Transforming Patent Strategy Using AI

Intellectual Property

Technology, Privacy, and eCommerce

Banner artwork by */ AI*

Cheat Sheet:

- **AI-powered idea capture.** Connected, enterprise-grade AI platforms can proactively surface patentable concepts across fragmented SaaS systems, replacing slow manual review with continuous, context-rich innovation monitoring.
- **From reactive to proactive IP Strategy.** Scheduled AI workflows can identify both past R&D activity and upcoming feature releases, enabling earlier engagement with engineering and marketing teams to strengthen invention pipelines.
- **Reduce duplication and portfolio sprawl.** Similarity analysis across invention disclosures

and prior-filed patents helps detect technical overlap early, minimizing redundant filings, continuation complexity, and unnecessary spend.

- **Governance, privacy, and privilege first.** Effective deployment requires zero data retention, strict permission controls, attorney oversight, and clear labeling to manage confidentiality, discovery risk, and global privacy obligations.

The challenge

In today's rapidly evolving technological landscape, traditional approaches to idea capture no longer keep pace with technological changes in how individuals ideate and create new inventions. Many organizations still rely on traditional idea harvesting methods: manual processes such as requiring engineers or developers to sift through documents, emails, chats and meeting notes to identify what is potentially patentable, then taking the time from their busy schedules to turn those into an invention disclosure. This delays the identification of valuable intellectual property and adds operational complexity.

The "Jigsaw Puzzle" problem

Many large enterprises continue to "discover" IP by searching technical review decks, getting plugged into product development processes to review product requirements documents, reviewing customer proof-of-concept materials and internal meetings notes, all with the goal to identify inventions. Occasionally different engineering and product teams even within the one organization use a multitude of varied tools and data sources which disparately house information in fragmented locations, such as messaging tools (e.g. Teams, Slack), workflow systems (e.g., JIRA, Asana), documentation repositories (Confluence, Sharepoint, GDrive, Box) and often a variety of commercial or sometimes homegrown systems. This makes it challenging to find "one source of truth" where all the innovation sits, but rather it is akin to identifying different slices of data across multiple sites, and then attempting to construct a 1000-piece jigsaw puzzle to finally reveal the full scope of innovation in the organization. This is the reality that many patent counsel find themselves in, particularly in large distributed organizations.

The value of visibility

The good news is that this jigsaw puzzle can be solved far more rapidly if all the pieces are laid out in a logical manner, face up, and oriented correctly for reassembly. An enterprise-level SaaS-connected AI platform with IP and privacy protections such as zero retention and zero training (sometimes referred to as Zero Data Retention) is a useful tool in surfacing technical concepts that serve as ingredients for inventions. A general enterprise retrieval-augmented generation (RAG) platform or integrated knowledge platform could be helpful also for such purposes, but may depend on the user to provide the appropriate knowledge sources. Another important factor is the level of context provided in the search results, as a general keyword search may be quite noisy, whereas a context-rich search result tailored to your business activity will generally provide higher quality input to the AI tool, increasing likelihood of higher quality output.

A connected AI platform can be configured to regularly monitor this SaaS tool activity and understand the nature and scope of the data within each, and proactively surface candidate ideas as they emerge in a regular and repeatable manner. Of course, to ensure privacy and confidentiality of the company data, such a solution should include zero data retention, and strong permissions and data governance to ensure that users can only see the data that they have native permissions to in each of the underlying SaaS tools. This ensures data governance, and also reassures privacy, legal, compliance, and other important stakeholders.

Non-connected AI tools may also be used to attempt to identify ideas and technical concepts, but such tools generally require data to be identified and fed into the tool directly, rather than collected via prompting of a connected AI tool. The non-connected tools put the burden of surfacing the ideas on the user, which can be quite time consuming and potentially miss or delay coverage of valuable concepts.

Tech stack factors influence the outcome

If you're in a modern cloud-first SaaS environment, the techniques outlined in this article will be relevant. The greater connectivity between your SaaS environments and your AI platform, generally the more visibility and utility you will derive from the AI tool.

If you are in a more legacy tech stack, it is often preferable to start with a narrower scope of connected applications (primarily technical data sources, and exclude HR systems or sources containing GDPR special categories of data), and focus on one initial use case (e.g. engineering data repositories and core technical communications channels) and connect those first to show initial success, before expanding the set of connected applications.

