informer
by entrinsik

WHITE PAPER

THE INFORMER PLATFORM

Informer *AI*

Data Security & Governance

How the Informer platform protects your data when you use AI, including in Informer GO. Provider choice, the query-not-data architecture, and governance the AI cannot be talked around.

ABOUT THIS DOCUMENT

This paper explains how Informer safeguards your data when you use its AI features, in both the Informer web application and the Informer GO companion app. It is written in two layers: a plain-language executive summary for business and procurement readers, followed by technical sections that give security reviewers the detail they need to evaluate the controls.

It describes *what Informer guarantees and how those guarantees are controlled*. It is intentionally not an implementation manual.

What is Informer GO? Informer GO is Entrinsik's companion client for Informer, available as a desktop app (Windows, macOS, Linux) and a mobile app (iOS and Android). It gives users a chat-first way to work with the same AI features described here, under the same permissions as the web application, and stores only your session credentials on the device, never your analytical data or conversation history, which remain in your Informer instance.

CONTENTS

Table Of Contents

OVERVIEW

00 Executive Summary

When you ask Informer's AI a question about your data, five things are true.

01 You choose where the AI runs.

Use Informer's managed AI service, bring your own provider key (BYOK), or run a model entirely inside your own network. Supported options include OpenAI, Anthropic (Claude), Google (Gemini), Amazon Bedrock, and self-hosted / local models.

02 Informer sends questions, not your raw data.

The AI receives your question and a description of your data's structure, not the data itself. It proposes a query; Informer runs that query inside your instance; only the results go back to the AI — results that may include specific field values you are authorized to see (Section 3) — so it can explain them.

03 The AI can only see what you can see, and can't be talked out of it.

Every query runs under your own identity and permissions. Data not shared with you is absent from the AI's view entirely. It has no knowledge it exists, and the limit is enforced beneath the AI, so it can't be bypassed by prompt-injection or jailbreaks.

04 Your data is not used to train AI models.

Under Informer's managed service, our agreements with each provider prohibit training on your prompts or results. With BYOK, handling is governed by your own provider agreement. With a local model, your data never leaves your network at all.

05

Your data stays in your instance.

Whether you run Informer on-premises or in a dedicated cloud environment, your analytical data and AI conversations remain in your instance. Only usage metrics (for billing) are shared with Entrinsik, never your data or conversations.

Your data stays where it lives. The AI works from questions and answers rather than your raw data.

PROVIDERS & BYOK

01 You Control Where The AI Runs

Informer does not lock you into a single AI vendor or a single trust boundary. You select the posture that matches your data-sensitivity and compliance needs.

OPTION	WHAT IT MEANS	WHERE YOUR DATA GOES
Managed AI	Informer operates the AI connection for you using Entrinsik-held provider keys. No setup required.	Your question, a description of your data, and query results are sent to the provider Informer manages (OpenAI, Anthropic, or Google) under Entrinsik's enterprise agreements.
Bring Your Own Key	You supply your own provider credentials. Informer uses your account to call the model.	The same minimal exchange, but to your provider account (OpenAI, Anthropic, Google, or Amazon Bedrock), governed by your agreement with that provider.
Self-hosted / local	You point Informer at a model running inside your own network.	Nothing leaves your environment. The entire question-and-answer exchange stays behind your firewall.

Because Informer's server can itself be deployed on-premises, a fully behind-the-firewall configuration is possible: a self-hosted Informer instance paired with a local model means your analytical data and your AI prompts never traverse the public internet.

CREDENTIAL PROTECTION

When you provide a key (BYOK), it is encrypted at rest and used only to authenticate to the provider you configured. Keys are never exposed to other tenants and are not embedded in client applications.

POLICY & CONTRACT

02 The Training Guarantee

A common and reasonable concern is whether an AI provider will learn from or retain your data. Informer addresses this through the trust boundary you choose.

MANAGED AI

Entrinsik's agreements

Entrinsik uses enterprise API tiers, not consumer products. Our agreements with each provider stipulate that your prompts and results are not used to train their models.

BYOK

Your agreement

Because the call is made on your account, the provider's handling of your data is governed by your own agreement with that provider. You own the relationship and the terms. Confirm your provider tier includes a data-handling / no-training commitment before deploying with sensitive data.

