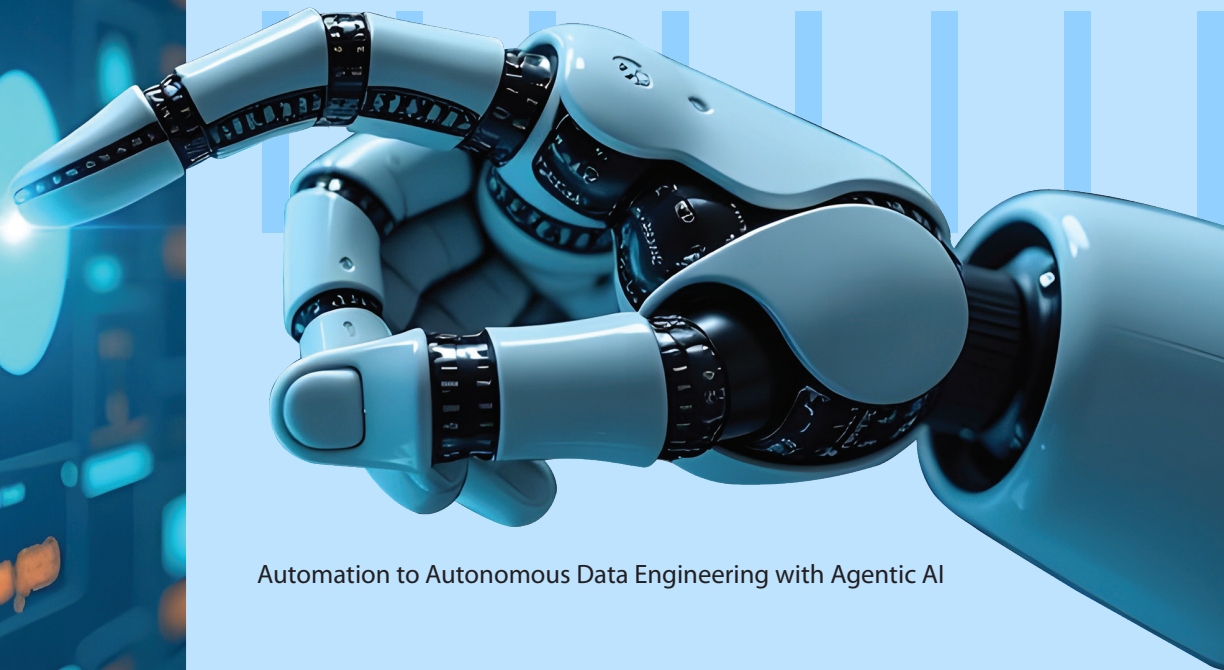




INFOSYS DE.AI - THE AGENTIC LEAP FOR DATA ENGINEERING



Automation to Autonomous Data Engineering with Agentic AI

An Evolution in Autonomy

Just as cars evolved from manual operation to cruise control (automation), driver-assist(assistance), and finally to fully autonomous vehicles (like Waymo), data engineering is on the same journey. We are moving from manually scripting pipelines to a future where autonomous agents manage the entire data lifecycle, from intent to production.

The Evolution of Data Engineering

Data engineering is evolving beyond manual scripting. Agentic AI introduces a new paradigm, moving from simple automation to intelligent, autonomous assistants that manage the entire data lifecycle.



Stage 1: Automation Driven

Rule-based and schedule-driven. Orchestrators (Airflow, cron) run pre-defined scripts. This is "automation," but it's brittle and 100% human-defined.



Stage 2: Assisted Automation

AI "co-pilots" help engineers write code, fix bugs, and draft documentation. This is an "assistant" model, reducing tasks but requiring human oversight.



Stage 3: Agentic for Personas

Specialized agents assist specific roles: an "Architect Agent" suggests designs, a "Test Agent" generates test cases, a "Dev Agent" writes pipeline code.



