



AI Agents

Beyond the prototype

Victor van den Broek

October 8th, 2025, Amsterdam



Introductions



Victor van den Broek

Senior Solutions Architect

17 YoE in Data+AI

Focus areas:

- Public sector / digital sovereignty
- GenAI
- Data+AI governance
- Angel investing



databricks

Cloud Data+AI platform

- > 20.000 customers
- > 4B ARR of which >1B in AI
- > 100B valuation



The rise of AI agents: Transforming work and society

The rise of the AI manager

AI agents will do the grunt work of coding

the high-performing digital teammates, the professionals who will define the future of work

How AI Agents Are Changing the Way We Learn

Agent hype is everywhere

AI agents: from co-pilot to autopilot

Can one person build a billion-dollar company now with AI agents?

The Rise Of The Hybrid Workforce: Humans And AI Working Together

INNOVATION

The Rise Of AI Agents And The Future Of Data

An agent workforce could be on the horizon



Decoding AI ROI Paradox: Why 85% of Projects Die in Production

88% of AI pilots fail to reach production but that's not all on IT

The tough task of making AI code production-ready

Between 70-85% of GenAI deployment efforts are failing to meet their desired ROI

Reasons Why Generative AI Pilots Fail To Move Into Production

CIOs' lack of success metrics dooms many AI projects

But,
agents could
still be
easier



Agentic AI can clearly be very successful and productive



Assistant



Agent Mode

Automate end-to-end analyses and workflows

✦ Generate a sales forecast

✦ Diagnose OKR performance

✦ Summarize and interpret a notebook

Code is *unambiguous*
100% verifiable
Human-in-the-loop
Fast feedback cycles



Beyond code...



General Intelligence fails for enterprise use cases

General Intelligence

Foundation models trained on a broad dataset
disconnected from your business data

vs

Data Intelligence

AI **connected** to your customer data and able to solve domain-specific problems



What is an Agent?

An Agent making actions by retrieving your data and calling your API ... an Agent is a LLM acting if it has a brain 



"Can you help me return my last order?"

"Done! Here is your shipping label..."



I should look up the recent order and our return policies.

1. Look up the user's last order
2. Query our policy documents

The return policy is <...>, order <id> fits within this policy.

Start our return process

I should let the user know I've started their return & send the shipping label.



Common Agent Design Patterns



By now, these can all be procured 'as-a-service'

