

What IT wants to know *before agents get write access.*

Agents that write into CRM, ads and website belong in every proper security review. This pack answers the questions upfront: permissions, data flows, GDPR, prompt injection, EU AI Act — plus a pre-filled questionnaire. Open, no form, ready to forward.

AS OF **OCTOBER 2026** IGROMO E.U. WWW.IGROMO.COM/EN/SECURITY-GOVERNANCE

01 Permission ladder

02 Data flows per tool

03 GDPR & data processing

04 Prompt injection

05 EU AI Act

06 Security questionnaire

AT A GLANCE

Six answers for *the busy reviewer.*

Runs in your accounts

Tools, knowledge base, hub and agent setup run in your environment. Models, enrichment and tools run on your own contracts.

Reading isn't writing

Agents read for the analysis. Write access is granted deliberately per tool and agent — never wholesale.

No go, no external effect

Outreach, publishing and budgets need a single go. Irreversible actions always do.

Complete change log

Every change with before/after, reason, approval and timestamp. Rollback where technically possible.

Data isn't instructions


Content from emails, chats and web pages is information, never a command. Rules only come from the knowledge base.

Stoppable at any time

Revoke an approved rule, pause an agent, remove access: the agent stops immediately.

Permissions are granted *step by step*.

Every agent starts at the bottom and climbs only as far as its job requires — set per tool, documented and revocable at any time.



Always a single go — even with an approved rule

- Messages to external people (outreach, sequences, emails)
- Publishing (landing pages, posts, ads going live)
- Budget and bid changes
- Anything that can't be rolled back


Typical starting configuration of the core agents

AGENT	READS	WRITES	MAX.
Signal Agent	CRM, web, intent	Scores and notes in the CRM (rule), alert to sales	L3
Angle & Content Agent	Knowledge base, existing assets	Drafts; publishing only after go	L2
Paid Media Agent	LinkedIn Ads, Google Ads, GA4	Optimizations, budgets, bids — per go	L2
Pipeline Reporting Agent	CRM, ads, analytics	Reports and dashboards in the hub	L1
Outreach Routing Agent	CRM, outreach tools	Routing, contact hygiene (rule); sequences after go	L3
Inbox & Ops Agent	Approved inboxes and channels	Tasks with an owner after go	L2
Product Truth Agent	Product docs, release notes, website	Flags on contradictions	L1
QA & Brand Guard Agent	Live pages, ads, emails	Findings; fixes after go	L2

Guide values. Which level an agent actually gets per tool is set and documented in the project. Many teams deliberately start with L0 and L1 only.

Where data comes from – *and where it goes.*

An agent run always follows the same path. The concrete data-flow document with your tools, fields and access is created in the project and handed over.



AREA	EXAMPLES	AGENTS READ	AGENTS WRITE	PERSONAL DATA
CRM	HubSpot, Salesforce	Companies, contacts, deals, stages, activities	Notes, scores, properties, workflows, deal stages	Yes — business contact data
Ads	LinkedIn Ads, Google Ads	Campaigns, cost, performance	Audiences, creatives, bids, budgets (budget always with a go)	Low — audience lists
Web & analytics	GA4, Search Console, Fiddler	Pages, traffic, search queries, company-level visits	— (read only)	Low — aggregated
Enrichment & outreach	Clay, Sales Navigator, HeyReach, Instantly	Company and role data, sequence status	Enrich lists; sequences only after go	Yes — business contact data
Website & CMS	Webflow, Vercel, Lovable, HubSpot CMS	Pages and content	Landing pages and content, only after go	No
Communication & work	Slack, Teams, Outlook, Gmail, Notion, ClickUp	Approved channels, inboxes, tasks	Tasks, summaries, drafts	Yes — message content
Model	Claude (default), others after testing	The context of the given run	—	Depends on the run — via your account
Knowledge base & hub	Your repo or workspace	Rules, angles, decisions, metrics	Updates after the owner's go	Minimal — aggregated

A selection of typical tools. Which tools, fields and inboxes an agent actually sees is set by you in the project — new tools are only added when the revenue logic calls for it.

Clear roles. *Your contracts.* Data stays where it is.

Because RevPrime® runs in your accounts, your data map changes less than you'd think: the data lives in the tools you've already reviewed.

