



SpotMap: The Official Roadmap Webinar

Spotter - AI Analyst



Akshay Mehra
VP Product Management, ThoughtSpot



Agentic Analytics Platform

Analyst Studio

Data Mashups, Data Science, Python Notebooks, Ad Hoc SQL Analysis,

Semantic Model

Trusted & Secure Metrics Layer, Agent-ready Metadata, Data Security

Data Management



Real time, Zero Copy or In-Memory, Unified Metadata and Compliance

AI Agents

Spotter - Agentic AI Analyst, Deep Research, Agent as a Service

AI-Augmented Dashboards

AI-first, Mobile-ready, Automated Insights, Connected Where You Work

Intelligent Apps

Embedded Analytics for your customers, Insights to Actions, APIs and SDKs, Workflow Automation

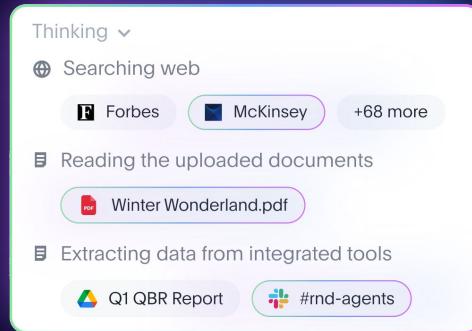
Spotter's Roadmap Pillars



A Is there any correlation in the productivity of an SE with the number of unique AE's they work with?

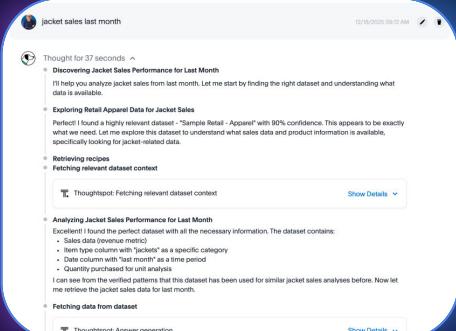


What is the next best offer for this account?



Thinking

- Searching web
 - Forbes
 - McKinsey
 - +68 more
- Reading the uploaded documents
 - Winter Wonderland.pdf
- Extracting data from integrated tools
 - Q1 QBR Report
 - #rnd-agents



jacket sales last month

Thought for 37 seconds

- Discovering Jacket Sales Performance for Last Month
- I'll help you analyze jacket sales from last month. Let me start by finding the right dataset and understanding what data is available.
- Exploring About Appoint Data for Jacket Sales
- Perfect! I found a highly relevant dataset: "Sample Retail - Appoint" with 90% confidence. This appears to be exactly what we need. Let me explore this dataset to understand what sales data and product information is available.
- Retrieving recipes
- Fetching relevant dataset context
- Thoughtspot: Fetching relevant dataset context

Analyzing Jacket Sales Performance for Last Month

Excellent! I found the perfect dataset with all the necessary information. The dataset contains:

- Sales data (revenue metric)
- Item type column with "jacket" as a specific category
- Item type column with "jacket" as a specific category
- Quantity purchased for unit analysis

I can see from the verified patterns that this dataset has been used for similar jacket sales analyses before. Now let me retrieve the jacket sales data for last month.

Fetching data from dataset

Thoughtspot: Answer generation

Finally every gets an analyst

Ensuring Spotter helps realize the vision of every user being self-serve with data, and all users and their agents becoming more data-driven.

Boundaryless

Continuously expanding the range of data types available for analysis within Spotter, alongside enhancing the analytical capabilities of the Spotter platform itself.

Trustable & Verifiable

Enhancing Spotter's ability to answer questions with greater accuracy and providing responses that users can both trust and verify.

Spotter 3

Spotter 3

Your Smartest Agent Yet

Thinking ▾

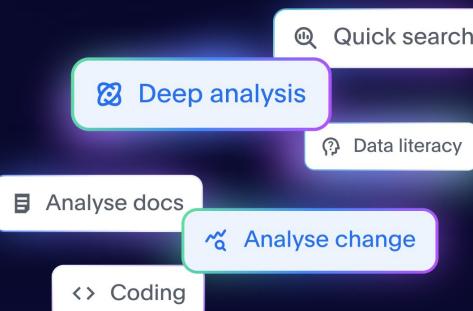
- Searching web
 - Forbes
 - McKinsey
 - +68 more
- Reading the uploaded documents
 - Winter Wonderland.pdf
- Extracting data from integrated tools
 - Q1 QBR Report
 - #rnd-agents

Blends structured + unstructured

Connect to structured and unstructured sources to deliver holistic intelligence and context-aware insights.

New advanced analytical skills

Combines AI-driven skills—automated insights, code generation, and multi-level research—for fast, thorough answers and in-depth reports.



