#FABCONSQLCON2026

FABCON | SQLCON

Microsoft Fabric COMMUNITY CONFERENCE | Microsoft SQL COMMUNITY CONFERENCE

ATLANTA MARCH 16 - 20, 2026

Sound off.
The mic is all yours.
Influence the product roadmap.

Join the Fabric User Panel

Share your feedback directly with our Fabric product group and researchers.

https://aka.ms/JoinFabricUserPanel

Join the SQL User Panel

Influence our SQL roadmap and ensure it meets your real-life needs

https://aka.ms/JoinSQLUserPanel

WELCOME

AI Agents, Power Platform and SQL

for Next Gen Automation

WHY YOUR AUTOMATION STRATEGY NEEDS TO **EVOLVE NOW**

INTRODUCTION

Hello,

Chike Eduputa

Head of Microsoft

Capgemini Invent UK

in in/ceduputa

AGENDA

In this session, we will cover:

01

The Treadmill Problem

02

What is an Agent

03

Where do Agents fit

04

The quiet hero

05

Trusted systems

CHAPTER 1

The Treadmill Problem

CHANGE OF CUSTOMER DETAILS PROCESS

01
Gather required identification documents.

02
Visit a branch.
Some companies accept forms by Post or Online

03
Complete a Paper form.
Some companies use Digital Forms

04
Verify identity with relevant identification documents.

05
Receive confirmation of updated details.

Sounds simple and straightforward to automate, right?

CHAPTER 1

The Treadmill Problem

CHANGE OF CUSTOMER DETAILS PROCESS

01

Customer identifies the need to update personal details.

02

Visit a branch, queues and speaks to a staff (typical). *Some banks accept by Post or Online*

03

Complete a Paper form. *Some banks have a Digital Form*

04

Verify identity with relevant identification documents.

05

Receive confirmation of updated details.

Staff greets customer and asks reason for visit.

Staff provides change of details form. *(hopefully the latest)*

Staff assists customer with completing the form. (*if required*)

Staff verifies identity documents.

Staff checks form is complete and signed.

Staff accepts the form and confirms it will be processed.

Staff issues a reference number and sends it for processing.

Paper form and ID are batch scanned and manually entered in internal systems

Data entry team reviews for legibility and completeness

Identity verification checks are performed

Business rules applied like fraud screening, address validation, duplicate records

Customer information file updated

Changes propagated to core banking systems downstream

Audit logs are created for compliance

Notification sent to customer confirming update

Branch Teller System used by staff to capture updates

Document Scanning System to digitise paper forms

Customer Information File to master customer profile

Core Banking System for Account Servicing

Fraud Detection System for risk screening

Address Verification Service to validate addresses

Data Integration system to sync updates across systems

Notification service to send SMS, email and/or post

Audit and Compliance System for regulatory tracking

CHAPTER 1

The Treadmill Problem

BANK CHANGE OF CUSTOMER DETAILS

01

Customer identifies the need to update personal details

02

Visit a branch, queues and speaks to a staff (typical). *Some banks accept by Post or Online*

03

Complete a Paper form. *Some banks have a Digital Form*

04

Verify identity with relevant identification documents.

05

Receive confirmation of updated details.

⚠ **Requires physical branch visit**

⚠ **Queues and waiting times**

⚠ **Manual form errors**

⚠ **Slow turnaround times**

Staff greets customer and asks reason for visit

Staff provides change of details form. *(hopefully the latest)*

Staff assists customer with completing the form if required

Staff verifies identity documents

Staff checks form is complete and signed

Staff accepts the form and confirms it will be processed

Staff issues a reference number and sends it for processing

Paper form is batch scanned and manually entered in internal systems

Data entry team reviews for legibility and completeness

Identity verification checks are performed

Business rules applied like fraud screening, address validation, duplicate records

Customer information file updated

Changes propagated to core banking systems downstream

Audit logs are created for compliance

Notification sent to customer confirming update

⚠ **Paper handling and storage**

⚠ **Manual data entry errors**

⚠ **Slow processing times**

⚠ **Fragmented system updates**

Branch Teller System used by staff to capture updates

Document Scanning System to digitise paper forms

Customer Information File to master customer profile

Core Banking System for Account Servicing

Fraud Detection System for risk screening

Address Verification Service to validate addresses

Data Integration system to sync updates across systems

Notification service to send SMS, email and/or post

Audit and Compliance System for regulatory tracking

CHAPTER 1

The Treadmill Problem

Now picture a 100+ years old Global Bank,

100,000’s	1,000’s	10,000’s	10,000’s
Transactions per year	Form types	Queues	Operations Staff

Paper forms

Manual entry

AI Builder + Power Automate Desktop struggled

Regulatory changes broke mappings

Scheme drift

Extraction quality collapsed

FRAGILITY

The pixel perfect fallacy - One UI change, total failure

CHANGE FAILURE

Every update becomes an innovation tax on the backlog

LOCALISATION FRICTION

Forms vary by region; bots can’t generalise

HERO CULTURE

One developer holds the keys; knowledge never scales

IDLE BOT ECONOMICS

Licensed bots sitting unused, cost without value

The Treadmill Problem

BANK CHANGE OF CUSTOMER DETAILS: THE RPA SOLUTION

Typical automation architecture

1. Branch scanner to a shared mailbox
2. Document repository (IBM Filenet)
3. OCR tool (AI Builder)
4. RPA tool (Power Automate Desktop)
5. Rules engine (Excel, Visio, Word)
6. Work Queues tool for Ops staff
7. Exception tool (Power App) for Ops staff
8. Legacy banking systems (20+ yr old IBM zOS Mainframe)
9. Notifications system (CCM)
10. Audit log repository

