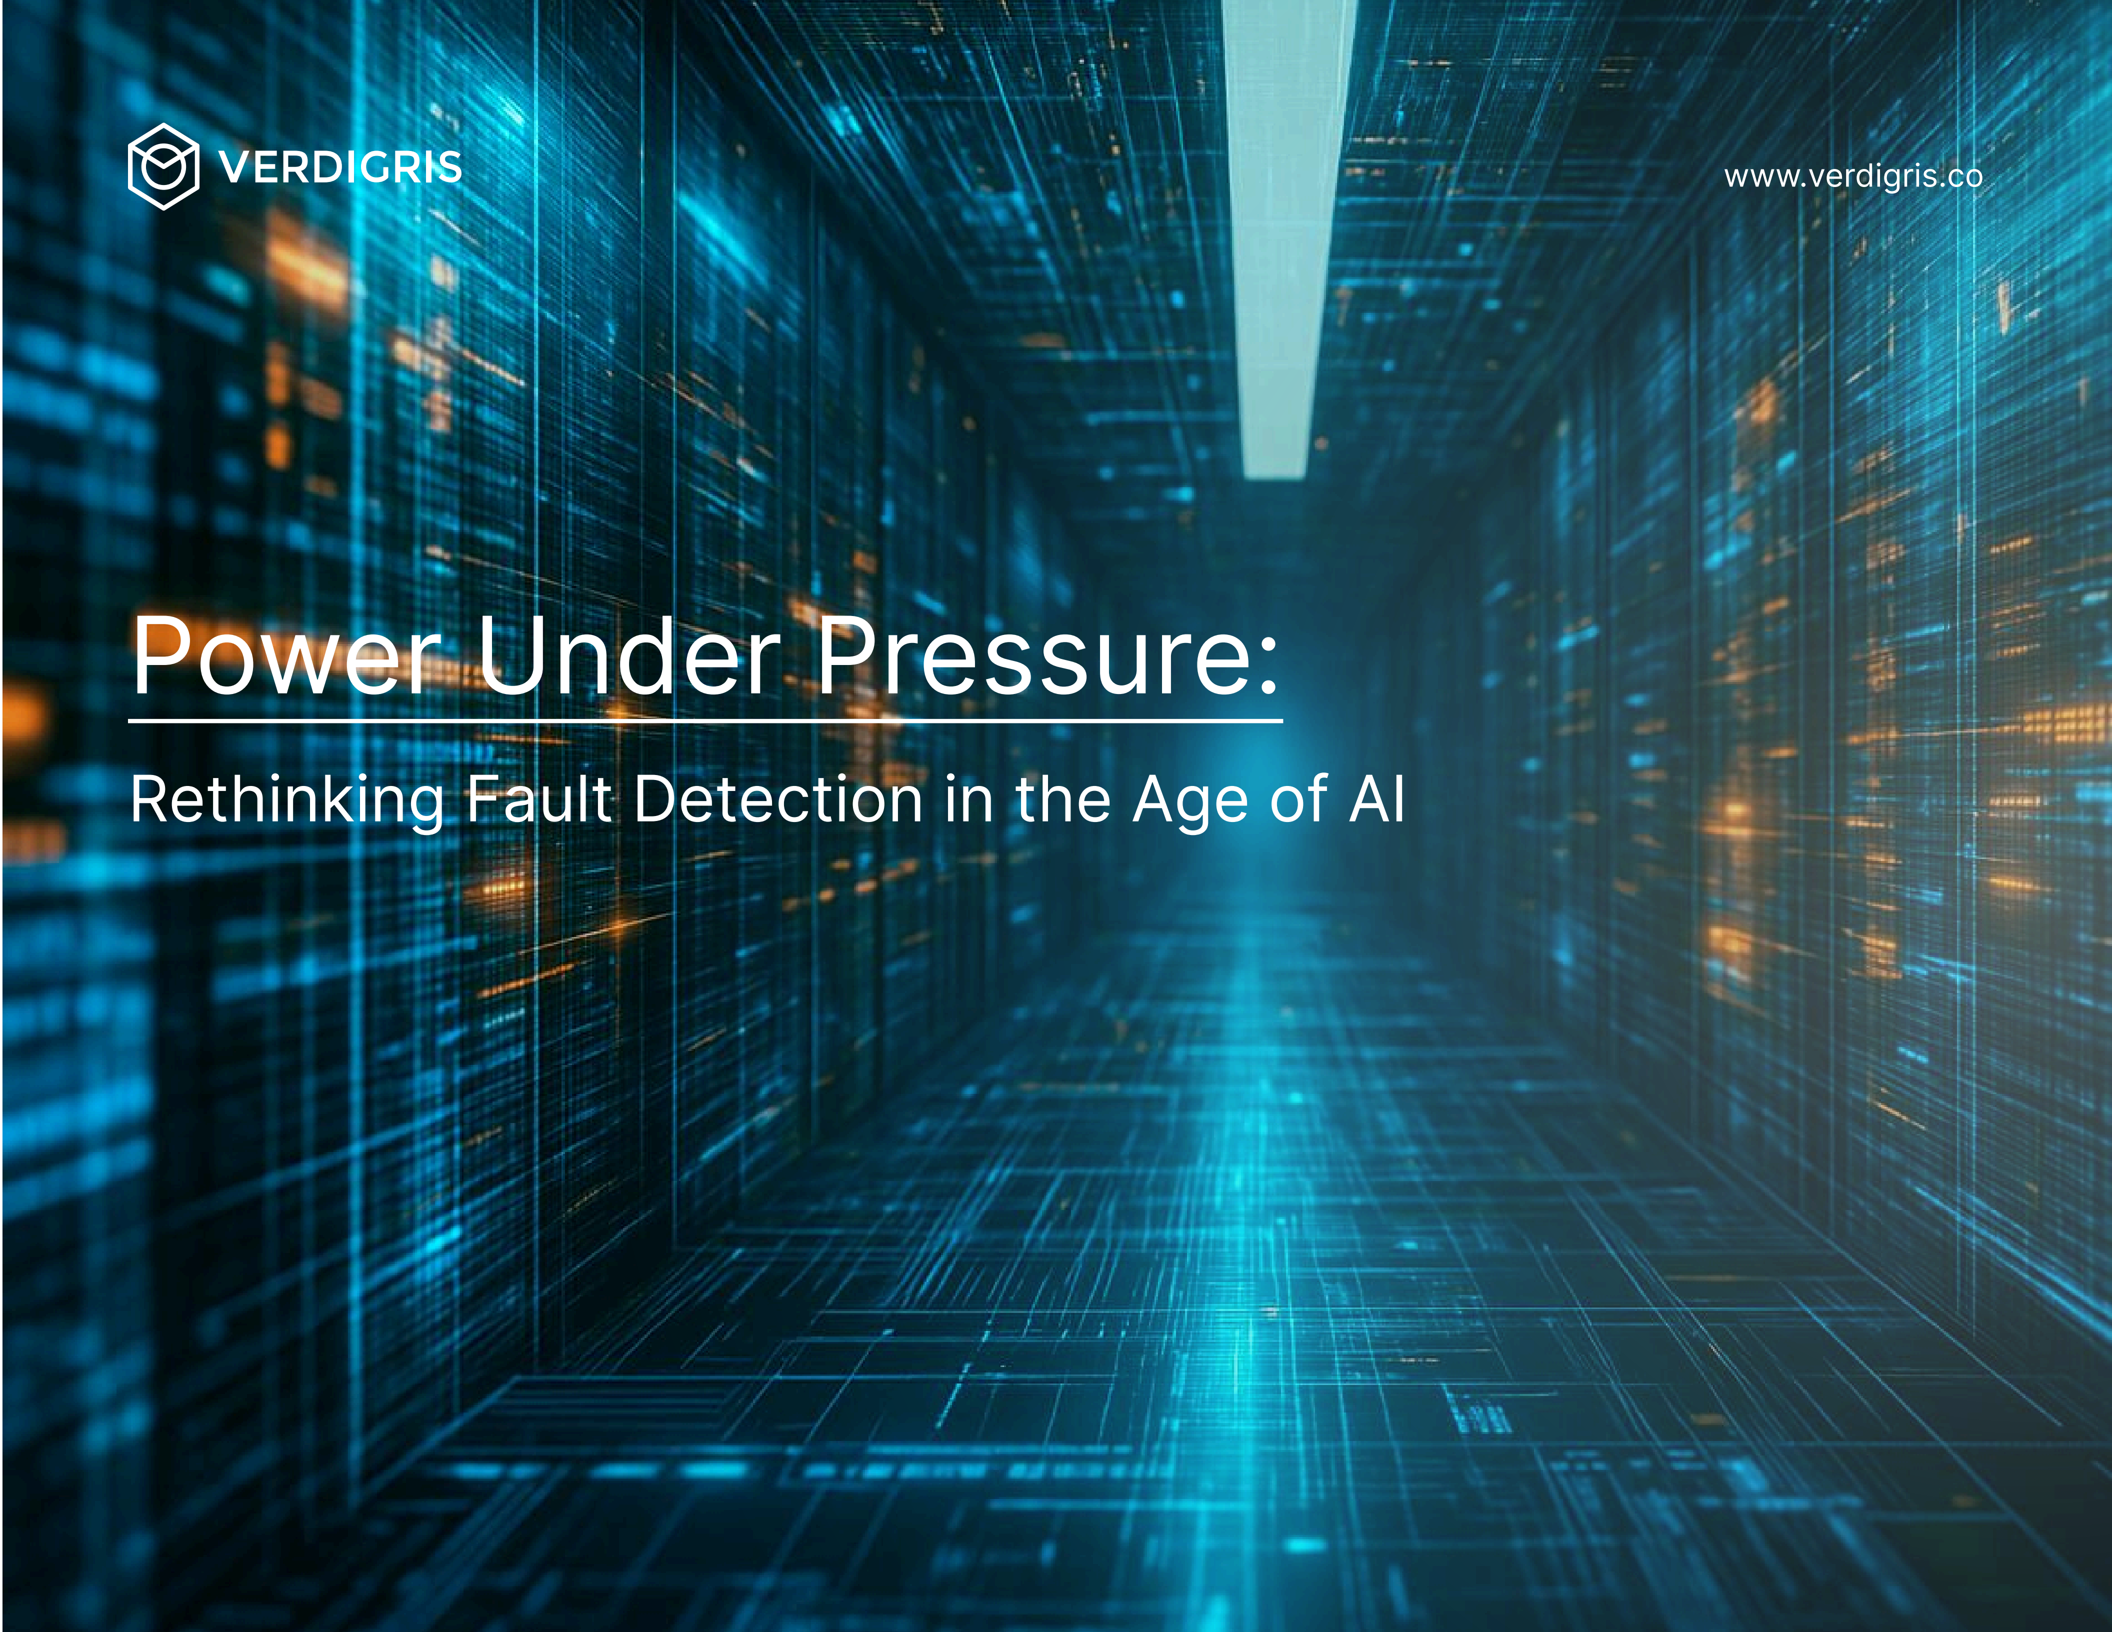

VERDIGRIS

www.verdigris.co

Power Under Pressure:

Rethinking Fault Detection in the Age of AI

Executive Summary

The rapid acceleration of AI workloads is placing unprecedented strain on data center infrastructure—especially among emerging neocloud operators built specifically for AI and high-density workloads, where electrical volatility is a first-order design challenge. Traditional fault detection and diagnostic (FDD) systems, built for an earlier era of predictable compute and static workloads, are falling short in this new reality.

