only after **two different people** approve — the first approver cannot give the second approval. Only **drafts** enter your systems, and the server re-checks approval rights on every request.

Automations, not features

A new request is solved with an automation before a screen is added. Morning brief · shift handover · weekly report are templates.

Read freely, write on approval

An automation only reads, within its creator's rights and scope. Writes are attached as drafts and go out only when a person approves.

The same screens on a phone

No separate mobile edition — the same screens fit a phone's width. Add it to the home screen as an app to receive push.

07 · Where work happens

⑫ Ask about every function in one window

Chat is not a separate feature but **the entrance to every function above**. With one sentence — “Show me unit 3's current state, whether this job's drawing is the latest, and the last repair video if there is one” — the AI **decides for itself what to look up**, calls the tools it needs in sequence and merges them into one answer.

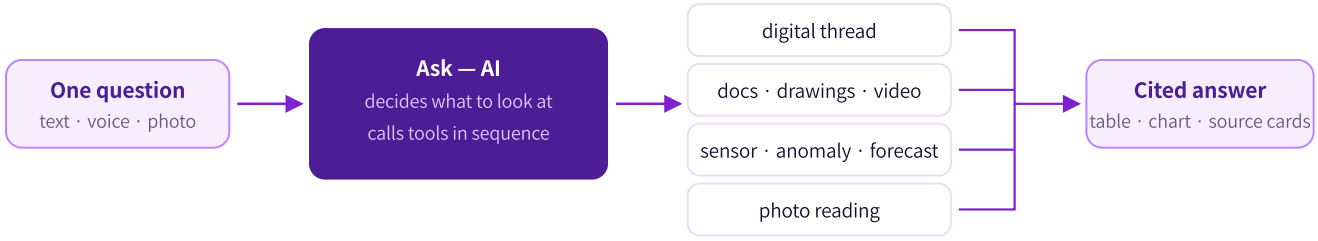




Figure 11. The answer to “Rank the assets in poor health by priority” — the top shows **how many tools were used**, followed by the ranking table and what to inspect first.

What comes with the answer

Evidence	How it appears in the answer
Documents	Source card — click a number such as [1] in the text to jump to that card
Drawings	Drawing card — “drawing no. · Rev · sheet”; click to open that sheet in the right-hand panel, with the conversation kept in place
Video · audio	Video card — click to play **that segment** in the right-hand panel. It also drafts work procedures from video
Digital thread	On connected installations — which record in which system the answer came from
Numbers · verdicts	“Used N tools” — expand to see which tools were called for the answer

WHY IT IS FAST AND CONSISTENT

It does not scrape everything; it calls only what the question needs. Anomaly scores, health indices and readiness are already computed, and the model is in-house, so there is no external round trip. And because **the same question follows the same procedure**, the judgement does not change with who is on shift.

07 · Where work happens

⑬ How often it is right — measured the same way every release

All values below are based on an **in-house evaluation set · heavy-industry demo data**, not measured on customer site data (public pipeline of the 2026.10 re-release 6). Scoring is done by machine, not by a judge model, and a release that falls short of the bar does not ship. During deployment validation we run the same method again on your data.

Chat AI

Item	Result	What was measured
Tool choice	**22 / 25**	On 25 floor questions, did it call the expected lookup tool (machine-scored)
Faithful to evidence	**6 / 7**	Tool numbers carried into the answer unchanged — key made by calling the same tool directly
Ontology distinctions	**14 / 16** median of 5 runs (13 – 15)	Keeps distinct axes such as equipment · operation · order apart
Daily briefing	**8 / 8**	Morning briefing invents nothing not in the data
Delivery calc. citing	**6 / 6** · invented numbers 0 2 / 2 when it must not call	For «if a part is N days late», only the engine's numbers are cited; no call when none is needed

System

Item	Result
Service response	All 6 MCP servers · 19 components respond
REST queries	174 GETs — 0 server errors (5xx)
Platform tool contracts	63 kinds — 0 with inputs · outputs changed from the previous release

Video · audio search (eval set v2 — 83 questions × 3 runs)

Item	Result	What was measured
Right video ranked 1st	**90%**	The right video is first in the results
Right video in top 5	**97%**	The right video is among the top 5
Scene-time error	**median 7.6 s**	Playback start vs. correct scene
«I don't know»	**95 – 100%**	Says it doesn't know for content not in the material — 0 video cards then

