Greyquill

v1.0 · May 2026

# The Greyquill Platform

## Architecture & Governance Brief.

What it is, where it runs, who owns the data,
and how the audit answer gets assembled.

For Heads of AI, CISOs, Internal Audit,
and the people who have to defend this to a regulator.

01 · WHAT THIS DOCUMENT IS

# What this brief is.

This is a short architecture and governance brief for the Greyquill platform. It is written for the people who have to approve AI workloads inside a regulated enterprise: the Chief Information Security Officer, the Head of Internal Audit, the IT Architect who has to fit a new platform into an existing operating model, and the procurement reviewer who needs to know what they are signing.

It is not a sales document. It does not list features. It answers the four questions that gate every regulated AI deployment: where does the platform run, where does the data sit, what is it built on, and what do you take with you when the engagement ends.

If you need more than this brief covers, the Greyquill team will sit with your security and architecture leads in a working session. Most enterprises move from this brief to a 90-minute review and from there to a defensible plan in under two weeks.

AT A GLANCE

# How the platform sits inside your environment.

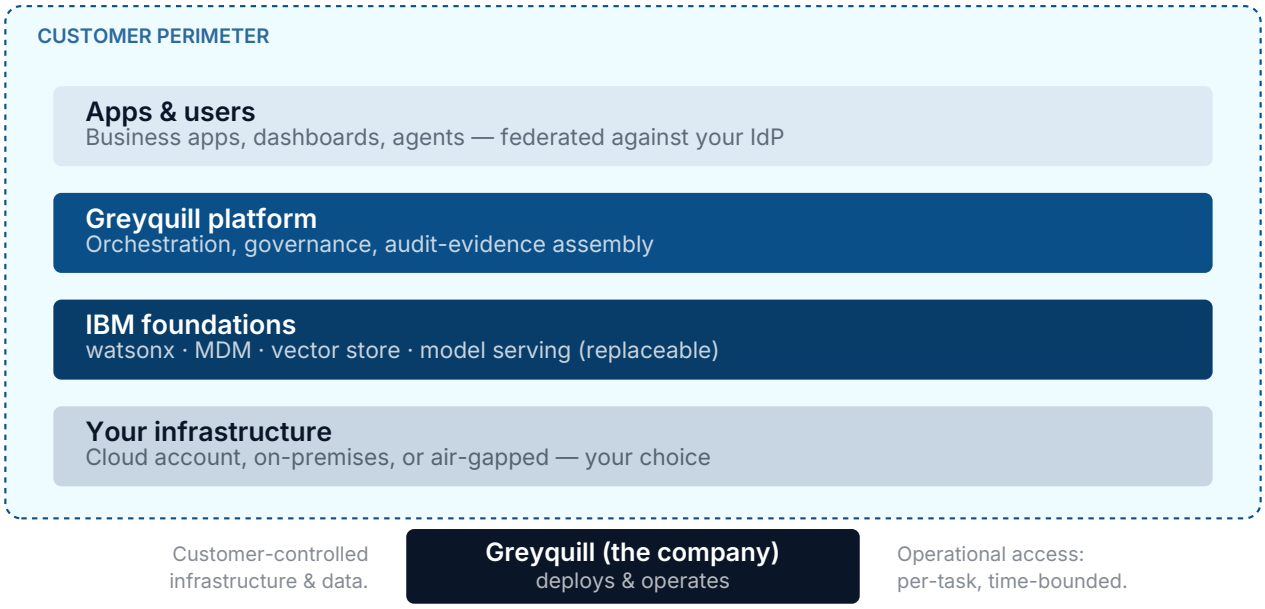


**Figure** · The platform is layered inside the customer perimeter. Greyquill, the company, deploys and operates from outside, with no standing access to production.

02 · DEPLOYMENT TOPOLOGY

# The platform runs where you need it to run.

Greyquill deploys inside your perimeter, on infrastructure you already control. The platform comes to your data, not the other way around. There is no default Greyquill-hosted tenancy that customer data passes through.

In practice this means one of three deployments, in order of how regulated buyers usually pick:

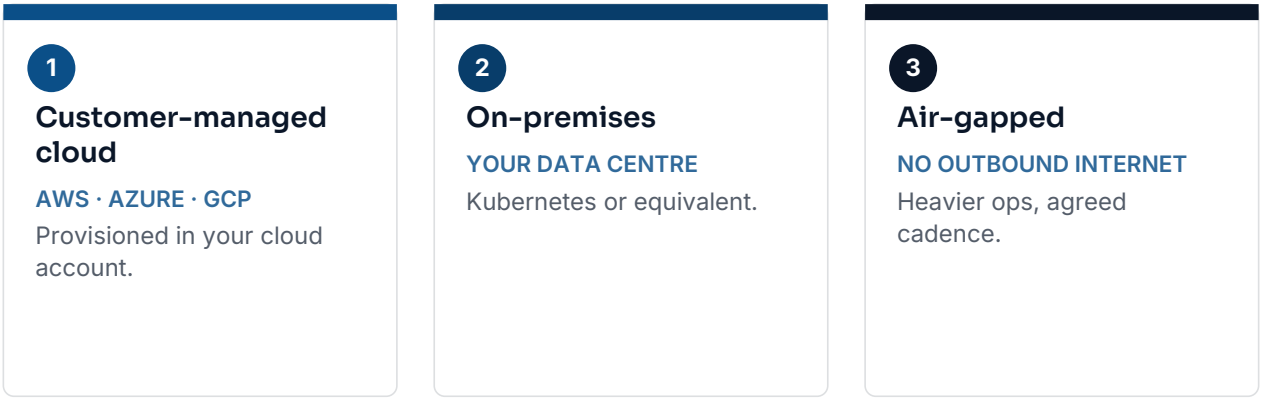


**Figure** · All three deployment modes share the same property: the platform lives inside your perimeter.

- **Customer-managed cloud:** AWS, Azure, GCP, or any other cloud account your team already owns and operates. Greyquill provisions and operates the platform inside that account. Your network policies, your IAM boundaries, your logging.
- **On-premises:** a customer-owned data centre, on Kubernetes or a comparable orchestrator. Used by buyers with sovereignty constraints, very strict data localisation rules, or existing on-prem investments they want to extend.
- **Air-gapped or restricted-egress:** where the platform must run without outbound internet. Supported with a slightly heavier operations model and a release cadence agreed up front. Used by some defence-adjacent and Tier-1 banking buyers.

Greyquill engineers participate in the deployment but do not require standing access to your production environment. Operational access is granted on a per-task, time-bounded basis through your existing access workflow.

---

03 · DATA RESIDENCY

## Your data never leaves your perimeter without your agreement.

This is the governance principle that drives every deployment choice above. It is worth stating plainly:

> **Customer data, by default, stays inside the customer perimeter. It is not copied, replicated, mirrored, or shipped to any Greyquill-controlled environment as part of normal operation.**

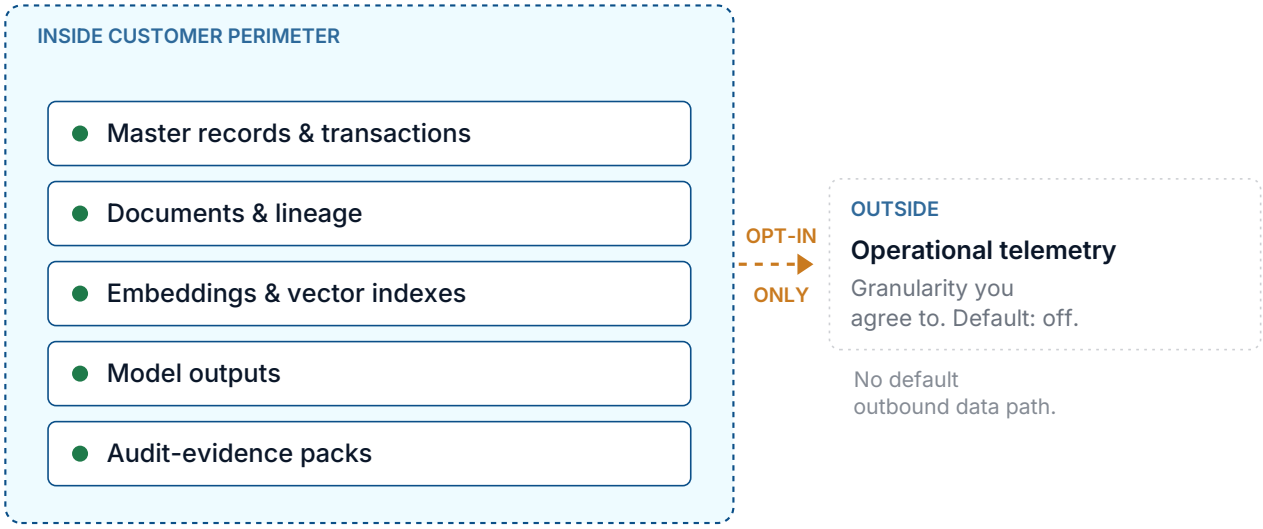


**Figure** · Five classes of data stay inside the perimeter. The only outbound path is opt-in operational telemetry, at the granularity you agree to.

Where Greyquill engineers need access to a sample of your data for an engineering task, that access is requested explicitly, time-bounded, and logged through your existing approval workflow.