CONTROLLER

You

You decide on the purposes and means of processing — and on every approval in the system.

PROCESSOR

igromo

Where igromo has access to personal data while building and running the system, a data processing agreement under Art. 28 GDPR governs it.

YOUR PROCESSORS

Model and tool providers

Anthropic, HubSpot, LinkedIn & co. run on your accounts and your contracts — just as they do today.

Data minimization

Agents get access to the fields, channels and inboxes their job requires — not to everything. The hub stores aggregated metrics and status instead of personal data.

No special categories

RevPrime® processes business contact data (name, role, company, business email, activities). Special categories under Art. 9 GDPR are not part of the setup.

No automated individual decisions

Scores and routing prepare decisions; the system makes no decisions with legal or similarly significant effect within the meaning of Art. 22 GDPR.

Legal bases

Account research and B2B outreach often rely on legitimate interest — the balancing test is yours. Your rules (e.g. on email marketing) live in the brand, claims & legal module of the knowledge base, and every agent follows them.

Third countries

Many tools in the stack are US providers. Transfers run on your contracts with those providers, e.g. based on the EU-US Data Privacy Framework or standard contractual clauses. Where providers offer EU data regions, they can be used.

Documentation for DPIA and records

Whether a data protection impact assessment is needed is your call. igromo provides the system description, data flows and measures as input for the DPIA and your records of processing.

Data subject rights

Because the data lives in your systems, your existing processes for access, rectification and erasure apply. Agents don't create shadow copies.

End of project

The system is handed over with documentation. igromo's access is removed and tokens are rotated. igromo keeps no copies of your customer data.

Security doesn't hinge *on the model* – it's in the architecture.

No model today is fully immune to prompt injection. So RevPrime® doesn't rely on the model spotting a manipulation — it limits what a deceived agent can do in the first place.

RISK	SAFEGUARD
Hidden instructions in emails, web pages or forms	Data isn't instructions: external content is treated as information and never becomes a rule unchecked. Rules only come from the knowledge base, with the owner's go.
Data exfiltration through outbound actions	No external effect without a single go. Recipient and content are part of the proposal. Write access only to the tools the agent needs.
Overly broad permissions	Least privilege per tool and agent along the permission ladder. Start read-only, more rights only once the agent has proven itself.
Wrong or outdated claims go live	Approved claims in the brand, claims & legal module. The QA & brand guard agent checks everything before and after go-live, the product truth agent matches against product docs.
Error cascades and runaway cost	Budget and bid changes always need a go. Routines run with thresholds from the approved rule and can be paused at any time.
Compromised credentials	Separate, narrowly scoped access per tool. No secrets in the knowledge base or prompts. Revoking or rotating access stops the agent immediately.

To be clear: even within approved routines an agent can make mistakes, such as a wrong CRM note. That's why these routines are narrowly scoped, every change is in the change log, and anything with external effect stays with the single go.

Where agentic marketing *sits in the AI Act*.

An assessment from an architecture point of view, not legal advice. It helps your legal team get to the right questions faster.

AUG 2024	FEB 2025	AUG 2025	FROM AUG 2026
Entered into force	Prohibited practices and the AI literacy duty (Art. 4) apply	Obligations for providers of general-purpose AI models apply	Further obligations, incl. transparency (Art. 50), apply step by step

For some obligations the EU has initiated adjustments and deadline shifts through the digital omnibus package. Your legal counsel checks the current status.

Risk class

Typical RevPrime® use cases — content, reporting, campaign monitoring, account research and company-level scoring — do not, in our assessment, fall under the high-risk areas of Annex III, such as employment, creditworthiness of natural persons or critical infrastructure. Use cases in those areas are not part of RevPrime®.

Transparency

Chatbots that talk directly to people must disclose that they are AI; synthetic images, audio and video — deepfakes in particular — must be labeled. With RevPrime®, messages to external people only go out after a single go. Your labeling rules live in the brand, claims & legal module and apply to every agent.

Oversight and traceability

Human approval, a change log with before/after and clear owners are formally required mainly for high-risk systems. RevPrime® has them built in anyway — because otherwise leadership won't trust the system.

Roles

You use the system in your company and are therefore usually the deployer. The model provider (Anthropic by default) is the provider of the general-purpose model. The allocation of roles is documented in the project.