Transcription · drawings

Item	Result — what was measured
Long training video	**12** of 12 utterances in a 22-min video transcribed — no spoken segment missed
Short training video	Text · timestamps of 5 segments in a 2-min video match **5 / 5**
Fit formulas	IT5 – IT8 · g · f deviations, across 3 – 500 mm, **0 mismatches** with published conventional values — 6 cells that deviate from the standard are «not calculated»
Drawing reading · revision compare	2 synthetic test drawings · 3 revision pairs — **every** answer-key item matched, 0 false positives (pp. 10 · 11)

We hand over the yardstick, not just the numbers. Answers come from formulas · schemas · operating records, so anyone running it at any time gets the same values. During deployment we rerun this evaluation with your equipment names · tags · documents and check the values on your floor together.

08 · Architecture

What runs where

The language model, digital thread, search, document understanding, time-series and vision engines each play their part, and all of them run **inside the customer's infrastructure**.

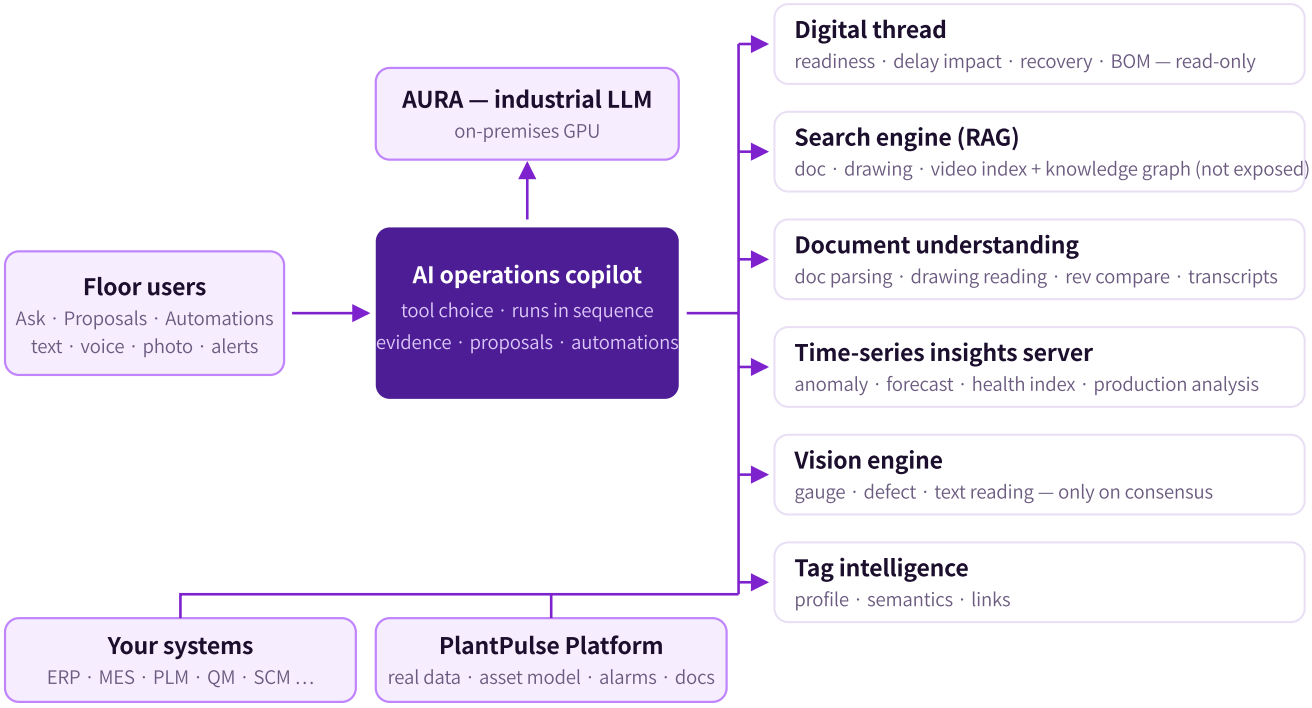


Figure 12. The AI keeps no data of its own. It reads what the platform has standardised and the records in your systems over **standard open interfaces**, so it sits on top of the existing data rather than rebuilding it.

Connecting other systems

Public API · MCP

Other systems can ask the AI questions (**/api/v1/chat**) and fetch today's task proposals. Over MCP it exposes **read-only** tools only (ask_plant · daily_tasks · list_plant_tools). Keys are issued in the admin console.

