

Revolutionizing Customer Service: Building Bricks, Not Crafting Code

An illustrative scenario demonstrating the art of the possible with
Databricks Agent Bricks

Author:

Eddie Edgeworth
CTO, Koantek



December 9, 2025



The Business Case for Agent Bricks

Introduction

Customer service is broken. Fifteen-minute hold times. Agents navigating seven systems for simple answers. Turnover rates pushing 40%. Traditional solutions, including chatbots, DIY frameworks, and months of LangChain development, haven't delivered enterprise-ready results.

The Call Center Challenge

ElectroHub's Critical Business Problem

Current State Pain Points

- TIME** 15-minute average hold times
- PEOPLE** 35% annual agent turnover
- SYSTEMS** 7+ disconnected systems for basic queries
- QUALITY** 45% first-call resolution rate
- COST** \$8.00 cost per interaction

Business Impact

- RATING** 2.8-star customer satisfaction
- REVENUE** \$2.4M annual lost revenue from abandoned calls
- TRAINING** \$6,000 training cost per new agent
- RISK** Regulatory compliance risks from inconsistent responses

THE ULTIMATUM

"Transform customer service in 8 weeks
or outsource the entire operation"

500

LOCATIONS

50,000

MONTHLY INTERACTIONS

8 WEEKS

BLACK FRIDAY APPROACHING

Enter Agent Bricks: Databricks' answer to enterprise AI agents. It transforms the platform's deep technical capabilities, including Unity Catalog governance, MLflow evaluation, Vector Search, and AI Gateway, into wizard-like solutions that teams can configure without coding. Companies can deploy production-ready agents in weeks, achieving 60-80% cost reductions while improving quality by up to 2x.

As we approach Black Friday, let's explore what's possible through an illustrative scenario. This story illustrates how Agent Bricks fundamentally transforms agent development, shifting from months of coding to weeks of configuration, and from ungoverned prototypes to enterprise-ready solutions.



The Scenario: Eight Weeks to Black Friday

Meet ElectroHub, our illustrative electronics retailer with 500 locations and 50,000 monthly customer interactions. Their challenge: transform customer service before Black Friday with an expiring call center contract and social media complaints trending.

Maria, Director of Digital Transformation, observed the current state: agents toggling between seven systems to answer basic questions, twelve minutes to locate a simple order, and inconsistent responses about return policies. The metrics painted the picture—2.8-star satisfaction, 35% monthly turnover, 45% first-call resolution.

The board's ultimatum: fix it in eight weeks or outsource everything.

Maria had a small team of three data professionals and limited AI experience. Traditional approaches would require months and a team of ML engineers. But with Agent Bricks, she saw a different path.

Traditional vs. Agent Bricks: The Fundamental Difference

<u>TRADITIONAL AGENT DEVELOPMENT</u>	<u>AGENT BRICKS APPROACH</u>
Write code for every component	Describe desired outcomes
Manual prompt engineering	Auto-optimization
Build evaluation frameworks	Built-in MLflow evaluation
Implement governance later	Unity Catalog from day one
Stitch together multiple tools	Unified platform experience
6-12 months to production	3-8 weeks to production
Requires ML engineering team	Business teams can build