Strengths

- Fastest way for manual, legacy process
- Works well where APIs are weak or absent
- Reduces repetitive rekeying by Ops teams
- Can be done incrementally system by system
- Lower change impact on core systems

Limitations

- Automates existing mess not transforming
- Fragile/brittle when layouts/fields change
- OCR accuracy inconsistent for handwriting
- Limited reasoning for ambiguous cases
- Hard to scale into an intelligent service
- Data remains fragmented

RPA is the right choice when:

Short-term productivity gains

Low disruption automation

Support for legacy systems with no APIs

A stop-gap to future transformation

CHAPTER 1

The Treadmill Problem

BANK CHANGE OF CUSTOMER DETAILS: THE AGENTIC SOLUTION

How might we reimagine the solution in an intelligent, orchestrated and auditable agentic workflow

DOCUMENT INTELLIGENCE COPILOT STUDIO POWER AUTOMATE SQL & FABRIC FABRIC DATA AGENTS AI AGENTS

What is an agent?

It is not a chatbot with better prompts and answers

It is not an RPA wearing a language model

It is *certainly* not uncontrolled autonomy

The Litmus test: *"If you remove the human from the chat, does the work still get done?"*
*If **yes** -> Agent, if **no** -> Chatbot*

APPLIES AI MODEL **REASONING**	CALLS **TOOLS**
INTERPRETS **INPUTS** TAKES **ACTION**	FOLLOWS **INSTRUCTIONS**

CHAPTER 2

What is an agent?

