




NLPatent Rebrands to Clerq with Agentic AI That Compresses
Patent Research from Weeks to 10 Minutes



NLPatent Rebrands to Clerq with Agentic AI That Compresses Patent Research from Weeks to 10 Minutes

Patent research that took your junior associates three weeks now runs in 10 minutes. The billable hour just got a lot more expensive to justify.

The News: Clerq Launches End-to-End Patent Automation

Toronto-based NLPatent [announced its rebrand to Clerq on August 18, 2026](#), alongside the launch of an agentic AI platform designed to execute patent research workflows from start to finish. The company claims its system compresses work that traditionally spans days or weeks into approximately ten minutes.

Two workflows shipped on day one: a rapid triage assessment for screening invention disclosures, and comprehensive patentability reports with feature-by-feature reasoning. The company has also announced expansion into invalidity and freedom-to-operate analyses.



NLPatent Rebrands to Clerq with Agentic AI That Compresses Patent Research from Weeks to 10 Minutes

The launch comes with strategic partnerships. [Clerq announced cross-referral integrations with RPX Corporation and Park IP](#), signaling a go-to-market strategy that positions the platform as infrastructure for existing IP service providers rather than a direct competitor. The company also hired Michael Chernoff as Director of IP Strategy to lead enterprise sales and product direction.

Clerq previously raised \$3 million in funding led by Draper Associates and Mighty Capital. The rebrand represents the company's transition from a patent search tool to a platform that executes complete research workflows autonomously.

Why This Matters: The Economics of Legal Work Just Shifted

Patent research sits at a peculiar intersection of high-value and high-tedium work. A typical patentability search at a major firm involves a junior associate spending 15-30 hours combing through prior art databases, reading dozens of related patents, and synthesizing findings into a memo. At associate billing rates of \$300-500 per hour, a single patentability opinion costs clients \$5,000-15,000 before a senior partner even reviews it.

Clerq's claim of ten-minute execution attacks the economics directly. If the output quality holds, the value proposition isn't "slightly cheaper patent research"—it's "patent research at a fundamentally different price point."

The immediate winners are obvious: in-house IP teams at technology companies who file dozens or hundreds of patents annually, and patent portfolio managers who need to evaluate large acquisition targets. A company evaluating a 500-patent portfolio for acquisition due diligence could theoretically run preliminary assessments in hours instead of weeks.

The losers are less obvious but more consequential. Junior associates at IP boutiques currently build expertise by doing exactly this kind of work. The traditional model uses patent research as training ground—associates learn the craft by doing the tedious work. If AI handles the tedious work, where do junior attorneys develop judgment?

The question isn't whether AI can do patent research faster. It's whether



the output is good enough that clients stop paying humans to do it at all.

Law firms face a strategic choice: adopt automation and compress billing, or watch clients bring the work in-house. [As SiliconANGLE reported](#), Clerq positions itself as enabling firms to “take on more work with existing teams”—but that framing obscures the more disruptive reality that clients can now bypass firms entirely for preliminary research.

Technical Architecture: How Agentic Patent Research Works

The term “agentic AI” has become overused in enterprise software marketing, but Clerq’s implementation reveals what the architecture actually means in practice.

Traditional patent search tools operate as retrieval systems. You input keywords, classifications, or patent numbers; the system returns matching documents. Semantic search improved this by understanding meaning rather than just keywords. But both approaches still require a human to synthesize results, identify relevant prior art, and write the analysis.

Clerq’s agentic approach chains multiple AI operations into a workflow that mirrors how a human researcher would work. The system:

- Parses an invention disclosure to extract novel claims and technical features
- Generates search strategies across multiple patent databases
- Retrieves and ranks relevant prior art
- Analyzes each reference against specific claim features
- Synthesizes findings into structured reports with feature-by-feature reasoning

The key architectural difference is autonomous decision-making during execution. When the system encounters ambiguous terminology in a disclosure, it generates interpretive alternatives and searches across all of them. When a promising reference appears, it follows citation chains to find related art. These aren’t pre-programmed rules—they’re emergent behaviors from a system designed to complete a goal rather than execute a fixed sequence.



The Reasoning Layer

Patent analysis requires more than retrieval—it requires legal reasoning about claim scope, obviousness, and anticipation. Clerq’s patentability reports include “feature-by-feature reasoning,” which suggests the system maps individual claim elements against prior art references and generates explanations for why each reference does or doesn’t anticipate specific features.

This is where the architecture gets interesting. Patent law has developed precise language around concepts like “teaching away” (when prior art suggests the opposite of an invention), “unexpected results” (when a combination produces surprising outcomes), and “motivation to combine” (when multiple references together suggest an invention). A competent patent analysis must engage with these doctrinal concepts.

The quality bar for agentic patent research isn’t “does it find prior art?”—it’s “does it reason about prior art the way a patent examiner would?”

Without access to Clerq’s actual outputs, it’s impossible to evaluate whether the system achieves this standard. The company’s claim of “feature-by-feature reasoning” suggests architectural awareness of this requirement, but the proof will emerge from real-world usage in prosecution and litigation contexts.

