



 **Human
Intelligence**

Challenge



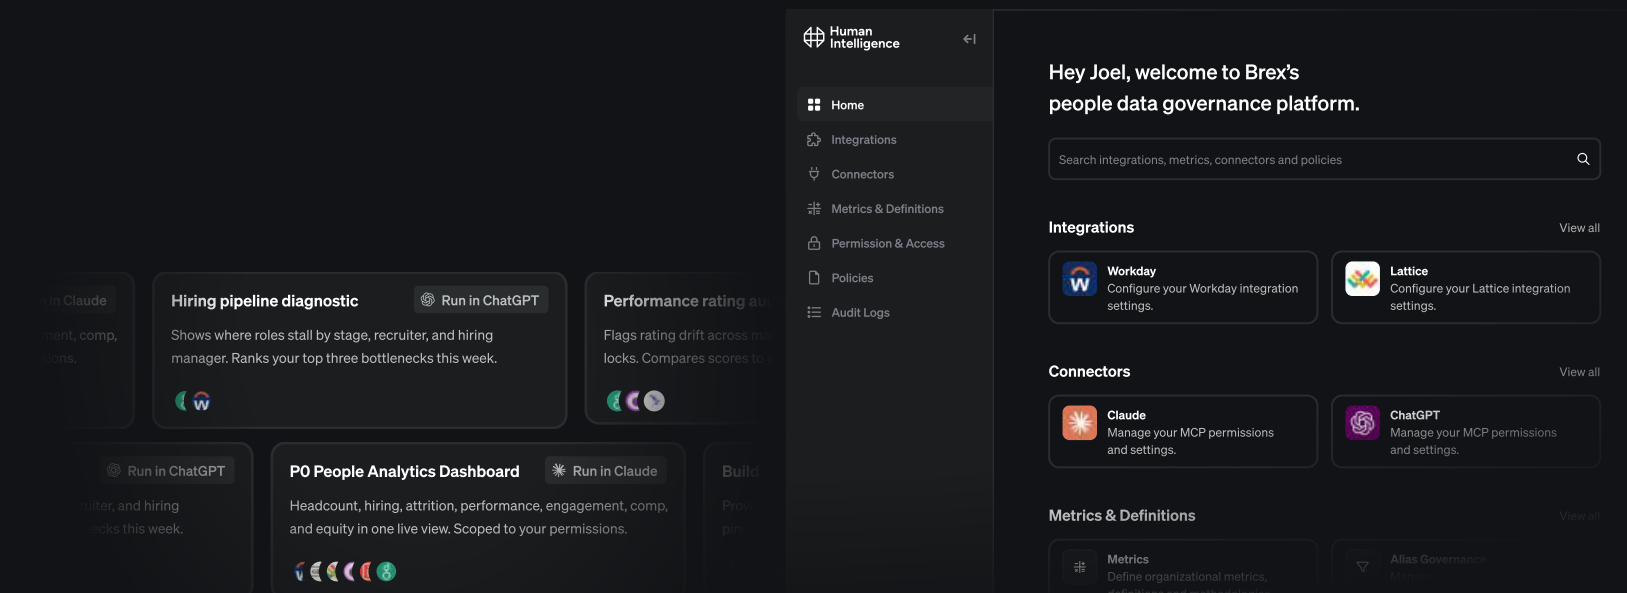
Move fast across the entire stack without ever compromising trust.

Human Intelligence is the governed data platform for people data. It unifies the enterprise's most sensitive systems (Workday, Lattice, Greenhouse, Culture Amp, Carta, Envoy, and 10+ others) into a single AI-ready model, with org-aware row-level security, policy management, metric governance, and a secure MCP that puts people data to work in AI tools without putting it at risk.

The bar is high because the data is sensitive. HI's promise — a single, governed, AI-ready model over an enterprise's most sensitive systems — only holds if every layer advances together. A faster connector is worthless without governance; natural-language querying is dangerous without permissions and audit.

But the existing model didn't scale. HI's current infrastructure, a purpose-built data-mart layer meant analytics leaned on pre-baked tables, so nearly every new customer question routed back through engineering. Adding sources, reshaping data, and answering novel questions were all bottlenecked on the same small team.

And speed couldn't come at the cost of trust. HI needed to land new connectors in days to support their customers' needs, stand up AI querying and agents, and expose data to AI tools via MCP while preserving row-level security, metric governance, and a complete audit trail.



