

Hello Anthropic

Global Creative Agency RFP

Competency Evaluation Questions

Priority Markets: FR, DE, KR, JP, IN

Fit and Focus

1. What are you genuinely strong at, and where are you adequate but not exceptional? What do you typically recommend clients go elsewhere for?

B2B tech creative is what we do. Not as a specialization we lean into for pitches — it's the only thing we do. Enterprise buying journeys, developer personas, the prosumer middle ground: we've built real depth there, and most brand agencies haven't.

We're also genuinely connected when it comes to media and creative. We run paid media for our clients, so we build creative knowing exactly how it'll be activated and measured. That changes what we make.

65% of our process is AI-assisted — meaning faster concepting, more variants, and smarter prep for multi-market adaptation. Not as a talking point, as a workflow.

Where we're adequate but not exceptional: high-budget cinematic film (we'll co-produce or point you to a specialist), pure consumer creative (we stretch into prosumer, not CPG), and markets where we don't have in-region staff — we'll tell you which ones rather than overstate presence.

What we'd send you elsewhere for: standalone brand identity or visual system design, PR and earned media, event or experiential production. We execute inside brand systems; we're not brand strategy consultants.

Fit and Focus

2. What mix of brand vs performance work do you do today, and where is the center of gravity?

We don't split them. Our model — brand-to-revenue — means creative is built from the start to work across the funnel. Performance creative disconnected from brand erodes long-term equity. Brand creative that ignores conversion mechanics is hard to justify.

Around 70% of our current engagements involve both brand and performance creative. As a B2B-focused agency, we typically see a natural skew toward performance-driven work, while ensuring brand building remains integral to long-term effectiveness.

For Anthropic specifically, we'd expect both: building brand familiarity in markets like Japan and Germany where enterprise trust takes time, alongside high-volume performance creative for developer and prosumer audiences where conversion is the immediate metric.

3. Who are your current or recent clients in tech, SaaS, or developer-facing products? Are you actively working with other AI companies, and how do you think about conflicts?

Recent and current: Databricks, Adobe, PayPal, ServiceNow, Palo Alto Networks, Cloudera, Box, Tenable, Autodesk. Multi-market work across developer, enterprise, and security audiences.

As you can see from our client list, we are actively working with a number of AI-led and AI-enabled organizations across enterprise technology, spanning data platforms, automation, and applied AI.

On conflicts: we maintain a register and disclose at proposal stage — before any confidential information is shared. If we're working with another AI company in an overlapping product segment, we'll name them and have a direct conversation about whether a genuine conflict exists. Where it does, we run separate account teams, separate file systems, and no shared personnel. Working with multiple AI companies isn't inherently a conflict given how segmented the space is — but we'll follow your lead on where the line sits.

Craft and Creative Process

1. Walk us through your end-to-end process from brief to delivered assets. What do you need from us to start, what do we get back at each stage, what's the typical timeline, and what most often slows things down?

What we need to start

A proper brief: audience, message hierarchy, proof points, tone, mandatories, and a clear stakeholder map showing who reviews, who approves, and who has final sign-off in each market. We also need access to brand assets, existing guidelines, and any campaign performance data worth factoring in.

Our localization partner also requires the English optimization strategy (so they can find the in-market equivalent, not just translate it), existing Translation Memories if you have them, and clarity on who's reviewing in each market — native speaker or English speaker. That last one changes the whole approach.

What you get back, and when

- **Weeks 1–2:** Brief alignment and strategic framing. We often push back on a brief at this stage — that's intentional. You get a written brief confirmation before any creative work starts.
- **Weeks 2–3:** 2–3 creative territories with strategic rationale, visual direction, and copy tone. Directional, not finished — the goal is to lock direction.
- **Weeks 3–4:** Refined concept(s) with initial format explorations.
- **Weeks 4–6:** Full asset development, market adaptations, final QA.
- **Delivery:** Organized asset package with naming conventions, technical specs, and handoff notes for your media team.