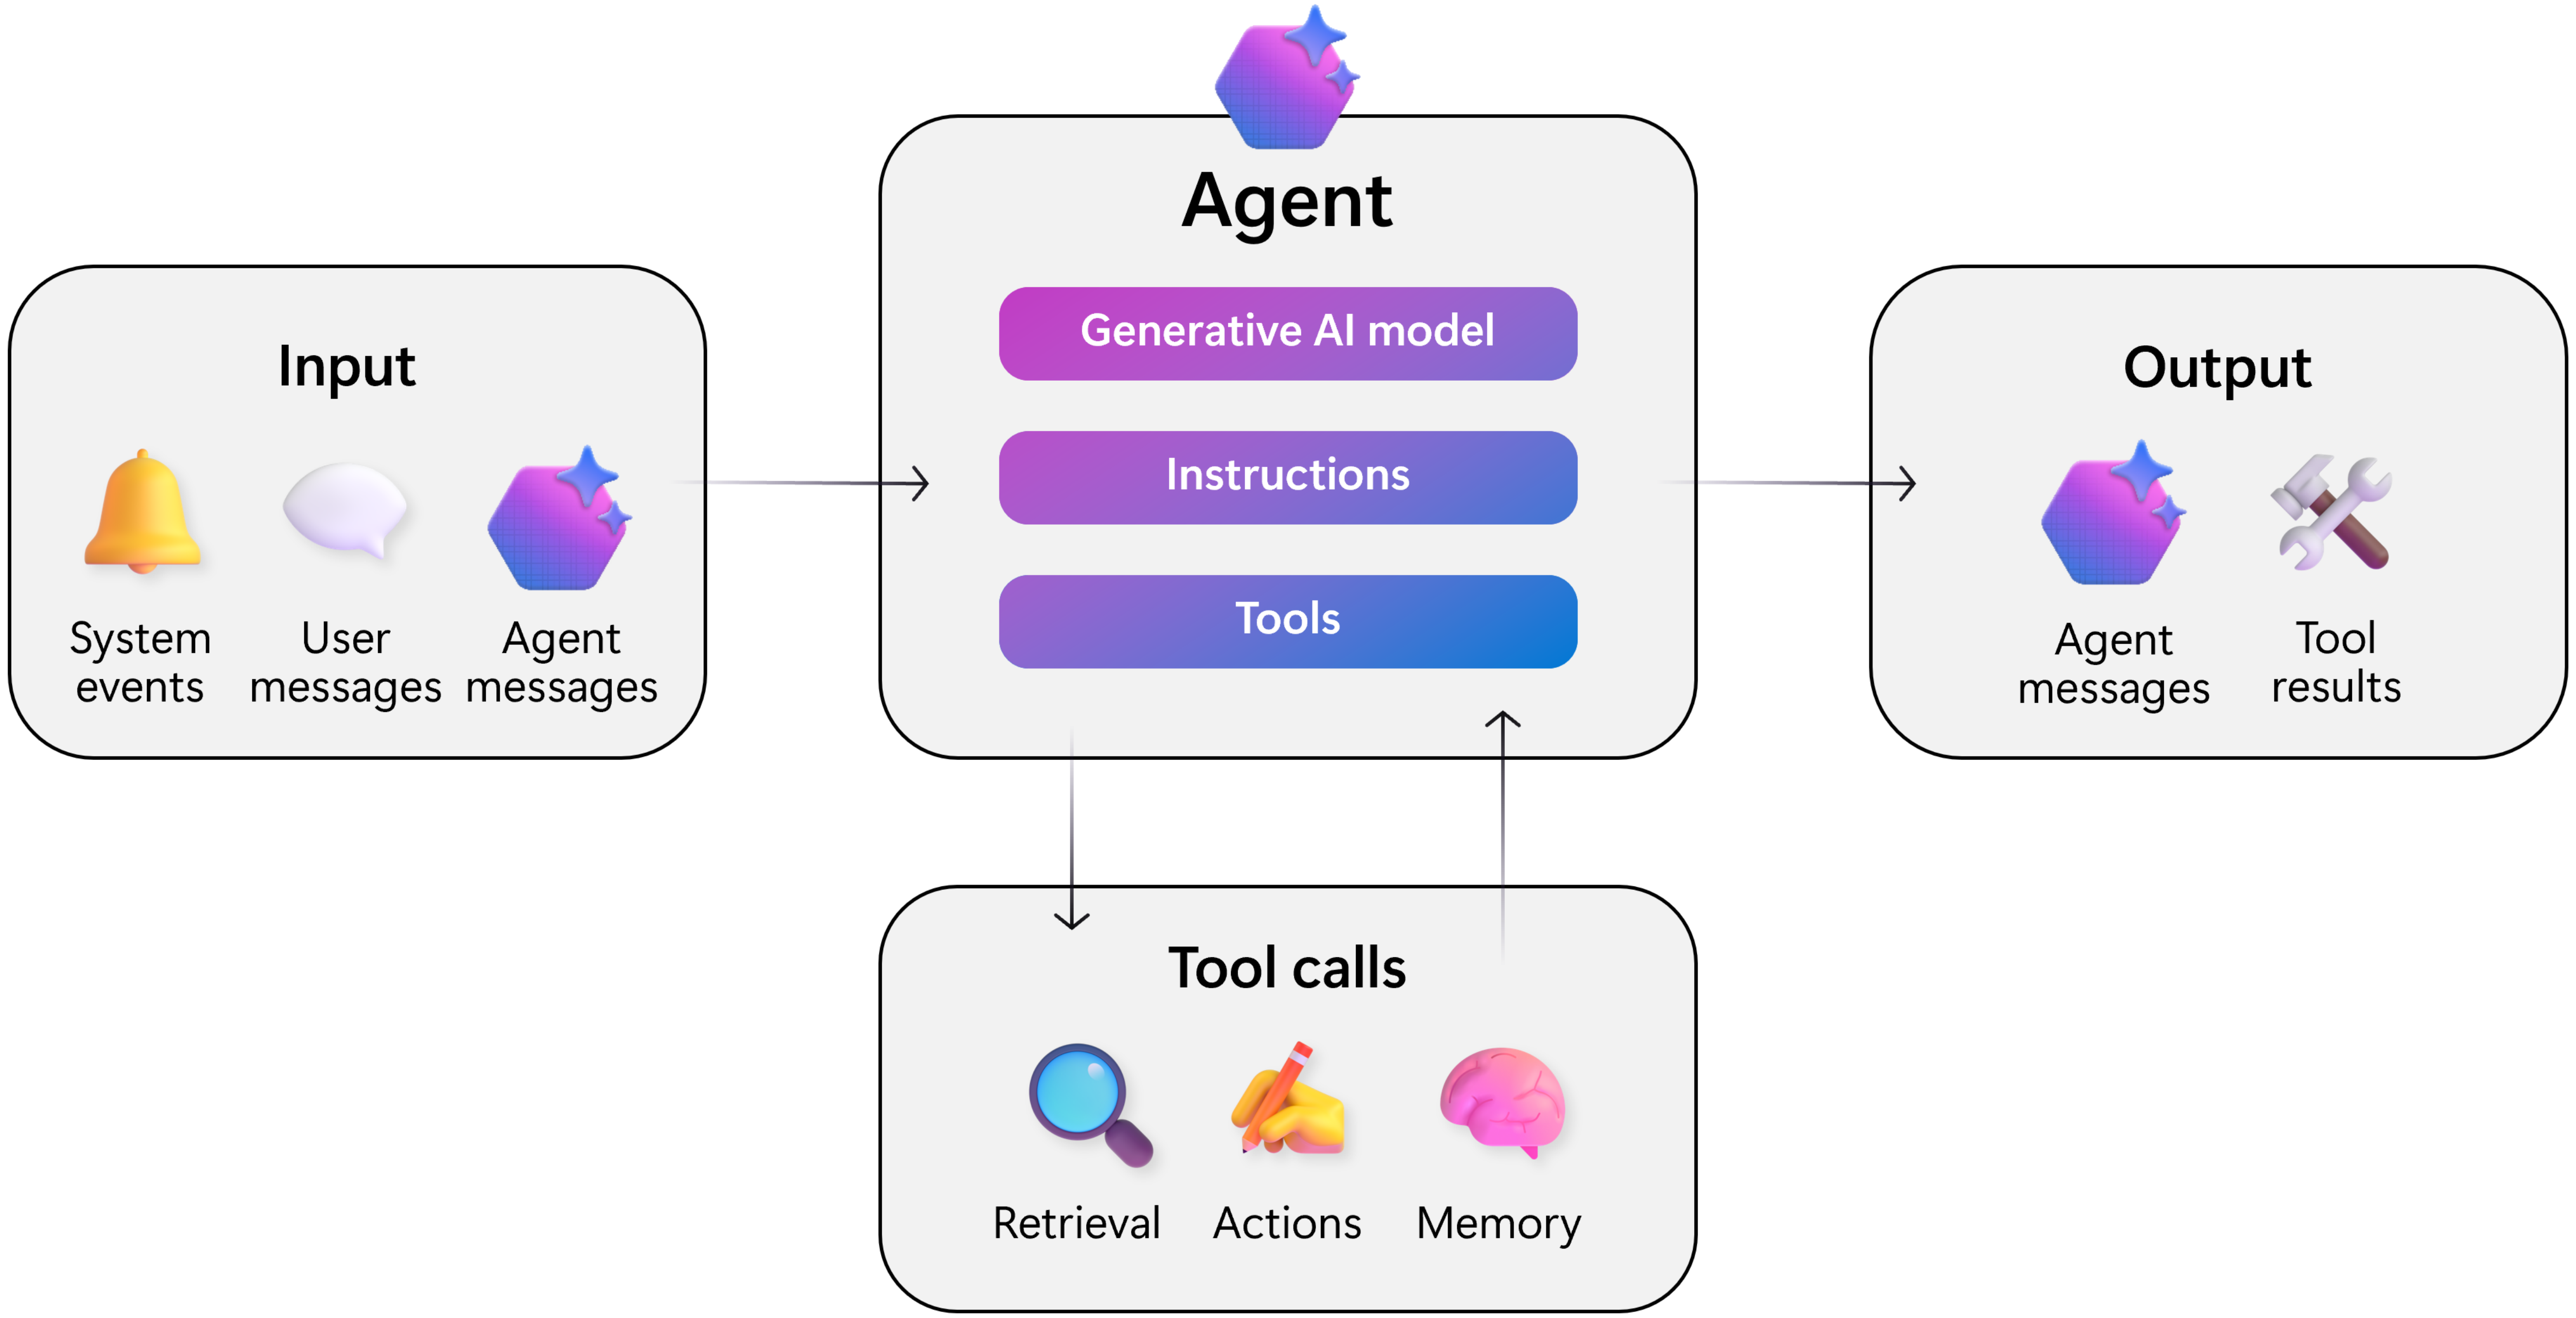

Source: Microsoft

What is an agent?