AI literacy

Art. 4 requires people working with AI systems to be sufficiently trained. Onboarding, playbooks and the documented handover are part of every project and contribute to this.

The usual questions. *Already answered.*

The questions that show up in almost every vendor and security review — with answers that hold for every RevPrime® project. Anything that depends on the specific setup is marked as such.

01 • COMPANY & ACCOUNTABILITY

Who is igromo?	igromo e.U., based in Heiligeneich, Lower Austria (EU), is the founder-led studio of Mario Maier: management consultancy and IT services from one source. Contact: hello@igromo.com.
Who has access to our systems?	Mario Maier personally — every engagement runs founder-led and senior-only. Other people only get access with your prior consent and under the data processing agreement.
Does igromo hold certifications like ISO 27001 or SOC 2?	No. igromo runs no platform of its own that stores your customer data: the data stays in your tools, which are covered by their own certifications and your contracts. igromo's measures are described in this pack; details are set out in the data processing agreement.

02 • ACCESS & IDENTITIES

How do the agents access our tools?	Through the tools' official interfaces (APIs, connectors, MCP) with dedicated, narrowly scoped access per tool — no shared personal logins. The rights an agent gets follow the permission ladder and are documented.
Who grants and revokes permissions?	You do. Access is created in your tools, and you can revoke it at any time. Without access an agent can neither read nor write.
Where are credentials and tokens stored?	Not in the knowledge base, not in prompts and not in the repo, but in the secret management of the respective runtime. Exactly where is set and documented per project.
How is igromo's own access secured?	igromo uses personal accounts in your tools with multi-factor authentication wherever the tool supports it. Your requirements, e.g. SSO or access only through specific accounts, are adopted.

03 • DATA & STORAGE

Where is our data stored?	In your tools and accounts. The knowledge base lives as versioned files (Markdown, YAML) in your repo or workspace; the hub stores aggregated metrics and status. Customer data is not copied into igromo's systems.
Which personal data do the agents process?	Business contact data from CRM and outreach tools (name, role, company, business email, activities) and content from inboxes and channels you approve. No special categories under Art. 9 GDPR.
Is our data used to train models?	The agents run on your own model account under your contract terms. Under commercial API and business contracts of the major providers, inputs are not used for training by default — check this in your contract with the provider.
How long is data retained?	The data stays in your systems and follows your deletion policies. How long the change log is kept is set by you in the project.

04 · AI & AGENTS

Which model do the agents run on?	Claude by Anthropic by default. The system is designed to be model-agnostic; other models and agent harnesses are possible. Every combination is tested and signed off for the specific deployment — switching isn't a single click.
Can agents act without approval?	Only within routines you have explicitly approved once, such as audience upkeep or CRM notes. Anything with external effect — outreach, publishing, budgets — and anything irreversible needs a single go.
What happens when an agent is wrong?	You reject the proposal, and the rejection with its reason goes into the decision memory. Executed changes are in the change log with before/after and can be rolled back where technically possible.
Can we stop an agent immediately?	Yes. You can revoke any approved rule, pause any agent and remove any access. After that the agent no longer acts.

05 · OPERATIONS & INCIDENTS

Is every action traceable?	Yes. The change log records every change with agent, before/after, reason, approval and timestamp. Agent status, pending approvals and anomalies are visible in the hub.
What happens in a security incident?	igromo informs you without undue delay, locks the affected access and analyzes the incident together with you. Deadlines and contact channels are set out in the data processing agreement.

06 · CONTRACT & EXIT

Is there a data processing agreement?	Yes, as soon as igromo has access to personal data. You can use your own template.
Who owns the system?	You do. Knowledge base, prompts, playbooks, hub and integrations run in your environment and are handed over with documentation. Usage rights are set out in the proposal.
How big is the lock-in risk?	Low. The knowledge base consists of open files, the system is designed to be model-agnostic and runs in your accounts. Your team can keep running it after the handover.

This pack describes the architecture and way of working of RevPrime®. It is not legal advice and does not replace a review by your data protection and legal counsel. The proposal and the data processing agreement are binding.

You have your own questionnaire? Send it over. I'll fill it in — usually before the first meeting with IT or data protection.

CONTACT

Mario Maier · igromo e.U.
Hasengasse 6, 3452 Heiligeneich, Austria
hello@igromo.com
www.igromo.com/en/security-governance