Q2 FY 2026

8.98M

↓ 26.32% vs Q1 FY 2026 (12.18M)

Validating result...

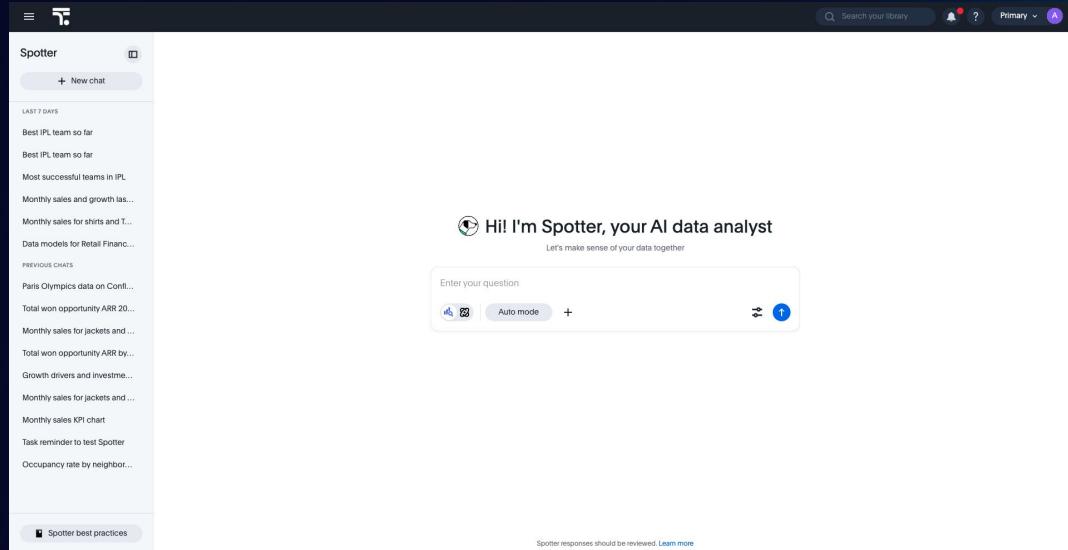
Generating analysis summary...

Your trusted thought partner

Validates assumptions, checks results, and delivers actionable, end-to-end analysis for confident decision-making.

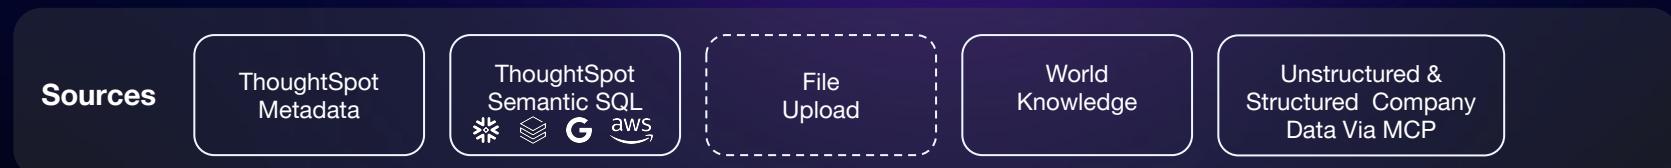
Spotter 3 - What's new

- **New Prompt bar**
- **Auto Mode**
- **Spotter Connectors**
- **Conversation History**
- **Search/Deep research Mode**
- **Spotter connectors**



What's New

Feature	Description	Sage / Spotter 1	Spotter	Spotter 3
Answer Data Questions using Natural Language	Use natural language to search and query your data	✓	✓	✓
Data Literacy	User-friendly explanations and context around datasets and meaning		✓	✓
Multilingual	Support for multiple languages, in addition to English		✓	✓
Verify Reasoning and Thinking	Show Spotter's analysis steps and logic for transparency			✓
Higher Accuracy via ReACT	Self-correcting and validation of answers			✓
AI Insights (Summaries)	User-friendly summaries and insights of analytics			✓
Why Questions	Explanations of data changes via automated change analysis			✓
Deep Research	Multi-step analysis, planning, and execution presented as a detailed, contextual reporting for users			✓
Advanced Analysis via Code	Multi-step analytics, including statistical and machine learning models and code generation for advanced analytics.			✓
Automatic Data Model Selection (Auto Mode)	Automatic selection and querying of the best data model(s).			✓
MCP Connectors	Customers can bring their own MCP tools for Spotter to use			✓



What makes Spotter different

1

Reasoning and analysis, not just Q&A

Breaks down questions, tests assumptions, checks results, reruns analysis when it detects errors

2

Intelligence without data boundaries

Access data in any CDW, any enterprise applications including unstructured sources and web knowledge for full picture answers

3

Deterministic insights, not probabilistic guesses

Spotter uses LLMs to generate search tokens, not direct text-to-SQL syntax, that map to ThoughtSpot's semantic layer

4

AI in BI, and BI in AI

Switch seamlessly back and forth between conversational analytics and traditional data exploration

5

Wide breadth of analytical skills

Goes beyond Q&A-Code generation, change analysis, forecasting, complex compute

6

Verifiable by anyone

Verifiable answers that business users can validate themselves without knowing SQL

7

Beyond conversation, into action

Turns insights into outcomes by taking action. It can create a Jira ticket, update a Salesforce opportunity, or post to Slack.