Where Agentic Systems Break

Understanding the failure modes of agentic patent research helps technical leaders evaluate the technology realistically.

The foreign-language problem: Significant prior art exists in non-English patents, particularly from China, Japan, Germany, and Korea. Machine translation of technical patent language introduces errors that compound through analysis. A system might miss relevant prior art simply because translation failed to capture the technical nuance.

The non-patent literature problem: Many patentability challenges arise from academic papers, technical standards, product manuals, and conference proceedings. These sources are harder to search comprehensively than patent databases, and coverage varies dramatically by technical field.



The claim construction problem: Patent claims use language in ways that diverge from ordinary meaning. “Comprising” has a different legal meaning than “consisting of.” A term might have an explicit definition in the specification that contradicts common usage. Agentic systems must navigate this legal-linguistic complexity.

The adversarial problem: Patent litigation is adversarial. The prior art that matters most is often the art that’s hardest to find—art that uses different terminology, exists in obscure databases, or requires creative interpretation to apply. Human researchers develop intuition for finding this art through experience. It’s unclear whether current AI architectures can replicate this adversarial creativity.

The Contrarian Take: What’s Overhyped and Underhyped

What’s Overhyped: The Ten-Minute Claim

Let’s be direct about the marketing: “ten minutes” describes system execution time, not total workflow time. A patent attorney still needs to review the disclosure, verify the search scope, evaluate the AI’s reasoning, and sign off on the analysis. The actual workflow might compress from three weeks to three hours, which is still transformative—but the headline number obscures the reality that human judgment remains in the loop.

More importantly, the ten-minute figure applies to initial research, not the iterative process that happens during actual prosecution. When a patent examiner rejects claims based on prior art the AI missed, someone still needs to reformulate claims, distinguish the art, or argue around it. The research phase is only one component of patent work.

What’s Underhyped: The Data Moat Potential

The underhyped story is the feedback loop. Every time Clerq’s system runs a search and an attorney modifies the results—accepting some prior art as relevant, dismissing other references as inapplicable—the platform accumulates training signal about what constitutes quality patent analysis in specific technical domains.

If Clerq processes thousands of patentability searches across semiconductor design,



antibody therapeutics, and machine learning systems, it accumulates domain-specific understanding that new entrants would struggle to replicate. The company with the most workflow data trains the best models, which attracts more users, which generates more data.

The strategic asset isn't the AI—it's the corpus of attorney-validated reasoning across technical domains.

RPX Corporation manages patent risk for over 500 member companies and controls one of the largest patent databases in the industry. Park IP specializes in patent portfolio analysis and technology transfer. These partnerships give Clerq access to transaction and analysis data that would take years to accumulate organically.

What Everyone Misses: The Prosecution History Problem

Patent analysis doesn't happen in a vacuum. Understanding how a patent was prosecuted—what arguments the applicant made to the patent office, what amendments narrowed the claims, what prior art was considered and distinguished—often matters more than the face of the patent itself.

Prosecution history estoppel prevents patent holders from claiming scope they gave up during prosecution. A prior art search that ignores prosecution history misses this context entirely.

[Law360's coverage noted](#) that Clerq plans expansion into invalidity and freedom-to-operate workflows. These applications require deeper engagement with prosecution history than basic patentability searches. The technical challenge of incorporating this context into agentic workflows is substantial and will determine whether Clerq can move beyond initial research into higher-value analysis.

Practical Implications: What Should Technical Leaders Do?

For Technology Companies with IP Portfolios

If your company files more than 20 patents annually, the calculus has changed.



NLPatent Rebrands to Clerq with Agentic AI That Compresses Patent Research from Weeks to 10 Minutes

Evaluate your current patent research costs—not just the external counsel bills, but the internal time spent on disclosure review, prior art review, and coordination with outside counsel.

Run a pilot. Pick ten recent invention disclosures and run them through Clerq’s rapid triage assessment. Compare the results to your existing workflow outputs. The goal isn’t to replace patent counsel—it’s to compress the feedback loop between inventor disclosure and counsel review.

The specific metric to watch: time from disclosure submission to first attorney feedback. If you can compress this from weeks to days, you accelerate your entire patent pipeline without adding headcount.

For Legal Tech Buyers

The broader lesson extends beyond patent law. Agentic AI architectures that chain research, analysis, and synthesis into complete workflows are emerging across legal practice areas: contract review, regulatory compliance, case research, due diligence.

The evaluation framework for these systems differs from traditional software:

- **Task completion vs. task assistance:** Does the system complete work products, or does it require substantial human assembly?
- **Reasoning transparency:** Can users trace how the system reached conclusions? This matters for professional liability.
- **Edge case behavior:** How does the system handle ambiguous inputs, conflicting sources, and incomplete information?
- **Integration architecture:** Does the system fit existing workflows, or does it require process redesign?

For Engineering Leaders Building AI Systems

Clerq’s architecture offers lessons for anyone building agentic systems in complex, knowledge-intensive domains.