Discovery: Five Problems, Five Agent Types

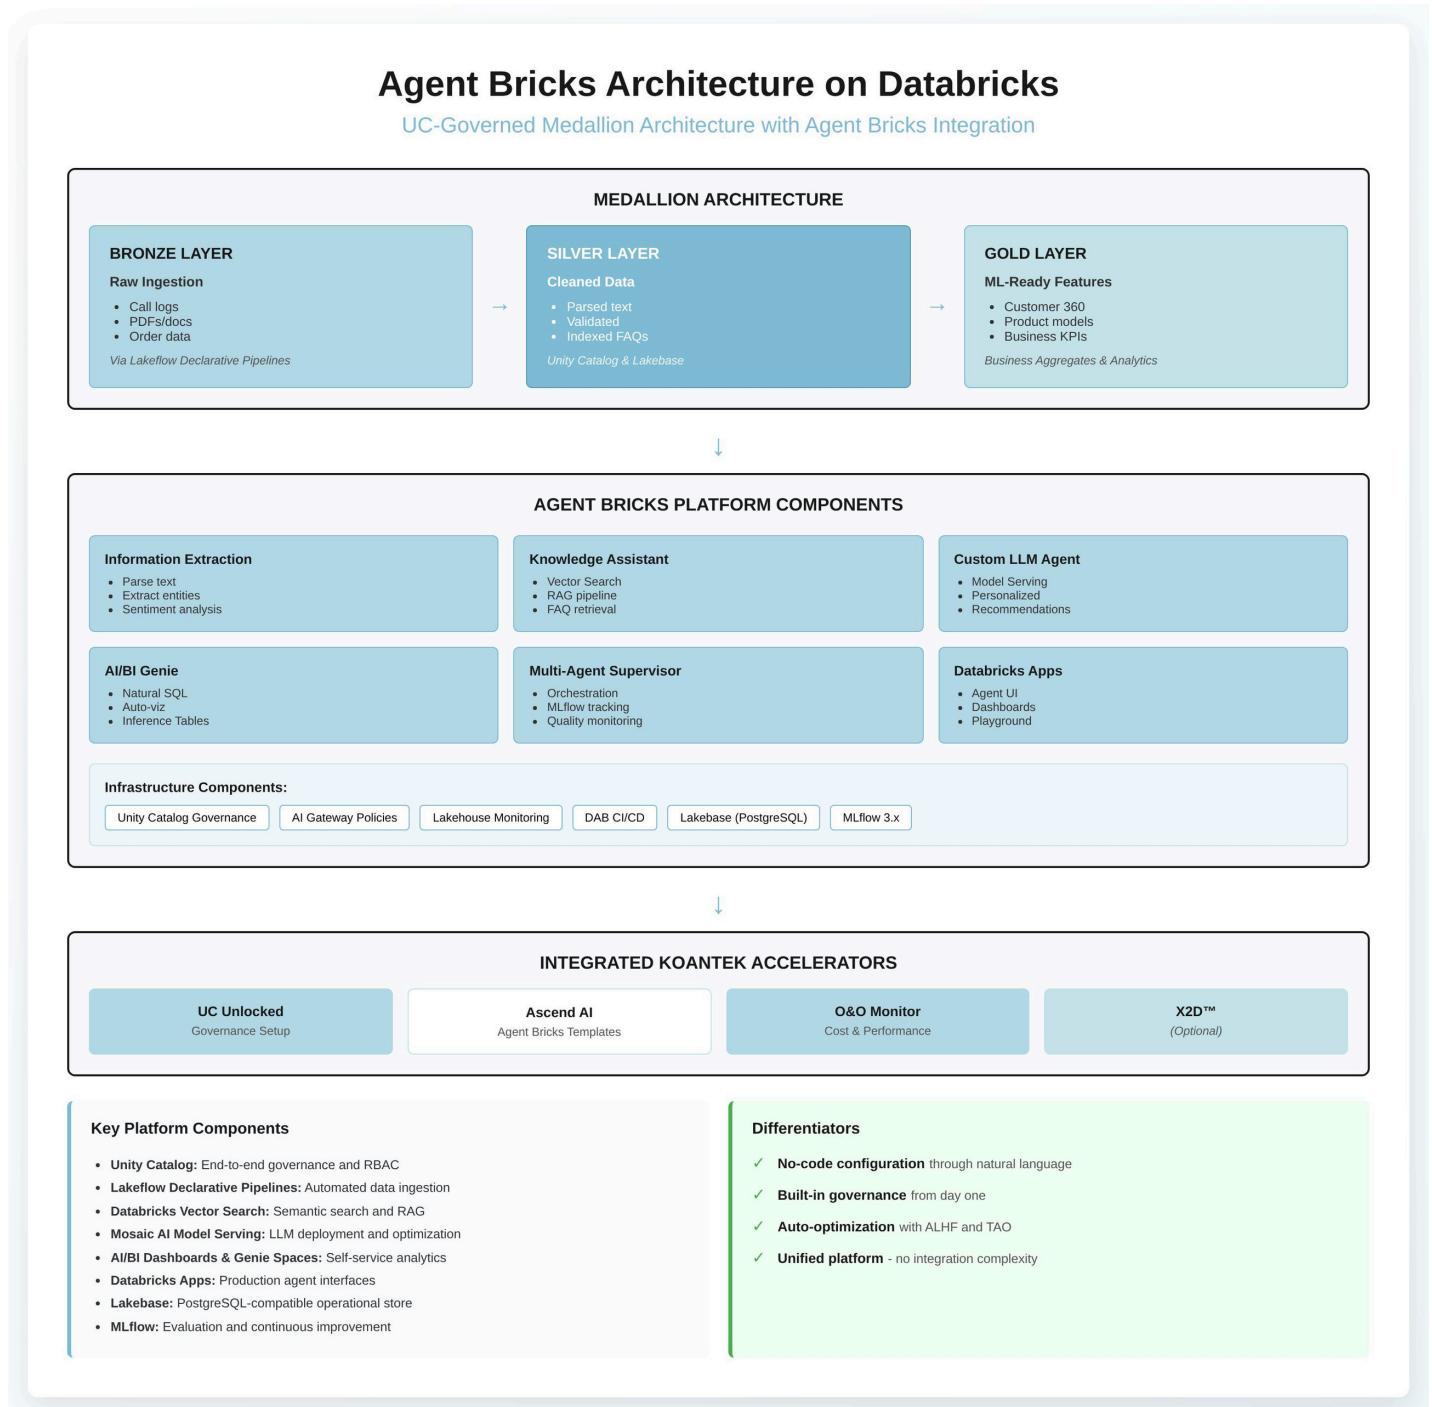
Maria's analysis revealed five distinct call categories, each requiring different capabilities:

1. **"Where's my order?"** (30%) – Database queries across multiple systems
2. **"How do I set this up?"** (25%) – Knowledge retrieval from 10,000+ documents
3. **"Can I return this?"** (20%) – Policy verification and purchase history
4. **"Is this compatible?"** (15%) – Technical specifications and cross-product knowledge
5. **"I have a complaint"** (10%) – Sentiment analysis and escalation



In lieu of a single monolithic chatbot, they realized they needed specialized agents that could be orchestrated together. Agent Bricks makes this possible through simple configuration and natural language descriptions.

Maria's decision was pragmatic: "We'll build one agent at a time, starting immediately."



Building with Agent Bricks: Configuration, Not Code

Day 1: Information Extraction Agent

The Challenge: Customers ramble—"I ordered that black TV thing last week, maybe Tuesday?" Agents waste 5 minutes extracting basic information from unstructured conversations.

The Solution: Maria's team uploaded call transcripts to Unity Catalog and configured the AI Parse Document function. Two-step pipeline: parse unstructured text, extract entities (Order ID, Product Name, Issue Type, Sentiment). No coding required—just configuration.

The agent learned context: whether customers said "Order ABC123" or "the Samsung TV I ordered last week," it extracted the right information.

Results:

- Call categorization accuracy: 92%
- Information extraction time: 5 minutes → 30 seconds
- Automatic priority routing based on sentiment

The screenshot displays the Unity Catalog interface. On the left, a navigation pane shows a tree structure of tables under the 'callcenter_demo' schema, with 'call_transcripts_parsed' selected. The main area shows the 'Sample Data' tab for this table. At the top, there are tabs for Overview, Sample Data, Details, Permissions, Policies, History, Lineage, Insights, and Quality. Below the tabs is a search bar and a 'Preview' button. A question prompt asks 'Ask your question about the sample data...' with three suggested queries: 'How many transcripts have errors?', 'What are common keywords in transcripts?', and 'Any correlation between paths and error_status?'. The main content is a table with two columns: 'path' and 'raw_parsed'. The 'path' column lists various PDF files, and the 'raw_parsed' column shows the corresponding JSON output from the AI parsing process, including fields like 'content', 'description', and 'type'.

path	raw_parsed
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_016_price_match.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_007_delivery_delay.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_018_bundle_discount.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_004_return_upgrade.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_010_printer_setup.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_019_defective_exchange.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_005_where_is_my_order.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_001_rambling_tv_order.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_002_angry_soundbar.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_017_complex_return.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_011_headphones_recommendation.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_003_tablet_recommendation.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_009_budget_tv.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_013_gift_ideas.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_008_wrong_item.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_020_multi_question.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_015_trade_in.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_014_cable_compatibility.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_006_smart_home_setup.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...
dbfs:/Volumes/callcenter_demo/default/call_transcripts/transcript_012_warranty_question.pdf	[{"document":{"elements":[{"bbox":{"coord":{"x":40.25,"y":147.4},"page_id":0},"content":"10/31/25, 3:23 AM","description":null,"id":0,"type":...