Digital thread

Connects ERP, MES, PLM, QM, SCM, APS, CMMS, EHS and the rest of the 11 families read-only through open MCP specifications (pp. 6–9). Recovery plans and maintenance requests pass only **as drafts** after approval, and it never controls equipment directly.

09 · Use cases

A day on the floor ① — supervisor's morning, maintenance's afternoon

Look at it through **people and moments**, not features. Drawn with demo data from a heavy-industry plant building a gas turbine (GT-2871). The AI answers with evidence, and only **drafts** approved by a person go into your systems.

U1 · SUPERVISOR — JUST BEFORE THE MORNING SHIFT

«What is blocked today, and which customer order is urgent?»

Today's work list → Blocked operations · reasons → Its drawing → Customer PO days of slack → Material expedite draft

Step	What the AI does — what the demo actually shows
Today's work	In Ask, “What's on today's work? Blocked first” — per operation: readiness verdict · blocking reason · customer PO days of slack · drawings to check, as a table
Why it is blocked	WP-HOT-01 precision 5-axis machining «blocked (drawing)» — design change ECO-GT-0142 moved blade BLD-SX-S1 from C → D, and the new revision is usable from 10-12
Open the drawing	Drawing card under the answer (DWG-BLD-SX-S1 rev D) → opens in the right-hand panel, with the conversation kept in place
Delivery slack	The customer PO lines the delay reaches and the days left until they are due, «at risk / slack»
Expedite draft	If material is the cause, the recovery plan's purchase-expedite draft — once **two** different people approve it in Proposals, one draft goes into SCM

U3 · MAINTENANCE — WHEN AN EQUIPMENT ANOMALY APPEARS

«What is wrong, and how did we fix it before?»

Anomaly · alarm → Diagnosis · evidence → Past action records · repair video → Maintenance request draft

Step	What the AI does
Anomaly · alarm	Continuous detection raises an «Anomaly to check»; once confirmed it becomes an incident — the same cause is grouped into one row
Diagnosis · evidence	The proposal detail carries the AI diagnosis with the sensor trend (anomaly span marked) · drawings · past cases. To dig further, «Ask about this»
Past cases	“How was it handled last time?” — past action records and manual procedures, and that segment of the repair video bound to the equipment (video enabled at installation)
Maintenance request	«Draft created on approval» in the proposal detail — a supervisor · maintenance · plant manager approves with a reason and one maintenance-request draft goes into MES; after the result is recorded, a recurrence check

09 · Use cases

A day on the floor ② — when a design change arrives, when a defect occurs

U5 · PRODUCTION ENGINEERING — A DESIGN CHANGE (ECO) ARRIVES

«Which jobs does this change block, and which revision is the floor drawing?»

Design change → Affected parts · jobs · milestones → Floor drawing revision → Old vs new revision compare

Step	What the AI does — demo: hydraulic manifold HMB-2040 C → D
Impact path	Ask about that ECO in Ask — affected parts → blocked jobs +N days → «delay this item adds» per milestone
Floor drawing revision	Work readiness «drawing» item — no approved revision at the workstation means «blocked — old drawing revision at the workstation»
Revision badge	✓ latest · ⚠ old revision · ⏳ change in progress on the drawing register — from the PLM-approved revision and open ECOs
Revision compare	C → D changes: title block (rev · ECO · date · approval), 4 dimensions (e.g. Ø32 H7 → Ø35 H7, zone D2), parts list O-ring P32 → P35, notes, changed spots marked

U7 · QUALITY — WHEN A DEFECT · NONCONFORMANCE OCCURS

«What did the drawing require, and where else was the same lot used?»

Nonconformance (NCR) → Part · lot · job · supplier → Required dim. · tolerance → Calibration · qualification

Step	What the AI does
Nonconformance	Reads QM nonconformance · inspection records — e.g. NCR-HV-0311 casting porosity, lot HT-VNE-S1-0921, job WP-HOT-01, supplier SUP-CASTING-A
Links	Ask by that number and the answer gives part · lot · job · operation · supplier together with the jobs it pushes back
Drawing requirement	“Bore requirement on this housing?” — dimension · tolerance and its place on the drawing; fit symbols expanded to upper · lower limits, marked «formula estimate»
Tools · people	Work readiness «tooling» · «staff» items — calibration · qualification valid or not

Auto-matching drawing inspection characteristics (balloons) against measurements is on the roadmap (p. 24).