Initial results from running these prompts may be noisy with a number of false positives and likely overbroad capture of technical information, but as your experience grows you can fine tune the prompts and thresholds to better suit your specific environment. Regardless, the prompt output in most systems will be significantly better than the old-fashioned manual review of ideas, and allows proactive conversations with engineering/R&D about potential areas for harvesting.

> The greater connectivity between your SaaS environments and your AI platform, generally the more visibility and utility you will derive from the AI tool.

In both scenarios (modern SaaS, and also legacy tech stack) the quality of your data is key. AI tools operate best with clean and fresh data, which companies should generally strive for, but often older environments benefit from data cleaning to see the full benefit of AI analysis.

In many companies, SaaS permissions and document sharing settings grow stale over time and as employees come and go, so it is important to keep those up to date. Legal teams can develop valuable partnerships with their IT and Security colleagues by collaborating on these areas.

It is important to select an AI tool that understands and respects the user's access permissions for any connected data source, and that does not grant broader access than a user's existing permission set. This permission structure should fundamentally be built at the core of the AI tool, not

a later “bolt-on” which could cause risk of incorrectly permissioned access.

Operationalizing AI: Practical Workflows

(Sidebar 1) Search for projects and ideas developed by Engineering/Product team

Example prompt for your AI/ops team:

[SOURCES] Review Slack, Jira, Confluence, Email, Gdrive, Asana and any other engineering and product repositories

[TIME PARAMETER] Focus on the time period from [DATE] forward.

[REQUEST] Identify all information relating to new ideas and feature development for [PROJECT NAME] and [FEATURE NAME] and any related terminology.

[DEFINITION] For the purpose of this search, define ‘innovative ideas’ as any product concept or feature which matches Criteria and Threshold

[CRITERIA] Positive mentions or feedback from a member of the engineering or product teams or from a legal team member or intellectual property practitioner, ideas with positive feedback coming from a customer, from Sales, from executives in the company.

[THRESHOLD] the idea should have elements of novelty and technical contribution

[OUTPUT] Provide a listing of all innovative ideas, with a particular focus on any innovative ideas that enhance user experience and user productivity and provide better context to users to do their work.

[IDENTIFY ACTORS] List out in a table format the directly responsible individual from the product team and from the engineering team for each innovative idea surfaced, and provide links to specific sources, document sections and conversations supporting this list of innovative ideas.

Output 1: A structured table of innovative ideas, potential inventors, sources

#	*Innovative idea*	*Short description (UX / productivity / context)*	*Product DRI (PM)*	*Engineering DRI*	*Key supporting sources & positive feedback*

The prompt above can be run one-off as needs be, or configured as scheduled agent to run monthly and capture ideas or concepts identified during the prior month, and provide this report to the patent team to start identifying inventions potentially worth filing as patents.

The prompt [DEFINITION] can be modified according to user preference to make it more selective

(narrowing the focus), and [CRITERIA] can be tailored by industry to focus on higher priority technical topics and [THRESHOLD] can be tailored according to how wide or narrow the funnel should be. For purposes of initial idea capture, 103 obviousness issues will not be considered as those require legal judgment by a human.

Caveat: the “innovative idea” list is a starting point, not a conclusion of actual novelty. But it does help narrow down the quantity of material the patent practitioner needs to review, providing a significant time saving. This permits greater focus on candidate ideas for discussion with engineers and product managers as potential inventions, saving valuable engineering and legal time, and increasing the likelihood of positive engagement from technical teams.

Assessing future idea pipelines

Above we have addressed the issue of capturing documented R&D projects (lookback invention capture), but what about future projects?

With the right prompting and data sources, the AI solution can be configured to track new features, architectures, and use cases that are planned to be shipped in the future by a certain user at a designated time. This output can include links to the relevant feature timelines, proposed announcements, any relevant marketing collateral being prepared (e.g. draft white papers, blogs, press releases, etc.)

(Sidebar 2) Future marketing announcements, product updates, and conference collateral can provide useful clues

Example prompt for your AI/ops team:

[SOURCES] Review Slack, Email, Gdrive, Asana, Figma, product roadmaps, company content hubs, blog and announcement drafts, and any other marketing repositories

[TIME PARAMETER] Look at the next 3 months forward from today.