On the localization side, our localization partner runs parallel: Localization Framework Analysis first (to determine the right route — AI, AI + human, copywriter), then Cultural Consultation reviewing imagery and visuals alongside copy before anything is written, then production by native-language digital marketing specialists, then final QA by their regional Brand Captains.

Typical timelines

- **Net-new concept, single market:** 4–6 weeks from a complete brief.
- **Adapting an existing global concept to a new market:** 2–3 weeks.
- **Completely net-new market launch (research through first deliverable):** 5–8 weeks depending on research depth and client review speed.

What most often slows things down

An incomplete brief is the single biggest cause of rework. After that: unclear approval authority, scope additions mid-production, and legal/compliance review cycles — particularly in Germany and Japan. We build those in and flag them early. We don't absorb them silently at the end.

Craft and Creative Process

2. How do you handle revision cycles, and what happens when feedback fundamentally changes direction mid-process?

Two structured revision rounds per stage, included in scope. We say this in the kickoff conversation, not buried in an appendix. A third round, or a direction change after concept approval, gets scoped and costed before we continue.

If feedback fundamentally shifts direction: we stop. We don't quietly absorb it and deliver late. We have a direct conversation about what's changed and why, re-brief collaboratively, and reset timeline and scope with your agreement. If we missed something in the original brief, we'll own it. If the shift comes from your side — new stakeholder, changed business priority — we agree on a fair path forward together.

On the localization side, our partner builds revision expectations into the intake upfront. Because Cultural Consultation happens before production starts, mid-process direction changes are rare. When they do happen, their Brand Captains evaluate whether it's a genuine cultural issue or a personal preference, and route accordingly before resuming.

3. How does AI factor into your creative workflow today — ideation, production, localization, something else? What's your policy on GenAI in client work around disclosure, human oversight, and IP risk?

AI is operational infrastructure across concepting, production, localization prep, and performance variant generation. It generates options faster; humans make the creative decisions. AI doesn't produce final market-facing copy without human review — that's a hard rule.

Our localization partner operates as an AI marketplace, deploying whichever solution generates the best output for each use case — including Anthropic LLM licenses. Everything runs inside their secure proprietary stack, ISO 9001 and ISO 18587 certified. ISO 18587 is the standard specifically covering AI/NMT-produced translation and localization — it mandates post-processing QC and qualified human reviewers. They provide signed IT compliance attestation confirming zero client data is used to train public models, and maintain full audit trails on post-edit distance per language pair so you can see exactly how much human refinement went in.

We disclose which elements involved AI assistance. No AI-generated creative leaves without human sign-off. We use commercially licensed AI infrastructure and don't use tools with known IP litigation exposure without your knowledge. If you have internal GenAI restrictions, we operate within them.

Craft and Creative Process

4. Show us a piece of work you're proud of and one that didn't land. What did you learn from the one that didn't?

Work we're proud of — OpenText Cybersecurity

OpenText had built a comprehensive cybersecurity suite through multiple acquisitions but had no unified market presence. The category was already crowded with fear-led, technical messaging that all looked the same. Buyer engagement was low and differentiation was essentially zero.


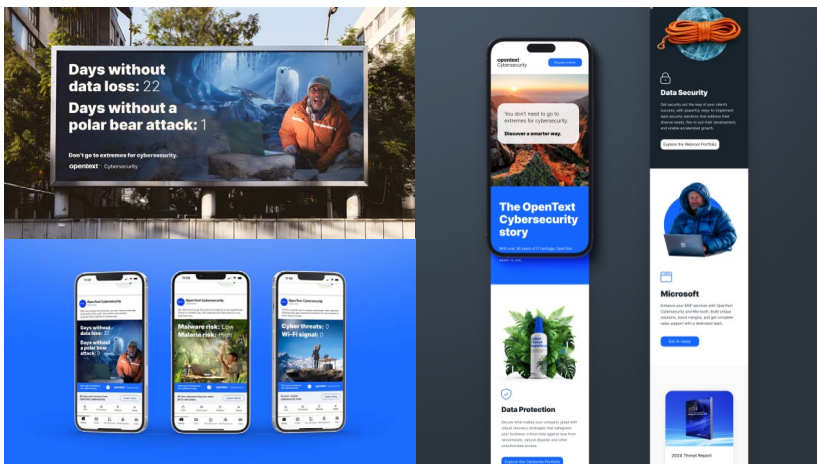
Our read was that the fear angle was the problem, not the solution. We reframed the brief entirely — instead of doubling down on threat and complexity, we looked for a way to make OpenText feel like the sane choice in an irrational category.