Other moments

When	You ask or click	You get
Logging a gauge	Gauge photo: “What does this pressure gauge read?”	The reading — only if 4 of 5 reads agree; rejected if blurry
Safety before work	“Walk me through the LOTO steps”	That segment for each step of the training video (video enabled at installation)
Before the plant manager's meeting	“Today's top 5 orders at delivery risk and why”	Late milestones · causes · recovery plans compared

09 · Use cases

What comes next — roadmap

The items below are **plans not yet in the product**. They build on the features in the preceding pages; timing and scope will be agreed together during deployment discussions.

Flow	What it would do	Foundation already in place
Operator job card	One card per job with the current-revision drawing · procedure video scene · tools · qualifications	Each element already works separately in chat
Delivery meeting	For the plant manager: risk per customer PO → cause chain (material PO · equipment · design change) → recovery plan → one report set	Recovery calculation · milestone delay · days of slack
Purchasing · supplier delays	Material PO delay → affected jobs · customer POs → expedite priority	Material PO delay impact · expedite drafts
Material movement · usage	Lot · serial → warehouse issue → AMR/AGV transport → consuming operation → finished-product serial, and in reverse	Timeline of warehouse (WMS) material movement records
Weld joint tracing	Per joint: welding procedure · welder qualification · conditions · NDT verdict on one line; a draft joint history sheet for shipment	Welder · inspector qualification checks, nonconformance links
Drawing standards	GD&T frame checks, fits → upper · lower limits, general tolerance · thread · surface roughness interpretation, citing standards documents the customer holds	Drawing reading · fit formula estimate
Inspection ↔ measurement	Drawing balloons (inspection characteristics) → first-article inspection sheet · measurement matching	Drawing reading · nonconformance links
Regulatory document package	A **draft** delivery · inspection document bundle from a single serial number	Revision compare · BOM check
Floor devices · source screens	Open an asset from a tablet · kiosk · QR code; offline viewing. Open a connected source system's screen directly, with context	Phone · home-screen app · Windows app
Permissions — next steps	SSO/LDAP integration	Role bundles (5 default + 8 advanced) · data scope · 4 document grades · audit screen · two-person approval · time-limited delegation · supplier accounts (p. 29)

Standards integration would work through rule logic and **citations from standards documents the customer holds** — no standard texts are shipped in the product.

10 · Comparison

How it differs from global industrial AI

In 2026 industrial AI is moving from copilots that answer to agents that act. The main offerings, side by side **from public information**.

Vendor · product	Strength	AI base · deployment (public info)	Recent moves
Siemens **Industrial Copilot · Intelligence Center X**	Engineering, operations and maintenance copilots for its PLCs, Teamcenter, Opcenter	Azure OpenAI (cloud); on-prem bundle on industrial PCs with NVIDIA NIM announced for operations	Eigen Engineering Agent GA (Apr 2026), nine copilots, industrial AI orchestration (Jun 2026)
Schneider Electric **AVEVA · Cognite Atlas AI**	Time-series historian (PI System) + industrial knowledge graph + agent workbench	Cognite: vendor-hosted SaaS by default, dedicated clusters and hybrid for large or regulated customers; Microsoft Azure AI partnership	Agreed to acquire Cognite for US$3.1B (30 Jun 2026); agentic manufacturing at Hannover Messe 2026
Rockwell Automation **FactoryTalk · Plex · Fiix**	PLC design copilot; agentic maintenance in MES and CMMS	FactoryTalk Design Studio is SaaS on Azure OpenAI	Design Plan · Build agents, maintenance agents, Plex Agent Composer
Honeywell **Experion Operations Assistant · Cognition**	Control-room operator assistance, alarm prediction	Google Cloud partnership (Gemini · Vertex AI)	Operations Assistant commercial launch (Mar 2026), Experion Cognition (Jun 2026)
Emerson **Ovation AI agents · Plantweb**	Power and water control, process asset optimisation	Deployment model not stated in public material	Five Ovation agents (alarms, predictive maintenance, root cause …)
Microsoft ecosystem **Factory Operations Agent etc.**	Manufacturing agents built by partners on Fabric and Copilot	Azure cloud	Accenture · Avanade “agentic factory”, GA planned for H2 2026

10 · Comparison

The path PlantPulse AI takes — five things

A customer-PO-centred digital thread and floor AI for air-gapped, engineer-to-order (ETO) plants. It targets plants where data cannot leave and each unit is built to order — heavy industry, shipbuilding, defence, power generation, aerospace.

