



POWER GENERATION

PRIVATE INDUSTRIAL AI FOR POWER PLANTS

Critical Infrastructure Protection, Reliable MW Capacity and
Coordinated Plant Knowledge

WHITE PAPER

CORE THESIS

Power plants do not need another isolated AI tool. They need a governed intelligence environment that connects asset condition, maintenance, operating context, generation capacity and expert knowledge - while keeping critical plant data and decision authority under local control.

SENSE	PROTECT	CONTEXT	REASON	ASSURE
Machine evidence	CIP governance	Plant knowledge	Local AI	MW capacity

PHOENIX Industrial AI
Made in the USA
+1 336 763 2882 | sales@phoenix-unifiedsystems.com



Executive Summary

Power generation is a capacity business. The value of reliability is not simply fewer alarms or better maintenance metrics; it is the ability to commit megawatts with confidence, avoid forced outages, plan maintenance at the lowest operational cost, and protect the revenue associated with available generation.

PHOENIX is designed around that outcome. It combines machine sensing, long-range wireless and edge collection, PHOENIX CIP for governed access and security, LabSource for local AI and knowledge, and AgentSWARMS for model and agent orchestration. Together, these layers create a coordinated plant knowledge experience rather than another isolated analytics application.

FROM ASSET HEALTH TO CAPACITY ASSURANCE

ASSET HEALTH -> AVAILABILITY -> RELIABLE MW CAPACITY -> PRODUCTION /
DISPATCH CONFIDENCE -> REVENUE PROTECTION

The opportunity

- Bring condition monitoring, maintenance history, operating state, process data, engineering documents and expert knowledge into one governed context.
- Run multiple local LLMs on customer infrastructure, avoiding external per-token inference charges for local workloads and keeping sensitive knowledge on-premises.
- Allow agents to select the best model for each task and hand critical conclusions to a second model or specialist agent for independent verification.
- Use PHOENIX CIP to control credentials, tools, data scope, audit, permissions and customer-initiated outbound HTTPS communication over port 443.
- Connect plant intelligence to the physical world using vibration, temperature, pressure, oil, electrical, GPS/BLE and other edge-collected sensor data.
- Turn the growing contextual database into reusable reports, prompts, internal applications, analytics and decision support.

Important security boundary

PHOENIX is positioned as a decision-support and intelligence layer, not a direct autonomous control layer for critical plant equipment. Read-only access, approved tools, scoped credentials, review gates and local governance can be used to keep AI separated from shutdown-capable or control-authority pathways.

PHOENIX CIP is a product architecture for Critical Infrastructure Protection. It can support a utility security program, but this white paper does not claim that PHOENIX by itself establishes compliance with NERC CIP or any other regulatory standard.



1. Why Power Plants Need a Different AI Architecture

A power plant already contains large amounts of valuable information, but it is usually distributed across separate systems, teams and time horizons. Reliability may understand vibration. Operations understands load and process behavior. Maintenance understands work history. Planning understands outages. Engineering holds design knowledge. Commercial teams understand the value of lost or available generation.

The AI problem is therefore not simply model intelligence. It is coordinated context.

Context	Typical sources	Decision value
Asset condition	Vibration, temperature, oil, current, speed, pressure, valve position, thermal data	What is degrading and how quickly?
Operations	Load, heat rate, starts, ramping, process conditions, operating mode	Is the behavior explained by how the unit is being operated?
Maintenance	CMMS work orders, PMs, backlog, labor, parts, outage scope	What intervention is already planned or required?
Engineering	Drawings, manuals, failure modes, limits, prior investigations	What does this equipment design allow or constrain?
Planning	Outage windows, dispatch expectations, redundancy, reserve requirements	When can work be done with least capacity impact?
Commercial	MW value, fuel/efficiency penalty, contract exposure, lost generation	What is the business consequence of the technical risk?

The gap with standalone AI

A model can be very capable and still make a weak plant decision if the evidence provided to it is incomplete. A vibration trend without load context can be misleading. A maintenance recommendation without outage context can be impractical. A production recommendation without equipment condition can increase failure risk. PHOENIX is designed to assemble the right evidence before asking a model to reason.



2. The PHOENIX Power Generation Architecture

The architecture is deliberately layered. Each layer has a different responsibility, allowing plant teams to separate sensing, security, context, reasoning and decision support.