Day 2: Knowledge Assistant

The Challenge: Customers ramble—"I ordered that black TV thing last week, maybe Tuesday?" Agents waste 5 minutes extracting basic information from unstructured conversations.

The Solution: Maria's team uploaded call transcripts to Unity Catalog and configured the AI Parse Document function. Two-step pipeline: parse unstructured text, extract entities (Order ID, Product Name, Issue Type, Sentiment). No coding required—just configuration.

For "My soundbar won't connect to my Samsung TV," the assistant identifies the specific model from purchase history, retrieves compatibility notes, and provides targeted instructions with manual references.

Results:

First-call resolution: 45% → 78%

Answer consistency: 95%

Document search: 8 minutes → instant

Agent training: 6 weeks → 2 weeks

The screenshot displays the configuration page for a knowledge assistant named "knowledge-assistant-product-faq". The interface includes sections for "Basic Info", "Configure Knowledge Sources", and "Optional" instructions. On the right, a "Test your Agent" chat window shows a user query: "what are the unique products under each category?". The assistant's response is: "Thank you for your question about unique products under each category. Based on the search results provided, I can identify various product types that are supported, but the documents don't contain a comprehensive categorized breakdown of unique products. The search results primarily contain troubleshooting and support information for different devices." The chat also shows a "Thinking" state and a cost breakdown: "2.88s to first token + 16.80s".



Day 3: Custom LLM Agent

The Challenge: "What's a good TV under \$500?" involves size, features, compatibility, inventory. Agents manually searched 200+ products, often missing optimal matches.

The Solution: Train Custom LLM on product catalog with prices, ratings, specifications. Configure through natural language: "Focus on quality and affordability," "Check inventory," "Consider purchase history for compatibility."

The agent provides nuanced recommendations—not just cheapest options, but thoughtful analysis: "Samsung Tab A8 (\$229): best overall value" vs "Fire HD 10 (\$149): budget-friendly, limited apps."

Quality scores improved from 56% to 69% through automatic optimization.

Results:

- Recommendation accuracy: 87%
- Cross-sell opportunities: +40%
- Customer satisfaction: 4.5/5 stars
- Recommendation time: 5 minutes → 30 seconds

The screenshot displays the Microsoft Copilot Studio interface for configuring a custom LLM agent. The main area shows the agent configuration page for 'product-custom-llm'. On the right, the 'Agent Configuration' panel is open, showing the 'Guidelines' section. The guidelines are as follows:

- Guideline suggestion:** Set clear guidelines to help your agent produce the right output. These will also be used to automatically evaluate quality.
- The response must avoid vague language, promotional tone, and recommendations based solely on popularity or low price.
- Formatting must clearly separate each product recommendation using bullet points or numbered lists for easy reading.
- Each product recommendation must include a specific product name and a concise justification referencing both quality and price.
- Recommendations must focus on products recognized for high quality and affordability, excluding luxury or premium-priced items.
- All recommended products must be currently available for purchase and not discontinued as of the latest market data.
- Recommendations must include at least three distinct brands or models to provide actionable choice for different user preferences.
- Each product recommendation must include a specific model name, release year (if available), and a clearly stated price range.
- Recommendations must avoid products with known reliability issues, poor build quality, or frequent negative feedback regarding
- Each recommended product must have a documented reputation for quality, supported by expert reviews or consumer ratings from
- Each product entry must present accurate model names, specifications, and features in a clear, consistent format suitable for

At the bottom of the panel, there is a 'Save and update' button.



Day 4: AI/BI Genie

The Challenge: "Where's my order?" requires querying three systems with different interfaces. Agents spent 5-7 minutes per lookup.