What	In the product today	Next plans
① **Fast to connect**	ERP · MES · SCM · PLM · QM · CMMS · EHS · WMS plug in via open-MCP-spec connectors, enabled only after **an automatic conformance test** passes. Your systems are only read	Measure and publish the time to connect a new system; a connector kit
② **PO to floor**	Customer PO line → project → work order → operation → equipment in one line; material POs bound to work orders show how far a delay reaches **into a customer PO's delivery slack**. The «Today's work» answer in Ask	Trace back from a defect · serial number to the customer PO
③ **Fully air-gapped**	Language model, search, transcription and drawing reading all run inside the plant; installed and operated from a signed air-gap bundle (amd64 · arm64)	—
④ **Fits floor people**	Three screens (Ask · Proposals · Automations), one-handed phone use, automations made in words, approval by work role, Korean transcription · drawing reading · video search	Tablet · kiosk · QR, field pilot metrics
⑤ **Into paperwork**	Drawing revision compare · BOM check · fit upper · lower limit estimates	A **draft** delivery · inspection document bundle from a single serial number

How it differs from other paths

General-purpose data platforms

They can hold anything, but become a build project that models the plant from scratch. PlantPulse ships with plant concepts (work order · operation · revision · qualification · permit) built in.

Equipment-centric industrial data platforms

They connect equipment and sensors well, but orders · jobs · drawing revisions must be added separately. PlantPulse takes the customer PO as its root.

AI from PLM · MES suites

Deep inside their own suite, but they break off on floors that mix other vendors' systems. PlantPulse connects through vendor-neutral spec connectors.

Cloud digital twins

They provide the framework, but data has to leave. PlantPulse runs everything, model included, inside the air-gapped network.

The comparison above describes product types in general terms and does not evaluate the performance of any specific company's product.

11 · Screens

Screen tour — automations · drawings · video · done

Every screen below was captured from a **system in live operation**.



Figure 13. Automations › Templates — start from a built-in automation such as condition watch · morning brief · shift handover memo · recurring fault causes, or create a new one from one sentence.

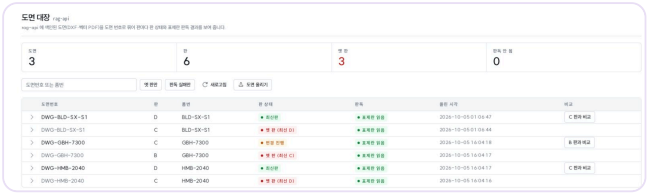


Figure 14. Admin console › Drawing register — grouped by drawing number, one row per revision with its status (latest · old revision · change in progress) and parse result.

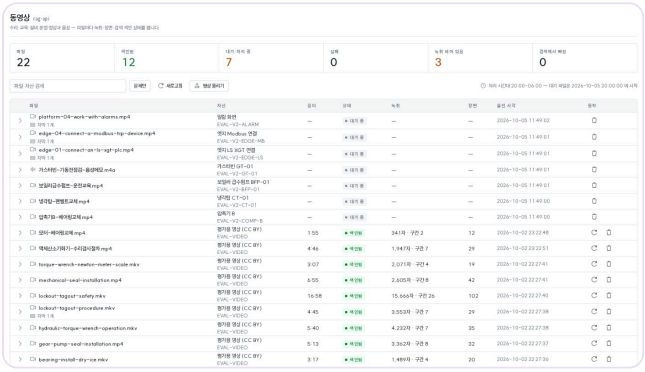


Figure 15. Admin console › Videos — transcript, scene and index status for each repair and training video or recording.



Figure 16. Proposals › Done — the history of executed actions and a one-line summary at the top: actions · time to close · passed check · recurred · awaiting result.

Closed work becomes the metric

Proposals that pile up unclosed are worthless. The top of the Done tab shows actions · time to close · **check pass rate** · **recurrence**, and the weekly report is produced by an automation as a PDF.

Screen design principle

The first screen is **Ask**, and there are three things you do — **Ask** · **Proposals** · **Automations**. Screens are monochrome; colour is used only for status dots (urgent · caution · normal). One primary button per screen. Without integrations, no empty slots or «connect to see this» remain.

