

Over the past decade, I've seen numerous B2B SaaS tools position their pricing in ways that create more questions than clarity. Now, with the rapid rise of AI visibility and large language model (LLM) observability solutions, pricing practices have become even more opaque. One term I keep hearing in the context of Fiddler's AI observability platform is the infamous "**quote-based pricing**." But what does that actually mean in practice? And how does it compare with more transparent tiered pricing models, such as those offered by newer players like Peec AI?

This blog post will dissect Fiddler's pricing model, demystify the concept of quote-based pricing, and place it in context with industry-standard approaches. Along the way, we will also highlight key features like AI search visibility vs classic SEO, prompt-level measurement and tracking, multi-LLM coverage and assistant benchmarking, and advanced metrics including share-of-voice, sentiment analysis, and citation tracking.

Understanding Fiddler's Pricing & The "Quote-Based" Model

First, let's unpack the pricing tiers typically referenced with Fiddler:

- **Fiddler Lite:** Entry-level access with basic AI explainability and monitoring features.
- **Fiddler Standard:** Adds advanced governance, bias detection, and model performance tracking.
- **Enterprise (Quote-Based):** Fully customizable with integration, scale, security, and support options.

You'll notice the Enterprise tier is not listed with a fixed price. Instead, Fiddler relies on what vendors call *quote-based* or *custom pricing*. This means, rather than publishing a straightforward rate card, pricing is tailored on a per-customer basis after consultation — with factors like model volume, user seats, data retention, and integration complexity influencing cost.

What Quote-Based Pricing Actually Means for Buyers

The appeal of quote-based pricing for vendors is obvious: It allows flexibility to serve everything from startups to large financial institutions, each with very different needs. However, from a buyer's perspective, there are some practical and measurable concerns worth highlighting:

1. **Lack of Price Transparency:** Without published prices or unit limits, it's impossible to benchmark Fiddler's cost-effectiveness before engagement.
2. **Hidden Scaling Constraints:** Many vendors include fine print limiting model queries, user seats, or API calls, but these are harder to spot with custom quotes.
3. **Potential Vendor Lock-In:** Complex contracts can disincentivize switching or scaling outside the initially quoted scope.
4. **Opaque ROI Assessment:** Without clear pricing tiers, measuring ROI and adopting measurable KPIs can be tricky.

So when you hear "quote-based pricing" from a vendor like Fiddler, ask yourself: "*What exactly breaks at scale?*" For example, is there a user-seat limit? How many LLMs or queries per month is included? And what data retention periods dictate cost?

Peec AI Pricing: A Transparent Alternative Benchmark

To put things in perspective, consider Peec AI, a newer entrant into the AI visibility and tracking space. Their pricing is tiered and clearly published:

Plan Monthly Cost (€) Key Features Starter €89 Basic prompt tracking, sentiment analysis, dashboard access Pro €199 Multi-LLM benchmarking, assistant monitoring, extended reports Enterprise Custom quote Custom integrations, SLAs, advanced governance, API access

Unlike Fiddler's quote-based model where all pricing for large customers is custom, Peec AI offers concrete, predictable pricing tiers addressing the needs of small and mid-market teams. This is especially valuable for organizations wanting better budget predictability and a clear path to scale.

Why AI Search Visibility Differs From Classic SEO

A common point of confusion when evaluating AI observability solutions is how "AI search visibility" differs from classic SEO metrics. Classic SEO focuses on optimizing website pages against keywords, backlinks, and content relevancy to rank well in search engine results pages (SERPs).

In contrast, AI search visibility measures how well AI assistants, chatbots, or language models surface relevant results or answers based on prompts—an inherently different mechanism. It involves:

- Measuring prompt success rates and completion quality.
- Tracking assistant response coverage against business KPIs.
- Monitoring model drift and bias in real-time or near real-time (something classic SEO does not address).


Tools like Fiddler and Peec AI emphasize this prompt-level measurement and tracking, which gives teams granular visibility into what queries work, which need refinement, and where AI assistants might be underperforming or biased.

Prompt-Level Measurement and Tracking: The Core of AI Observability

Unlike web analytics or classic metrics, prompt-level observability means measuring individual model inputs and outputs at scale. This includes:

1. **Prompt Attribution:** Understanding what prompts generate specific outputs, enabling targeted improvements.
2. **Performance Tracking:** Measuring precision, recall, sentiment, and compliance at the prompt-level.
3. **User Behavior Correlation:** Linking prompts with user actions or feedback for outcome-driven updates.