PHYSICAL WORLD	Turbines, generators, pumps, fans, motors, gearboxes, transformers, cooling systems, balance-of-plant equipment and distributed sensors.
EDGE / SX4	Local acquisition, preprocessing, buffering and edge collection for high-resolution and distributed machine data.
PHOENIX CIP	Identity, vault, credentials, permissions, approved tools, facility scope, audit, data governance and controlled outbound communication.
LABSOURCE	On-prem AI workspace, contextual database, local LLMs, RAG, documents, engineering knowledge, coding and reporting.
AGENTSWARMS	Routes tasks to the best model or specialist agent, assembles context, invokes tools and can request independent verification.
PHOENIX UI	Turns technical evidence into reliability, outage, capacity, maintenance, lifecycle and revenue decisions.



3. PHOENIX CIP for Critical Infrastructure

Power generation creates a stronger requirement for governance than a normal enterprise AI deployment. The AI environment must know not only what it can access, but what it must never control.

Core design principles

- Customer-controlled, on-premises AI infrastructure for sensitive plant knowledge and local inference.
- Approved system connections rather than broad unmanaged credentials.
- Read-only access patterns for critical operational systems where write capability is not required.
- Secrets and credentials held in a governed vault rather than embedded in agent prompts or code.
- Tool-level permissions that define exactly which queries, files, APIs or systems an agent may use.
- Facility, asset and role scope so access can be constrained to the user and task.
- Auditable records of the request, context used, model selected, tool calls, response, verification and human disposition.
- Customer-initiated outbound HTTPS communication over port 443 for approved external exchanges, avoiding unsolicited inbound plant connections.

CONTROL STAYS WITH THE PLANT

The objective is not autonomous control of generation equipment. The objective is to give engineers and operators better evidence, faster context and higher-confidence recommendations while preserving existing control-system authority and human approval.

A practical data-access pattern

1	2	3	4
AI requests context	CIP authorizes scope	Local tools retrieve data	AI returns evidence + recommendation

4. Local Multi-Model AI Instead of a Single Model Bet

Different plant tasks need different kinds of intelligence. Code generation, document interpretation, reliability reasoning, numerical analysis and independent review do not necessarily belong to the same model. LabSource can host multiple local LLMs and let AgentSWARMS choose the best fit by task, latency, context size, specialization and verification requirement.

Task	Primary AI role	Verification / guardrail
Python / HTML application development	Coding-focused local model	Test agent + second-model review before release
Equipment diagnosis	Reasoning model with plant context	Independent reliability-review model +



Task	Primary AI role	Verification / guardrail
		evidence check
Manuals / drawings / reports	Document or vision-capable model	Source citation / retrieval verification
Fast triage / classification	Smaller low-latency local model	Escalate uncertain cases to stronger model
Executive capacity assessment	Reasoning model across health, operations and planning	Cross-check numbers and assumptions against retrieved sources

For local models, the recurring economics shift from external per-token inference charges toward owned or leased compute, electricity, cooling and system administration. That can make high-volume internal AI workflows more predictable while preserving the option to use approved external models for specific tasks when policy allows.



5. From Asset Health to Reliable MW Capacity

The most important change in the PHOENIX power-plant model is that asset health is not the final KPI. It is an input into reliable generation capacity.

HEALTH	AVAILABILITY	CAPACITY	PLAN	VALUE
What is degrading?	Will it run?	How many MW are reliable?	Can we meet dispatch / outage commitments?	What revenue or cost is exposed?

A reliable-capacity calculation can consider

- Nameplate and current operating capacity.
- Health state of generation-critical and balance-of-plant assets.
- Redundancy and single-point-of-failure exposure.
- Known derates, operational restrictions and environmental limits.
- Current maintenance backlog and planned outage scope.
- Availability of labor, parts and specialist support.
- Expected operating mode, starts, stops, ramps and load profile.
- Efficiency or heat-rate penalties associated with degraded equipment.

EXECUTIVE QUESTION

Not: How many critical alarms do we have? Instead: What generation capacity is at risk, why is it at risk, when does it become material, and what is the lowest-cost action that protects the plan?

6. Example: Turning a Vibration Alert Into a Capacity Decision

Consider a critical induced-draft fan motor with rising vibration. A conventional monitoring workflow may raise an alarm and create a maintenance task. A contextual PHOENIX workflow goes further.