The insight: businesses were already going to absurd lengths trying to navigate cyber threats. We decided to show that.

The campaign — *Cybersurvivalists* — dramatized those extremes through character-led films set in surreal environments: arctic tundras, dense jungles, over-the-top survival tactics. The tone was deliberately funny in a category that had forgotten humor existed. OpenText was positioned as the smarter, simpler alternative to all of it.

We used GenAI to concept and prototype the campaign rapidly — detailed visualizations that got stakeholder alignment before a single dollar went into production. The final films were shot in LA and post-produced in-house. The platform was built to scale: video, social, display, and partner channels, across SMB and enterprise segments, across geographies.

Results: 4x lift in account engagement, 1.4 million views, 29% account engagement rate, and eight industry awards including Dotcomm, CreativePool, and Marcom Awards.



Craft and Creative Process

4. Show us a piece of work you're proud of and one that didn't land. What did you learn from the one that didn't?

Work we learned from — Zoho UK

We ran a high-profile brand awareness campaign for Zoho in the UK, spanning OOH across London Underground and major transit sites, alongside premium digital placements in the FT, Guardian, Times, and Economist. On paper, a strong, high-impact campaign. It underperformed.

Two things went wrong, and we own both of them.

First, the creative direction became overly influenced by a single client-driven message: “value.” While grounded in a real strategic point, it was overused across headlines and supporting copy to the point where it lost distinctiveness and impact. We had concerns during development but didn't push back strongly enough. Our role is to advocate for what will resonate with the audience, not align to stakeholder preference — and in this case, we could have been more direct, earlier.

Second, the campaign was designed and executed as brand awareness, but success was ultimately measured on lead generation. That misalignment meant the campaign was set up to underdeliver against expectations from the outset.

Following launch, paid social testing confirmed low engagement and limited cut-through. We reset to concept stage, refined the creative, and introduced pre-launch creative testing before relaunch — allowing us to identify the strongest-performing route with greater confidence.

What we changed: success metrics are agreed in writing before creative development begins. We are more direct, earlier, when client-driven creative decisions risk diluting impact. And we now prioritize pre-launch testing — particularly where there are differing stakeholder views — so we resolve uncertainty before going live.

Local Market Capability

1. For each of FR, DE, KR, JP, IN — do you have in-market creative talent or run these from a regional hub? If hub-based, how do you ensure cultural and linguistic accuracy?

Just Global — hybrid model

- **France:** Hub London. Native French creative and copy via specialist resources.
- **Germany:** Hub London. One of our most active B2B tech markets. Native German creative and copy expertise on the account.
- **Japan:** Hub Singapore. Japan isn't just a language challenge — tone, hierarchy, and visual conventions all differ materially from Western norms.
- **South Korea:** Hub Singapore. Korean B2B tech audiences have specific expectations around brand authority and technical credibility.
- **India:** Hub Singapore or Americas. India is internally diverse — the right approach depends on whether we're targeting a national enterprise audience, a developer segment, or regional.

Our localization partner — physical offices in all five markets

In-market staff aren't traditional linguists. They're native-language digital marketing specialists: bilingual copywriters, paid media people, data analysts. Each market also has a dedicated Brand Captain — an in-market authority who reviews all final assets for cultural resonance and brand voice before anything goes back to the client.