Stage 4: Completely Autonomous

A system of agents manages the full lifecycle based on human "intent." It can self-heal, self-optimize, and handle new requests from discovery to publishing.



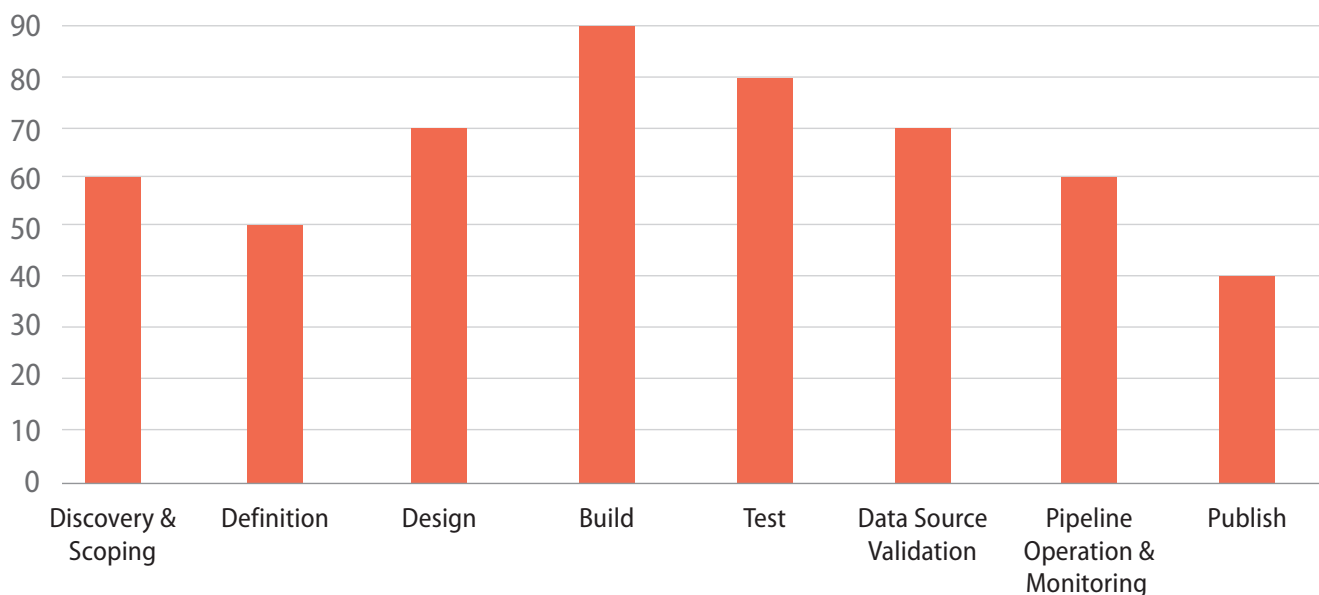
The Bottleneck: Disconnected Data Engineering and Technology Evolution