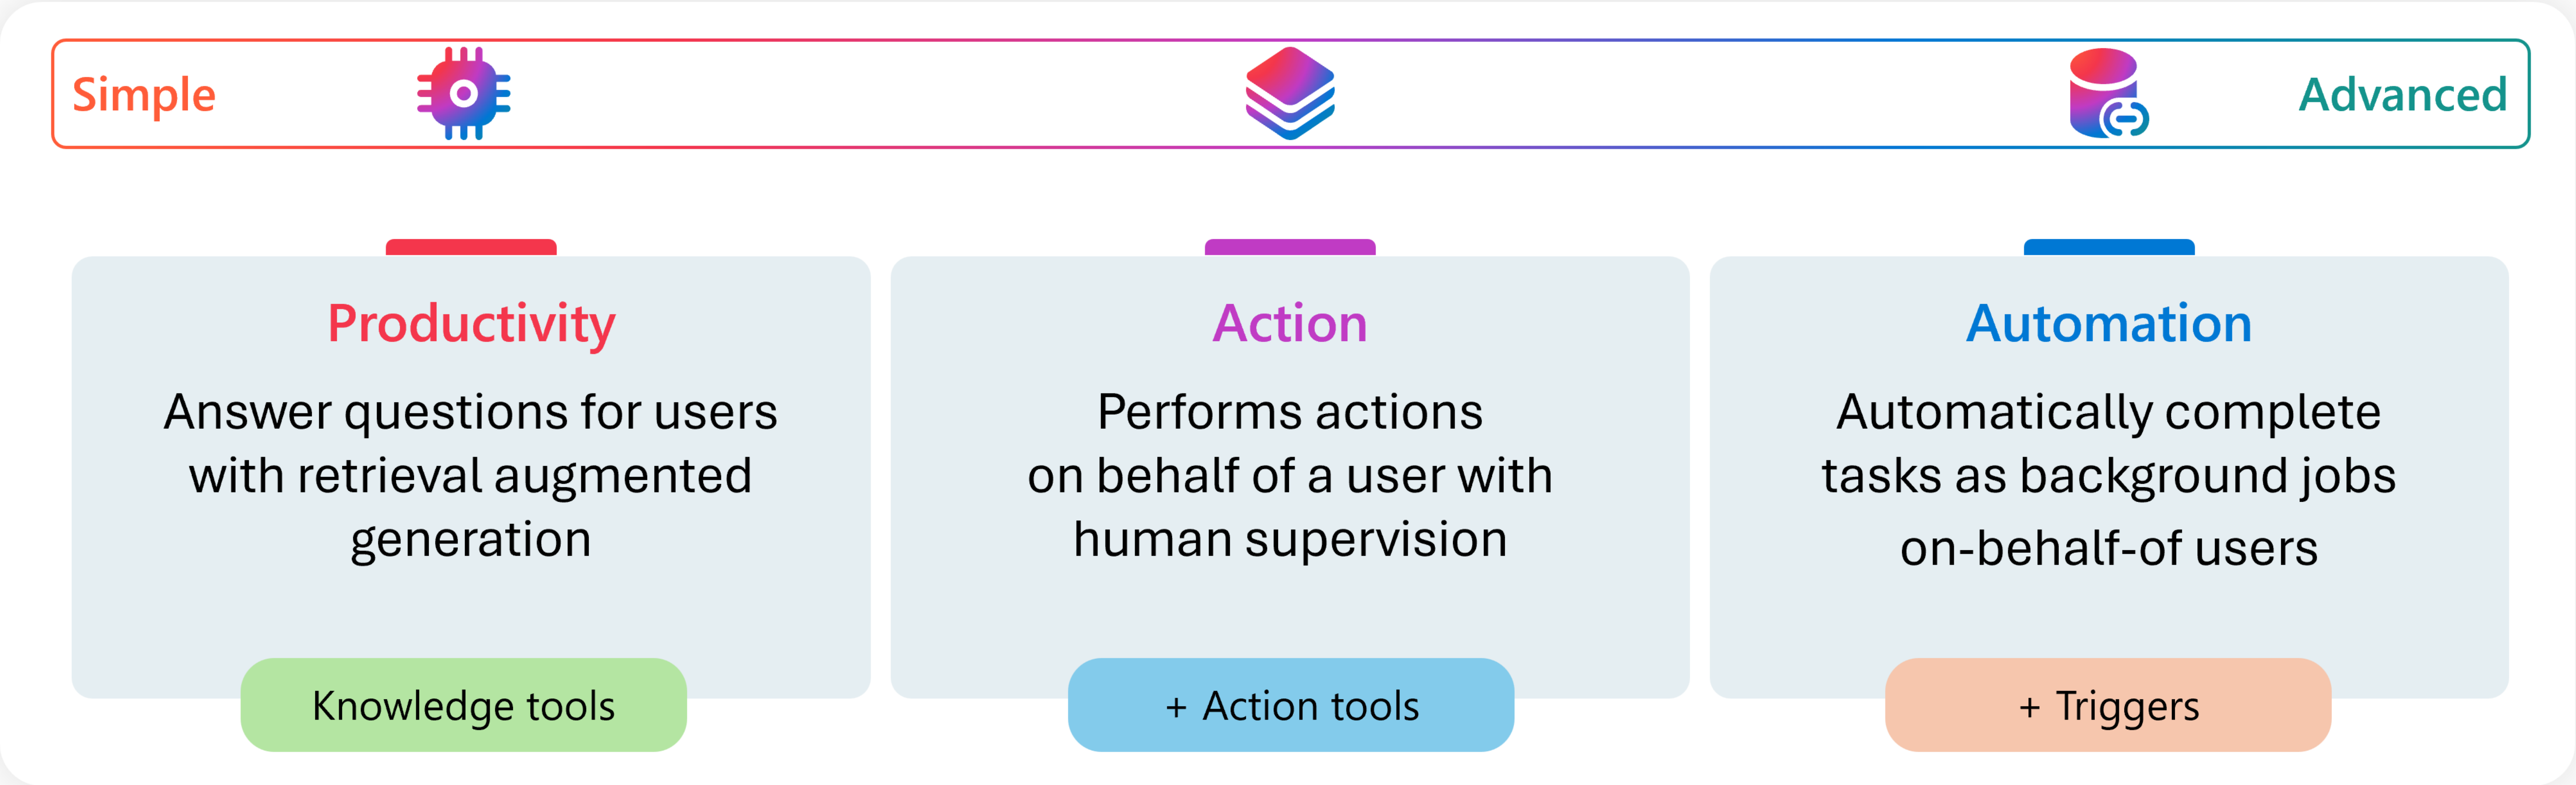

Simple

Advanced

Productivity

Answer questions for users with retrieval augmented generation

Knowledge tools

Action

Performs actions on behalf of a user with human supervision

+ Action tools

Automation

Automatically complete tasks as background jobs on-behalf-of users

+ Triggers

Source: Microsoft

CHAPTER 2

What is an agent?