Local Market Capability

2. When you take on a new market, what's your research process for building an understanding of the local audience and competitive landscape? How long from kickoff to first deliverable in a net-new market?

Our approach combines Just Global's B2B marketing intelligence with specialist in-market expertise from our localization partner. Together, we move from fragmented data to actionable local insight — quickly and with confidence.

Our research process, by phase

1. Market and Competitive Intelligence

Just Global leads on market dynamics, growth trends, digital investment, intent signals, and competitive messaging and whitespace. Our localization partner enhances with AI/search visibility analysis — what buyers in-market are actively asking across search and AI answer engines — and local content gap identification.

2. Audience and Buyer Understanding

Just Global leads on ICP and buying group mapping, persona roles and pain points, creative strategy by funnel stage, and media habits. Our localization partner enhances with local audience nuance through social listening and in-market behavioral signals.

3. Local Market Validation

Just Global leads cultural and linguistic briefings with in-market specialists, and stakeholder interviews with regional marketing and sales teams. Our localization partner leads social listening, sentiment analysis, and in-market qualitative/quantitative panels (5–10 participants) to test concepts pre-transcreation.

4. Synthesis into Connected Intelligence

Led by Just Global, informed by both partners: priority audiences, competitive positioning, content and messaging opportunities, and channel and media recommendations.

5. Creative Concept Testing (optional)

Together with our creative testing partner, we validate creative effectiveness before scaling — testing emotional response and effectiveness benchmarks, measuring attention and engagement, identifying the strongest-performing concepts, and comparing against category norms.

Timelines

- **Just Global — research and orientation:** 2–3 weeks; first adapted creative concepts: 4–6 weeks from kickoff.
- **Adapting existing global creative:** 2–3 weeks to first deliverables.
- **Full net-new market launch (localization partner):** 5–8 weeks total (3–4 weeks research, ~2 weeks production once the brief is set). Panel review adds 1–2 weeks if required.
- **Creative concept testing (if required):** adds 1–2 weeks.

Local Market Capability

3. How do you approach copy for non-English markets — translate, transcreate, or write locally from scratch? Who owns quality control for cultural fit and language accuracy, and where are they based?

We don't use a single approach. The right answer genuinely varies.

- **Brand and campaign copy (headlines, CTAs, taglines):** transcreation. The meaning and emotional intent must survive the language shift. Starting from an English source is a constraint, not a process.
- **Markets with significantly different communication conventions — Japan especially:** local copywriting from scratch. We brief a native copywriter on strategic intent and let them write independently.
- **Product descriptions, spec content, legal copy:** translation with native review. Precision and proximity to source matter more than creative adaptation.

Our localization partner formalizes this as a Localization Framework Analysis — they evaluate the source content and recommend the exact route before any production starts: AI-only, AI + human post-editing, human translation, transcreation, or copywriter + SEO.

Quality control: Just Global's Creative Director reviews all market adaptations against strategic intent. Final linguistic and cultural sign-off is with a native-language specialist in or from the relevant market — non-English copy doesn't leave without it. Our localization partner's Brand Captains own the final QC layer in-market. For Germany, Japan, and South Korea, legal and compliance review time is built into the schedule proactively.

Local Market Capability

4. When a local market team disagrees with global creative direction, how do you handle it? Give a specific example.

Both sides are usually partly right. Local teams know their audience. Global teams know the brand system. Our job is to facilitate a well-informed decision — not pick a side or let whoever argued most recently win.

In practice: we document both the global rationale and the local objection formally. Our localization partner's Brand Captains evaluate whether the objection is a genuine cultural or linguistic issue, or a subjective preference. That distinction matters — it determines what happens next. We then facilitate a structured conversation between global brand owner and local stakeholder. We advocate for local adaptation when we believe it will materially affect performance. We advocate for global consistency when we believe the objection is preference-based.