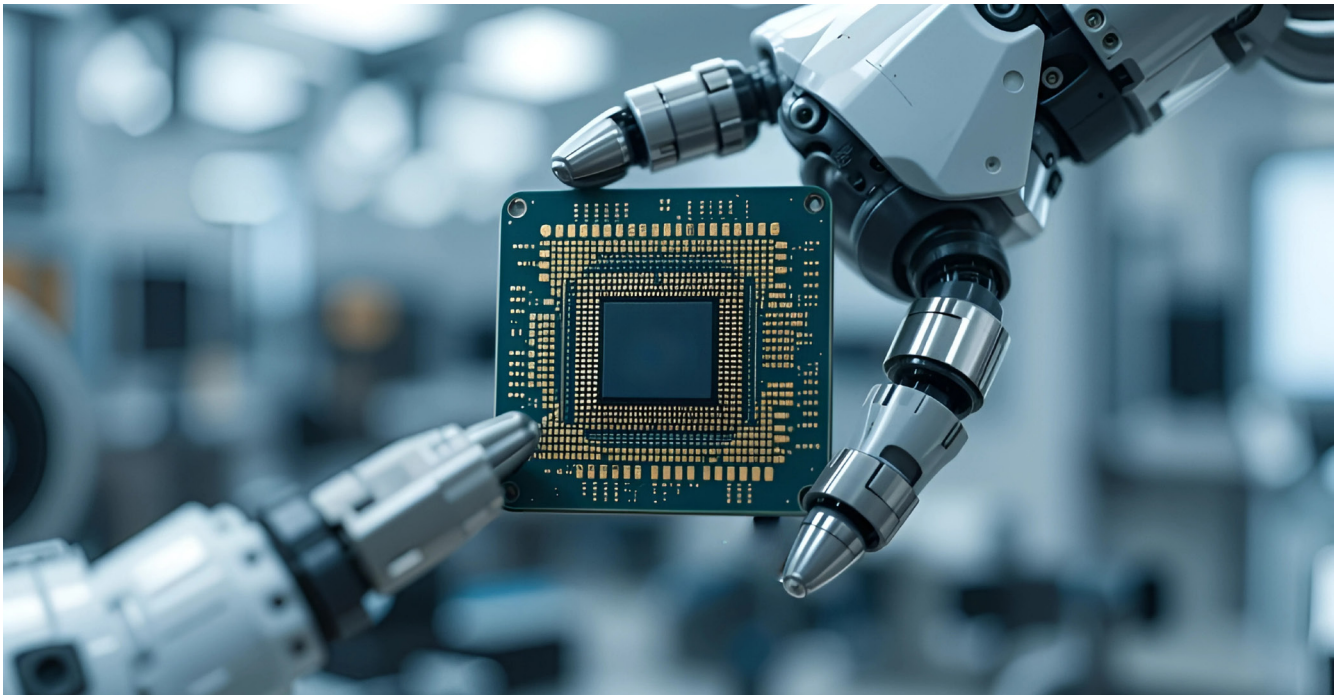
The conventional data engineering lifecycle is linear and labor-intensive. While data *platforms* have continuously evolved—from warehouse appliances to cloud platforms, data platforms for AI, and now unified platforms—the core *process* of engineering data

has remained a manual bottleneck. This disconnect is where value is lost.

- **High Effort:** 'Build' and 'Test' phases are notoriously time-consuming.
- **Reactive Operations:** 'Operate' is often reactive, fixing breaks rather than preempting them.
- **Rigid Design:** 'Design' is a static blueprint that's difficult to change.

