MarcoPolo: Data Security & Architecture Overview

The Challenge with Agentic AI

When AI agents need to work with your data—querying databases, pulling from CRMs, analyzing files—security becomes critical. How do you give an AI the access it needs without exposing credentials, leaking sensitive data, or losing visibility into what it's doing?

MarcoPolo solves this by providing a **secure, dedicated workspace** where AI agents can operate on your data safely.

Your Workspace: A Computer in the Cloud for Your AI

Think of MarcoPolo as provisioning a **dedicated computer for your AI agent**—a secure environment where it can connect to your data, run analysis, and persist its work.

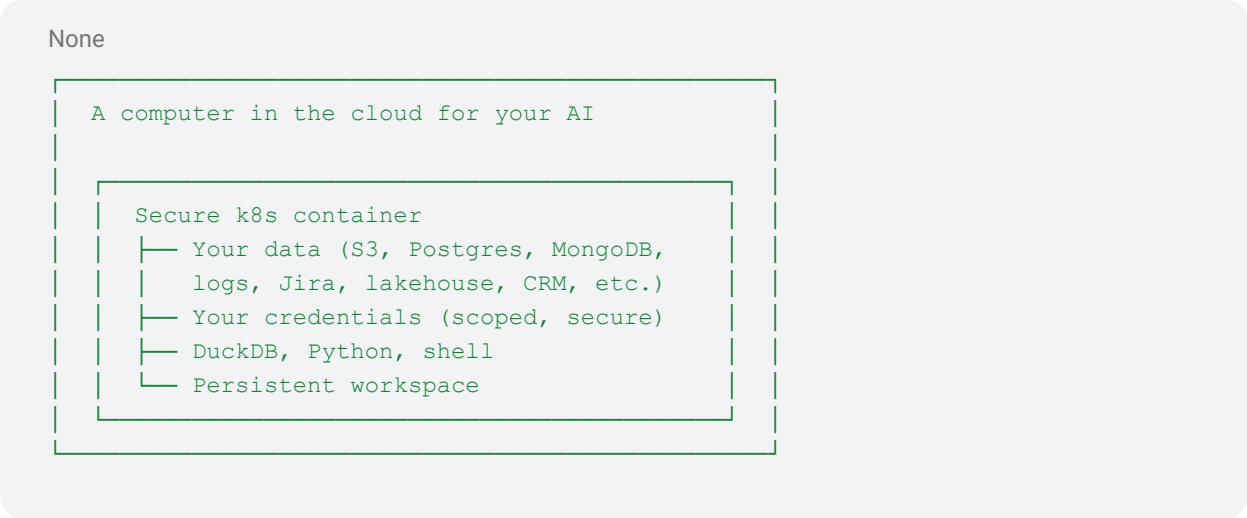

For example, when you use Claude with MarcoPolo, Claude gets its own workspace where it can write queries, download files, run Python scripts, and cache results—all in an isolated environment dedicated to you.

Persistent across sessions — The workspace stays available between conversations. Cached results, downloaded files, and generated artifacts persist so the AI can pick up exactly where it left off with full context of previous work.

Fully isolated — Each workspace is a dedicated Kubernetes container. No other users or AI agents can access your environment.

Data Integrity

MarcoPolo never stores your source data. The platform operates as a secure passthrough layer:

- Queries execute directly against your systems (databases, APIs, cloud storage)
- Results are cached in your workspace for continued analysis
- Your source data is never replicated, warehoused, or used for training

Your data stays where it is. MarcoPolo connects to your existing infrastructure—we don't migrate or centralize your data.

Credential Security: The Privilege Boundary

A key concern with agentic AI: if the AI can access your databases, can it also see (and potentially leak) your credentials?

MarcoPolo implements a strict privilege boundary that **prevents the AI from ever accessing credentials.**

None

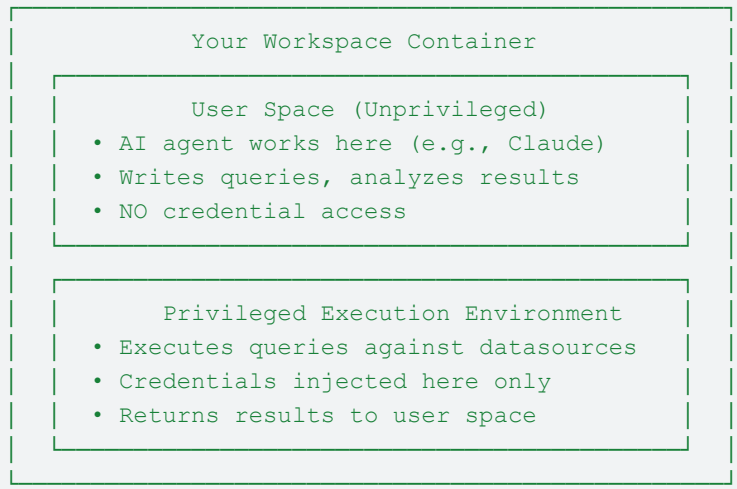

How it works: The AI writes queries in unprivileged user space. Those queries are handed off to a privileged execution environment that injects credentials, runs the query against your datasources, and returns only the results. The credentials themselves are never exposed to the AI—or to you during normal operation.

This means even if an AI agent were compromised or behaved unexpectedly, it cannot access or exfiltrate your credentials.

Transparency: Not a Black Box

A common concern with AI agents: "What is it actually doing?"

MarcoPolo provides complete transparency—everything the AI does is visible and auditable.

Full access to your workspace. All queries, scripts, and artifacts generated during sessions are saved. You can browse the filesystem just like your own computer:

- `/workspace/queries/` — every query written
- `/workspace/downloads/` — files pulled from your datasources
- Any scripts, analysis outputs, or generated files

You see every operation. The AI surfaces the exact SQL, API calls, and commands it executes. Nothing runs in the dark.

Audit trail built-in. Conversation history captures every tool call and query. Combined with workspace file access, you have complete visibility into what happened and when.

Common Questions

Q: Can the LLM or MarcoPolo see my data? A: The LLM only sees query results you explicitly request—it has no background access to your systems. MarcoPolo infrastructure doesn't access your workspace contents.

Q: Is my data used for AI training? A: No. Customer data never enters any training pipeline.

Q: Could the AI leak my credentials? A: No. The AI operates in unprivileged user space and never has access to credentials. Credentials exist only in the privileged execution layer.

Q: Can I see what the AI did? A: Yes—in your conversation history and as files in your workspace. Full transparency.

Q: What persists between sessions? A: Your workspace persists—cached results, generated queries, downloaded files. This gives the AI continuity and context. Your source data stays in your systems.

For technical architecture details or compliance documentation, contact support@marcopolo.dev