The automation and Done screens were captured on staging on 2026-10-12, the drawing register and videos on 2026-10-05.

11 · Screens

Screen tour — administration and control

Adopting AI also means **managing AI**. Operators must be able to see and adjust what it answers from, by what criteria it calls something abnormal, and how often it was right.

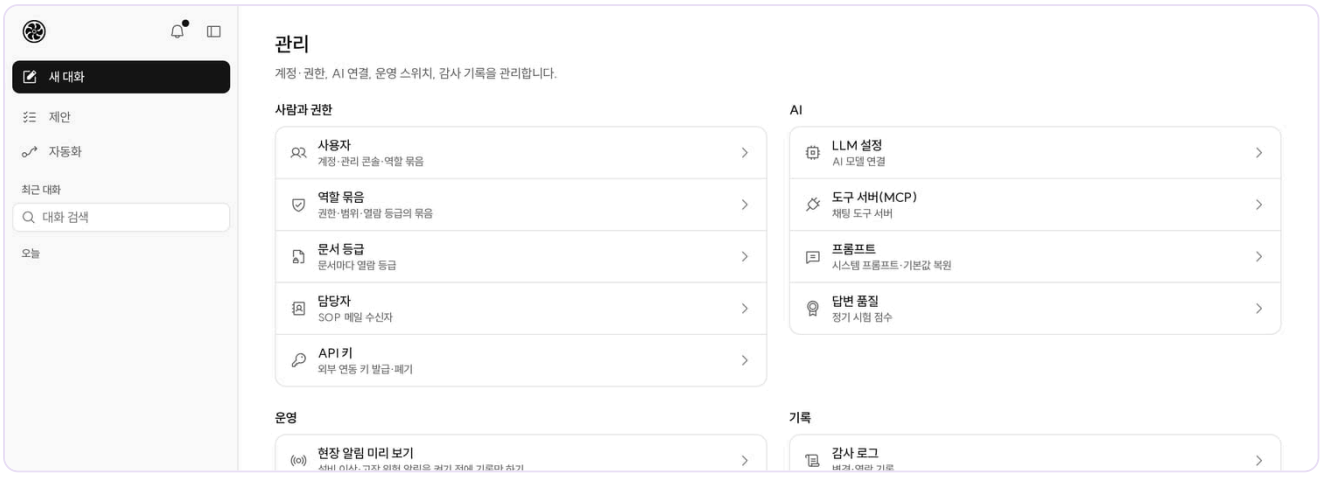


Figure 17. Admin — four groups: people and permissions · AI · operations · records. It sits at the same address as the copilot and opens only for administrators.

What operators adjust

Item	Content
Users · permissions	Role bundles · data scope · document grades — who sees and approves what
Detection criteria	Asset criticality grades and alert criteria. Anomaly watching is moving to automations — the floor says in plain words what to check and when
Guidance wording	Diagnosis and proposal wording tuned to floor vocabulary — edit history is kept automatically
Model connection	Choose between the on-premises model and an external API, and confirm with a connection test
API keys	Keys for the public API and MCP — the value is shown only once, right after issue
Document · drawing checks	Document parse check · drawing parse test · drawing register — checked before indexing, old revisions filtered out
Digital thread	Manage connector registration · conformance tests · bindings on the integration map
AI accuracy	How often the AI is right, from the floor's dismiss reasons · accept/reject · effect records and scheduled test scores
Branding	Customer name · logo · CI colours · login screen changed from the admin screens (white-label)

> **Not a black box.** The floor accepts or rejects AI proposals and records their effect, and dismisses false alarms with a reason. Those marks feed accuracy and threshold tuning. Accept and reject are records only — they never change equipment or schedules.

12 · Security

Data sovereignty · security

The most common point where industrial AI projects run aground is not performance but the **security review**. PlantPulse AI is designed to pass it.

A structure where data does not leave

Component	Placement
Language model (AURA)	Runs on a GPU in the customer's infrastructure — prompts never go outside
Time-series insight server	Runs in-house; models ship in the image and work offline
Document search engine	Opens no ports to the outside; reached only through the search API — documents never leave
Source data	Stays on the platform — the AI makes no copies
Your systems	The digital thread only reads and keeps keys · references instead of source data — the only write is a single approved draft
Drawings · video · photos	Originals and indexes stay on your servers, transcription uses the in-house model too — photos are discarded after reading

In other words, a configuration in which **neither process data, documents nor question text leaves the site** is possible. Every function is complete in closed and air-gapped environments.