Traditional Lifecycle: Manual Effort





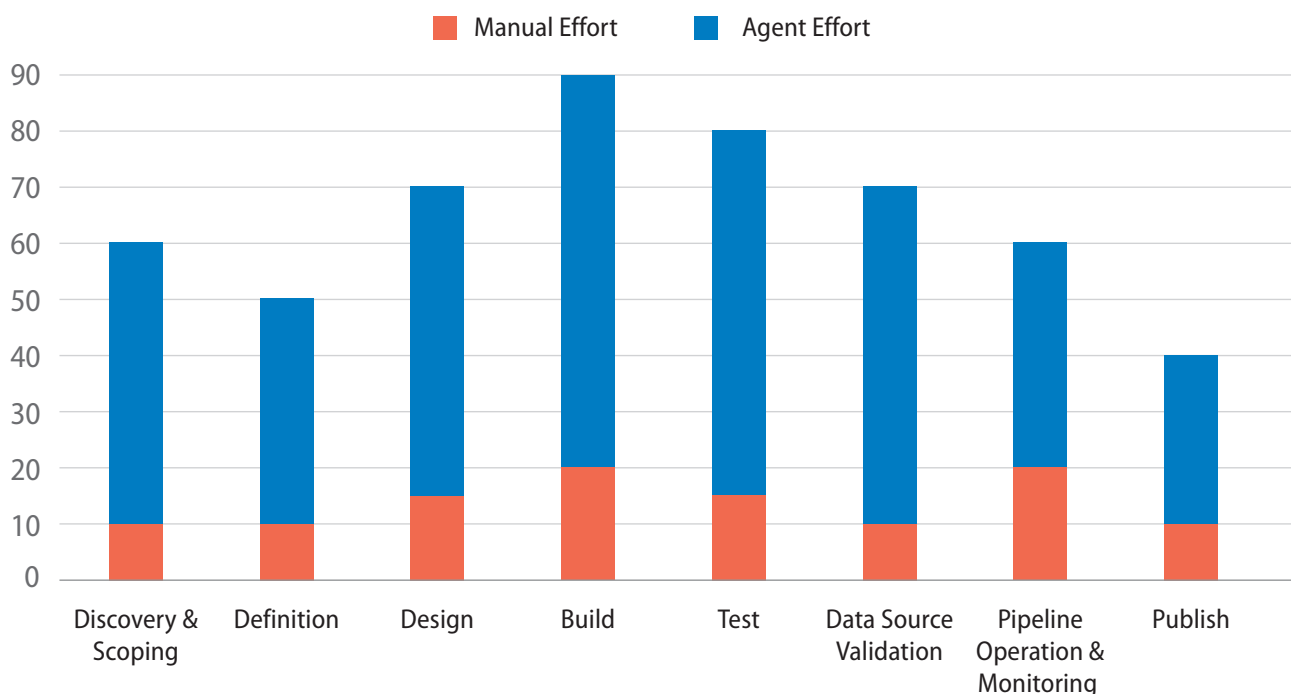
The Solution: The Agentic DE Lifecycle

Agentic AI transforms the lifecycle by offloading the majority of effort to intelligent agents. Human engineers transition to a strategic role, defining high-level goals and validating outcomes.

The agent handles the asynchronous execution of discovery, design, building, and validation.