[REQUEST] Identify all information relating to new feature releases by searching drafts of web page updates, blogs, white papers, conference announcements.

[DEFINITION] For the purpose of this search, define ‘new feature release’ as any planned announcement relating to new company products, services, or offerings, and also any enhancement of existing company products, services, or offerings.

[OUTPUT] Provide a table listing of all new feature releases, with a particular focus on features that add new capabilities and/or provide measurable outcomes to users of company products. Include citations to source documents relevant to the new feature releases.

[IDENTIFY ACTORS] List out in a table format the directly responsible individual from the marketing team and the product team for each new feature release surfaced, and provide links to specific document sources, document sections and conversations supporting this list of new feature releases.

[IDENTIFY SCHEDULE] For each new feature release, provide an estimate of when that new feature release will occur, and explain the support for that estimate with references to the documentation identified.

Output 2: A structured table of upcoming features/capabilities, individuals, timing, and sources

#	New feature / capability	What it adds	Est. release timing	Product DRI	Marketing DRI	Key supporting sources

As before, the prompt above can be run one-off as needs be, or configured as scheduled agent to run monthly and capture all ideas identified during the prior month, and provide this report to the patent team to start identifying inventions potentially worth filing as patents, and also to the broader legal team to review branding, trademarks, etc.

Caveat: marketing announcements often are not sufficiently detailed or enabling to meet 112 thresholds, but often provide a useful starting point for discussions with engineers and developers building the features. Think of them as an “early warning radar” rather than an actual identification of invention.

The overlap challenge: managing decentralized innovation

In large and decentralized organizations, similar customer problems often arise independently in different regions or business units. Without a way to effectively track and manage innovations, engineers may independently recreate and re-submit essentially the same patent idea, leading to duplication of work. This can also cause multiple overlapping patent submissions and even patent filings, which sometimes are addressed using continuation and continuation-in-part (CIP) applications, causing a tangled patent portfolio.

For example, in a particular company:

- Team A describes “invention A” in a disclosure meeting, which is recorded and stored in a secure folder in the company’s cloud instance.

- Team B independently creates “invention B” and discloses it in a presentation to management that is stored in a shared folder, whereby “invention B” overlaps with “invention A.”

- Team C, in another product group, creates a use case which could alternatively use Team A and/or Team B’s invention and is getting ready to deploy it using industry engineering and

workflow tools, duplicating effort.

This scenario can create a maze of patent continuations and incremental “add-on” filings, complicating portfolio coherence and straining budgets. And this is even assuming that the IP practitioner is aware of their duplicative or overlapping filings.

A connected platform can track these documents and events across the disparate systems, and if set up to run as a scheduled agent can alert the IP practitioner when a new disclosure substantially overlaps with prior filed ones thus saving both time and money, and potential headache for the IP team.

Caveat: any continuation or continuation-in-part decisions should be made by a human as part of a legal analysis, and typically will include analysis of written description support, prosecution history estoppel, and potentially of future enforcement posture.

(Sidebar 3) Compare new invention ideas against prior-filed patents summaries

Example prompt for your AI/ops team:

[SOURCES] Review [FOLDER] containing prior-filed patents and patent families, and [IDEAS] list of inventions being considered for patenting.

[TIME PARAMETER] Look at any patents filed in the past [TIME PERIOD] from today.

[REQUEST] For each invention, compare the technical features to the prior-filed patents, and calculate a percentage of technical overlap between the invention and the prior-filed patents..

[DEFINITION] For technical features of each invention, examine the structure and arrangement of technical elements of the invention and any technical steps required to implement the invention. For technical overlap, conduct a similarity analysis between the invention and the technical content and elements of the prior filed patents,

[OUTPUT] Provide a table listing out each invention, the degree of technical overlap with any prior-filed patents, and explain the specific technical elements or steps of the prior-filed patent which overlap with the technical features of the invention.

Optional: [IDENTIFY ACTORS] List out in a table format the individuals noted on the invention, and any individuals noted as inventor on the prior-filed patents which have overlap with the invention.

Output 3: A structured table of invention idea, approximate technical overlap, overlapping patent families, explanation of overlapping elements/steps

#	Invention	Approx. technical overlap (0–100%)	Key overlapping patent families	Explanation of overlapping

	(within TIME PERIOD)	**technical elements / steps**

Caveats: This output is machine-generated so is not a legal opinion and is not privileged, it is simply a statistical similarity analysis. However, it does provide a useful overview of which inventions may have technical overlap of technical features or technical steps with prior filed patents, providing a relevant starting point for IP practitioner review. For larger portfolios, this can save a significant time burden on the IP practitioner, and can also assist with finding “needle in a haystack” prior inventions within the company’s portfolio.