The Solution: Connect Unity Catalog tables (orders, shipping, inventory) to AI/BI Genie. Natural language queries automatically convert to SQL. "Where's my soundbar ordered last week?" instantly generates the appropriate multi-table query.

The Genie maintains context—when customers follow up with "Will it arrive before Saturday?", it references the previous query. Complex requests like "Show tablet sales trends" generate both data and visualizations automatically.

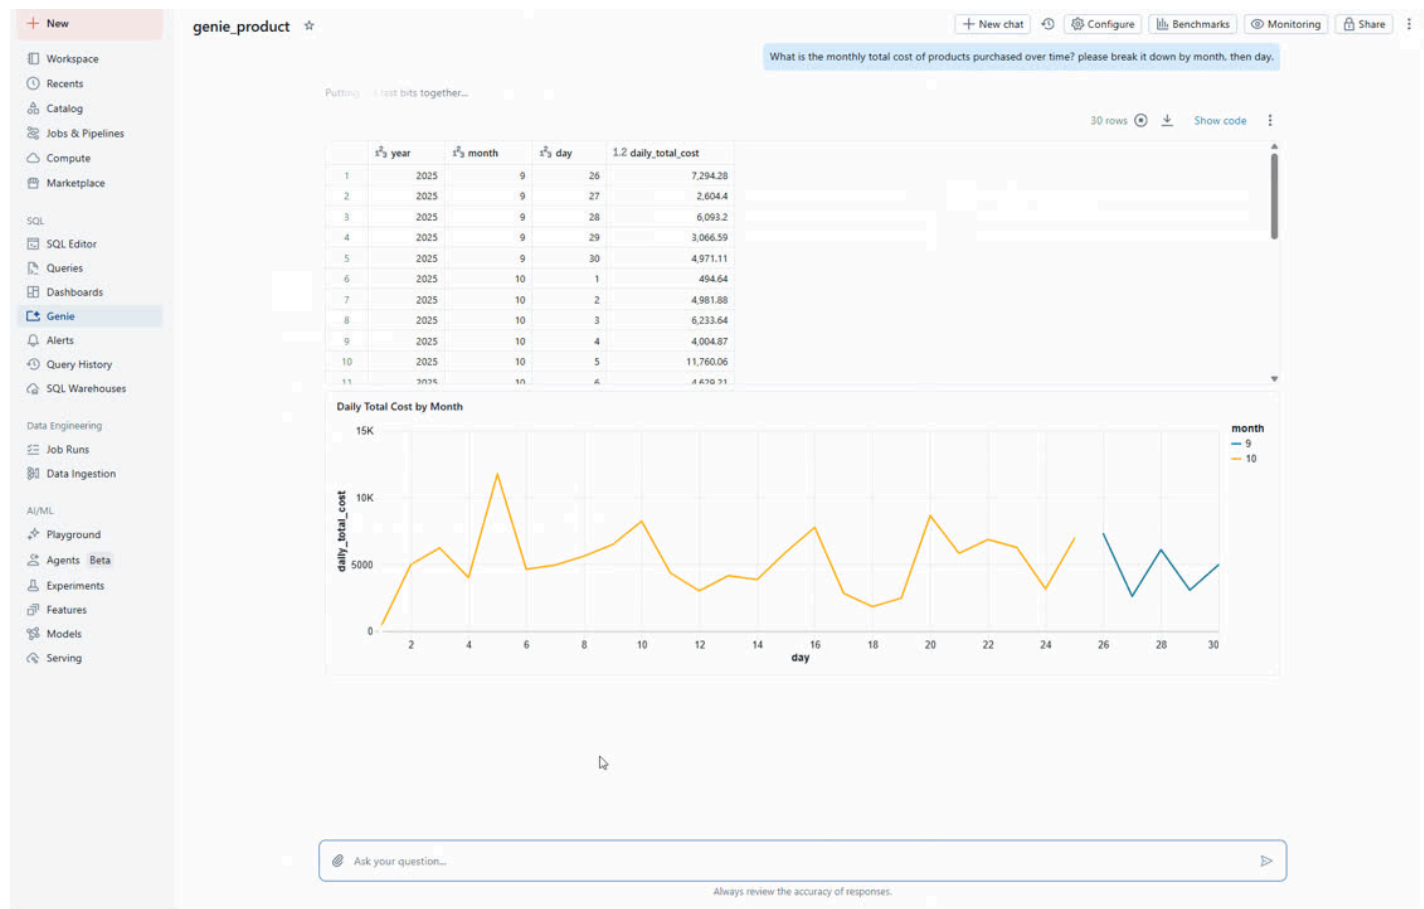
Results:

Order lookup: 5 minutes → 3 seconds

Query accuracy: 94%

Auto-generated charts for trend analysis

Zero SQL training required



Day 5: Multi-Agent Supervisor

The Challenge: Real service requests aren't linear. "My soundbar is broken. I want to return it for the next model up if it's in stock and can be delivered by Friday." This traditionally required three agents and 45 minutes.

The Solution: The Multi-Agent Supervisor orchestrates specialized agents. Configuration defines routing: simple questions to the Knowledge Assistant, data queries to the AI/BI Genie, and recommendations to the Custom LLM. Complex scenarios trigger coordinated responses.

For the broken soundbar scenario, the Supervisor seamlessly orchestrates:

Extract order details (Information Extraction)

Verify return policy (Knowledge Assistant)

Check inventory (AI/BI Genie)

Generate recommendations (Custom LLM)

Process return and new order (AI/BI Genie)

One conversation, under 5 minutes.

Name	Problem type	Endpoint	Last modified
product-custom-llm	Custom LLM	t2t-6882c987-endpoint	October 29, 2025 3:48 PM
multi-agent-callcenter	Multi-Agent Supervisor	mas-75ab8642-endp...	October 27, 2025 12:16 PM
knowledge-assistant-product-faq	Knowledge Assistant	ka-6083e90b-endpoint	October 25, 2025 4:48 PM
customer_extract	Information Extraction	kie-d4cf82b6-endpoint	October 24, 2025 4:03 AM

Results:

Complex resolution: 45 minutes → 8 minutes

Single-interaction resolution: 82%

Agent handoffs: -75%



Enterprise Readiness: Built-In, Not Bolted On

Governance Through Unity Catalog

Agent Bricks leverages Databricks' Unity Catalog for enterprise-grade governance:

- Agents access order history but never credit card numbers
- Role-based access controls are enforced automatically
- Full audit logging for compliance
- PII masking in all logs and analytics

The AI Gateway routes all LLM calls through monitored, optimized channels. Every interaction is tracked, measured, and continuously improved.

Continuous Optimization

Unlike traditional approaches requiring manual tuning, Agent Bricks automatically improves:

- Auto-generated evaluation: Synthetic test data validates edge cases
- MLflow integration: Tracks quality, cost, and latency metrics
- ALHF (Agent Learning from Human Feedback): User feedback drives improvements
- TAO (Test-Adaptive Optimization): Real usage patterns optimize performance

The platform identified cost optimizations that Maria's team would have never found—achieving better accuracy with models 10 times cheaper than the initial configurations.

The screenshot displays the Agent Bricks interface for a request. The request is for "Audio" and the response is "Audio Product Recommendations - Quality & Affordability Focus". The response includes four product recommendations: Audio-Technica ATH-M40x (2014), Sennheiser HD 280 Pro (2003), JBL Flip 6 (2021), and Soundcore Liberty 3 Pro (2021). The Assessments section shows a "product_availability_accuracy" assessment with a "No" result and a detailed rationale.

