

The WildRaven Technical Reality Assessment™

A forensic, seven-day diagnostic that tells you what is real, what is smoke, and what it will actually cost — before you deploy capital or bet a roadmap on it.

Large technology decisions fail when marketing claims outpace technical reality. The Reality Assessment is a “bullshit detector” run by one senior architect — not a team of junior analysts. In seven days you get a Red / Yellow / Green verdict, the evidence behind it, and a specific plan. Not a deck full of frameworks.

What This Solves

A vendor demo looks flawless. A pitch deck promises 100x scale. An internal team swears the AI “just needs a few more weeks.” None of it survives contact with production — and by the time you find out, the check has cleared or the quarter is gone. The Reality Assessment closes the gap between what you have been *told* and what actually exists in the code, the architecture, and the vendor's track record.

Who Engages WildRaven

- ✓ Investors and acquirers validating a technology asset before a Series B, growth round, or M&A close
- ✓ Operators about to sign a mid-six or seven-figure AI, media, or complex software contract
- ✓ Boards and CEOs watching an AI initiative burn cash with nothing shipped
- ✓ CTOs who inherited a system they cannot yet trust and need an independent second set of eyes

The Four-Pillar Diagnostic Framework

Every engagement is scored across four pillars. Each pillar is where real technology bets are quietly won or lost.

01 · Technology Reality Check

Does the product actually function as described — or is it a “**manual-in-the-middle**” demo, a fragile wrapper around someone else's API, or slideware with a login screen? I separate proprietary IP from rented third-party dependencies and establish what the company genuinely owns versus what it merely resells.

02 · Architectural Scalability

Can the current infrastructure survive 10x or 100x growth, or does it fall over at the first real load? I hunt for hidden bottlenecks, single points of failure, and the technical debt that turns a “quick integration” into an 18-month re-architecture.

03 · AI Edge-Case Reliability

Where AI is involved, I test behavioral consistency and prompt fragility, and I look for the thing most teams skip entirely: a real evaluation and fallback harness. An AI feature that works in the demo and hallucinates in production is not a feature — it is a liability with a launch date.

04 · Vendor & Implementation Truth

I forensically audit vendor roadmaps and past performance to expose **the Gap** — the distance between what was promised in the proposal and what was actually delivered. Recurring slippage, moving goalposts, and “proprietary” hand-waving all leave fingerprints.

The Red / Yellow / Green Verdict

Every pillar and the engagement as a whole receives a plain-language rating so a board can act in minutes:

- **GREEN** Sound. Proceed — the technology substantiates the claims.
- **YELLOW** Fixable with defined risk. Proceed only with specific conditions, remediation, or a renegotiated price.
- **RED** Material misrepresentation or structural failure. Do not proceed as-is — renegotiate, restructure, or walk.

The Seven-Day Process

01

Discovery Call · 15 min, complimentary

We confirm the situation, determine fit, and assign the right tier (Standard or Complex). No pitch.

02

Deep Technical Review · Days 1–5

Access to architecture docs, codebase samples, vendor contracts, and key people. I personally review everything — no junior analysts — and map the system to your business and money model.

03

Written Assessment · Days 5–6

Current-state analysis, critical risks, time and cost to fix, and a 90-day de-risking outline with a clear

04

Executive Briefing · Day 7

A 60-minute walkthrough with Q&A and a decision framework built for your deal team and operators.

What You Receive

- ✓ **Forensic Assessment Matrix** — a prioritized register of technical risks and vulnerabilities, each rated Red / Yellow / Green
- ✓ **Systems Recovery Roadmap** — if the project is failing, a precise, sequenced list of corrective actions to get it to production-grade reliability
- ✓ **Truth-Verification Summary** — hard findings on system stability, scalability, latency, and operational viability
- ✓ **Executive Summary** — a one-page decision memo written for boards and partners, not engineers
- ✓ **90-day de-risking outline** — what to fix first, what it costs, and what “good” looks like

Scope of a Standard Engagement

To keep the seven-day timeline honest, a Standard engagement covers one primary product or platform, up to five core services or components, up to three code repositories, and up to three stakeholder interviews (60 minutes each). Conducted remotely. Larger, multi-system, or multi-entity targets are scoped as Complex.

Why It Works

One senior mind. The person reviewing your system designed and rescued systems for Stanford, Cisco, ESPN, NASA, and the SEC — personally, start to finish. No hand-offs to associates.

Forensic, not diplomatic. My only job is to protect you. If the honest answer is “kill it,” you will hear it in plain English before your next invoice.

Actionable. You leave with a specific plan and real numbers — not an abstract strategy deck.

Investment

From \$7,500 for a Standard assessment. **From \$12,500** for Complex, multi-system engagements. Scoped precisely on the discovery call.

The guarantee: if the assessment is not worth what you paid, you get a full refund of your fee. I can offer that because in twenty-five years the answer has never been in doubt.

Ready to Get the Truth About Your Technology?

Every engagement starts with a complimentary 15-minute discovery call — no pitch, no obligation.

keenan@iinc.tv · 310.570.0774