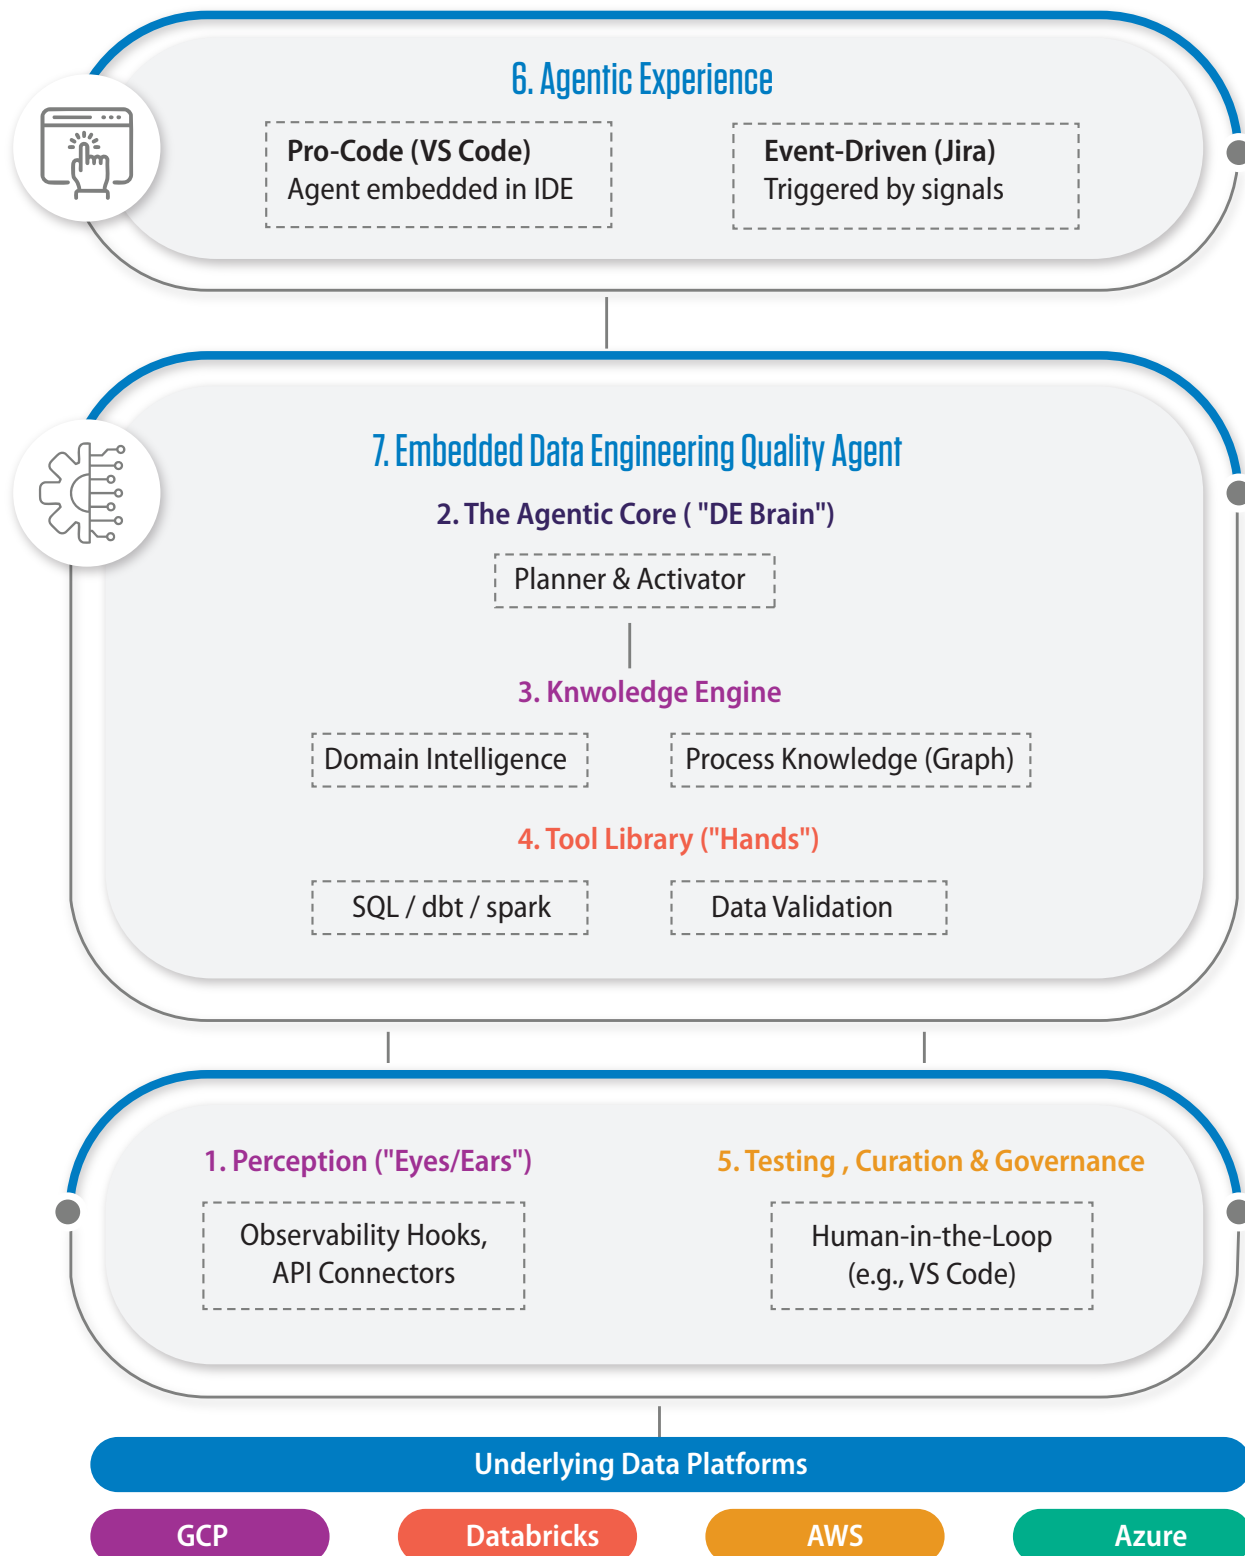
- **Human-in-the-Loop:** Manual effort shifts to goal-setting and validation.
- **Autonomous Execution:** Agents autonomously perform the heavy lifting.
- **Continuous Adaptation:** The system can self-heal and adapt to schema drift or source changes.