Permissions and audit

- **Role bundles · data scope · document grades** — 5 default + 8 advanced bundles, per-person scope, four grades up to controlled, every view on the **audit screen**.
- **Two-person approval · delegation** — no self-approval; two approvers for supplier expedites.
- Every tool the AI calls is written, with its arguments, to the **tool-call ledger**. Secret-looking arguments are masked on screen and in storage.
- Tools that change things are blocked by default — they require **human approval**.
- Work orders and recovery plans the AI writes are **drafts** — never registered in CMMS or MES automatically.
- URLs the AI passes on cannot reach private-network or metadata addresses, and stored credentials are encrypted with AES-256-GCM.

Trust in the answers

If it does not know, it does not answer

Low-confidence tag interpretations stay out of diagnosis; an empty document search returns “not found”; a meter reading needs 4 of 5 reads to agree on value and unit, or it is “unreadable”.

It shows its evidence

Document answers show their source document; numeric answers keep, under “How this answer was made”, the tools called and what they returned.

We recommend a security review before deployment. The configuration above is what the product supports; the actual placement (model location, network segregation level, log retention) is agreed during deployment to match your policy.

13 · Requirements

Installation · hardware requirements

For an AI product, **where it runs** is the requirement. PlantPulse AI does not assume a large GPU cluster.

Three configurations

Configuration	Description	Suits
On-premises (recommended)	The model is deployed on the customer's own GPU hardware. Data and prompts never leave the site	Closed networks · no data export
In-house gateway	Connects to a model-serving gateway the company already operates	Enterprise AI infrastructure in place
External API	Connects a commercial model API	Where security policy allows

Hardware

Language model

The language model (aura-35b, NVFP4), embedding, reranker and time-series engine share **one GPU**. The measured configuration is the **GB10 class** (arm64, 128 GB unified memory).

Server GPU configuration

A single 96 GB-VRAM server GPU (amd64, e.g. AWS g7e) is also supported. Its GPU figures are converted estimates, to be confirmed by measurement before deployment.

Exact specifications are sized from **tag volume, concurrent users and document volume**. Closed-network installs use a **signed air-gap bundle** — after import its signature (cosign) and hash are verified, and an unverifiable bundle never counts as a pass.

Prerequisite

> **PlantPulse Platform is assumed.** The AI answers from the asset model, time series, alarms, ontology and documents the platform has standardised. Deploying the AI without the platform leaves no data structure to ground the answers in, and it becomes **no different from a generic chatbot**.

14 · References

References

An AI built on data gathered over 15 years on manufacturing floors.

162 projects delivered	**34** sites in operation	**642,021** connected sensors
100 TB+ data under management	**15 years** manufacturing-floor experience	**6** awards

Representative deployments

POWER · HEAVY INDUSTRY

Doosan

ICT-converged wind-farm operations system — combined with a digital-twin framework to improve availability through turbine health tracking and predictive maintenance.

FOOD & BEVERAGE

Lotte Chilsung

Smart-factory platform and real-time OEE system for production lines.

SHIPBUILDING · OFFSHORE

Hanwha Ocean

Real-time pipeline for large-scale shipyard OT data.

LOGISTICS · ELECTRONICS

Samsung Electronics

OT data platform for a next-generation logistics centre.

PUBLIC · TESTBED

AX testbed

Selected as the builder of an industrial digital transformation (AX) testbed.

RESEARCH · TESTBED

KERI

Industrial data platform for the Korea Electrotechnology Research Institute's smart-factory testbed plant.