AI/BI Genie

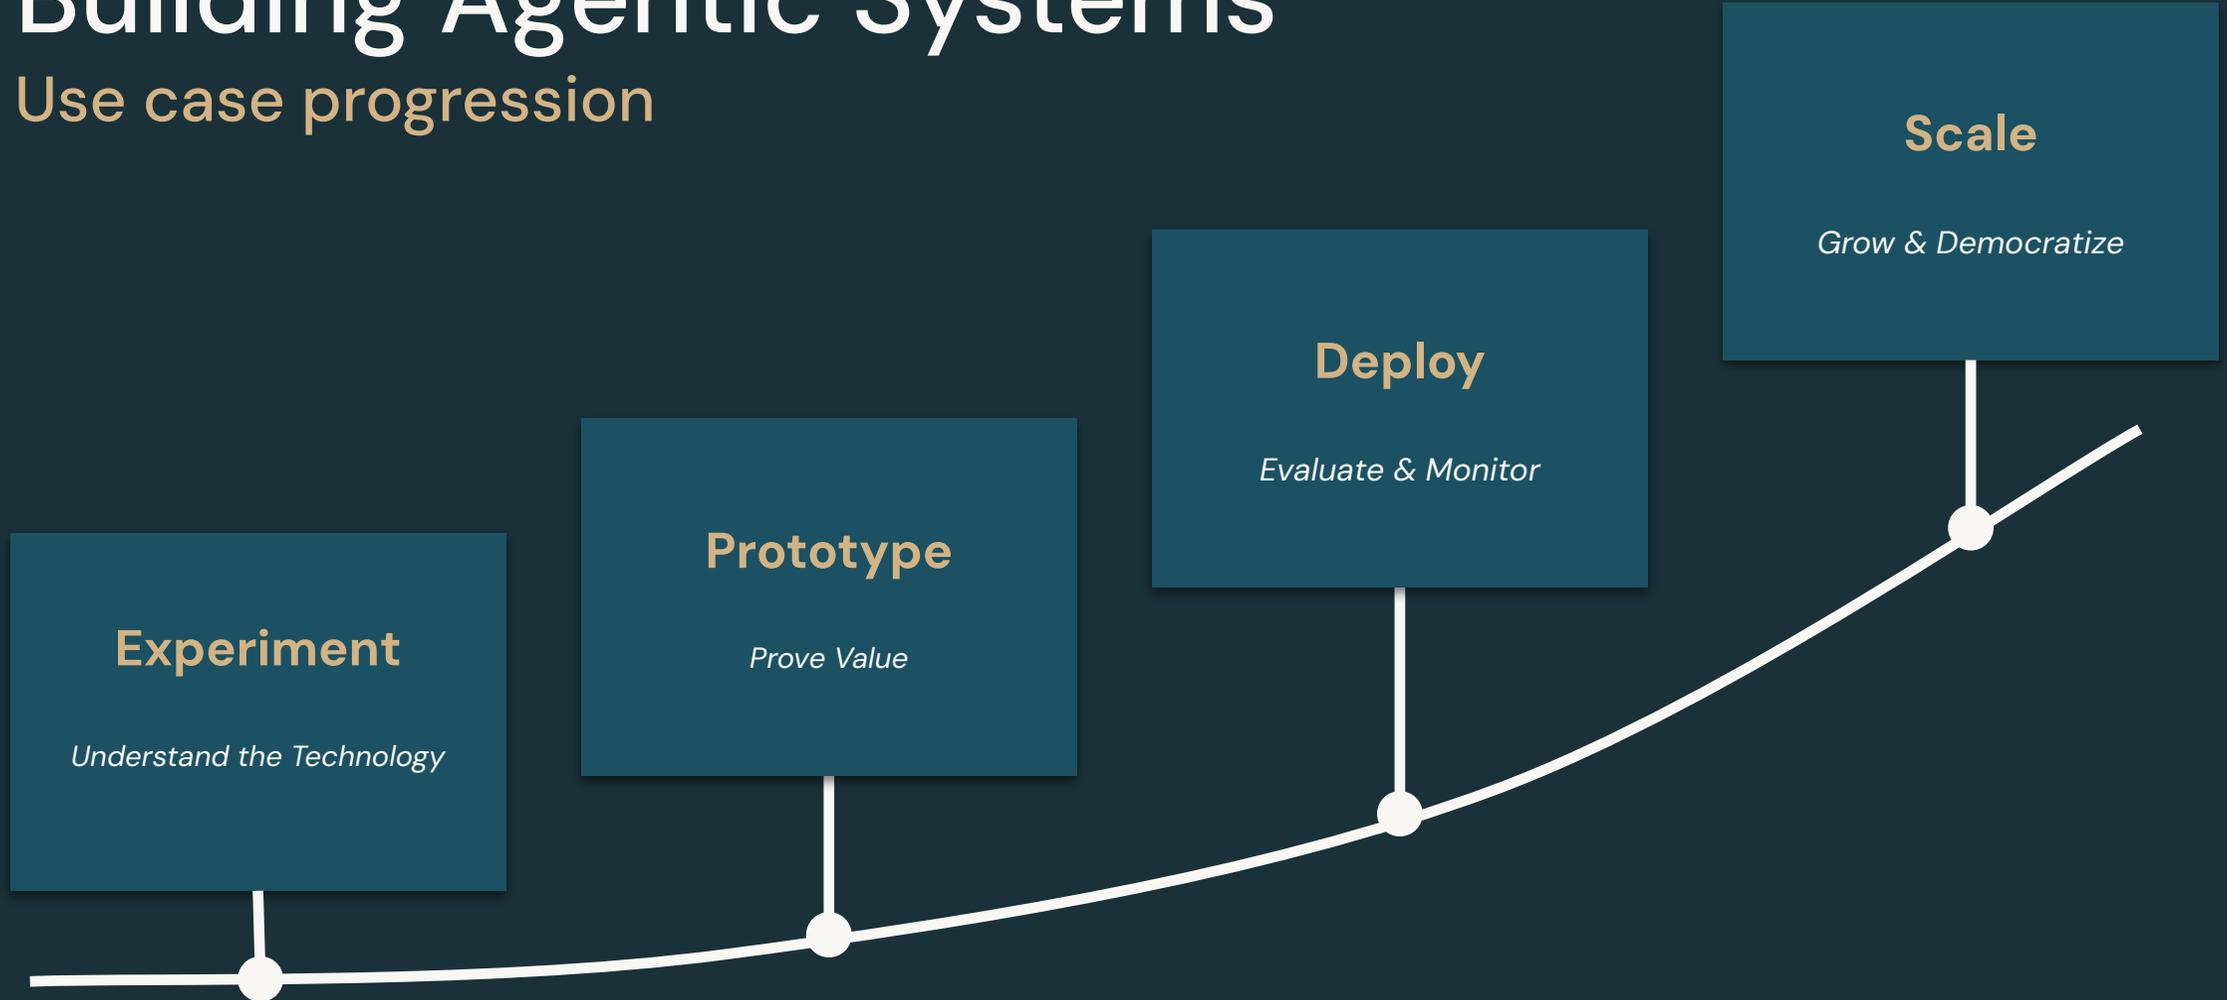


Agent Bricks



Building Agentic Systems

Use case progression



AI Pipeline

Agent Framework

Rapid Development

Evaluation Datasets