A specific example — global engineering firm recruitment campaign

The US team was attached to a “top secret/NDA” creative direction for a global recruitment campaign. Our local teams in Japan and India told us it would actively deter top talent — treating corporate secrecy as a punchline doesn't convey prestige in those markets. The Brand Captains documented this as a genuine cultural breach, not a preference. We brought it to US HQ, made the case, and got agreement to pivot the localized creative toward prestige and innovation. It launched successfully. The disagreement was resolved through data, not opinion.

Brand Stewardship and Measurement

1. How do you balance global brand consistency with local flex? How do you build creative systems rather than one-off assets?

The difference between a campaign and a creative system is that a system tells every market what they can change and what they can't — and hands them the tools to do it without coming back to us every time.

In practice: we define the core concept and visual language globally. We explicitly document the flex points — photography style, copy tone, color emphasis, format mix, sometimes message hierarchy. We define the fixed elements that don't change regardless of market. And we deliver a Creative System Guide as a formal deliverable — not internal documentation, but a working document your in-house team, local market teams, and any future agency partners can actually use.

For this engagement specifically, building the system in Phase 1 means every market added in Phase 2 gets onboarded into something that already exists. We extend it, we don't rebuild it.

Brand Stewardship and Measurement

2. How do you measure creative effectiveness, and how do performance insights feed back into the next round of concepts? Share an example where data changed your creative direction.

We measure at two levels: output (individual asset performance) and outcome (pipeline and revenue contribution). The connection between them is where the work gets interesting.

We also look at creative through two complementary lenses: concept impact (macro) and asset performance (micro). At a macro level, we assess how a suite of assets built around a single concept performs collectively in driving engagement, pipeline generation, differentiation, and long-term brand impact. At the micro level, we analyze how individual executions perform in-market and what specific variables are driving results.

We track engagement (CTR, video completion, scroll depth), conversion (form fills, downloads, trial signups), and brand metrics where measurable (aided recall, message association, consideration shift). We run a monthly or bi-monthly creative review where media data is reviewed alongside creative variables — looking for patterns across assets and markets, not individual data points. Findings come back as hypotheses to test, not conclusions to implement.

An example: in a campaign targeting Global Systems Integrators, we tested two near-identical creative variants where the only change was background color — shifting from one brand color to another. Despite everything else remaining constant, the revised version delivered a clear uplift in engagement and click-through rate. We scaled this learning across the campaign and applied it to subsequent activity targeting the same audience. Small, data-led optimizations at the asset level compound over time, strengthening not just individual campaigns but future creative effectiveness within that audience and market.

Brand Stewardship and Measurement

3. How do you think about the handoff with an in-house media team?

We ask for your media team's technical specifications, naming conventions, and format preferences at briefing — not at delivery. Assets are built to your specs, not ours.

We provide a Media Team Briefing Document with every delivery: intended placement, performance context, creative notes relevant to optimization. We offer a brief alignment call at the start of each campaign cycle so your media team understands creative intent before activation. And if certain assets underperform in activation, we want to know — that feedback loop is built into the engagement as a formal cadence, not an informal conversation.

Team, Operating Model, and Scale

1. What would our team look like — dedicated, shared pool, or hybrid? Name the specific people who'd be on the account, their seniority, and their % allocation.

We propose a hybrid model: a dedicated core team with flexible specialist support to scale with demand. This ensures consistency, deep account knowledge, and strategic continuity, while giving you the agility to flex resource around peaks in activity, market launches, or campaign bursts.

Just Global — dedicated core team

- **Client Services Leader:** David Neal Dedicated. Senior client relationship, strategic direction, delivery oversight.
- **Senior Creative Director:** John Langstone Dedicated. Creative strategy, concept development, quality sign-off across all markets. 20% allocation.
- **Associate Creative Director:** Ben Trimmerger Dedicated. Creative strategy, concept development, execution. 50% allocation.
- **Art Director:** Desiree Jue Sinyork Dedicated. 70% allocation.
- **Copywriter:** Matt Mundy Dedicated. 50% allocation.
- **Motion / Video Specialist:** Tom Budd 50% allocation.
- **Integrated Production Director:** Katie Ryder Dedicated. 60% allocation.

