IGNITION
DESIGN LABS

TECHNICAL WHITE PAPER

AIOps
A Network That Explains Itself

Multi-signal anomaly detection, correlation and LLM-ranked root cause — replacing monitoring that only tells you something already broke.

Monitoring tells you what already happened

By the time the alarm fires, the user has already noticed.

As IT infrastructure grows more complex, the inflow of log data, system metrics and operational signals becomes unmanageable. Traditional monitoring systems lack the intelligence to analyse multi-dimensional data in real time, and they are fundamentally reactive: they respond to problems after those problems have occurred.

The result is familiar to anyone who has run a network. An alert fires. It is one of hundreds. Somewhere in a flood of logs is the line that explains it, and an engineer will spend the next two hours finding that line.

> **The distinction that matters.** Detecting an anomaly is easy — thresholds do that. Explaining it is hard. IDL's AIOps platform is built around the second problem, because the second problem is the one that costs money.

What makes IDL's approach different

Rather than focusing narrowly on wireless, IDL takes a **multi-signal** approach spanning networking, cloud services and system performance. Logs, metrics and external signals are learned *together*, in a single unified model, so that one signal can explain another — instead of arriving as a stack of disconnected alarms that a human is left to correlate.

One model, many signals

The Holistic Log Model.

At the core is a model trained across several dimensions at once:

- Log rates, sequences and structure
- Anomalies in numeric and discrete fields
- Time-based patterns and seasonality — the daily and weekly cycles that make "unusual" a moving target

Data is collected in real time across servers, networks and applications, with metrics sampled frequently enough to capture live performance. Log data is sampled at ten-minute intervals — 144 samples a day, 1,008 a week — which is what allows diurnal and seasonal cycles to be modelled rather than smoothed away.

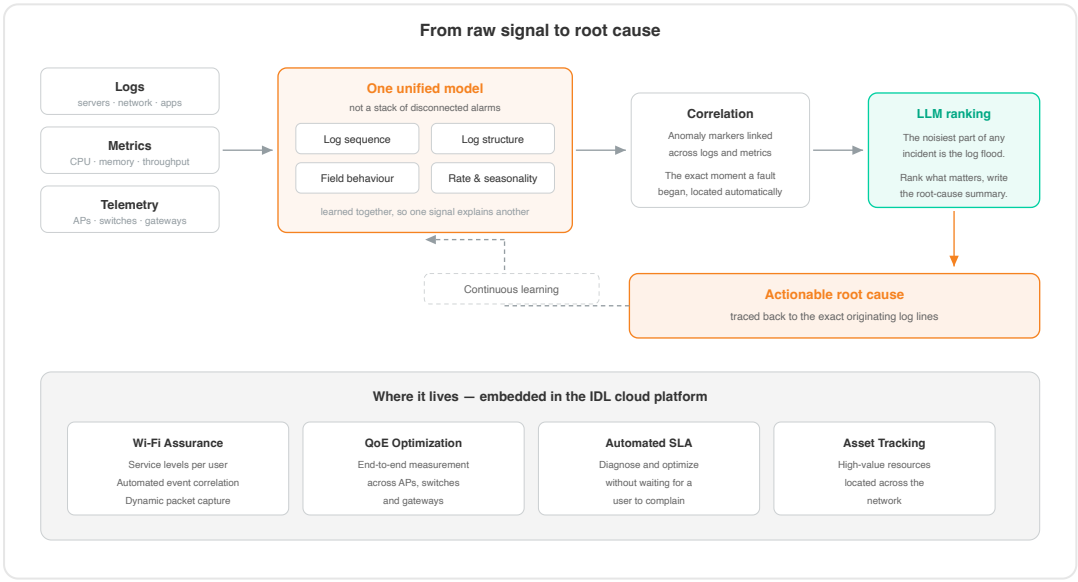

From raw signal to actionable root cause

How anomalies are actually found

Five detectors, one verdict.