Inputs

- request: Audio

Outputs

- response: Audio Product Recommendations - Quality & Affordability Focus
 - Audio-Technica ATH-M40x (2014)
 - Price Range: \$99-\$119
 - Justification: Professional-grade studio monitor headphones with excellent build quality and neutral sound signature. Widely praised by audio professionals and consumers for durability and accurate sound reproduction. Offers 90% of the performance of higher-end models at significantly lower cost.
 - Sennheiser HD 280 Pro (2003)
 - Price Range: \$80-\$100
 - Justification: Closed-back dynamic headphones designed for professional use with excellent isolation and sound quality. Consistently rated highly for comfort and durability. Provides studio-quality sound without the premium pricing of flagship models, with over 10 years of proven reliability.
 - JBL Flip 6 (2021)
 - Price Range: \$80-\$130
 - Justification: Portable Bluetooth speaker with IP67 waterproof rating and 12-hour battery life. Delivers surprisingly good sound quality for its size and price, with JBL's reputation for durability. Offers features typically found in \$200+ speakers at significantly lower cost.
 - Soundcore Liberty 3 Pro (2021)
 - Price Range: \$80-\$100
 - Justification: True wireless earbuds with active noise cancellation and 32-hour battery life. Receives consistent 4+ star ratings for sound quality and build reliability. Provides features comparable to \$200+ competitors while maintaining Anker's reputation for quality at affordable pricing.

Assessments

- accurate_strength_limitation_assessment: Yes
- brand_variety_coverage: Yes
- budget_friendly_product_details: Yes
- product_availability_accuracy: No
 - Span: root_span
 - Feedback: No
 - Rationale: The guideline states that all recommended products must be currently available for purchase and not discontinued as of the latest market data. The response includes products like the Audio-Technica ATH-M40x, Sennheiser HD 280 Pro, JBL Flip 6, and Soundcore Liberty 3 Pro. However, the Sennheiser HD 280 Pro was released in 2003, and while it may still be available, the guideline does not specify a requirement for the release date. Therefore, without current market data confirming the availability of each product, we cannot definitively say that all products comply with the guideline. Since the guideline is not fully satisfied due to the lack of confirmation on the current availability of the products, the result is 'no'.

Black Friday: The Results

Eight weeks from inception to production. The timeline:

Weeks 1-2: Information Extraction + AI/BI Genie

Weeks 3-4: Knowledge Assistant + Custom LLM

Weeks 5-6: Multi-Agent Supervisor

Week 7: Load testing

Week 8: Black Friday launch

Black Friday Success Story

50,000 Interactions Handled with Excellence

Performance Transformation

Metric	Before	After	Impact
Resolution Time	12 minutes	3 minutes	75% reduction
Cost per Interaction	\$8.00	\$0.80	90% reduction
First-Call Resolution	45%	78%	73% improvement
Customer Satisfaction	2.8 stars	4.7 stars	68% increase
Agent Turnover	35%	8%	77% reduction

How Agent Bricks Made It Possible

- ✓ 5 specialized agents built in 5 weeks
- ✓ 3-person team (no ML engineers required)
- ✓ Natural language configuration vs months of coding
- ✓ Built-in governance and compliance from Unity Catalog
- ✓ Continuous improvement through MLflow evaluation

Business Impact

- SAVINGS** \$1.8M annual savings in operational costs
- AUTOMATION** 85% of queries fully AI-handled
- RELIABILITY** 99.9% uptime during Black Friday surge
- NPS** Industry-leading NPS improvement

THE KOANTEK DIFFERENCE

"From crisis to competitive advantage in 8 weeks"

Partner with Koantek to Transform Your Customer Experience

Ascend AI Accelerator

Pre-built Agent Bricks templates

Proven Methodology

Fortune 500 deployments

Guaranteed Results

80% automation | 60% TCO reduction

Contact: info@koantek.com | koantek.com/ascend-ai



November 24th: 50,000 interactions flooded in. The results validated everything:

Performance Metrics:

- 85% of queries are fully AI-handled
- Resolution time: 12 minutes → 3 minutes
- Customer satisfaction: 2.8 → 4.7 stars
- Cost per interaction: \$8.00 → \$0.80
- System reliability: 99.9% uptime

Business Transformation:

Metric	Before	After	Impact
Cost per interaction	\$8.00	\$0.80	90% reduction
Agent turnover	35%	8%	77% improvement
First-call resolution	45%	78%	73% improvement
Agent satisfaction	6.2	8.9	44% increase

The human impact proved equally significant. Agents now focused on high-value interactions—complex consultations, relationship building, empathetic service recovery. One agent summarized: "The AI handles the routine. I handle what matters."

The board's response was immediate: "How quickly can we expand this to all customer touchpoints?"