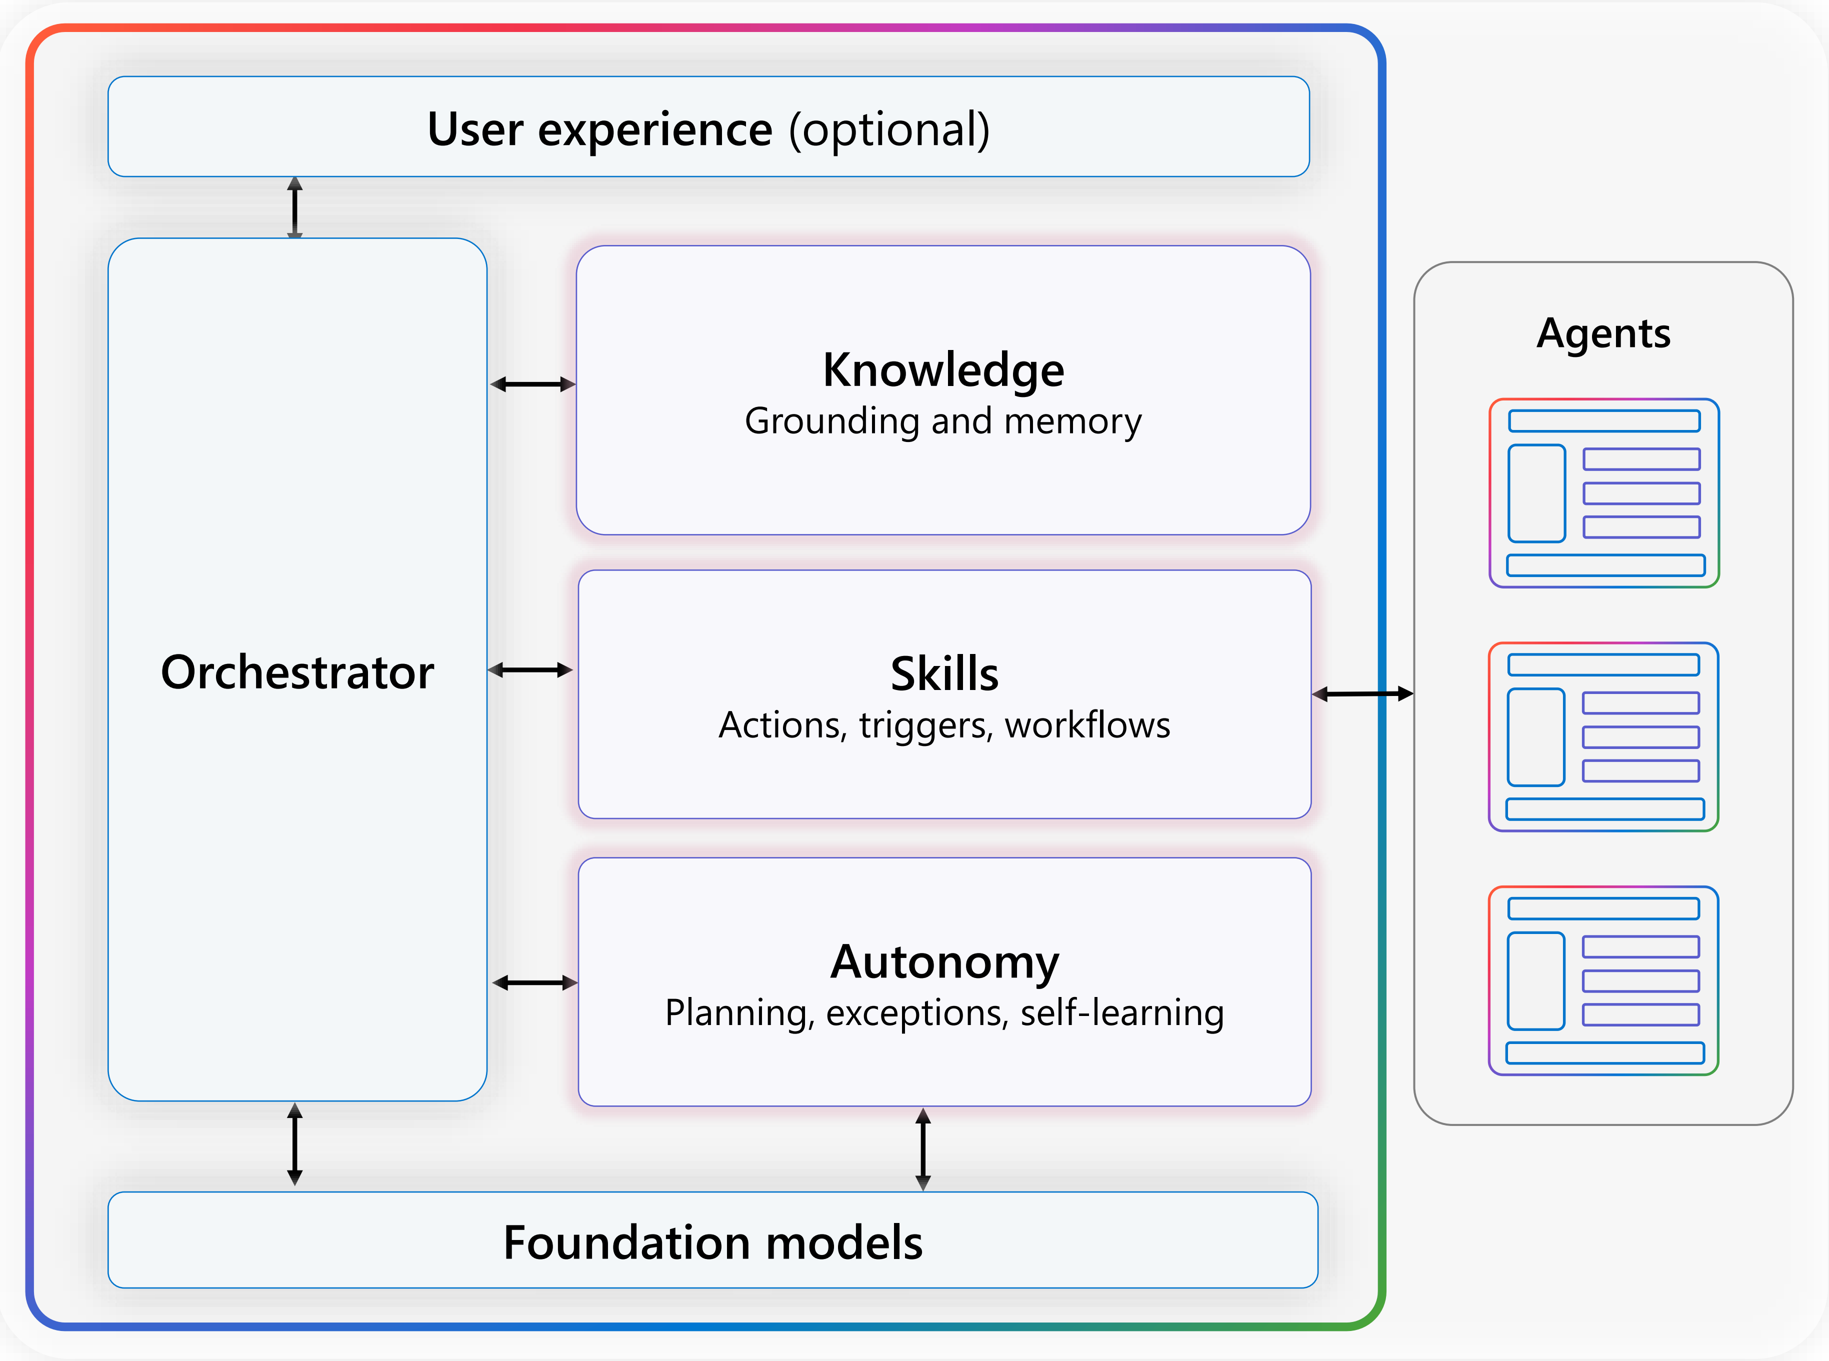
User experience (optional)
Orchestrator
Knowledge
Grounding and memory
Skills
Actions, triggers, workflows
Autonomy
Planning, exceptions, self-learning
Foundation models
Agents

Source: Microsoft

What is an agent?