Human Feedback

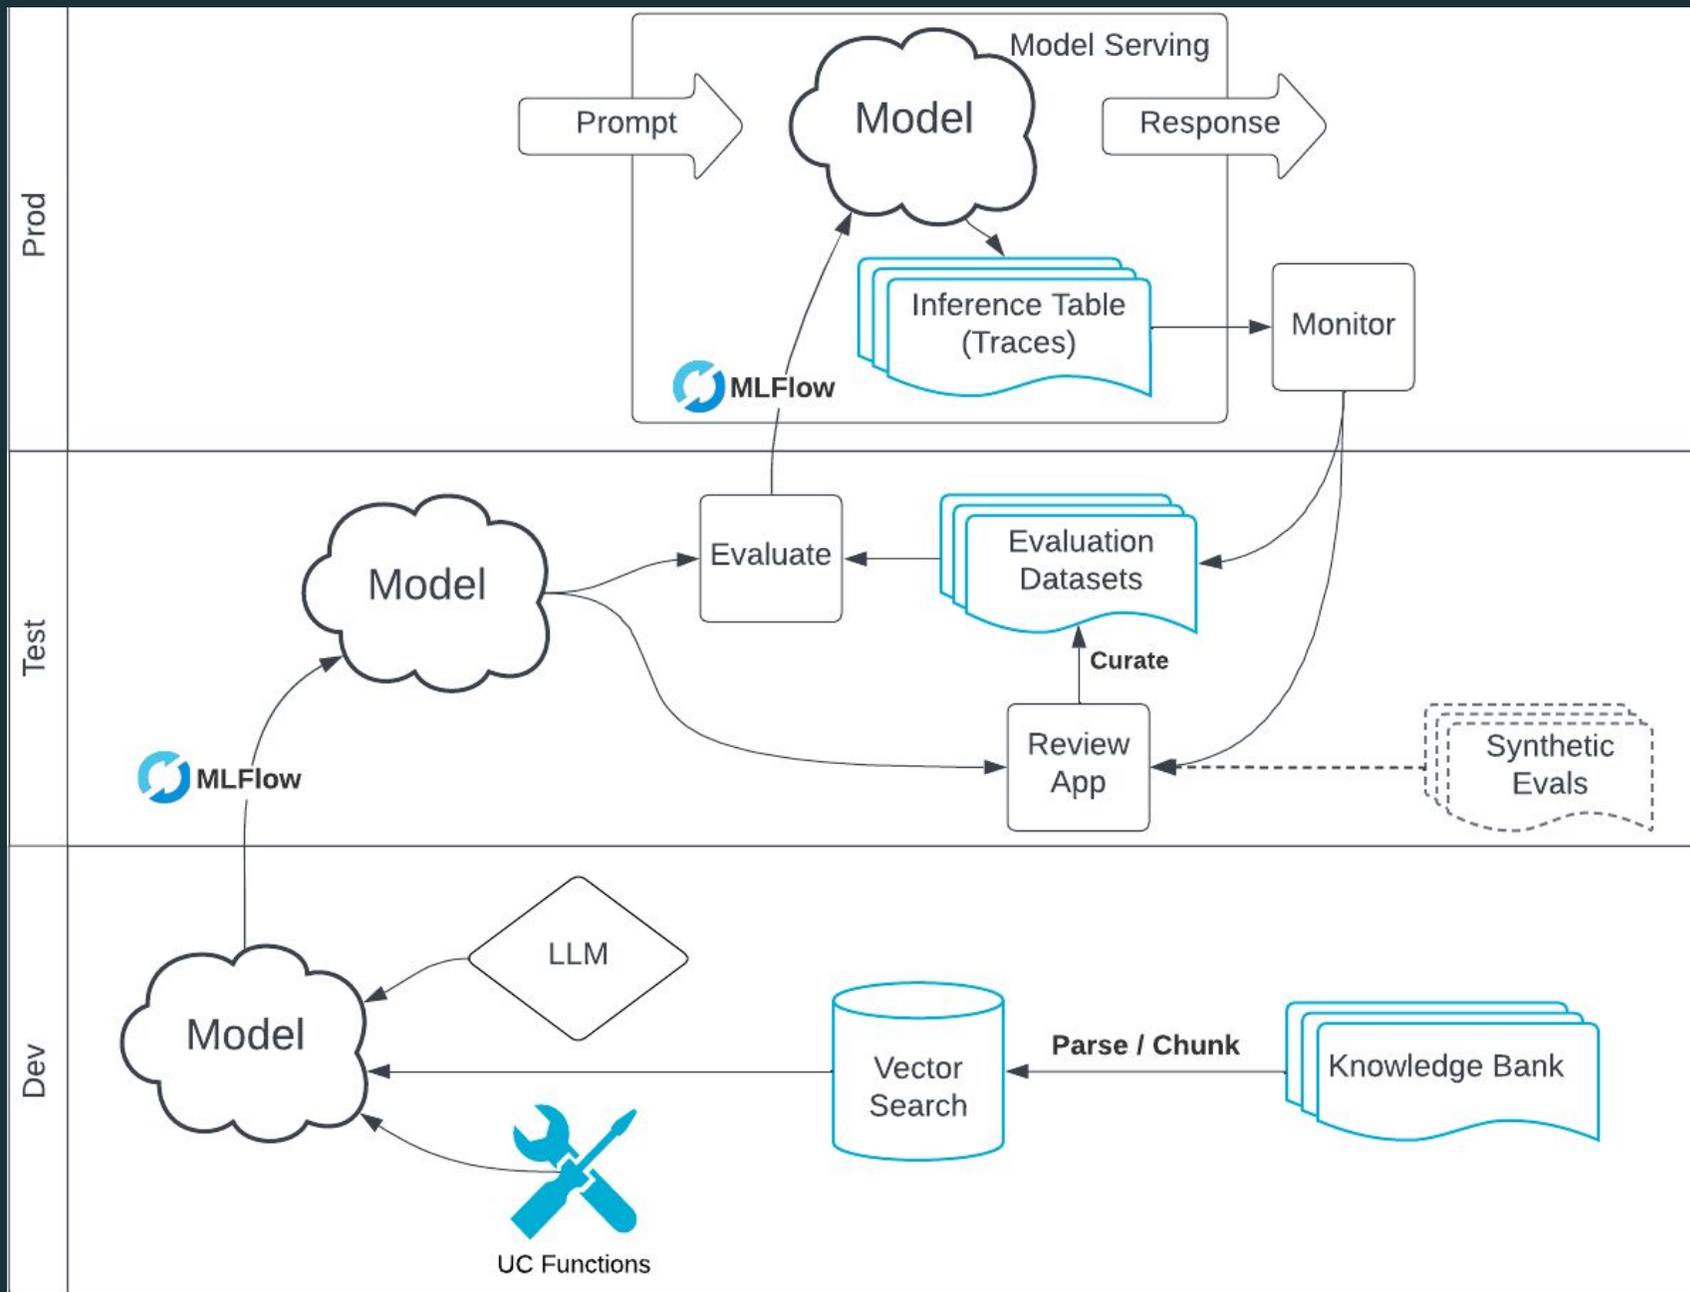
Automated Evaluation

Deployment

A/B Testing

Auditing/Tracing

Monitoring





Prepare data

Data ingestion

ML features

Vector index