LOCAL MODEL

No third party

There is no external provider involved. Data handling is entirely within your control, and your data never leaves your environment.

This guarantee is a *policy and contractual* control. It works alongside, not instead of, the technical fact that Informer sends as little data as possible in the first place, described next.

03 How Informer AI Works With Your Data

This is the heart of Informer's data-protection model, and it is deliberately simple. It applies to all of your data, whether it lives in a database, a file or spreadsheet, a web or REST API, or another connected source.

When you ask a question, Informer gives the AI two things: your question and a description of the shape of your data: the names and types of the fields available to you — scoped to your permissions, so if you can access 20 of a Dataset's 50 fields, the model is given those 20 and the rest are never named to it. The model uses this to reason about what to ask for. It does not hand the model your data. The AI responds by proposing a query. Informer's engine runs that query inside your instance, under your identity and permissions, and returns only the results, frequently a summary or an aggregate. These queries execute automatically as read-only, permission-bound operations; you are not prompted to approve each one, and because every query runs under your own identity, automatic execution cannot widen your access. A complex question may take several such turns in a bounded loop: the model proposes and refines queries until it can answer, each one independently checked against your permissions, with the loop itself capped. The AI then phrases a clear, written answer from what it received.

The model is a query author, not a data consumer. It works from descriptions and from the answers to the specific queries it runs, across a bounded number of turns.

What this means in practice

EXCHANGED WITH THE MODEL

- ✓ Your question
- ✓ A description of your data's structure (field names & types)
- ✓ Specific query results, usually aggregated

NOT SENT

- ✗ A bulk copy or export of your data
- ✗ Data outside your own permissions
- ✗ Another tenant's data

If a question genuinely needs record-level detail, the AI can ask Informer to return it, and Informer will. What comes back is still only data you are permitted to see (Section 4), and only the records that answer the question. Bulk transfer of raw data is avoided by design: Informer favors compact, aggregated results and limits the size of each exchange. Minimal data sharing is a deliberate default and a permission boundary, not a claim that the model is technically sealed off from your records.

A NOTE OF TRANSPARENCY

Informer does not alter or obscure the values inside a field that a user is authorized to see. If a question legitimately requires data from an accessible field, the actual values may be included in the query results sent to the model. Exposure is therefore governed by your access configuration: which fields and records each user (and the AI acting on their behalf) is permitted to reach. Section 7 describes how to use this control deliberately.

04 Governance: The AI Inherits Your Permissions

Informer is built for multi-tenant, role-governed enterprise use. Its AI features do not sit beside that governance model. They operate inside it.

The access model the AI inherits

Informer's data access is explicit and administrator-configured, and it exists independently of AI: it is the same model that already governs dashboards, data views, and ad hoc queries. A few principles define it:

- **Nothing is visible by default.** A user reaches a Datasource or Dataset only by owning it or having it shared with them, directly or through a Team. Assets that were never shared do not appear at all.
- **Sharing is read-only and graduated.** A shared Datasource is granted at a chosen level (full, restricted, or dataset-only), which determines both the query tools and the fields a user can reach. Sensitive fields can be marked restricted and must be granted individually.
- **Row and field limits travel with the data.** When a Dataset is shared, an administrator can attach a row filter so the recipient sees only permitted records, and can defer field visibility to the Datasource's restricted-field rules. These limits apply to every report built on that Dataset.

This model is organized around Teams and roles and is configured once, for the whole platform. Because it already decides what each person can reach, the AI does not need, and does not have, an access model of its own: when it queries on your behalf, it reaches exactly what your configuration permits, and nothing else.