Other possible use cases

Although the similarity analysis shown in the third example is intended to find overlap within a given company’s own patent portfolio, if the [FOLDER] contains a set of different items (e.g. third party marketing collateral) and the [IDEAS] contains a set of the company’s granted patent claims, then the similarity analysis can be used to search for potential areas of overlap that may be interesting to an IP practitioner for other purposes.

Caveat: This is not a formal prior art analysis and should not be used as a substitution for one, but can help narrow the aperture and focus more clearly on a smaller set of documents.

Privilege, discoverability, privacy, and governance

Given that the use cases described herein may be subject to discovery at a later state, the prompts and actions described should be created by and overseen by attorneys as part of a legal analysis. For strongest protection, an attorney needs to define the scope and purpose of each specific query, and the attorney should review and curate the output in a timely manner. The logs and results of these prompts should be treated like legal documents, and protected appropriately with access controls and confidentiality labeling. This can be added to the output documents using a prompt like: “Label generated output as being company confidential, created at attorney direction.” However, to strengthen privilege, the output so-labeled should be used solely for legal analysis purposes, and not for metric reporting, data operations, or business use cases. Even with such markings, a legal risk remains that a machine-generated analysis will not withstand privilege scrutiny, but this should be weighed as a business decision versus the benefit for improving the pipeline of inventions.

However, to strengthen privilege, the output so-labeled should be used solely for legal analysis purposes, and not for metric reporting, data operations, or business use cases.

Those use cases can (if so desired) be separately labeled as “company confidential business information” to show a clearer distinction between the legal use and the business use.

Further, where the scope of content includes employees under privacy regimes such as GDPR, care

should be taken to ensure that the privacy rights of employees are respected, and this is particularly so in distributed multinational companies. As a mitigation, prompts can be specified to replace individual names with more general departments or team names instead. It is important to ensure that AI governance practices are followed, and the AI tool you select should have AI guardrails and governance at the core. The activities described in this article are internal business processes for IP management and patent strategy, which are generally considered low-risk under the E.U. AI Act. To meet transparency obligations, outputs from prompts should be labeled as machine generated.

> It is important to ensure that AI governance practices are followed, and the AI tool you select should have AI guardrails and governance at the core.

Conclusion

This article outlines many of the benefits that a connected AI platform can provide to patent practitioners (particularly those in-house) to help:

- Identify potential invention ideas scattered across existing technical data repositories.
- Identify potential invention ideas that may be announced in upcoming marketing events or announcements.
- Identify potential areas of overlap between invention ideas in a company's existing patent portfolio.

Utilizing such a tool can offer significant time savings, efficiency improvements, and enhanced visibility into the IP development process within the broader organization. It is particularly beneficial for decentralized organizations where visibility into different business units, functions, and geographies may be limited, and where the IP team may be thinly staffed. Such an AI solution increases the impact of the IP team and also helps generate valuable data-backed metrics on portfolio strategy and structure, invention activity, and help support IP budgets, headcount, tooling, and other priorities for the IP team. By enabling IP practitioners to focus their attention on the most strategic and valuable work and spend less time just searching and parsing through information, the AI solution raises the overall effectiveness of the team.

Join ACC for more AI insights!

Michael T. Moore

Vice President, Head of Legal, and Company Secretary

Glean Technologies, Inc.

Michael T. Moore is Vice President, Head of Legal, and Company Secretary at Glean Technologies Inc., where he leads global legal, privacy, intellectual property, and AI governance. He has previously held senior in-house leadership roles at a cybersecurity startup, IQVIA, Pure Storage, and Rambus, spanning product counsel, patent strategy, privacy, cybersecurity, and regulatory compliance. Moore is an IAM-300 recognized patent strategist and holds multiple IAPP certifications, including CIPP/US, CIPP/E, CIPP/C, CIPM, CIPT, and AIGP. He is a frequent author and speaker for the Association of

Corporate Counsel and Practising Law Institute on intellectual property, privacy, cybersecurity, and AI governance.