04 · WHAT IT IS BUILT ON

## Built on IBM foundations. Replaceable parts by design.

Greyquill is an IBM Business Partner. The platform is built on top of the IBM stack that regulated enterprises already trust for production workloads. This gives the platform two properties that matter to your security and architecture teams:

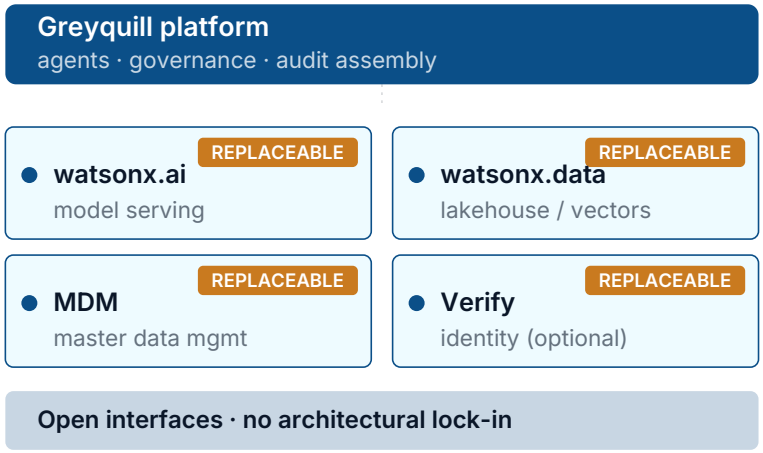


**Three properties.**

**Recognised.**
IBM components your security team has already approved.

**Replaceable.**
Swap any component for an equivalent your standards prefer.

**Subscribable.**
Buy IBM directly, under your own enterprise agreement.

**Figure** · Default IBM components are recognised by enterprise security teams and replaceable through open interfaces. No architectural lock-in to a specific SKU.

- **Recognised foundations.** Your security team is not being asked to evaluate a stack of unfamiliar components. The IBM components Greyquill uses have already been reviewed, accepted, and deployed inside enterprises like yours.

- **Customer-subscribable components.** Where it makes sense, you can subscribe to the underlying IBM components directly and have them installed under your own enterprise agreement. The platform is designed to compose with components you already own, not duplicate them.

Equally important: **the IBM components are replaceable.** Where a customer has a strong reason to use a different equivalent (a different vector store, a different model serving runtime, a different identity provider), Greyquill works against open interfaces. The default is IBM because it works for enterprises and is recognised by their security teams. The exception is anything else, supported.

---

05 · IDENTITY AND ACCESS

## Bring your own identity provider.

Greyquill does not run its own authentication system for your users. The platform integrates with whatever identity provider your organisation already uses. This is not optional. It is the only supported model.

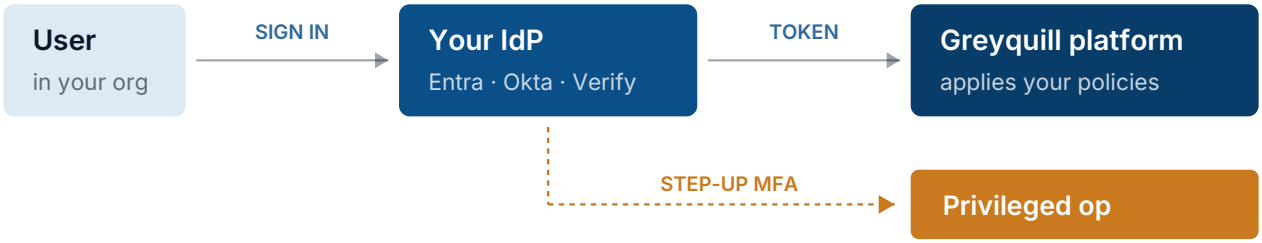


Greyquill maintains no parallel user directory. Your IdP is the source of truth.

**Figure** · Sign-in flows through your IdP; the platform reads claims and applies your policies. Privileged operations require a step-up event from the same IdP.

In practice this means:

- **Federation against your IdP** (Microsoft Entra ID, Okta, Ping, Google Workspace, IBM Verify, or any standards-compliant SAML or OIDC provider).
- **Your users, your groups, your conditional access policies.** Greyquill does not maintain a parallel user directory. The platform reads identity and group claims from your IdP at sign-in and applies them.
- **No shared secrets stored long-term.** Service-to-service trust uses short-lived tokens issued by your IdP or by infrastructure you control.

Privileged operations (changing a governance policy, promoting a model into production, modifying an audit trail) require a step-up authentication event that uses your IdP's existing strong-auth mechanism. Greyquill does not invent its own MFA.