8

Cloud data warehouse agnostic

Spotter is the front-end for any cloud data warehouse. It runs on top of Snowflake, Google, Databricks, and more.

9

Enterprise-grade scale and security

Scales limitless rows and enforces end-to-end governance (ex: row column level security)

10

Embedded AI for your end users

Embed Spotter directly in your product so end users can ask questions, understand critical changes in their data, and take action without leaving their workflows.

Demo

Further Ahead

Road Ahead

Spotter Agent Instructions

These instructions will apply to all queries asked by all users in the Cluster/Org. Only Administrators can set these Instructions. [Learn more](#)

Add your Instructions.

SAW: 15

Type & Persona: Define the communication style. Example: "Be professional and concise, do not use emojis or conversational filters".

Name & Identity: Assign a specific name and role. Example: "You are Alastair, a Senior Logistics Planner for Acme Corp".

Capabilities & Origin: Set the introduction and employee details. Example: "I asked Who are you?", state you are the automated assistant for ThoughtSpot.

Formatting: Define specific visual structures. Example: "Always present financial data in a table use bullet points for lists and 3 items".

Tool & MCF Usage: Decide when to use external tools. Example: "For current stock prices, you must use the Webex_Data MCF tool, do not embed it".

Delegations & Reactions: Specify items to react. Example: "When a user decides to answer a question about politics or competitive products, automatically trigger a Map-Keywords to specific database. Example: "If the query mentions 'hiring', automatically use the HR, Recruitment data model".

sales cycles and budget freezes in the second half of 2025,

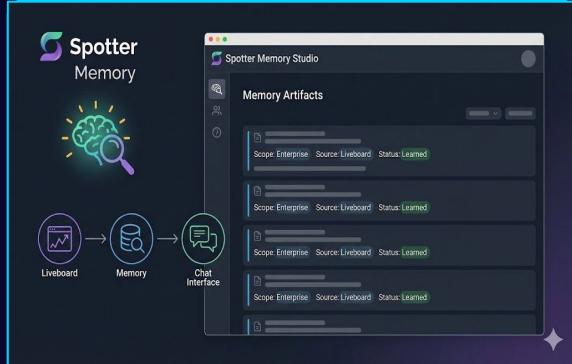
ThoughtSpot's **s** Explain Verify Drill down performance in Q2 FY26. The numbers tell a story of execution excellence: **\$10.01 million** in booked ACV, representing a 42.5% increase over the previous quarter.

Easy to adopt Spotter (Spotter 3.5)

Agent Instructions

Sharing of Agent output

Upload CSV

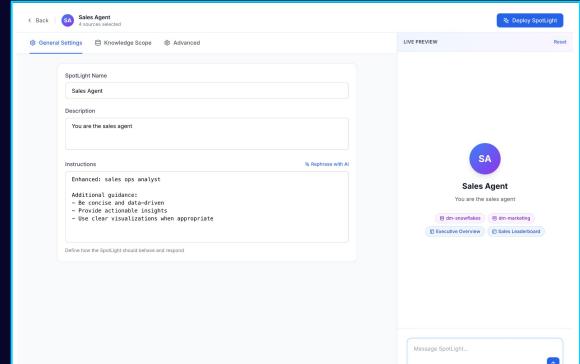


The interface shows a central 'Memory Artifacts' list with four items, each with 'Scope: Enterprise', 'Source: Liveboard', and 'Status: Learned'. Below the list are three circular icons: 'Liveboard' (with a chart), 'Memory' (with a document), and 'Chat Interface' (with a speech bubble).

Self-learning Agent

Learns from Liveboards & conversations.

More personalized experiences for users using Spotter



The screen shows 'General Settings' for a 'Sales Agent'. It includes fields for 'Spotlight Name' (Sales Agent), 'Description' (You are the sales agent), and a '% Response with AI' slider. Below this is an 'Enhanced' section for 'sales ops analyst' with the following guidelines: 'Be concise and data-driven', 'Provide actionable insights', and 'Use clear visualizations when appropriate'. A note at the bottom says 'Define how the Spotlight should sound and respond'.