Localization partner — dedicated leadership plus scalable in-market specialist pool

- Aron Simmons, Director of Strategic Accounts: US-based. Dedicated point of contact, owns the overall client relationship and day-to-day US communication. 50% allocation.
- Louise Hornberg, Director of Business Development (EU): EU-based. Time zone coverage and regional commercial alignment. 50% allocation.
- Colt Foutz, Digital SME & Production Orchestration Lead: Oversees global production workflows, Localization Framework Analysis, and AI integration. Time tied to project execution.
- In-market Brand Captains and Specialists (FR, DE, JP, KR, IN): Managed pool of native-language digital marketing specialists. Each market has a dedicated Brand Captain providing consistent third-eye QA on all brand deliverables.

Team, Operating Model, and Scale

2. What monthly volume can you sustain across 5 markets simultaneously, and how does that scale if we add markets?

Across 5 markets simultaneously, we sustain 10–12 net-new concepts and 40–50 market adaptations per month on static. For video and motion, 4–5 net-new concepts per quarter with 15–18 cut-downs or adaptations per month. For performance creative, 6–8 copy sets per market per month — each covering headline variants, body copy, and CTAs for a single campaign. Those are real numbers, not ceiling figures. They hold on a complex brief.

Adding a market follows a defined onboarding process — research, local resource confirmation, creative system extension — with a 4–6 week lead time to first deliverables. We can absorb one new market per quarter without affecting quality or timelines in existing markets.

The constraint isn't production capacity, it's coordination: brief alignment, local review cycles, and QA overhead. We resource ahead of expansion, not in response to it, and we'll be transparent if a proposed timeline requires additional resourcing before committing.

On the localization and production side, volume capacity is effectively uncapped. Executing hundreds of assets per month across 5 markets simultaneously is standard operating procedure. Adding a market doesn't require onboarding a new regional agency — our partner already has offices and specialist networks in over 100 cities globally. They spin up a local Brand Captain, run a Market Intelligence sprint, and route the new volume through an existing global production engine.

Team, Operating Model, and Scale

3. What's your turnover rate on accounts like ours, and what happens if the creative lead leaves mid-engagement?

Our account team turnover rate for engagements of this scope is around 20% annually — broadly in line with the industry, and something we manage proactively rather than reactively.

Good senior people get opportunities. What we control is what happens when they do. Account knowledge is documented and centralized, not stored in one person's head or inbox. Briefs, decisions, creative rationale, and campaign history live in shared systems accessible to the full team. Onboarding a new team member to an active account typically takes 6 weeks without disrupting delivery.

If the creative lead leaves mid-engagement: you're notified immediately, there's a structured handover with overlap, and no active project is disrupted without your knowledge and agreement. We don't wait until a replacement is confirmed to tell you — you hear it first. The people named in this proposal are the people on your account, and if that changes for any reason, you'll know before it happens.

On the localization and production side, turnover for strategic engagements of this scale sits under 5%. Brand knowledge lives in a centralized platform — Localization Frameworks, tone of voice guidelines, cultural consultation decisions, Translation Memories, and glossaries all documented and accessible. If a core team member transitions, production continues without disruption because the in-market Brand Captains and the underlying platform remain constant.

Team, Operating Model, and Scale

4. Walk us through your pricing model. How does it differ for video vs static, and for net-new concepts vs adaptations of global assets? What does adding a new market cost in time and money?

We work on a retainer model for engagements of this scale, covering dedicated team time, a defined monthly volume of deliverables, and a regular reporting cadence. The retainer is stated as outputs, not hours. You know what you're getting each month, not just who's on the clock. We provide a rate card that breaks out video vs. static clearly.