Solution

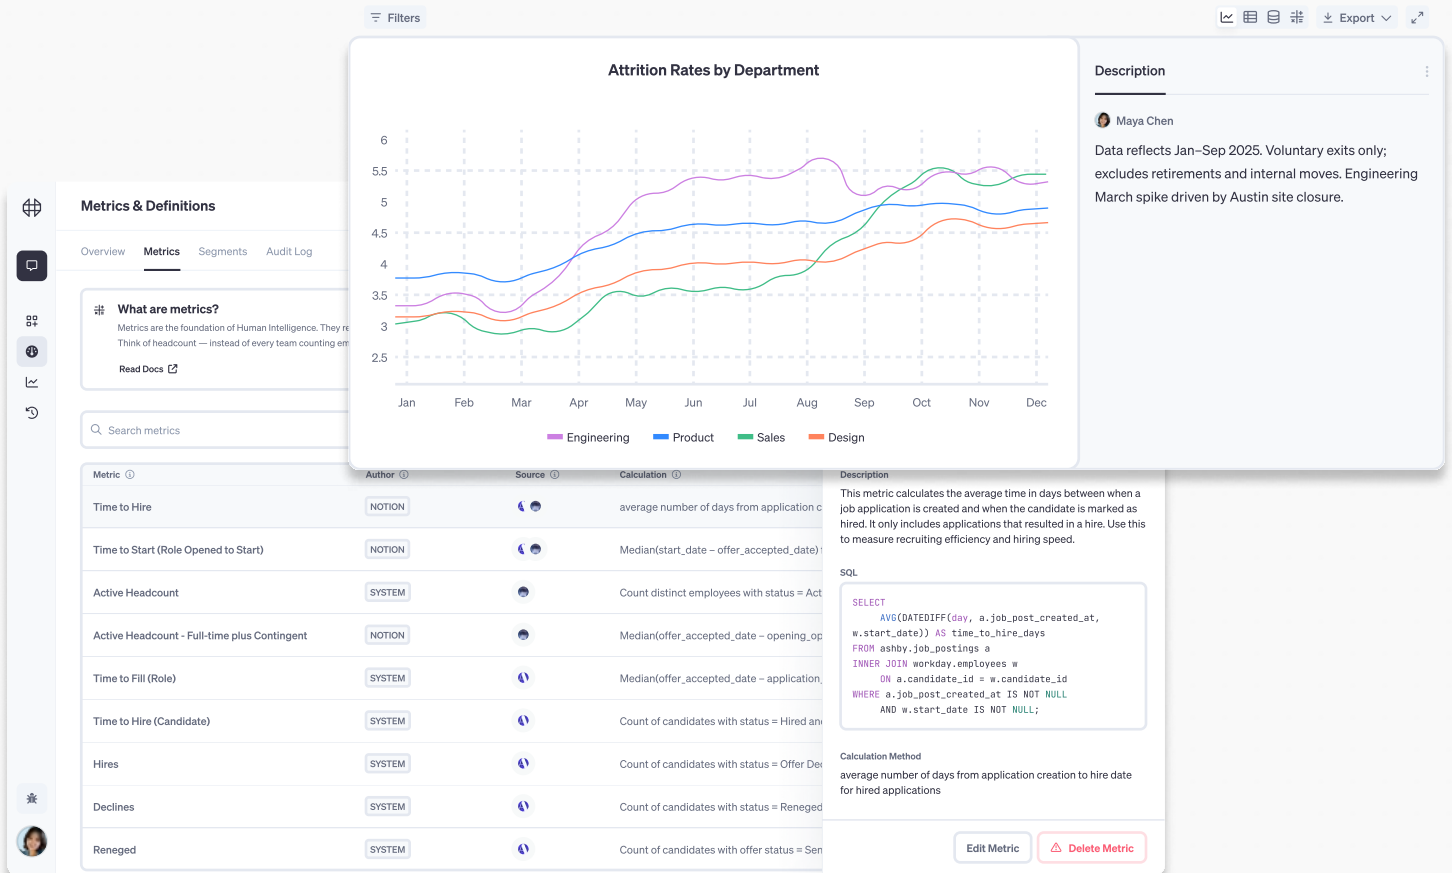


Galaxy embedded as a forward-deployed team alongside HI's team, shipping production code across every layer of the stack over the course of the partnership. Our work spanned from data modeling to custom AI implementations:

DATA OAuth connectors for Greenhouse, Leapsome, Slack, Envoy and Carta, ready for production in < 1 week. Built the canonical entity-based model that replaced HI's purpose-built mart layer, an entity-resolution engine mapping a canonical entity across every source, and a denormalized graph where queries traverse 5+ systems in one view.

AI Talk to your data via natural language with metric-catalog matching, plus chat-to-charts and a dashboard copilot joining 5+ systems in one view. LLM agents for metric management and chart creation made engineering-gated workflows self-serve. MCP shipped to Slack and Notion in production, with every agent call moved onto auditable API endpoints.