Customers we have worked with

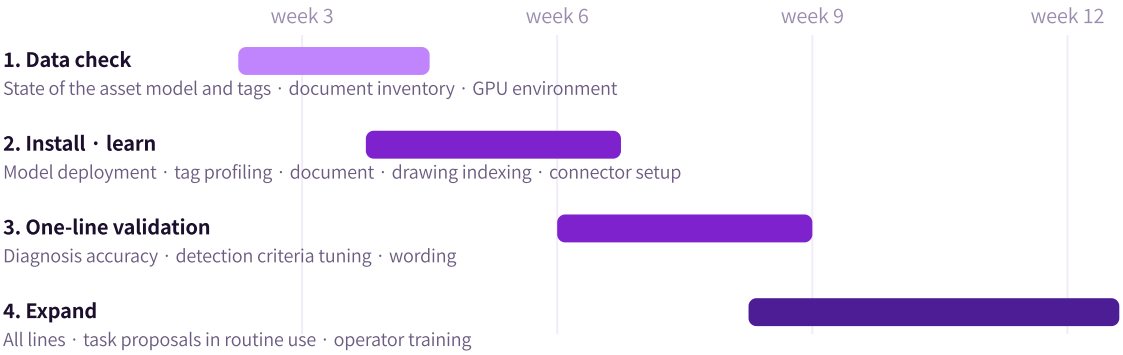
Samsung Electronics · Samsung C&T · LG · Doosan · Doosan Mottrol · Hanwha · Hanwha Ocean · Hanwha Aerospace · Hanwha Engine · Lotte Chilsung · Lotte Chemical · Lotte Fine Chemical · Lotte Wellfood · Amorepacific · Hyundai Rotem · Hyundai Mobis · HD Hyundai Heavy Industries · Hyosung · Hansol · SK Chemicals · CJ Logistics · KAI · KEPCO · KT · Denso · Carrier · Seegene · Shilla · Winia · STX · Iljin · KERI · KEIT · Ministry of Trade, Industry and Energy and others.

AI deployment cases and quantified outcomes are provided individually through your sales contact.

14 · Deployment

Deployment plan and schedule

AI answers only as well as the data is prepared. So we recommend deploying in the order **data check → start narrow → widen**.



What to confirm in validation

- When asked with our equipment names and tags, **does it understand**
- Do the anomalies it flags **match the floor's judgement** — checked against correct / false-alarm marks on the AI Accuracy screen — and can the criteria be adjusted when they don't
- Do answers **carry evidence** — which data, which document · drawing · video
- Does work readiness **correctly flag blocked operations from our MES · ERP records**, and does it read our drawings accurately
- In a closed network, **is every function complete**

> **Honest expectations are the condition for success.** AI does not work magic on unorganised data. In the stage-1 data check we first agree on **what is possible now and what is not yet**.

14 · Next steps

Decisions to make together

Where to run the model

On-premises GPU, in-house gateway or external API — whichever fits security policy and budget.

Where to start

Picking the one line where downtime costs most gives the strongest validation.

Which systems to connect first

Work readiness rests on MES. On top of it, decide in what order to add PLM · QM · SCM · CMMS · EHS.

What to include as evidence

Which of equipment manuals, work standards, drawings, maintenance and training videos to use as evidence.

Who receives the results

Assignment rules for task proposals and the notification channels (email · web push · Slack).

Seen as a product family

The four PlantPulse products form one continuous flow. AI is the last stage, and the more complete the earlier stages, the more accurate it becomes.

Product	Role
PlantPulse Edge	**Collect** — brings data out of the equipment
PlantPulse Platform	**Process** — stores, standardises and turns it into KPIs
PlantPulse AI	**Intelligence** — connects systems, reads documents · drawings · video, judges the floor
PlantPulse Studio	**Create** — the floor builds the screens it needs

NEXT STEP

30-minute demo · data-check workshop

The fastest way is to see natural-language queries on a live instance. Share some of your equipment names and tag conventions and we will demonstrate **in your vocabulary**.

PlantPulse® AI

Plant operations AI copilot · On-premises / air-gapped

Product documentation docs.plantpulse.io · Document Rev. 2026.10

PlantPulse® **AI**

PLANT OPERATIONS COPILOT

GET IN TOUCH

Start your industrial AI transformation today.

Consultation, demo requests, quotations — whatever you need, leave us a message. A specialist replies within one business day, and we can schedule a live demo on the real product right away.

PHONE
+82 1588-5268

EMAIL
webmaster@kopens.com

INQUIRIES · DEMO REQUESTS
kopens.com/contact

COMPANY
kopens.com

PRODUCT
plantpulse.io

DOCUMENTATION
docs.plantpulse.io

THE PLANTPULSE® FAMILY — DEPLOYED TOGETHER, ONE STACK FROM THE FLOOR TO AI

Edge	Platform	AI	Studio
Industrial protocol edge gateway	Enterprise IIoT / DataOps platform	Plant operations AI copilot	AI application builder

KOPENS, Inc.

© 2026 KOPENS. All rights reserved. PlantPulse® is a registered trademark of KOPENS, Inc.

Document Rev. 2026.10 · Figures and screens are as of the date of writing and may change without notice.

kopens®