Your Path Forward: From Possibility to Production

This scenario demonstrates what's achievable with Agent Bricks. Maria's team of three built five production agents in eight weeks—not because they were AI experts, but because Agent Bricks transformed Databricks' deep technical capabilities into configurable solutions.

The Platform Advantage

Traditional Approach:

- 6-12 months timeline
- Team of ML engineers
- Custom code for prompts, evaluation, and governance
- High failure risk

Agent Bricks Approach:

- 3-8 week timeline
- Small team, no AI expertise required
- Configuration through natural language
- Built-in enterprise readiness



The difference: Agent Bricks unifies what others stitch together (Unity Catalog governance, MLflow evaluation, Vector Search RAG, AI Gateway optimization) into a single, wizard-like experience.

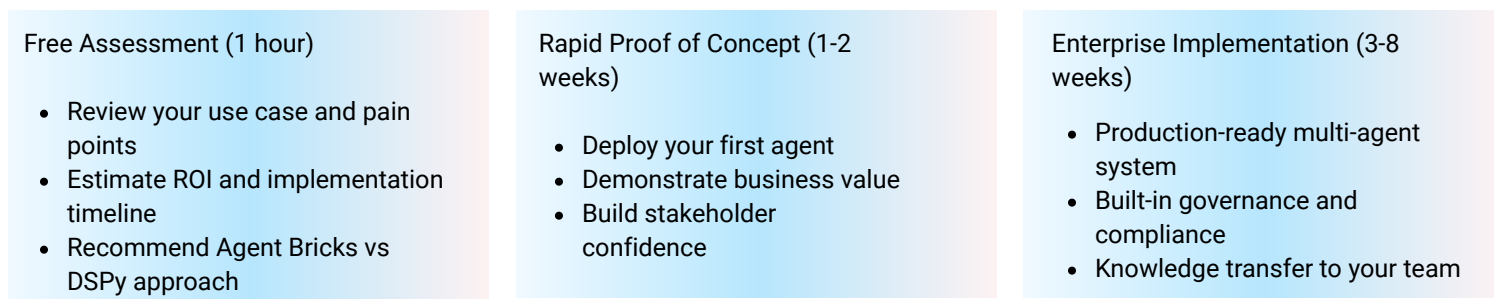
Your Implementation Roadmap



Take Action: Partner with Koantek for Accelerated Success

Your busy season is approaching. While competitors debate approaches, you could be building with proven expertise.

Partner with Koantek—Your Fast Path to Production:



Additional Resources:

Agent Bricks documentation: databricks.com/agent-bricks

Andrew Ng's Agent Bricks workshop for hands-on learning

Databricks Community Slack [#agent-bricks](#)



About Koantek: Your Partner in Enterprise AI Transformation

Koantek is a Databricks Premier Partner specializing in enterprise AI agent implementations and data platform migrations. Our Ascend AI Accelerator framework has helped Fortune 500 companies achieve production-ready AI agents in weeks, not months.

Our Proven Track Record:

- 80% automation of routine customer interactions
- 60% reduction in total cost of ownership
- 90% compliance coverage from day one
- First governed agent deployed in under 3 weeks

Two Paths to Success:

- Agent Bricks Fast Path: Rapid deployment through low-code configuration
- DSPy Code Path: Custom, Pythonic agents for complex requirements

Both approaches leverage unified governance, evaluation frameworks, and enterprise-grade MLOps best practices.

Enterprise-Ready Solutions: We provide pre-built Policy Packs for regulated industries (HIPAA, Basel II/III), ensuring your AI agents meet compliance requirements from day one. Our agents are treated like enterprise software—versioned, evaluated, and promoted behind rigorous quality gates

Ready to Transform Your Customer Service?

Schedule a consultation to discuss your specific use case and see how Koantek's Ascend AI Accelerator can transform your customer experience:

Web: koantek.com/ascend-ai-brickbuilder

Email: sales@koantek.com

LinkedIn: linkedin.com/company/koantek

About the Technology

Agent Bricks is powered by the Databricks Data Intelligence Platform, combining declarative development, automatic optimization, and built-in governance through Unity Catalog. Learn more at databricks.com/agent-bricks.