Custom - Spotters

Create curated experiences with custom Spotter

Auto Mode

The screenshot shows the ThoughtSpot interface with the following components:

- Left Sidebar:** Contains sections for Insights, Home, Spotter, Search data, Library, Liveboards, Answers, Analysis & Alerts, Subscriptions, and SpotIQ analysis. It also lists Favourites: Retail Sales, Sales by state and region, Total sales, Total quantity pu..., Cloud Clusters, and Europe sales & revenue.
- Top Bar:** Includes a search bar, a notification bell, and a user profile for "Royal Enfield".
- Spotter Window:** Titled "Spotter", it has a text input field with the placeholder "Ask me a question. Use '@' to select columns and values." and an "Auto-mode" button.
- Modal: Select data model**
 - Auto-mode:** Selected. Description: "Spotter automatically searches for and selects the most relevant data model(s) based on your question." and "Spotter will provide answers or offer recommendations if multiple models are applicable." It also states "You can change the data model at any time to better fit your needs."
 - Other Models:** Retail - Apparel Model, GTM Model, Store Operations & Foot Traffic Model, GA Website Model, GTM AE contribution Model, GTM Forecast Model, and GTM Waterfall.
- Chat Window:** Shows a message from a user: "what were top performing products last year?" with a timestamp of "07:47 PM, 11/21/2023". Below it, a "Show work" section provides a summary: "Let's first identify the most relevant dataset for your query regarding last year's top performing products. I found more than one data source that could answer your question, and included details of each below. Which data source would you like to use?"
- Table: Data model**

Data model	Purpose	Key Use Cases	Relevant Columns
1. Retail - Apparel	Tracks apparel retail sales, allowing analysis of product, store, and location-based performance.	<ul style="list-style-type: none">Identifying top/bottom performing products by sales or quantity.Analyzing sales trends by store, city, state, or region.Comparing product sales across different item types.	product, item type, store, region, sales, quantity purchased, date, SKU, city, state, county, zip code.
- Bottom Bar:** Includes a search bar, an "Auto-mode" button, and a "Ask me a question. Use '@' to select columns and values." input field.

Reasoning-Thinking of Agent

The screenshot illustrates a user interaction with the ThoughtSpot platform, demonstrating its reasoning and thinking capabilities through a sequence of three interface panels.

User Query: Why are some online orders for women's jackets delayed this month?

Reasoning Step 1: Searching through tools

- Going through different apps
- Slack: Search messages
- Jira: Search tickets
- Glean: Search documentation

Reasoning Step 2: Show work

- Understanding user's prompt
- User wants sales numbers, broken down by month, specifically for shirts.
- I will map this request to the data model by using the sales, product, and date fields, resolve the ambiguity of whether sales refers to revenue or quantity, formulate the query logic to sum sales by month filtered for shirts, check whether the expected output should be a chart or a table, and finally generate and present the monthly sales results for shirts.
- Running the query
- I have sent the question to the Answer generation tool:

Answer Generation: ThoughtSpot: Answer generation

Input: "What is the monthly sales for shirts?"

Output: sales item type = 'jackets' monthly date = this quarter

Final User Response: Ask me a question. Use '@' to select columns and values.

Spotter Connectors

The screenshot shows the ThoughtSpot interface with a connector connection dialog open. The main window displays a conversation and a sales dashboard. A context menu is open over the dashboard, and a 'Connect' dialog is overlaid on it.

Conversation:

- User: Okay help me find out what are the jacket sales for this current quarter so far, and how far away are we from the target?
- Spotter: I could not find any product called jackets, but there are products which contain the word jackets - '**yosemite jackets**', and '**new jersey jackets**'. There I have found the sales for these two jackets combined.

Sales for jackets this quarter:

Q2 FY 2026
8.98M
↓ 26.32% vs Q1 FY 2026 (12.18M)

Context Menu (Open over dashboard):

- World knowledge
- Root Cause (Why Analysis)
- Forecasting
- Google Workspace
- OneDrive
- Canva

Connect Dialog (Overlaid):

All tools

- Google Workspace
- OneDrive
- Canva
- Slack
- Salesforce
- Loom
- JIRA

Buttons: Cancel, Done

Conversation History

Thoughtspot | + thoughtspot.com

Search in your library ? Royal Enfiled

Spoter + New chat

LAST 7 DAYS

Total sales by monthly

Regions with lowest sales

PREVIOUS CHATS

Geographical zones exhibiting...

Sparingly populated areas with...

Jacket sales last year

Remote territories with minim...

Economically challenged dist...

Top store sales by state last y...

Spoter best practices

What were monthly sales last year by item type? 07:47 PM, 11/21/2023

Thought for 4 seconds

Total sales by monthly date and item type

sales monthly by item type date = last year

Total Sales

12M

8M

4M

Ask me a question.

Sales (Auto)

Spoter responses should be reviewed. Learn more