DEFENSE IN DEPTH, AND THE AI RUNS INSIDE EVERY LAYER

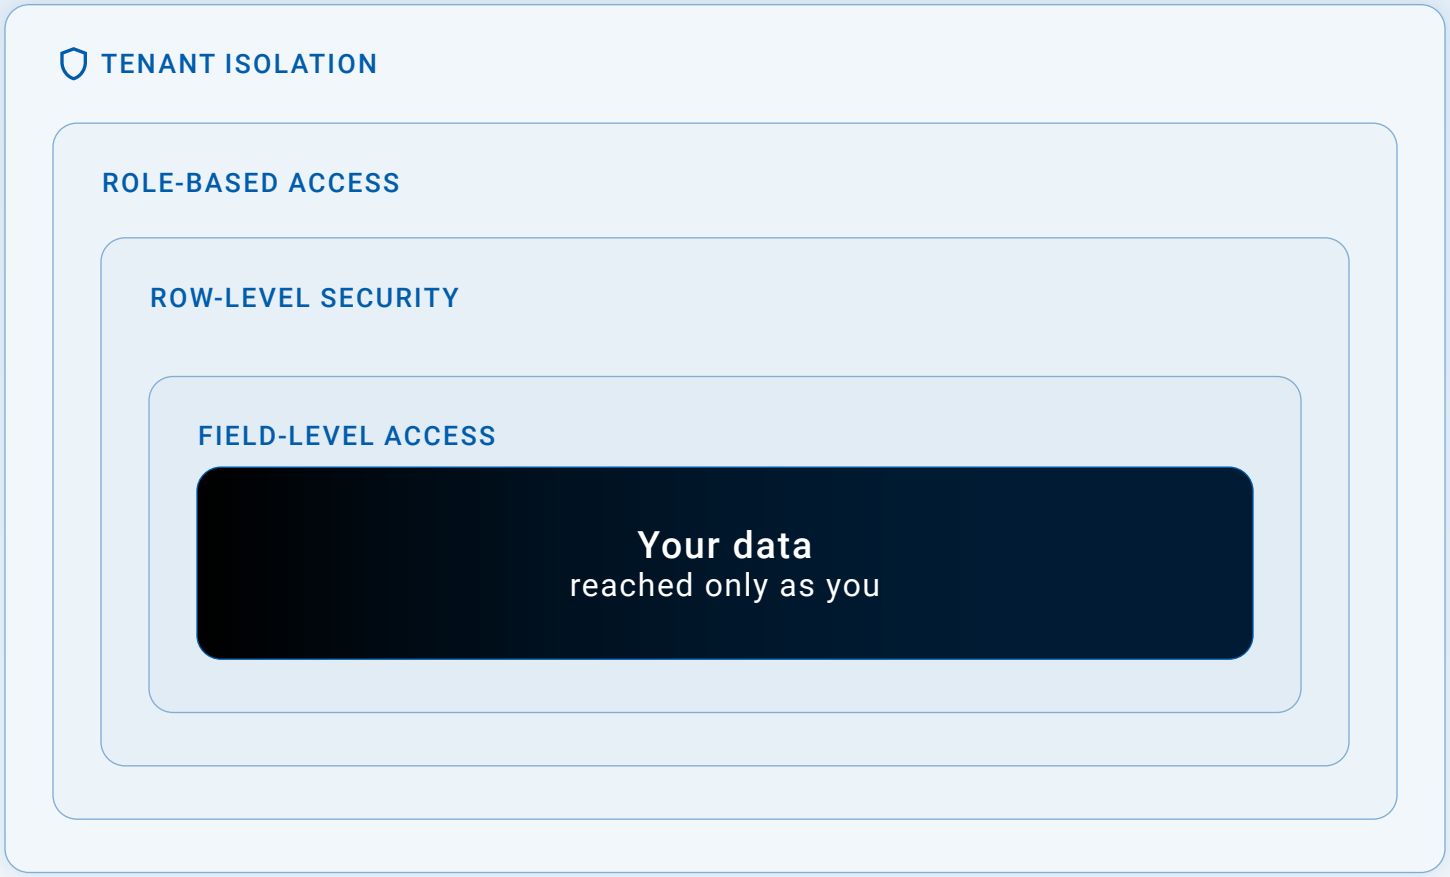

KEY GUARANTEE

When the AI runs a query on your behalf, it runs as you, subject to every layer above, with no elevated identity and no bypass path.

Each read is automatically scoped to what the requesting user is entitled to. The AI is bound by the same scoping.

So the AI can only query Datasets you can already access; within those Datasets it can only see the fields you can see; and its results reflect the same row-level filtering your own queries would. You configure access once, and the AI inherits it. There is no separate AI permission model to maintain.