- **Static vs. video:** Net-new static campaign concepts run \$20–30k — strategy through to final assets. Net-new video runs \$45–65k for script through to delivered master, reflecting the additional production resource required for motion graphics, voiceover casting and recording, subtitling, and timing QA.
- **Net-new concepts vs. adaptations:** Market adaptations of existing global assets are typically 25–35% of concept cost depending on the degree of localization required — a straightforward resize and copy swap sits at the lower end; a market requiring cultural reworking of imagery, copy, and format mix sits at the higher end.
- **Adding a new market:** One-time onboarding cost of \$20–30k covering research, local resource setup, and creative system extension, with a 4–6 week lead time to first deliverables. Ongoing monthly cost for that market is itemized against agreed volume and added to the retainer transparently.

On the localization side, AI-assisted translation for technical and low-emotion assets is priced per word. Transcreation, cultural consultation, and net-new local copywriting are scoped by hourly effort based on the specific content. Video adaptations carry higher cost and longer turnaround — motion graphics replacement, voiceover, subtitling, and timing QA all require specialized effort. Adding a new market includes a one-time Market Intelligence cost (AI visibility, social listening, cultural consultation) plus ongoing per-word or hourly rates for that language pairing.

Nothing outside the agreed scope starts without being flagged and quoted first.

Legal, IP, and Information Handling

1. Who owns the work product and underlying files? How do you handle talent releases, music licensing, and stock rights across markets?

Intellectual Property and Work Product Ownership

All deliverables produced under our engagement are treated as work made for hire and transfer to the client upon full payment. This includes all campaign assets, creative files, and underlying source materials scoped within the Statement of Work. We retain ownership only of our pre-existing proprietary tools, methodologies, and frameworks, for which the client receives a license as embedded in the deliverables. Our agreements include a full IP assignment as a backstop for any work that does not qualify as work made for hire under applicable law.

Talent Releases, Music Licensing, and Stock Rights

We structure our agreements to clearly allocate rights clearance responsibilities between parties at the outset of each engagement. For content and assets we produce or procure on your behalf, we take responsibility for securing the necessary rights. For client-supplied materials, the client warrants that all required clearances are in place.

2. How do you handle data privacy compliance across GDPR, Japan's APPI, Korea's PIPA, and India's DPDP?

Our standard Data Processing Addendum provides full coverage for GDPR, UK GDPR, and Swiss FADP, including EU Standard Contractual Clauses, a UK IDTA Addendum, 24-hour breach notification, Article 32 security measures, and data subject rights obligations. For US-based data flows, our CCPA Addendum covers CCPA/CPRA and applicable state privacy laws across Colorado, Connecticut, Virginia, and Utah.

For APAC jurisdictions, we implement jurisdiction-specific data processing terms as part of our engagement setup. This includes cross-border transfer controls and purpose-limitation requirements under Japan's APPI, granular consent frameworks required under Korea's PIPA, and data fiduciary obligations under India's DPDP Act. We monitor regulatory developments — particularly as India's DPDP implementing rules are finalized — and update applicable terms to maintain compliance.

Legal, IP, and Information Handling

3. How do you handle confidential pre-launch product information — access controls, storage, NDAs?

We treat pre-launch product information with the highest level of confidentiality controls.

Contractually: our Mutual Non-Disclosure Agreement establishes a five-year confidentiality obligation from the date of first disclosure, prohibits copying of confidential materials, and requires certified return or destruction of all confidential information upon request. Team members with access are individually bound to confidentiality obligations, and we maintain liability for any unauthorised disclosure by our personnel, agents, or subcontractors.

Operationally: access to pre-launch materials is restricted on a named-personnel, need-to-know basis. Materials are stored in encrypted environments with access logging. We do not permit pre-launch content to be processed through unauthorised third-party tools or platforms. For particularly sensitive launches, we are prepared to execute project-specific confidentiality protocols — including defined access lists, storage requirements, and destruction milestones — as an exhibit to the Statement of Work.