Build agents

GenAI models

Classical ML models

Functions & tools

Deploy agents

Agent serving

MLOps/ LLMOps

Lineage

Evaluate agents

LLM judges

Peer labelling

Tracing

Govern agents

AI guardrails

Credentials

Usage tracking

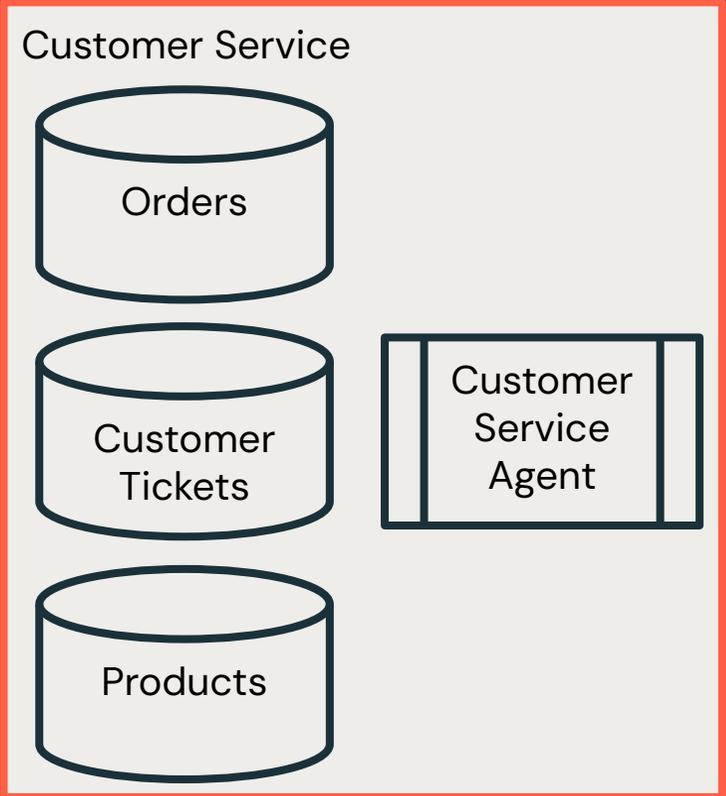
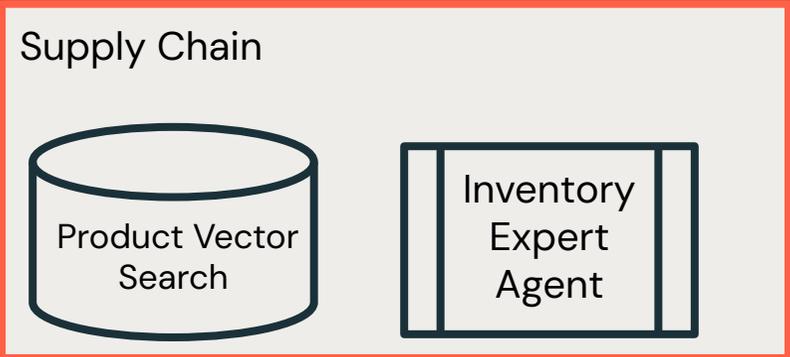
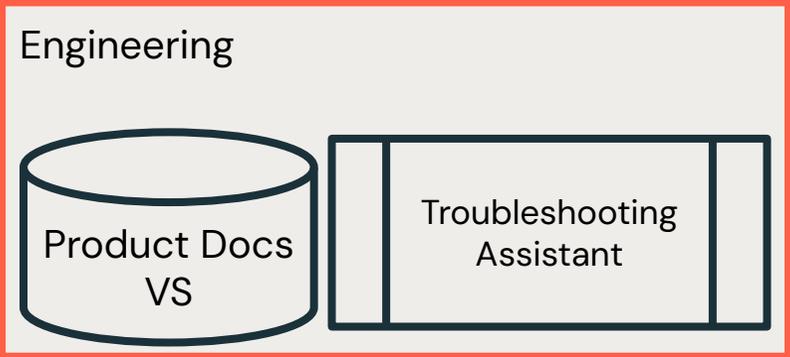
Rate limits

Beyond the single agent...