Where each layer is enforced. Tenant isolation lives in the data stores themselves. In PostgreSQL (metadata and governance data) every connection runs under that tenant's own database role, governed by database-native row-level-security policies; in the analytical store (Elasticsearch / OpenSearch) each tenant uses its own account, limited to its own namespaced indices. Tenants share infrastructure by default; isolated instances of the whole stack are available as a paid option. Within a tenant, row and field limits are enforced centrally in Informer's query layer, applied to every Dataset query before it runs, beneath every feature including AI.

The AI acts only as you, and can't be prompted around it

Datasets, Ad hoc Queries, and Datasources are shared to specific users and teams. You see an asset only if it has been shared with you, directly or through a team you belong to. The AI operates strictly inside that same sharing. There are two ways to limit what an AI can see, and they are worlds apart:

✕ BEHAVIORAL GUARDRAIL

Instructed to withhold

The model is given the data, then told not to reveal certain parts of it.
The boundary depends on the model following instructions.

✕ Vulnerable: a crafted prompt or jailbreak can coax it into ignoring those instructions.

✓ STRUCTURAL ENFORCEMENT

Incapable of reaching

The model is never given anything outside your permissions, and every query is enforced beneath the model, at the data layer, under your identity.

✓ Immune: there is no instruction to override, and the data simply isn't reachable.

Informer uses structural enforcement. If an asset isn't shared with you, the AI doesn't know it exists. There is nothing for a cleverly worded prompt to ask for, and nothing for the model to "discover." No prompt-injection, jailbreak, or social-engineering attempt can make it return, or even acknowledge the existence of, data you are not authorized to see.

WHY THIS MATTERS

The effect is that the AI model itself sits outside the trust boundary for data access. The unpredictability that language models are known for is not a path to your data, because the model is never the thing enforcing the limit.

Attacks this design defeats

ATTACK VECTOR	WHY IT FAILS IN INFORMER
Prompt injection or jailbreak	Access is enforced beneath the model at the data layer; instructing the model cannot grant data the requesting user lacks.
Asking for unshared or another tenant's data	Unshared assets are absent from the model's view; it cannot name or reach what was never shared with you.
Privilege escalation or confused deputy	The AI has no identity or credentials of its own; it acts within your authenticated session and is never elevated.
Bulk data exfiltration	Results favor aggregates and each exchange is size-limited, so bulk export is not the default path.
Coaxing the agent across many steps	The query loop is bounded and every step is independently permission-checked, so additional steps cannot widen access.

05 Extending The AI: Toolkits And Integrations

Informer's AI can do more than query your Informer data. It can be extended, deliberately and under your control, with toolkits (including MCP) and integrations to external services.

Built-in tools stay inside your permissions

The AI's built-in tools run on the Informer server as the requesting user, through Informer's own permission-checked APIs. Everything in Section 4 applies to them unchanged: they act only as you, bound by tenant isolation and row- and field-level access. The default toolset adds capability, not reach.

Toolkits and MCP: reaching external tools

Beyond the built-in set, you can give the AI additional tools: servers that speak the open Model Context Protocol (MCP), or custom API tools. Where they run is explicit and separated:

- **Local tools run on your own device.** A local ("stdio") MCP server runs inside the Informer GO app on the user's machine, under that user's own operating-system permissions. It never runs on Informer's servers.
- **Remote tools run server-side as outbound calls** to the endpoint you configure.

Adding or editing a toolkit requires the elevated Data Wizard role; toolkits are tenant-isolated, private to their owner by default, and shared deliberately.

AN EXTERNAL TOOL IS A BOUNDARY YOU CHOOSE TO CROSS

An external tool reaches a system Informer does not govern, so its access is bounded by the credentials you configure for it (below), not by your Informer row- and field-level permissions. Enabling one is a deliberate decision to let the AI exchange data with that system. The built-in experience stays inside your permissions; reaching beyond it is always your choice.