This white paper explores the forces reshaping the operational demands of data centers, how legacy fault detection tools are being outpaced, and what facility and operations leaders must consider to future-proof their infrastructure. Along the way, we share a practical readiness checklist and call attention to industry trends that suggest a fundamental shift in how we manage electrical risk.

In this guide, you'll discover:

✅ How AI workloads are changing the behavior of electrical infrastructure

✅ Why traditional fault detection systems fail in AI-driven environments

✅ What to look for in modern, AI-ready fault detection solutions

The AI Workload Effect on Electrical Infrastructure

The emergence of generative AI, large-scale model training, and real-time inference has created a new profile of power consumption: dynamic, bursty, and dense. These workloads break many of the assumptions underlying traditional data center design and operations.

- **Dynamic Power Consumption:** AI training jobs can create rapid power draw fluctuations, with spike durations under 30 seconds.
- **High-Density Deployment:** Some AI clusters consume 3–5x more power per rack compared to traditional server infrastructure.
- **Harmonic Distortion:** Power electronics in GPUs introduce high-frequency harmonics that degrade equipment and elevate thermal risk.
- **Coordination Complexity:** Fast power shifts now impact cooling, power redundancy, and mechanical infrastructure.

For neocloud operators—AI-native providers architected around GPU-intensive compute—the challenge is magnified. Power densities exceeding 200 kW per rack and rapid provisioning cycles mean even brief transient faults can trigger cascading thermal and electrical risks.

These changes demand more than monitoring—they require a **new model of electrical intelligence.**

The Limitations of Traditional Fault Detection Tools

Most fault detection systems in use today were designed for a different time: one where server loads were steady and equipment behavior was relatively predictable. As a result, many tools fail to deliver the early warning or diagnostic clarity needed for today's operating environment.

Common Limitations:

- **Slow Sampling:** 1-minute data resolution misses fast-transient faults.
- **Shallow Visibility:** PDU-level insights are insufficient for detecting circuit-level issues.
- **Generic Alarms:** Alerts lack cause, context, and correlation—often resulting in alarm fatigue.
- **Limited Historical Data:** Without detailed historical logs, trend and root-cause analysis is compromised.
- **Reactive Operation:** Most tools alert after the fault, not before.

For neocloud operators, legacy EPMS and BMS platforms often lack the high-frequency sampling and visibility required for modern AI loads, forcing a patchwork of tools that still miss transient signals.

💡 **Did You Know?**

Short-duration power disturbances—measured in milliseconds to seconds—are both common and operationally disruptive for sensitive, high-density loads. But because legacy monitoring and surveys rarely break out sub-second incidents, their contribution to downtime is likely under-reported.

Rethinking Fault Detection in the Age of AI

Imagine a neocloud operator scaling GPU clusters across multiple regions — as workloads spike unpredictably, sub-minute faults can cascade before traditional systems even register an anomaly.

Facility leaders must evolve their thinking from "monitor and alert" to "predict and prevent." This means reassessing how fault detection tools serve uptime, capacity planning, and safety in AI-era data centers.

Key questions to ask:

- **Granularity:** Are we monitoring at a resolution that matches fault behavior?
- **Context:** Do alerts help us understand and resolve the issue?
- **Correlation:** Can we connect events across circuits and systems?
- **Visibility:** Do we understand what's happening at the edge of distribution, not just the center?
- **Adaptability:** Can our system evolve as workloads and infrastructure evolve?

💡 **Analyst Insight**

Gartner expects 40% of AI data centers to be constrained by power availability by 2027, underscoring why electrical reliability is now a primary risk vector.

Four Characteristics of Next-Generation Fault Detection

Diagnose, Not Just Detect

Traditional FDD
Most systems rely on threshold-based alarms that indicate something is wrong—high load, unusual voltage, excessive harmonics—but provide little information about why it happened or how to fix it. The result is often alarm fatigue, delayed investigation, and finger-pointing between teams.