Governance Risk of Current Methods

Isolated Implementations

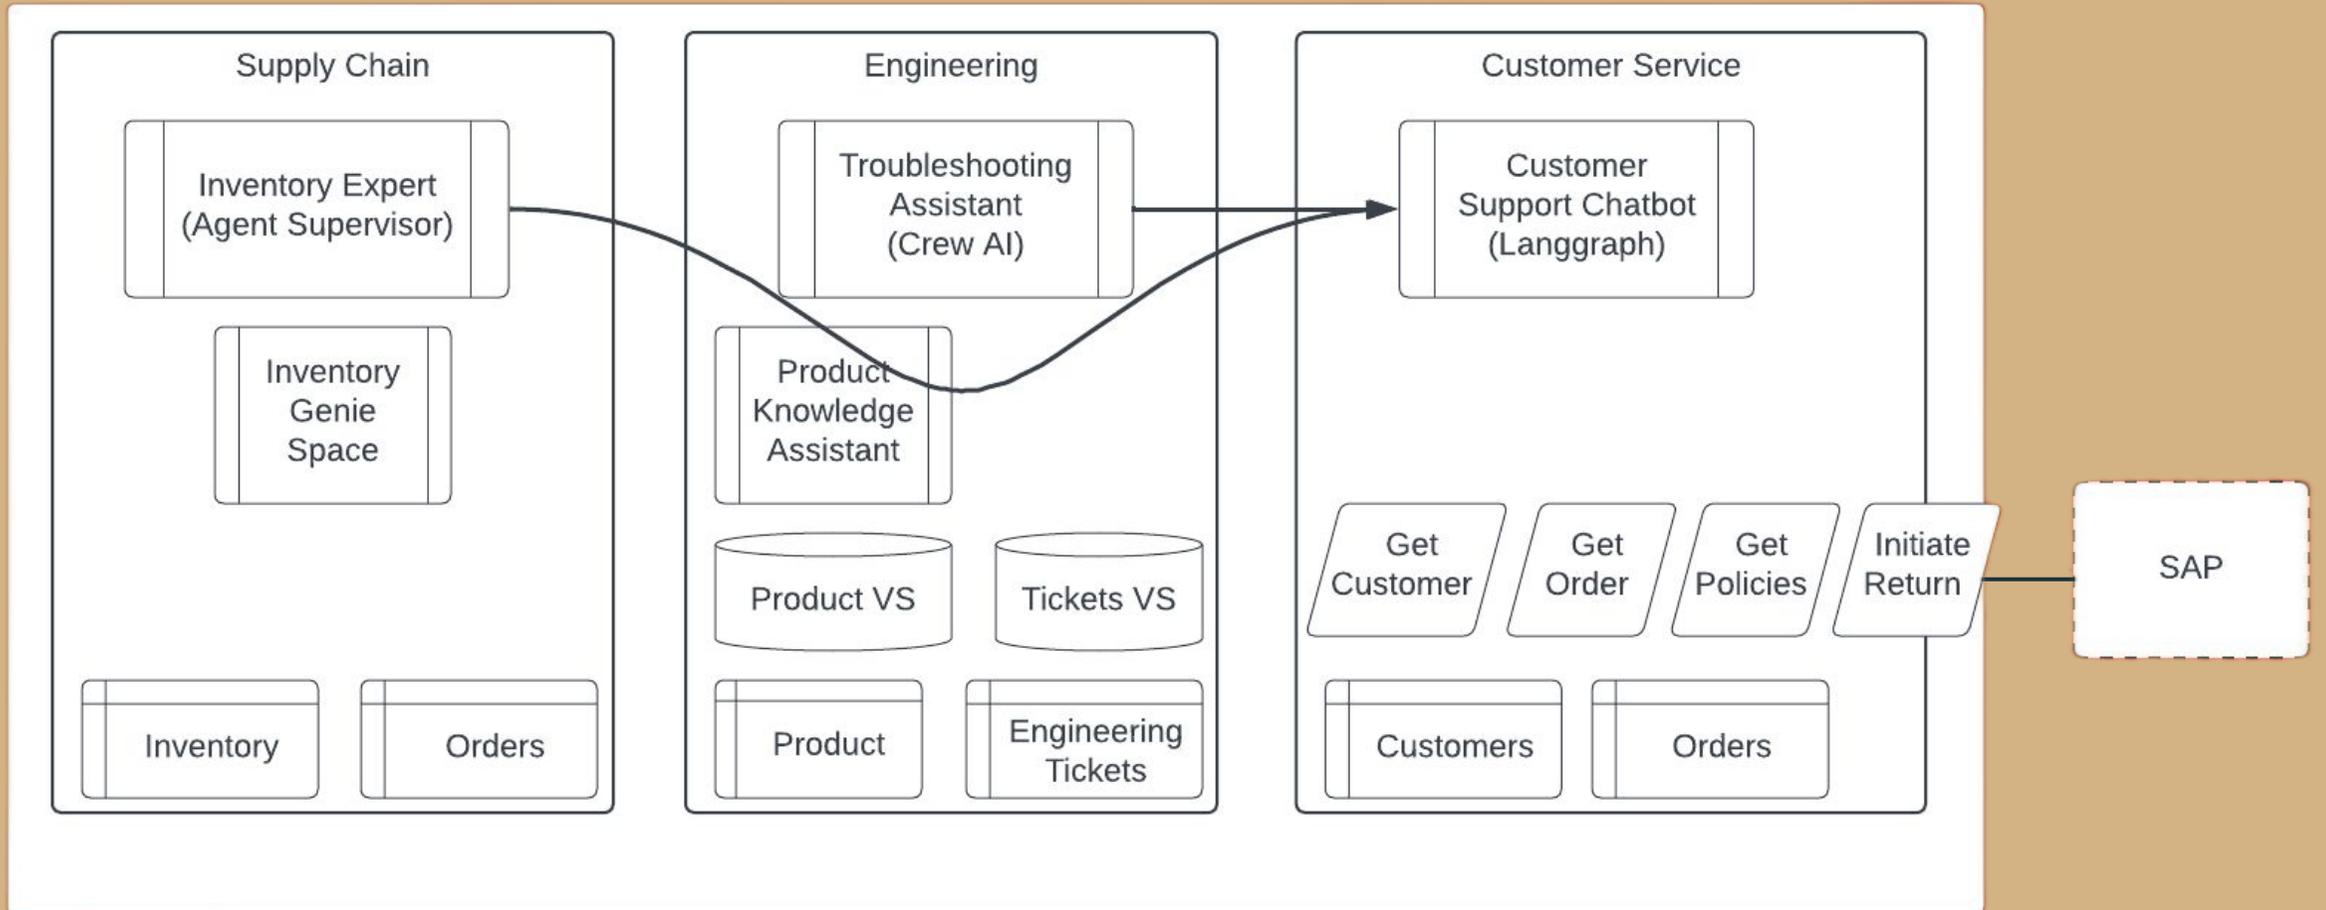


- Duplication of Efforts
- N Sources of Truth
- Change Management



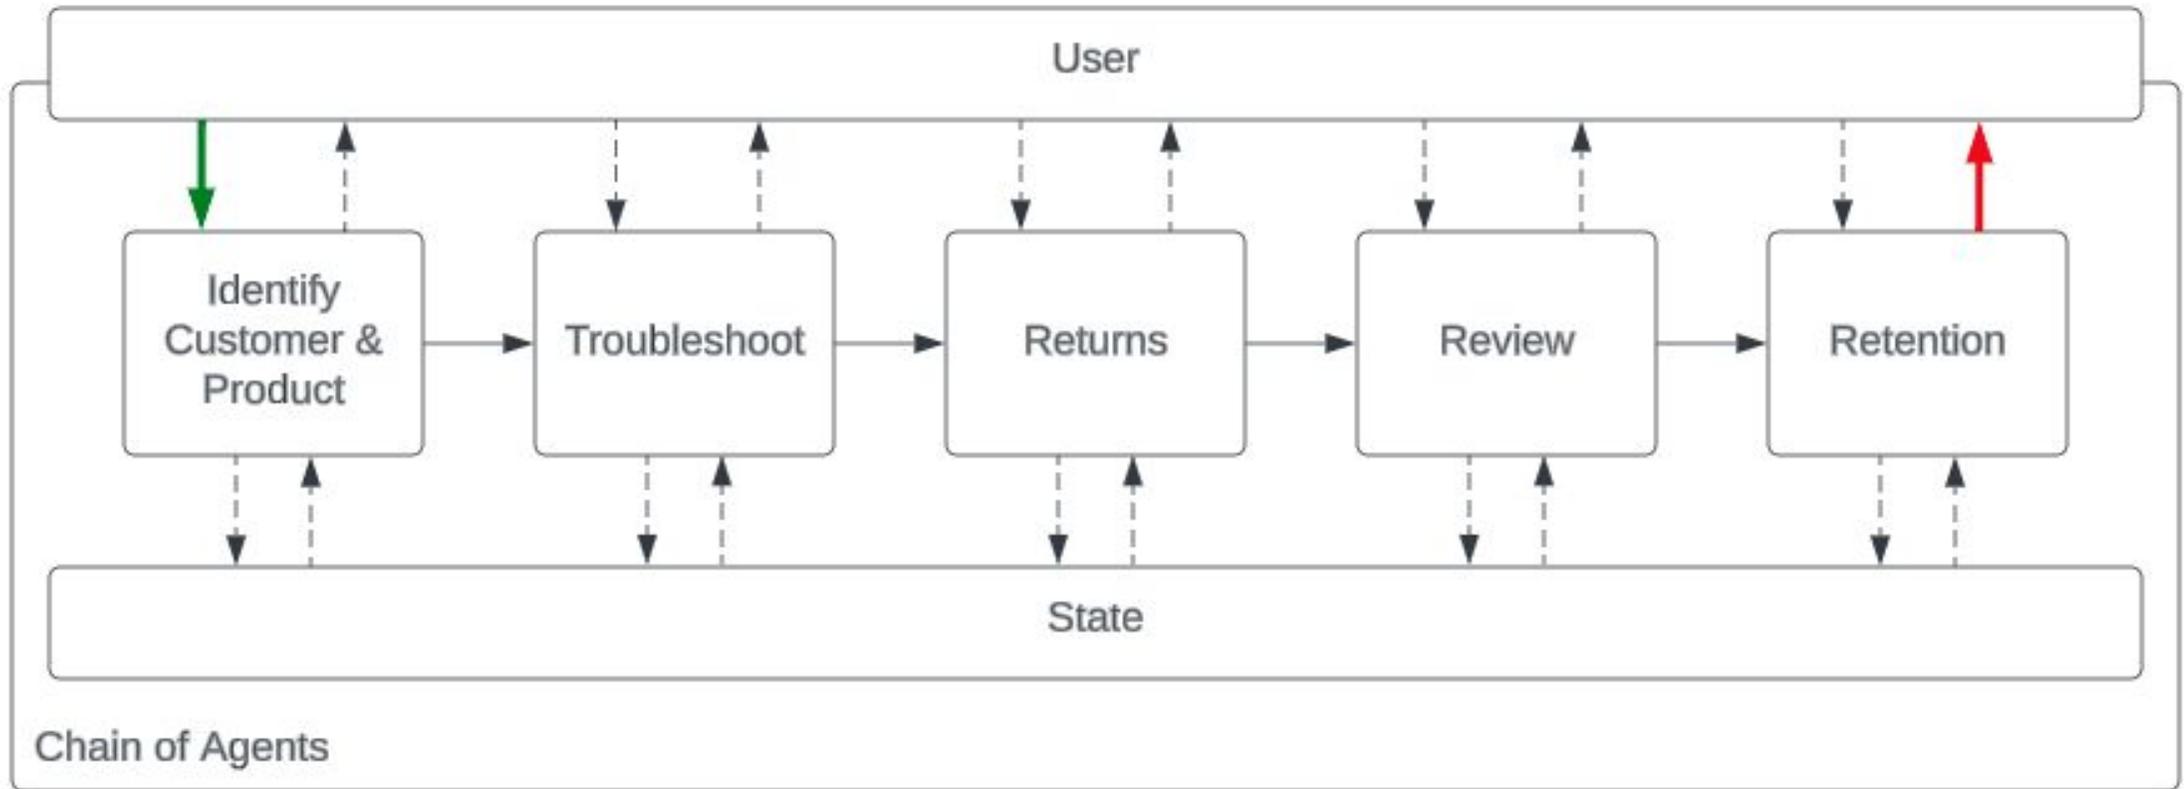
Data and Agent Reuse

Governed, provisioned and secured



Agent Chain

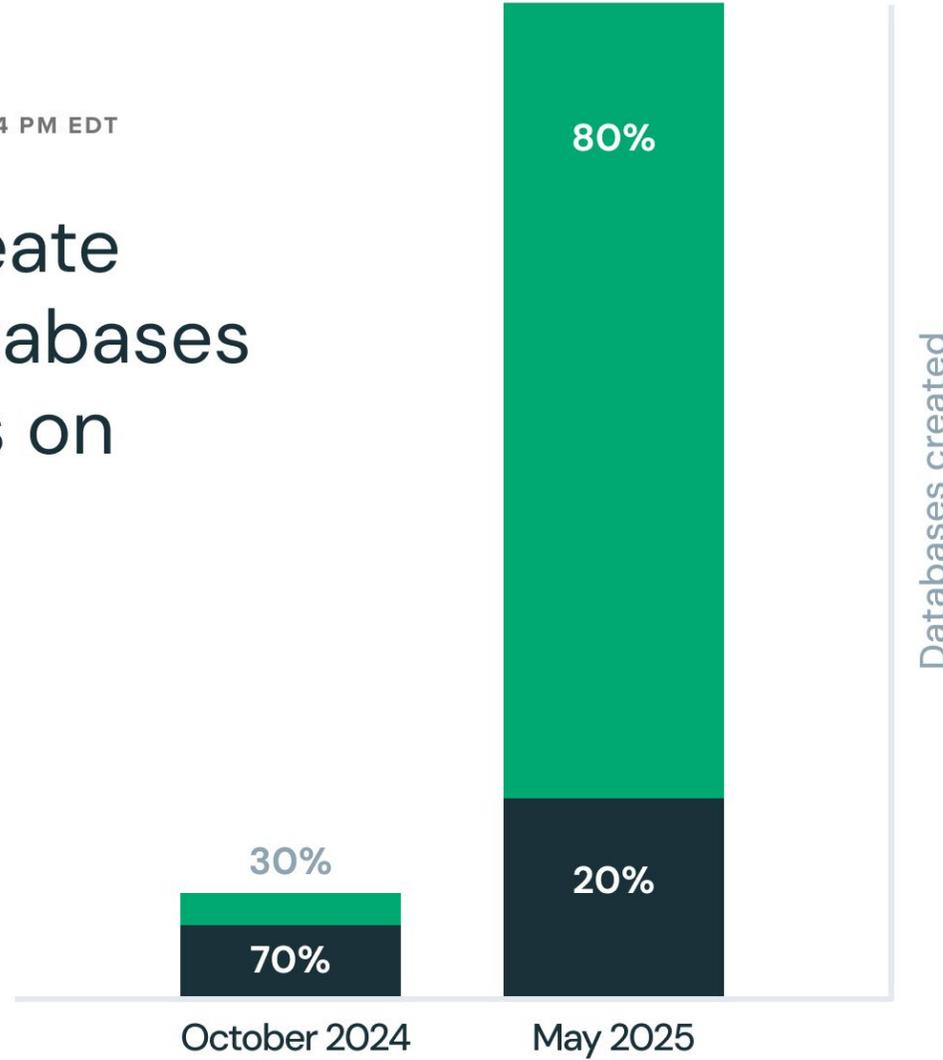
Agentic System



Databricks is buying database startup Neon for about \$1 billion

PUBLISHED WED, MAY 14 2025•6:00 AM EDT | UPDATED WED, MAY 14 2025•2:14 PM EDT

AI agents create **4x** more databases than humans on Neon



Multi-agents, now you need to...

➤ Integrate auth/SSO

➤ Define governance and access controls around data and models

➤ Manage app infrastructure/containers

➤ Integrate into a VPC

➤ Deploy a database to manage agent state

➤ Integrate with observability tools

➤ Configure audit/access logs

➤ Upgrade packages/models



Unified governance is key
to making AI agents
(and multi-agents)
successful at scale

(And you can do all of this on Databricks today)