Integrations: how external access is authenticated

Toolkits and other features connect to external services through integrations, which is where any credentials live; the tools themselves do not hold them. Each integration uses one of two authentication models, chosen when it is configured:

- **Per-user (OAuth)** is the default for services like Google Drive, OneDrive, Gmail, and Microsoft email. Each user completes their own sign-in and consent, and Informer stores a token for that user only. When the AI uses the integration, it reaches only that user's own account; one user can never act through another's connection.
- **Shared**. A single connection, configured once and used across the tenant. This is the model for service-style connectors authenticated by an API key, and an option where a shared account is intended.

You decide which model fits each service, matching it to how that data should be reached.

HOW STORED CREDENTIALS ARE PROTECTED

The credentials and tokens Informer holds for an integration are encrypted at rest with AES-256 and are never returned to the browser; the OAuth code-for-token exchange is completed server-side, not in the client. Configuring an integration requires the Data Wizard role; with a per-user integration, each user connects their own account, and their tokens remain private to them.

06 Supporting Security Controls

Beyond the AI-specific data flow, Informer provides the controls expected of an enterprise platform.

Independent attestation (SOC 2 Type 2)

Entrinsik maintains a SOC 2 Type 2 attestation covering the Security Trust Services Criteria, examined by the independent firm Sentry Assurance, LLC for the period September 2024 through August 2025. The examination found Informer's controls suitably designed and operating effectively throughout that period. The full report is available to customers and prospects under NDA. Entrinsik's next SOC 2 Type 2 examination covers January 2026 through August 2026, with the report expected in September 2026.

Encryption

Connections are protected with industry-standard TLS (HTTPS), including sessions secured with hardened cookie settings. Sensitive credentials (AI provider keys, Datasource passwords, and Integration tokens) are encrypted at rest using AES-256 with per-value initialization and standard key derivation. The encryption key is held in a protected key file on the host, separate from the application configuration and the database; it is guarded by the host operating system's access controls (and can be combined with disk- or volume-level encryption) and supports key rotation. This application-level encryption is scoped to secrets: AI conversation content and analytical Dataset data are not encrypted field-by-field by the application — their protection at rest relies on database- and disk-level encryption, applied per your deployment standards.

Authentication & single sign-on

Informer integrates with your identity provider and supports SAML, LDAP / Active Directory, and OAuth / OIDC, in addition to native authentication and multi-factor authentication. Users authenticate through your existing controls.

AI usage auditing

Every AI interaction is recorded for governance and billing: who initiated it, the model and provider used, the feature or mode it was used in (such as a specific Assistant or an embedded app), the volume of usage, and whether it ran on a managed or customer-supplied key. This gives administrators clear visibility into how AI is being used, independent of the conversation content itself.

Conversation storage, deletion & retention

AI conversations are stored inside your Informer instance, partitioned per tenant and subject to the same access controls as the rest of your data. By default, a chat's messages live for the life of that chat: when a user deletes the chat, its messages are permanently (hard-) deleted with it — there is no soft-delete grace period in the application. There is no indefinite out-of-the-box archive of conversation content. The usage audit above is kept separately (metadata, not message text), so administrators retain a record that AI was used even after a conversation is removed. On Informer Cloud, a deleted conversation may still exist in encrypted backups until it ages out of the backup retention window (see Cloud hosting, below). If your compliance program requires long-term retention of conversation content, that can be accommodated as a deliberate, opt-in arrangement rather than the default.

Incident response & breach notification

Entrinsik maintains an incident-response process and notifies affected customers as soon as possible in the event of a confirmed security incident. Entrinsik's security controls are independently attested under SOC 2 Type 2 (above); contact your Entrinsik representative for security documentation.

Cloud hosting, availability & resilience

Informer Cloud, the managed offering, runs on Amazon Web Services in United States availability zones with 24/7 physical security, under a 99.5% availability commitment. Customer data stays in the United States and is replicated in real time to a separate US region, with failover typically under an hour. Backups run nightly on the host and daily to offsite storage with a 7-day retention window, and restores are tested periodically.

Automated monitoring and intrusion detection alert a 24/7 operations team, and each customer runs in a sandboxed, per-tenant environment. When you offboard, Entrinsik returns a final copy of your data and then irreversibly removes it. *Self-hosted and on-premises deployments run on your own infrastructure, under your own controls*.