Modern Fault Detection
Diagnostic systems correlate signals across circuits, phases, and time to suggest likely root causes. They prioritize actionable signals and reduce noise, empowering operators to move quickly and confidently during critical moments.

Why It Matters
For neocloud operators managing globally distributed AI clusters, automated root-cause clarity shortens time-to-resolution across fleets where every minute of downtime can idle millions in compute.

Predict, Not Just Monitor

Traditional FDD
These systems only react to violations once they've occurred—like breaker trips or sustained overcurrent. They offer no mechanism to identify patterns or conditions that commonly precede faults.

Modern Fault Detection
Predictive systems analyze continuous data streams to identify early warning signs—load drift, voltage imbalance trends, or recurring harmonics—well before failure occurs. They enable teams to schedule preventive maintenance or adjust capacity before incidents happen.

Why It Matters
As AI workloads push electrical systems harder, proactive fault avoidance becomes more valuable than rapid response. Predictive insights help operators shift from firefighting to foresight.

Four Characteristics of Next-Generation Fault Detection *(continued)*

Integrate Broadly

Traditional FDD
Data often lives in isolated systems—DCIM, BMS, or EMS tools that don't share context with planning, energy, or sustainability platforms. This limits the usefulness of fault data beyond the control room.

Modern Fault Detection
Fault detection becomes part of a broader electrical intelligence layer, integrating with capacity planning, sustainability reporting, and digital twin models. This allows operators to use fault history to inform design upgrades, procurement, and resource allocation.

Why It Matters
For AI-native and neocloud environments, integration ensures electrical intelligence scales seamlessly from rack to region, unifying site operations and global reliability analytics.

Learn Over Time

Traditional FDD
Events are isolated, and past faults are rarely analyzed unless there's a major outage. Static rules govern detection, with no ability to revisit historical data as new patterns emerge.

Modern Fault Detection
These systems operate like a historian—storing detailed, high-resolution data and enabling retrospective analysis. As new fault conditions are discovered, they can be applied backward across months of data to identify hidden issues and recurring risks.

Why It Matters
With AI, electrical risks are no longer static. The ability to learn and reanalyze past behavior ensures continuous adaptation and greater resilience in the face of evolving workloads.

Together, these capabilities define **the foundation of AI-ready electrical intelligence.** They reduce downtime, uncover hidden risks, inform smarter planning, and empower teams to operate with greater confidence in a landscape where power dynamics are constantly shifting.

Future-Proofing Your Data Center Starts with Smarter Fault Detection

The AI era is rewriting the rules of reliability, pushing electrical systems to adapt in real-time. Traditional fault detection tools, while useful in static environments, are no longer sufficient. Operators who succeed will be those who evolve their infrastructure monitoring to be faster, deeper, and more adaptive.

To begin this shift, organizations must evaluate their fault detection systems not just on uptime history, but on future-readiness.

Is Your Fault Detection Solution Ready for AI Workloads?
10 Questions for Data Center Ops Leaders

The AI era is rewriting the rules of reliability, pushing electrical systems to adapt in real time. Traditional fault detection tools, while useful in static environments, are no longer sufficient. Operators who succeed will be those who evolve their infrastructure monitoring to be faster, deeper, and more adaptive.

This 10-question checklist will help you assess whether your current fault detection setup is truly ready for AI infrastructure—or leaving you vulnerable to power faults, missed capacity, and potential SLA penalties.

✅ **Score yourself honestly. If you're answering "No" to more than 3, it's time to re-evaluate your tools.**

For each question, select: Yes | No | Not Sure

1. **Can your system detect faults that occur in less than 60 seconds?**

   ⏱️ AI power spikes are often brief but critical. Minute-level data is no longer enough.