Agentic Lifecycle: Effort Rebalanced



Capability Blueprint: 'Lego Blocks' for an Agentic Data Engineering System

Agentic Data Engineering Architecture



1. Perception & Interface

The “eyes and ears” of the agent. This includes API connectors to data sources and observability hooks to monitor data quality and pipeline status.

2. The Agentic Core (The “Brain”)

The central planning and reasoning engine. It decomposes complex requests (e.g., “Build a sales dashboard”) into a series of executable steps.

3. Knowledge Engine

Provides context. Includes Domain Data Intelligence (metadata, business rules) and Process Intelligence (pipeline logs, best practices).

4. Tool Library (The “Hands”)

A curated set of functions the agent can call: schema discovery, data validation, running SQL queries, code generation, and pipeline deployment (dbt, Airflow).

5. Testing, Curation & Governance

The human-in-the-loop component, often integrated into a known IDE like VS Code, for validating agent-generated code, approving data publications, and providing feedback.

6. Agentic Experience

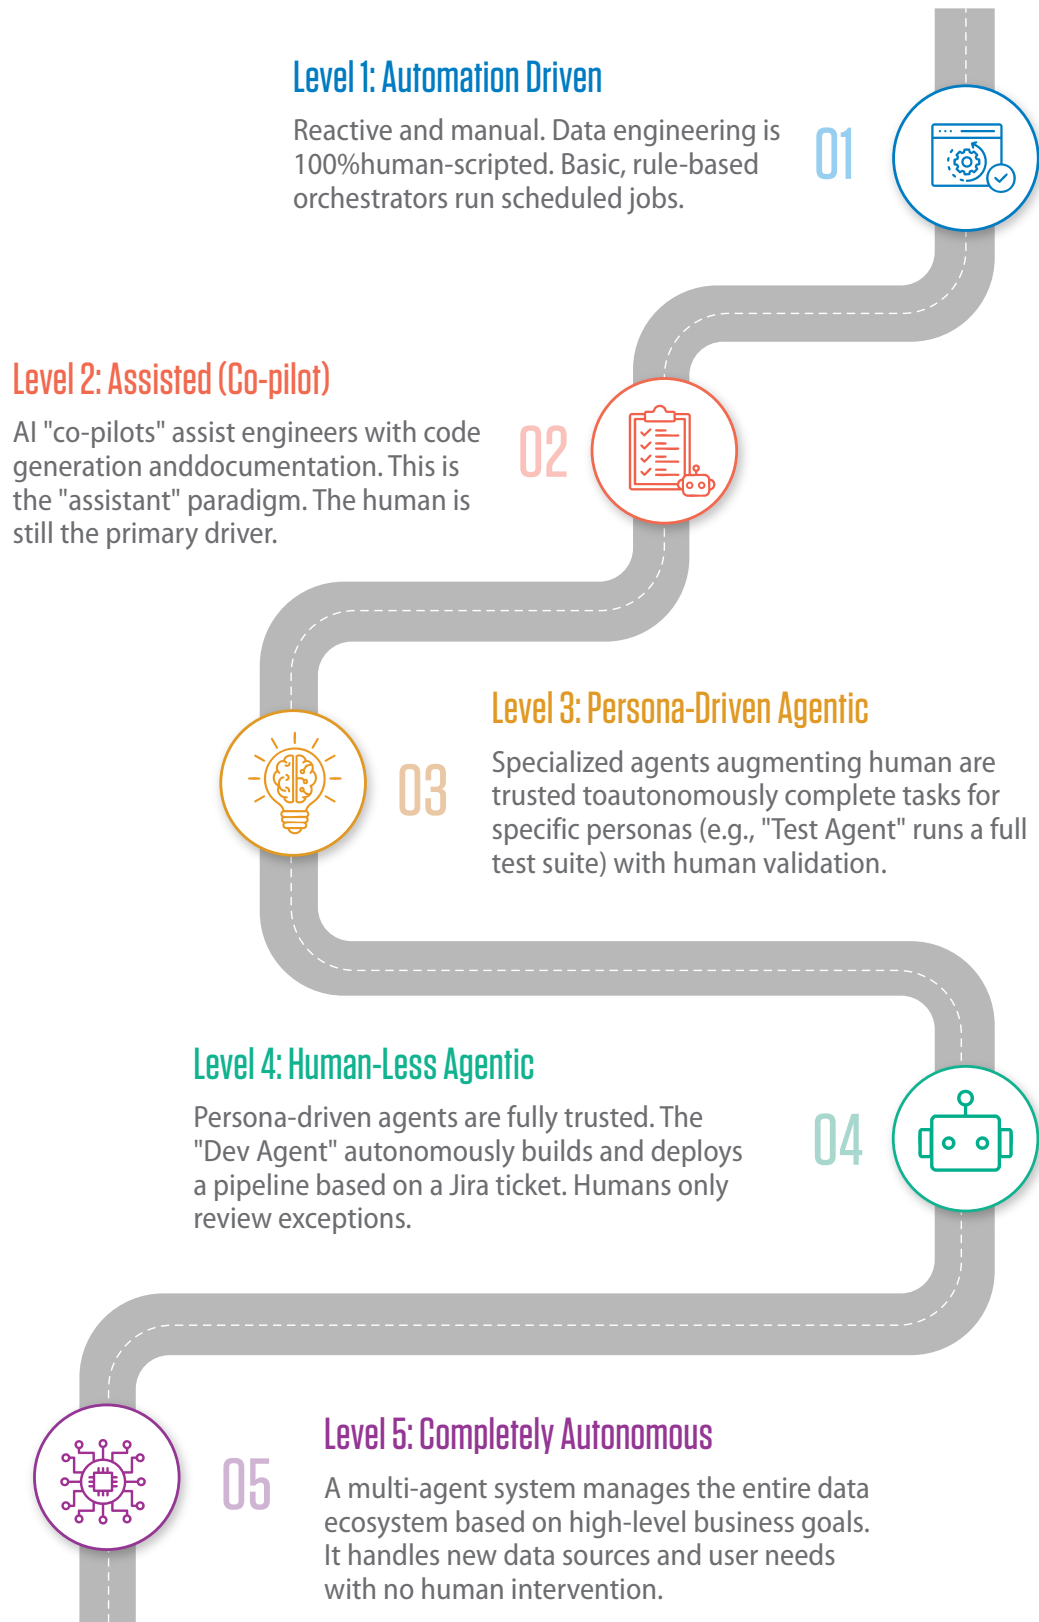
How humans interact with the agent. Includes “pro-code” experiences (embedded in VS Code) and event-driven triggers (a change in a Jira story triggers an agent).

7. Embedded Data Engineering Quality

An overarching layer of non-negotiable standards enforced by an adjudication agent. This includes Code Completeness, Correctness, Security by Design, Compliance by Design, and adherence to engineering Best Practices.

