TRUE NORTH PROGNOSTICS

High-Confidence Transformer Early Warning

A blind-validation proposal using existing utility historian data — no new sensors, no thresholds, and no operational risk during evaluation.

Utility validation version

April 2026

The utility problem: failures are rare, but consequences are severe

- HV transformer failures can drive emergency capital spend, outage duration, safety exposure, and public scrutiny.
- Conventional indicators often confirm risk after degradation has already progressed.
- The hard decision is not whether data exists — it is whether the evidence is strong enough to act early.

The risk lives between first indication and defensible action.

Earlier	Lower	Better
warning window	unnecessary alarms	capital timing

Why today's monitoring can still be late

Common current approach

- DGA and oil testing are critical, but they tend to become informative after chemical evidence appears.
- Threshold alarms are useful for protection, but they are not optimized for earliest detection.
- Single-signal monitoring can miss multi-signal behavioral drift.

What is hidden in the data

- Load, weather, seasonal and day/night cycles mask smaller degradation signatures.
- Signals arrive at different sampling rates and data quality varies by source.
- Early degradation may appear as subtle correlated changes before it becomes obvious.

Our position: an independent early-warning layer

- AI-MSET analyzes existing historian and SCADA data to learn expected transformer behavior.
- It is multivariate: signals are evaluated together, not as isolated thresholds.
- It is initially positioned as complementary to DGA — an earlier indication layer, not a replacement claim.
- The evaluation goal is simple: prove earlier, useful warning on data the utility already owns.

No new sensors
No operational changes
No automatic trips

How existing data becomes more useful

Normalize

Remove load, ambient, seasonal and day/night effects

→

Align

Convert multi-rate historian streams into a uniform analysis view

→

Repair

Handle missing values, drift, sensor disturbance and data quality defects

→

Infer

Use correlated signals over time to extract higher-fidelity system behavior

The method does not depend on one perfect sensor. It improves confidence by using relationships across many imperfect signals.

Detection is based on behavior, not thresholds

- Expected behavior is learned from historical operating data under comparable conditions.
- A first alert is generated when correlated signals deviate from expected behavior in a statistically meaningful way.
- Repeated and correlated alerts are used to separate isolated data noise from developing equipment risk.
- Outputs are confidence-weighted and designed to support engineering judgment.

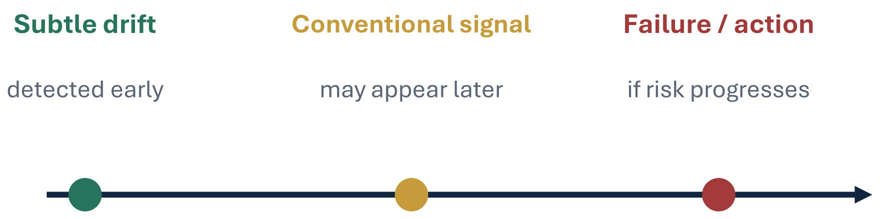

An alert must answer: data problem, sensor problem, or asset problem?

Data anomaly

Isolated bad value, missing data, communication issue, or historian defect

Sensor issue

Repeated abnormality tied to one instrument or measurement channel

Asset degradation

Correlated changes across equipment-relevant signals over time

The operational value is not just early warning — it is knowing what kind of action is warranted.

The right first test is blind validation

- The utility provides three anonymized transformer datasets.
- Any dataset may contain a failure or developing issue; only one may; none may; or all may.
- True North is not told if or when any transformer failed, had abnormal DGA, or received major maintenance.
- We return a dated event log and classification before ground truth is revealed.

Blind results are more credible than vendor claims.

Dataset	Possible truth state	TNP response
A	failure / no failure / data issue	dated event log
B	failure / no failure / data issue	dated event log
C	failure / no failure / data issue	dated event log

What we return before the utility reveals ground truth

Dataset	Finding	First detection date	Confidence	Classification
A	Developing equipment issue	YYYY-MM-DD	High	Asset behavior
B	No actionable issue	—	—	Normal
C	Repeated sensor disturbance	YYYY-MM-DD	Medium	Sensor / data

- The utility scores true positives, true negatives, false positives, missed known events, and lead time.
- This avoids hindsight bias and prevents the model from being tuned to known failures.
- A successful test creates the evidence needed for a low-risk shadow pilot.

Metrics utilities can act on

Lead time	Days/weeks between first detection and DGA, maintenance action, or failure
False alerts	Alerts per transformer-year, not abstract lab denominators
Missed events	Known failures or major degradation events not flagged
Actionability	Whether the output correctly separates data, sensor, and equipment causes

Published low-error methods are useful; utility-grade proof comes from blind fleet data.

How results become operational decisions

Detect — early behavioral deviation → **Classify** — data, sensor, or asset → **Prioritize** — engineering review → **Plan** — maintenance or monitoring → **Record** — audit trail and outcome

- The system supports engineering judgment; it does not bypass utility procedures.
- Outputs can feed work management systems, dashboards, or existing asset-health workflows.
- The initial deployment can run in shadow mode with no operational consequences.

Business case: risk avoidance, not incremental monitoring

- Avoiding one high-consequence transformer event can justify predictive capability many times over.
- Earlier confidence helps utilities schedule outages, reduce emergency crews, and protect reliability metrics.
- Better risk timing supports inventory strategy and capital discipline under long transformer lead times.
- High-consequence risks such as fire, collateral damage, and public scrutiny become easier to manage proactively.

**The benefit is not "more alarms."
The benefit is fewer surprises and more defensible action.**

Reliability	Capital	Safety
SAIDI / outage planning	planned vs emergency	risk reduction

Designed for trust, governance, and auditability

What utilities need

- Confidence-weighted outputs, not black-box directives.
- Evidence that can be reviewed by engineering, operations, audit, insurance, and regulators.
- Clear distinction between equipment risk and bad telemetry.

How the system supports it

- Event logs with timestamps, signal context, and classification.
- Integration into existing work management and asset-health workflows.
- Validation path from historical blind test to shadow mode before operational adoption.

Deployment path: low-risk by design

1

Blind historical benchmark

Three anonymized datasets; labels withheld; scored after results are submitted

2

Shadow-mode pilot

Live historian feed; alerts compared against existing process without operational action

3

Operational workflow

Confidence-weighted outputs routed to engineering review, planning, or work management

What we need from a utility partner

- Three anonymized transformer historian datasets, preferably 12–36 months each. Longer datasets, i.e. 36+ months are better to block out the diurnal and seasonal temperature variations but analysis can be done with shorter datasets.
- Available signals such as load, current, voltage, oil temperature, winding temperature, ambient temperature, tap position, alarms, status points, etc.
- Ground truth withheld until after the dated event log is submitted.
- After submission: DGA history, oil test results, inspections, maintenance records, and known events for scoring.

Data format

CSV export from existing historian systems is sufficient for initial evaluation.

Operating principle

No decisions are made from the test. The first objective is evidence, not deployment.

Evaluation discipline

Any result is useful: success, failure, or partial performance all clarify where the method belongs.

Strategic takeaway

From reactive confirmation to earlier, defensible action

- Utilities already have much of the data needed to evaluate early transformer risk.
- The next step is not a sales claim — it is a blind benchmark that proves whether the method finds real risk without being told where to look.
- If validated, the result is a practical early-warning layer that improves planning, reliability, capital discipline, and risk governance.

Proposed next step: approve a blind historical validation using three anonymized datasets.