---

2. **Are you monitoring at the circuit level (not just PDU or panel)?**

   🔍 Faults often originate deep in the distribution tree. Granularity matters.

---

3. **Does your fault detection solution surface voltage imbalance trends over time?**

   ⚡ Imbalances are a leading cause of transformer failure and go unnoticed without phase-level tracking.

---

4. **Can it identify harmonic distortion from AI workloads?**

   🎛️ GPU clusters introduce harmonics that most systems ignore—until damage occurs.

---

5. **Does your system offer root cause insights, not just alarms?**

   🧠 Alerts are useless without actionable context. Are your teams still guessing?

Is Your Fault Detection Solution Ready for AI Workloads?

10 Questions for Data Center Ops Leaders *(continued)*

6. **Can it diagnose nuisance tripping and prevent recurrence?**

   🛑 False trips hurt uptime. You need analytics that connect cause and effect.

---

7. **Are you able to run retrospective fault detection on historical data?**

   📈 When new fault types emerge, can you scan the past to surface risks?

---

8. **Does your system correlate power anomalies across circuits?**

   🔗 Complex faults rarely happen in isolation—are you seeing the big picture?

---

9. **Can it help you plan capacity safely for high-density AI clusters?**

   📊 You need more than thresholds—you need insight into what your infrastructure can handle.

---

10. **Are power faults still showing up as surprises?**

    💥 If you're still saying "we didn't see that coming," your fault detection solution isn't doing its job.

Tally your Yes answers:

- 8–10 Yes → You're in great shape. Your fault detection solution is AI-ready.
- 5–7 Yes → You're holding steady, but may be exposed to new fault types.
- 0–4 Yes → Your system likely isn't keeping up with AI demands. You're operating with risk.

About Verdigris

VERDIGRIS

Founded in Silicon Valley in 2011 by former semiconductor design engineers, Verdigris set out to solve energy challenges that traditional systems couldn't. Today, Verdigris powers the energy intelligence of mission-critical data centers for Fortune 500 leaders, delivering precision and efficiency to meet growing power management demands.

Verdigris' hybrid edge-cloud architecture combines precision sensing with AI-powered diagnostics, enabling early warning systems that uncover electrical risks before they become outages. By delivering real-time insight into electrical behavior, Verdigris helps data center operators unlock stranded capacity, avoid downtime, and make smarter, faster infrastructure decisions.

Backed by visionary investors like DCVC, Stanford StartX, and Verizon Ventures, Verdigris has earned recognition as a pioneer in energy intelligence. Named one of Fast Company's Top 10 Most Innovative Companies in Energy and honored as Frost & Sullivan's Technology Innovation Leader in North American AI-based Energy Management and Automation, Verdigris continues to shape the future of energy intelligence.

Power Fault Detection for the AI Era

Verdigris' approach to fault detection helps data centers prevent downtime, ensure reliability, and focus on the right issues at the right time. Advantages include:

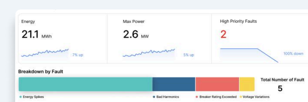

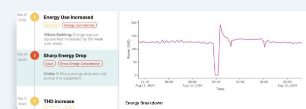

✅ **Granular & High-Fidelity Data** – Sub-secondly, circuit-level monitoring captures what others miss.
✅ **Retrospective Analysis** – Apply new fault detection algorithms to historical data, surfacing hidden risks.
✅ **AI-Ready Insights** – Find patterns, pinpoint root causes, and recommend fixes—not just alerts.

VERDIGRIS

A truly modern fault detection solution should offer:

✅ Deliver circuit-level visibility and intelligent context to focus on the issues that impact reliability.
✅ Enable retrospective fault analysis to uncover hidden risks in historical data.
✅ Go beyond alerts to provide root causes and proactive guidance that prevent downtime.

Ready to rethink your approach to fault detection?

See how Verdigris helps AI-native and neocloud operators achieve predictive reliability at scale.

🔗 **Schedule a Free Consultation**