1. Detect	High-resolution vibration indicates a rising component-specific trend.
2. Contextualize	Agent retrieves motor speed, load, fan operating point, temperature, recent work orders and prior alignment history.
3. Explain	Reasoning agent separates operating-condition effects from likely mechanical degradation.
4. Quantify	Capacity agent determines whether loss of the fan derates the unit, by how much, and whether redundant equipment



	can carry the load.
5. Plan	Maintenance agent checks outage windows, parts availability and required labor.
6. Verify	Second model reviews the diagnosis, assumptions and evidence before the recommendation is surfaced.
7. Decide	PHOENIX UI presents risk, MW exposure, recommended timing, supporting evidence and approval path.

The result is not merely a failure prediction. It is a coordinated maintenance-and-capacity decision with a defensible evidence trail.

7. Building a Coordinated Plant Knowledge Experience

One of the largest sources of plant intelligence is often not stored in a database at all. It lives in the experience of operators, maintenance supervisors, reliability engineers, OEM specialists and long-serving technicians. LabSource is designed to capture that knowledge as part of normal work rather than as a separate knowledge-management project.

EXAMPLE EXPERT CORRECTION

"That 1x rise is normal above this load on Unit 2. Check inlet-vane position and process draft before calling imbalance." Once reviewed and approved, that correction can become reusable context for future investigations.

What the contextual database can retain

- User question, facility, unit, asset, component and time range.
- Data sources retrieved and exact evidence supplied to the model.
- Agent and LLM selected for the task.
- Tool calls, calculations, code and report logic.
- Primary conclusion and independent verification result.
- Engineer comments, corrections, acceptance or rejection.
- Subsequent maintenance action and operational outcome.
- Reusable domain rules, prompts and approved examples.

This creates a compounding effect: the environment becomes more useful as the plant uses it, because new work is performed against a richer history of plant-specific decisions and outcomes.



8. Let Plant Teams Build on the Knowledge Base

The same governed context used for analysis can support local application development and reporting. Domain experts do not need to wait for a large enterprise software project every time a new operational need appears.

LOCAL APPS Build controlled Python / HTML tools such as outage readiness views, pump performance calculators, condition dashboards or inspection workflows.	CONTEXTUAL REPORTING Drag-and-drop report blocks can reference assets, work orders, condition, capacity, expert knowledge and recommendations rather than hard-coded isolated queries.
AI SUGGESTIONS Prompt suggestions can be generated from current asset, unit, role and operating context so users are guided toward high-value questions.	REUSABLE INTELLIGENCE Approved calculations, agents, reports and applications become part of the plant knowledge system instead of disappearing inside one-off analyses.

9. A Practical Deployment Path

A power-plant AI program does not need to start with every system connected. A controlled rollout can prove value while the governance model matures.

Phase 1 - Govern

Deploy LabSource and PHOENIX CIP; establish identity, vault, permissions, approved models and audit.

Phase 2 - Connect

Connect a limited set of read-only sources such as CMMS, historian extracts, documents and PHOENIX sensor data.

Phase 3 - Contextualize

Map units, systems, assets, components, operating states, maintenance relationships and capacity impact.

Phase 4 - Prove

Choose 2-3 high-value use cases: rotating equipment diagnosis, outage readiness, capacity-at-risk or engineering knowledge retrieval.

Phase 5 - Scale

Add approved agents, reports, local apps, more data sources and cross-model verification workflows.

Phase 6 - Institutionalize

Capture expert corrections and outcomes so plant knowledge continuously improves.



Initial use cases with executive relevance

- What generation capacity is at risk over the next 7, 30 and 90 days?
- Which maintenance actions protect the most MW for the least outage impact?
- Which critical assets are degrading faster than their current production plan assumes?
- Which derates are mechanical, process-driven or maintenance-driven?
- Which known plant issues are being repeatedly rediscovered because expert knowledge is not reusable?
- What work can be safely deferred, and what work cannot, when capacity and revenue exposure are considered together?

Conclusion: Protect the Plant, Grow the Knowledge, Assure the Capacity

The strategic opportunity for AI in power generation is larger than predictive maintenance. The objective is to create a secure intelligence layer that understands the equipment, the operating environment, the maintenance plan, the plant's engineering history and the business value of available generation.

PHOENIX combines physical sensing, governed connectivity, local multi-model AI, contextual knowledge, agent orchestration and decision support so plant teams can move from isolated information toward a coordinated knowledge experience.

PHOENIX POWER GENERATION PROPOSITION

Protect critical infrastructure. Keep intelligence local. Connect plant evidence. Preserve expert knowledge. Turn asset health into reliable MW capacity - and reliable capacity into stronger operational and revenue decisions.