Roadmap to Organizational Maturity

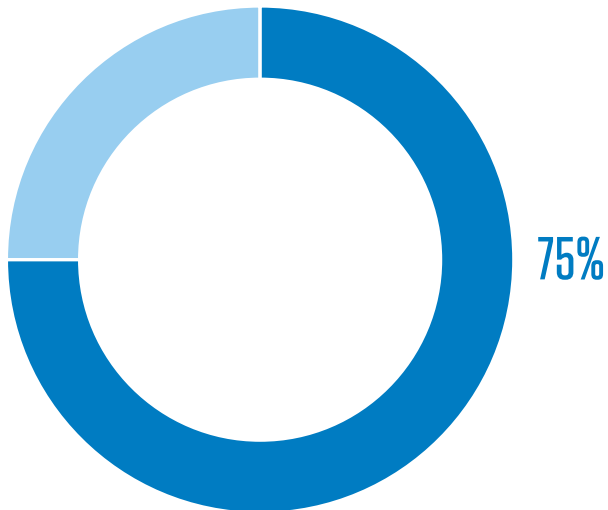
Adopting agentic AI is a journey, not a single step. Organizations typically progress through five stages of maturity, building capability and trust at each level.



The Transformative Impact

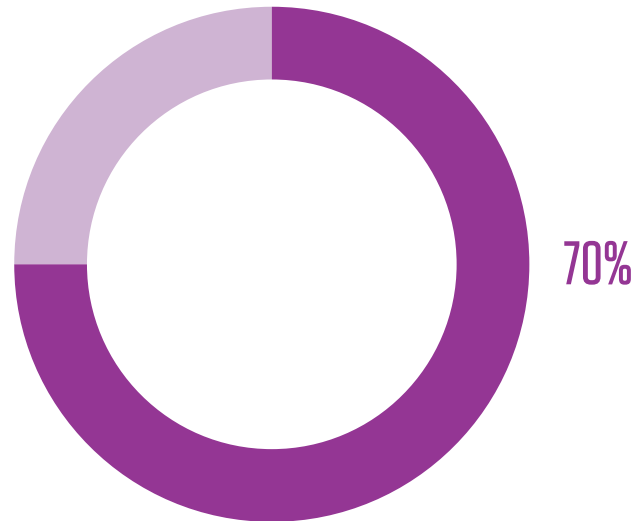
The shift to an agentic model unlocks profound business value by freeing your most skilled engineers from low-level maintenance to focus on high-impact strategic work.

Faster Development



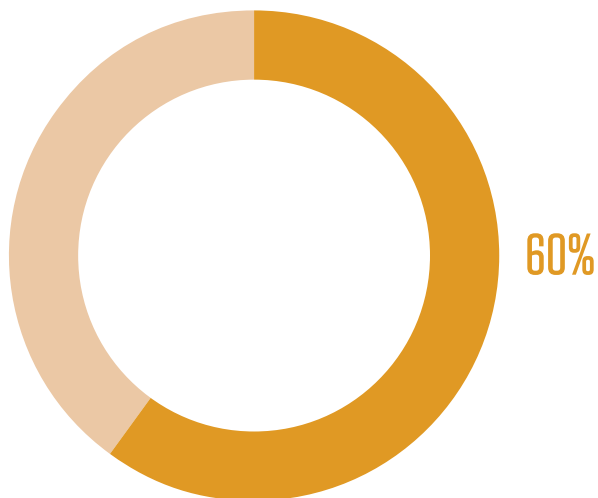
Reduction in pipeline development and deployment time by automating design, build, and test phases.

Fewer Errors



Reduction in data quality issues and pipeline failures through autonomous validation and self-healing.

Cost to Value



Improvement in cost-to-value ratio by automating high-effort, low-value tasks and focusing resources on innovation.

Authors

Rajan Padmanabhan

Unit Technology Officer, Data Analytics and AI, Infosys

Vishnu Sankar

Principal Technology Architect, Data Analytics and AI, Infosys



For more information, contact askus@infosys.com

© 2025 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.