The workflow decomposition matters: Breaking patent research into discrete stages (claim parsing, search generation, prior art retrieval, relevance ranking, synthesis) allows each component to be optimized independently while maintaining



coherent end-to-end execution.

Domain expertise shapes architecture: Patent analysis requires specific reasoning patterns—claim construction, obviousness analysis, anticipation evaluation—that don't emerge automatically from general language models. The system architecture must encode these patterns.

Human-in-the-loop positioning: Clerq positions attorney review at the end of workflows, not throughout. This design choice trades real-time human oversight for speed, betting that output quality is high enough to make retrospective review sufficient.

The Competitive Landscape: Who Else Is Building This?

Clerq isn't operating in isolation. The legal AI market has grown crowded, and patent-specific tools are proliferating.

Established players like Clarivate (which owns Derwent Innovation), LexisNexis (which owns PatSnap through recent acquisition), and Google Patents have massive prior art databases and are integrating AI capabilities. These incumbents have distribution advantages—they're already embedded in law firm workflows—but they've approached AI as enhancement to existing search tools rather than workflow replacement.

Startup competitors include Patent Bots (focused on automated patent drafting), Specifio (AI-assisted claim writing), and IP.com (prior art databases with AI search). Each addresses different parts of the patent workflow, creating potential for integration or competition depending on where Clerq expands.

General legal AI platforms like Harvey, CoCounsel (from Casetext/Thomson Reuters), and Spellbook are building horizontal capabilities that include patent research. These platforms trade domain depth for breadth—they're less specialized than Clerq but serve multiple legal workflows.

The strategic question for Clerq is whether to remain a patent-specific platform or expand into adjacent legal workflows. The partnership model with RPX and Park IP suggests a strategy of depth: become essential infrastructure for patent-specific



work rather than competing broadly.

What Changes in 6-12 Months

Quality Benchmarks Will Emerge

The current marketing claims are difficult to verify. Within a year, practitioners will have enough experience with Clerq and competing platforms to assess actual quality. Expect academic studies comparing AI-generated patentability reports against attorney-generated reports, litigation outcomes based on AI-identified prior art, and patent office statistics on examination outcomes for AI-assisted applications.

The first public failure matters enormously. If a significant litigation loss traces back to AI-missed prior art, it will slow enterprise adoption across the industry. Conversely, if AI-identified prior art invalidates a major patent, it validates the technology dramatically.

Pricing Models Will Shift

Clerq's current pricing isn't public, but the economics force certain model choices. If the platform charges per search, it's betting on high volume at moderate prices. If it charges subscription-based enterprise licenses, it's betting on deep integration with fewer, larger customers.

The more interesting development will be how law firms price AI-assisted work to clients. Billing ten hours for work that took ten minutes creates obvious ethical tension. Expect new pricing models: flat fees for AI-assisted research, hybrid arrangements that separate human judgment from AI execution, or efficiency-based pricing that shares automation savings between firm and client.

Specialization Will Deepen

Patent law is specialized, but patent analysis within specific technical domains is even more specialized. A patentability analysis for a CRISPR gene-editing technique requires different expertise than analysis for a 5G antenna design or a financial derivatives trading algorithm.

Vertical specialization is the next frontier. Platforms that develop deep



expertise in pharmaceutical patents, semiconductor patents, software patents, or mechanical patents will outperform general-purpose tools in their domains. The company that assembles the best corpus of validated analyses in biologics patents, for example, will capture that market segment.

Regulatory Response Begins

Patent offices are conservative institutions, but they're not oblivious. The USPTO has already issued guidance on AI-assisted invention disclosure. Rules around AI-assisted prosecution are inevitable.

Questions regulators will address: Can AI-generated patentability analyses satisfy duty of candor obligations? Do practitioners have independent verification obligations for AI-identified prior art? What disclosure requirements apply when AI systems identify or analyze prior art?

These regulatory developments will shape how platforms like Clerq are used in practice and will likely create new market requirements around auditability and explanation.

The Structural Change Beneath the News

The Clerq announcement is a datapoint in a larger transformation. Agentic AI systems are emerging across professional services—legal, financial, medical, engineering—wherever knowledge work follows identifiable patterns that can be decomposed into component tasks.

The pattern is consistent: identify a workflow that consumes junior professional time, decompose it into tractable subtasks, build AI systems that chain those subtasks into end-to-end execution, and position the output for senior professional review rather than junior professional execution.

This pattern challenges the apprenticeship model that has defined professional services for centuries. Junior professionals learn by doing supervised work. If AI does the work, supervision becomes review of AI output—a fundamentally different skill set.

The companies that navigate this transition successfully will be those that figure out how to develop senior professionals without the traditional junior proving ground.



NLPatent Rebrands to Clerq with Agentic AI That Compresses Patent Research from Weeks to 10 Minutes

The platforms that enable this transition—Clerq among them—are building infrastructure for a professional services economy that looks very different from today.

The ten-minute patent search is the visible symptom of a deeper change: AI systems are no longer just research tools, they're junior colleagues that produce work products for review rather than merely gather information for analysis.