Minimal client-side footprint

The Informer GO companion app stores only the credentials needed to maintain your session on the device. Your analytical data and conversation history are not cached on the device; they are served from, and remain in, your Informer instance.

CONFIGURATION

07 Shared Responsibility: Your Controls

Informer provides the boundary and the enforcement; you decide how to apply them. The most important levers are below.

CONFIGURE ROLES & FIELD ACCESS DELIBERATELY

Because the AI inherits each user's permissions, the cleanest way to control what the AI can reach is to control what each user can reach. Scope role, row, and field access to match your data-sensitivity.

CHOOSE THE RIGHT TRUST BOUNDARY

Match the deployment and provider posture (managed, BYOK, or local model) to the sensitivity of the data in question. For the most sensitive workloads, a self-hosted instance with a local model keeps everything inside your network.

REVIEW WHAT YOU EXPOSE TO AI

Be intentional about which Datasets and fields are made available to AI-assisted features, particularly where they contain regulated or sensitive values. Because field values are returned as-is rather than masked, the way to keep a sensitive value from reaching an external provider is to restrict access to that field, or to run a local model so values never leave your network. There is no separate "visible to users but hidden from AI" control; by design there cannot be one, because the AI acts as the user or team, never as an identity of its own.

MANAGE RETENTION

Delete conversations in line with your own data-retention policies, or arrange long-term retention if your compliance program requires it.

SUMMARY

08 Trust Posture At A Glance

CONCERN	INFORMER'S POSITION
Choice of AI provider	You choose: managed, BYOK (OpenAI, Anthropic, Google, Amazon Bedrock), or self-hosted / local.
Behind-the-firewall option	Yes. Self-hosted Informer with a local model keeps all data and prompts in your network.
Is my raw data sent to the AI?	Not in bulk. The AI gets your question and your data's structure; Informer runs the queries the AI composes (within your permissions) and returns their results, favoring aggregates. Your data is not copied out wholesale..
Can the AI see data I can't?	No. Unshared data never enters the AI's view; its queries run under your identity with your row- and field-level permissions.
Can a crafted prompt or jailbreak bypass that?	No. Access is enforced beneath the AI at the data layer, not by instructing the model, so it cannot be prompted, injected, or jailbroken around.
Can the AI use external tools or services?	Only if you enable them. Toolkits (including MCP) and integrations extend the AI beyond Informer; they require the Data Wizard role and are an explicit opt-in. Built-in tools stay bound by your Informer permissions.
Whose account does an integration use?	For per-user integrations (the default for Google Drive, OneDrive, Gmail, Microsoft email), only the requesting user's own connected account; shared integrations use one configured connection. Stored tokens are AES-256 encrypted and never returned to clients.
Is my data used to train models?	No (managed: contractual; BYOK: per your provider agreement; local: never leaves your network).
Where does my data live?	In your Informer instance (on-prem or dedicated cloud).Only usage metrics are shared, for billing.

Are conversations stored, and for how long?	Inside your instance, per-tenant and access-controlled. A chat's messages are hard-deleted when the chat is deleted (on Informer Cloud, copies may persist in backups until they age out); indefinite retention of content is available as an opt-in arrangement, not a default.
Is AI usage audited?	Yes. Each interaction records who used AI, the model and provider, the mode it was used in, and usage volume, kept separately from conversation content.
Encryption	TLS in transit; AES-256 at rest for credentials and secrets.
Authentication	SAML, LDAP / Active Directory, OAuth / OIDC, native auth, and MFA.
Independent attestation?	SOC 2 Type 2 (Security), examined by Sentry Assurance, LLC for September 2024 to August 2025. Full report available under NDA.
Availability (Informer Cloud)?	99.5% availability commitment, with real-time failover to a separate US region (typically under an hour) and daily offsite backups at 7-day retention.
Where is Informer Cloud hosted?	Amazon Web Services, in United States availability zones. Self-hosted deployments run on your own infrastructure.

IN CLOSING

Conclusion

Adopting Informer's AI capabilities does not expand your data-exposure surface.

You choose the trust boundary. Your data stays in your instance. The AI works from questions and answers rather than from your raw data, and every query it runs is bound by the same permissions that govern your users, permissions it cannot be prompted around. Combined with enterprise-grade encryption, authentication, auditing, and the option to run entirely behind your own firewall, Informer lets you bring modern AI to your data without surrendering control of it.