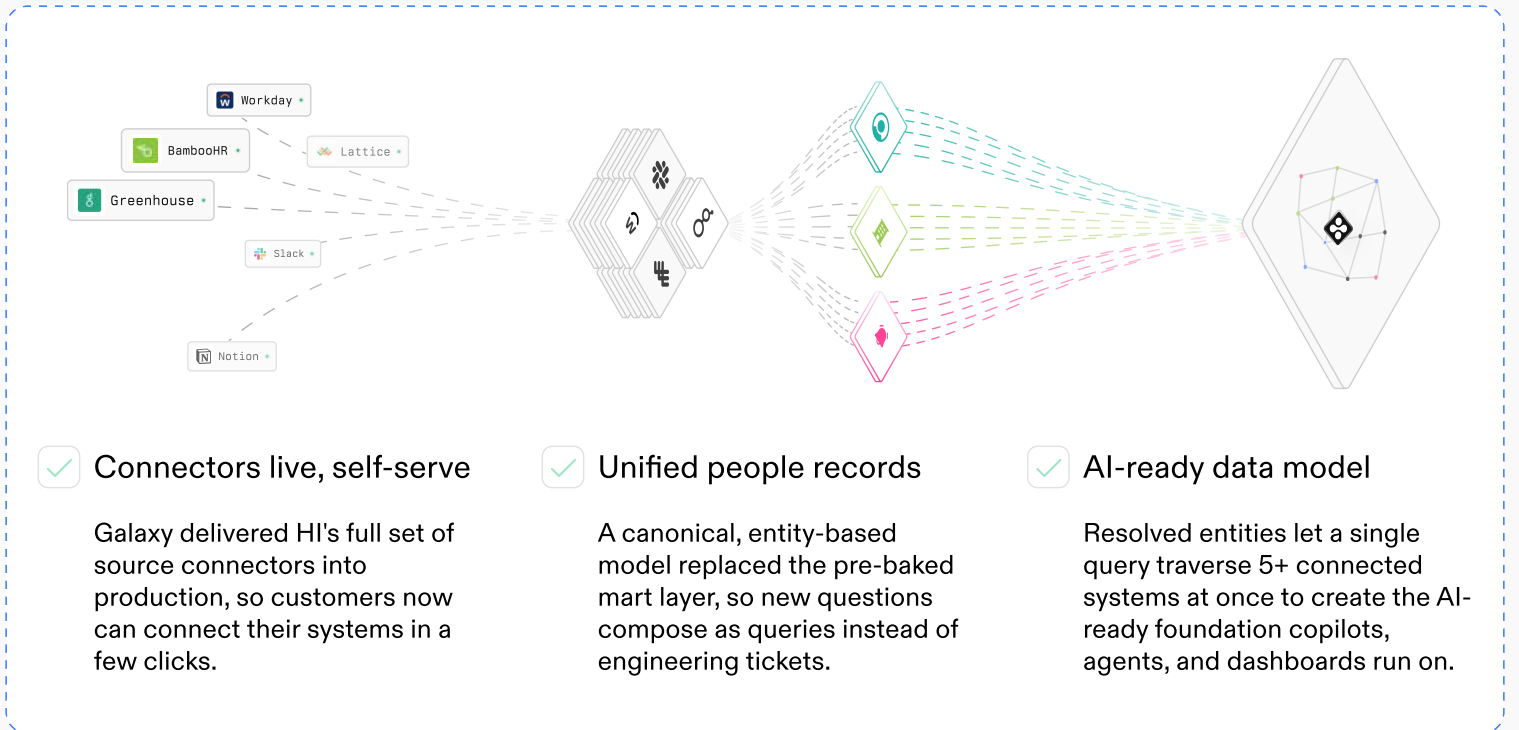
TRUST Metric governance with per-organization labels, definitions and field-level rules — HI's most defensible surface. Tag-based permissions let customers control who sees what data, and a full audit log plus LLM observability (Langfuse) covers every resource mutation.



Result



Reliable data infrastructure secure by default and configurable to every customer.



✓ Connectors live, self-serve

Galaxy delivered HI's full set of source connectors into production, so customers now can connect their systems in a few clicks.

✓ Unified people records

A canonical, entity-based model replaced the pre-baked mart layer, so new questions compose as queries instead of engineering tickets.

✓ AI-ready data model

Resolved entities let a single query traverse 5+ connected systems at once to create the AI-ready foundation copilots, agents, and dashboards run on.

Galaxy as your competitive advantage

By embedding as a forward-deployed team, Galaxy shipped to production daily across every layer of HI's stack: data, AI, and trust. Each round of customer feedback folded straight into the next iteration, so HI validated real usage in live environments instead of waiting on long build cycles. The result is a platform customers can extend themselves: new connectors in days, governed metrics they define on their own terms, and AI features running against production data.

+5

New customers live in less than 6 weeks

< 1 week

Time to ship a new source connector

+\$500k

Committed revenue from completed customer POCs



"The Galaxy team shipped alongside us every day and owned problems end to end. We got working software in front of customers in weeks, and we came out of it with a platform our own customers can extend themselves."

Peter Johnston
Founder & CEO, Human Intelligence



AI-readiness starts with *data readiness*

Custom-built for your stack, not the market.

`hello@getgalaxy.io`