- **Log sequences** — bidirectional transformers identify contextual anomalies in sequences of log lines, catching the log that is wrong *given what came before it*, not merely rare.
- **Log structure** — every log is parsed and tokenized by type. Histograms of those tokens are compared continuously against historical distributions; divergence produces a drift score.
- **Numeric fields** — gauge-type values are identified statistically, then density-based clustering isolates outliers. Clusters are normal behaviour; what falls outside them is not.
- **Discrete fields** — low-cardinality fields such as severity are tracked for changes in distribution.
- **Log rate** — the event rate itself is monitored as a signal, because a sudden silence is as diagnostic as a sudden flood.

A score is not an answer

Models produce scores — loss, drift, outlier percentage. On their own these are useless to an engineer. Effective troubleshooting requires retrieving and correlating the actual logs, metrics and thresholds behind the anomaly. That is the subject of the next section.

From symptoms to cause

Correlate, locate, rank, explain.

Correlation

When an anomaly occurs, the platform correlates markers from both logs and metrics to identify the relationships between signals — which ones are related, and which are affecting each other. The higher the anomaly score, the more severe the problem.

Normalization then sharpens the rising and falling trends, and the exact moment a fault began is located automatically — so the logs retrieved are the ones from the moment it started, not the moment somebody noticed.

Ranked by an LLM

The noisiest part of any incident is the log flood. A large language model ranks logs by relevance and severity — based on anomaly scores, context and rules — producing a filtered, prioritized list rather than an archive. For metric symptoms, the platform retrieves the metric name, its outlier percentage and its deviation from normal, and reasons about cause and effect across them.

> **The output is a summary, not a dashboard.** Interactive dashboards present root causes dynamically, with the reasoning made explicit. Engineers can add and adjust anomaly markers themselves rather than living inside a fixed set someone else chose.

Traceability

Every signal traces back. A log sequence returns to its original line by sequence ID. A log type returns the logs matching that structure at a given timestamp. A numeric or discrete value returns the logs that carried it. The model never asks to be trusted on faith.

It keeps learning

Networks change. A model trained once is a model that is already wrong.

The platform operates in continuous learning cycles. Log patterns are learned first without consuming GPU resources, assessing log volume and density to identify which data deserves priority. GPU-based training then targets the highest-priority logs. Inference runs continuously, with reports produced at ten-minute granularity.

When new data significantly alters the operational environment, models retrain. Reinforcement learning continuously adjusts detection in place, so accuracy improves with exposure rather than degrading with drift.

Where it lives

AIOps is not a separate product to integrate. It is embedded in the IDL cloud management platform:

Wi-Fi assurance

Service levels per user, automated event correlation, dynamic packet capture, policy configuration, guest access.

QoE optimization

End-to-end measurement using data from access points, switches and gateways; automated SLA measurement and AI-driven adjustment.

Asset tracking

High-value resources located across the network.

Across every solution

The same engine serves SD-WiFi, SD-Access, SD-Branch, SD-Edge and MDU deployments.

What it is for

Not to produce more alerts. To produce fewer, and to explain them — so that operations moves from reactive firefighting to prediction, and the cost of running a network stops scaling with its size.

About Ignition Design Labs

Ignition Design Labs (IDL) is a market leader in next-generation carrier-grade convergent Wi-Fi and 4G/5G wireless access technologies.

IDL has distinguished itself in the telecommunications sector most notably through its partnership with a Tier-1 operator in East Asia serving over 100 million subscribers — a collaboration that produced one of the world's largest integrated Wi-Fi and cellular systems, with over 400,000 access points across the capital area.

Beyond convergence, IDL addresses the user-experience problems of large-area Wi-Fi coverage: scalability, cognitive radio management, and integration with 4G/5G. It applies AI to network management and IoT security so that customers gain a better experience while spending less on maintenance.

IDL is headquartered in Silicon Valley, with R&D in Taiwan.

Related solutions

- **SD-WiFi** — carrier Wi-Fi offload inside the 3GPP core
- **SD-Edge** — vCPE and uCPE: a fully virtualized, software-defined carrier edge
- **SD-Access** — a zero-trust campus fabric without replacing your switches
- **SD-Branch** — networking, security and SD-WAN converged into one gateway
- **MDU** — one SSID for the whole property, a private network for every unit
- **AIOps** — from reactive monitoring to a network that explains itself
- **Security** — secure gateway and DPI engine

> Learn more at **ignitiondl.com/solutions**