APPENDIX A

A Provider & Deployment Options Compared

The table below compares the security-relevant characteristics of each way you can run Informer AI. This comparison concerns only the external trust boundary: where your data is processed and who holds the keys. The in-product governance described in Section 4 (tenant isolation, role-based access, row- and field-level security, and the guarantee that the AI only ever acts as you) is identical across every option. Choosing a provider changes where data is processed, not whether your permissions are enforced.

OPTION	MODELS	KEY & ACCOUNT	DATA PROCESSED	"NOT USED TO TRAIN" BASIS	TYPICAL FIT
Managed AI LEAVES TO PROVIDER	OpenAI, Anthropic, Google	Held & operated by Entrinsik	The managed provider's cloud	Entrinsik's enterprise agreements with each provider	Fastest to adopt; general enterprise use, no setup
BYOK (direct) LEAVES TO PROVIDER	OpenAI, Anthropic, Google	Your own provider key	Your own provider account's cloud	Your own agreement & tier with that provider (confirm it includes a no-training commitment)	You already hold an enterprise AI contract
BYOK (Bedrock) YOUR AWS	Models offered through Bedrock	Your AWS credentials	Your AWS account, in the region you choose	Your AWS agreement; content stays in your AWS boundary	AWS-centric orgs; data residency in your region
Self-hosted / local NEVER LEAVES	A model you operate	None (no third party)	Entirely within your own network	Not applicable; your data never leaves	Highly sensitive data; behind-the-firewall / air-gapped

Across all options, only a question, a description of your data, and specific query results are ever exchanged with the model, never your raw data (see Section 3). For the managed-AI option, prompts and results are processed in a United States provider region for Informer Cloud deployments, and in the provider region closest to your deployment for on-premises installations.

APPENDIX B

B Glossary

Managed AI	Informer's turnkey AI service, in which Entrinsik operates the connection to the AI provider on your behalf using Entrinsik-held credentials.
BYOK (Bring Your Own Key)	A configuration in which you supply your own AI provider credentials, so AI requests run on your own provider account under your own agreement.
Self-hosted / local model	An AI model you run inside your own network, so that AI requests and responses never leave your environment.
Dataset	A queryable, indexed collection of records in Informer, built from a Datasource. Datasets are the primary unit the AI queries.
Datasource	A connection to a system of record (database, file, spreadsheet, web / REST API, or other source) from which Datasets and Ad hoc Queries are built.
Ad hoc Query	A reusable query a user builds against a Datasource, shareable like a Dataset. The AI inherits the same access to an Ad hoc Query that the user has.
Tenant / multi-tenancy	A tenant is a single customer's isolated environment within Informer. *Multi-tenancy* keeps each tenant's data strictly separated, with no path from one tenant's data to another's.
Row-level security (RLS)	Access control that limits which *records* (rows) a user can see, applied automatically to every query.
Field-level access control	Access control that limits which ***fields*** (columns) a user can see. In Informer a field is either visible to a user or excluded entirely; values are not partially obscured (see *masking*).
Masking	Transforming a value to partially obscure it (for example, showing a card number as ****1234). Informer controls access by hiding whole fields, not by masking individual values.

Prompt injection / jailbreak	Techniques that attempt to manipulate an AI model, through cleverly crafted input, into ignoring its instructions. Because Informer enforces data access beneath the model rather than by instructing it, these techniques cannot widen what the AI is able to reach (see Section 4).
Aggregate / aggregation	A summarized result computed over many records (a count, sum, average, or ranking) rather than the underlying raw rows. Informer favors returning aggregates to the AI.
Schema / metadata	A description of the structure of your data (the names and types of fields), as distinct from the data values themselves.
Audit log	A record of activity. Informer's AI usage audit captures the metadata of each AI interaction (who, which model, what mode, and how much) for governance and billing.

---

informer
by entrinsik

This document describes Informer's data-security controls and guarantees. For contractual terms, provider agreements, or deployment-specific configuration guidance, contact your Entrinsik representative. © Entrinsik, Inc.