Fiddler Standard and comparable tiers aim to provide these capabilities with dashboards that track prompt performance, bias detection, and alerting. However, scaling to thousands or millions of prompts requires clear pricing that includes query volume caps, retention windows, and user controls—all often buried in quote-based pricing models.



Multi-LLM Coverage and Assistant Benchmarking

The modern AI stack rarely runs on a single LLM. Enterprises often integrate multiple language models—OpenAI’s GPT, Anthropic, Cohere, or proprietary models—to drive different applications. Observability platforms must therefore:

- Support multi-LLM monitoring with unified dashboards.
- Benchmark assistants’ performance against each other by domain, purpose, or user segments.
- Enable A/B testing and longitudinal analysis of model drift and quality.

Fiddler claims multi-LLM coverage with enterprise-grade scalability, but these features are often gated behind the enterprise quote. Peec AI, by contrast, begins multi-LLM benchmarking at the Pro tier (€199/month), making it measurable and accessible for mid-market teams.

Share-of-Voice, Sentiment, and Citation Tracking in AI Governance

Beyond monitoring raw prompt inputs and outputs, AI observability platforms are increasingly incorporating “share-of-voice” and sentiment analysis. These metrics come from marketing and PR analytics but adapted to AI assistants:

- **Share-of-Voice:** How often does an assistant or LLM generate a particular answer versus competitors or alternative models?
- **Sentiment Tracking:** Does the AI’s content skew positive, negative, or neutral in tone? Is there any harmful bias present?
- **Citation Tracking:** Can the assistant explain or reference sources where answers are derived, crucial for compliance and trust?

These advanced visibility features are integral to AI governance frameworks but are often summed up in marketing fluff as “AI governance support.” Realistically, implementations vary widely and are almost always confined to higher-priced, custom enterprise tiers like Fiddler’s “quote-based” offering or Peec AI’s enterprise plan.

What Breaks at Scale?

As a 10-year B2B SaaS analyst who’s reviewed countless martech and AI observability tools, I always ask: “*What breaks at scale?*” For platforms like Fiddler, the key scaling challenges typically include:

1. **Data Volume Limits:** How many prompt events and model responses can be ingested and retained before additional fees apply?
2. **User & Role Management:** Does the system support granular access controls and audit logs for compliance?
3. **Export & API Access:** Can teams export raw data or integrate observability data into broader analytics stacks?
4. **Latency & Refresh Cadence:** Are monitoring dashboards truly real-time, or is there a 15-60 minute lag?
5. **Customization & Integration:** Does the pricing include connectors to data lakes, MLOps tools, or governance platforms?

In many cases, these capabilities are only fully unlocked via a quote-based enterprise engagement — meaning you lack firm pricing visibility until committing significant resources.

Summary Table: Comparing Fiddler and Peec AI Pricing Models and Features

Feature / Pricing Aspect	Fiddler Lite	Fiddler Standard	Fiddler Enterprise (Quote-Based)	Peec AI Starter (€89/mo)	Peec AI Pro (€199/mo)	Peec AI Enterprise (Quote)
Pricing Transparency	Published	Published	Published	Custom quoting, opaque	Published	Published
Prompt-Level Tracking	Basic	Advanced	Full, customizable	Basic	Advanced	Full, customizable
Multi-LLM Support	Limited	Yes	Yes, with SLAs	No	Yes	Yes
Share-of-Voice & Sentiment	No	Some	Comprehensive	Basic sentiment	Yes	Full suite
Export & API Access	Limited	Limited	Full	No	Yes	Full
Access Controls & Governance	Minimal	Enhanced	Enterprise-grade	Minimal	Improved	Enterprise-grade

Final Thoughts

When evaluating AI visibility platforms like Fiddler, especially at enterprise scale, understanding what “quote-based pricing” actually entails is critical. Unfortunately, many vendors use this as a black box, leaving buyers unclear on measurable limits and true total cost of ownership.

Transparent pricing, with published tiers and documented usage caps, can significantly improve [generative engine optimization](#) vendor comparisons, budgeting, and feature adoption planning. Newer players like Peec AI highlight the benefits to buyers and marketplaces by offering clear starter and pro tiers before moving into custom enterprise engagements.

Ultimately, what breaks at scale are almost always the granular limits on prompt volume, user seats, governance features, and integration flexibility — all of which should be explicitly quantified rather than hidden behind buzzwords and custom quotes. Demand clarity on these points ahead of any sales engagement to make informed, measurable decisions about your AI observability investments.