---

06 · AUDIT EVIDENCE FLOW

# The audit answer is built as the system runs.

The single design choice that separates Greyquill from a model-serving platform with governance bolted on: every action the platform takes is logged into an evidence record at the moment it happens, not stitched together at month-end.

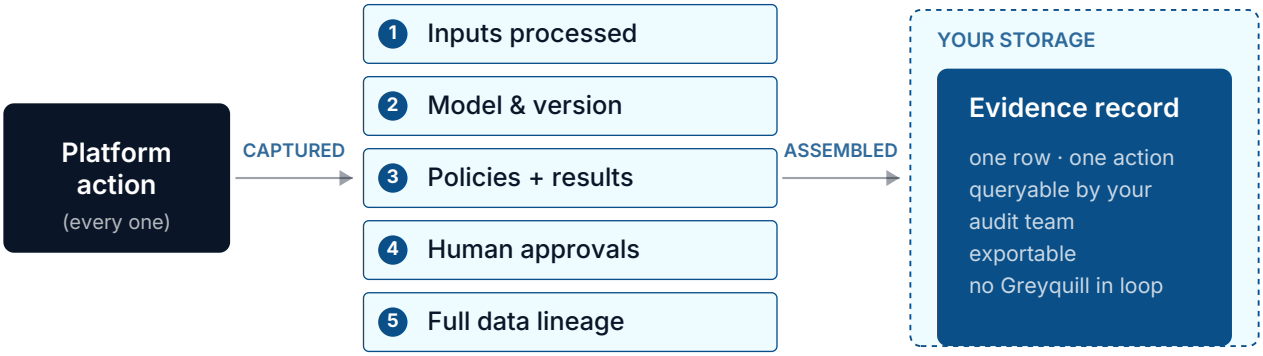


Regulator's question becomes a query, not a project.

**Figure** · Five fields are captured at the moment of each action and assembled into one evidence record per action, in storage you control.

For every piece of work the platform does, you can recover, in one query:

- Which inputs were processed.
- Which version of which model produced the output.
- Which governance policies were evaluated, and what each one returned.
- Which human approvals, if any, were attached.
- The full data lineage from source system to output.

That evidence record lives in storage you control. It is queryable by your audit team without involving Greyquill. It is exportable to whatever audit tooling your organisation already uses. When a regulator asks a question, the answer is not a project; it is a query.

---

07 · WHAT YOU TAKE WITH YOU

# Your configuration is yours.

If the engagement with Greyquill ends, your organisation does not lose what was built. The artefacts that matter (your master data model, your governance configuration, your agent definitions, your audit evidence history, your trained model artefacts) are yours, in your environment, in formats you can use without Greyquill.

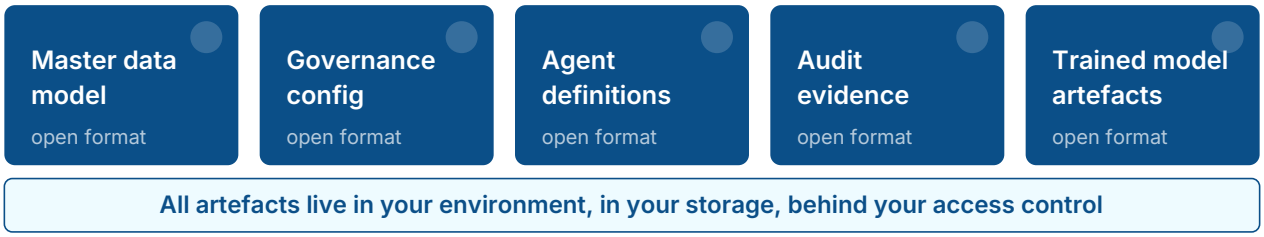


**Figure** · Five artefact classes, all stored in your environment in open formats, behind your access control.

This is not a contractual promise grafted onto a proprietary platform. It is the architecture: the platform stores customer-specific configuration in open formats, in customer-owned storage, behind customer-owned access control. There is no Greyquill-side database that your configuration depends on.

The practical test: if your team imagines, for a moment, that Greyquill the company ceased to exist tomorrow, would your AI workloads keep running? The answer the platform is designed to give is **"yes, until you choose to change them."**

---

08 · WHAT HAPPENS NEXT

## Where we go from here.

The fastest way through is a 90-minute working session with your security, data, and AI leads. Greyquill brings a senior architect and a senior governance lead. We work through the four questions this brief opens (deployment, data, identity, audit) against the specifics of your environment, and produce a written deployment posture document the next day.

From that posture document, an engagement plan typically takes another week to agree. Two weeks from first conversation to a defensible plan is a realistic expectation, not a slogan.

To start, send this brief sideways inside your organisation to whoever has to approve it. When the questions come back, mail them to **amarnath@greyquill.io** and we will answer them in writing before any meeting.

---