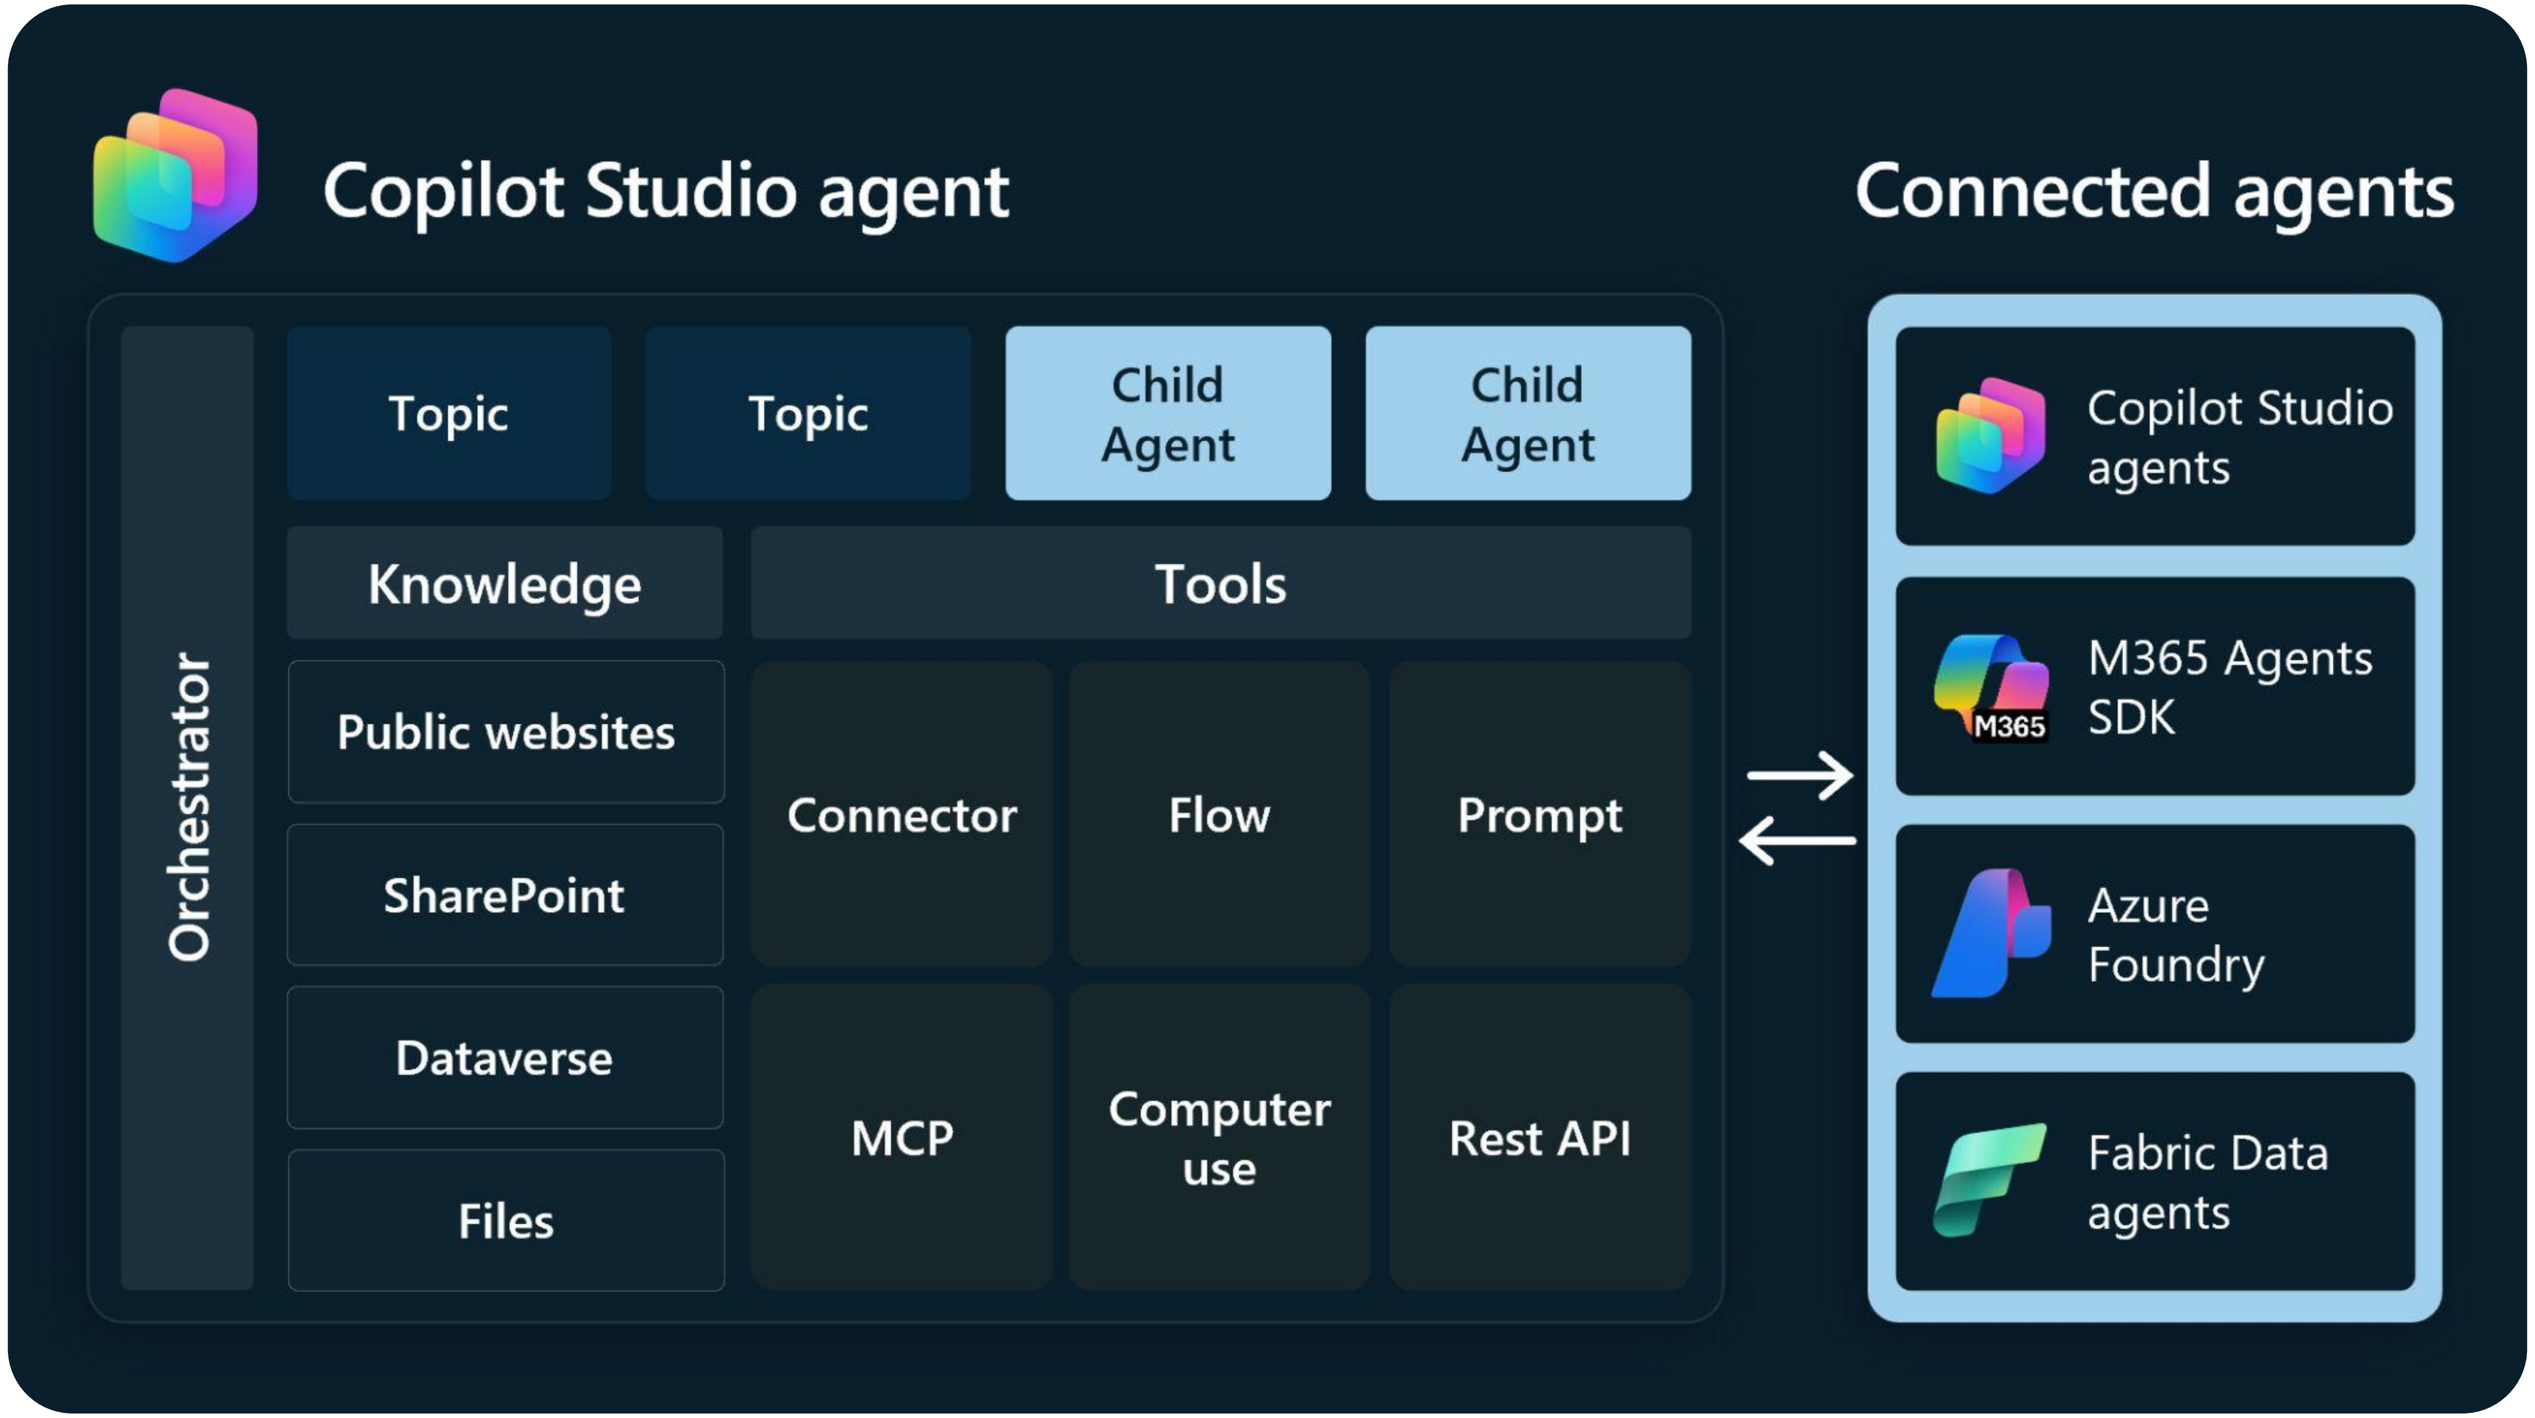

Copilot Studio agent

Connected agents

Orchestrator

Topic	Topic	Child Agent	Child Agent

Knowledge	Tools		
Public websites	Connector	Flow	Prompt
SharePoint			
Dataverse	MCP	Computer use	Rest API
Files			

- Copilot Studio agents
- M365 Agents SDK
- Azure Foundry
- Fabric Data agents

Source: Microsoft

Common Agent Patterns

Pattern	Description	Example
Sequential	Focused, sequential tasks	Market research analysis
Coordinator	Route requests to the right agent	Customer support triage
Parallel	Simultaneous specialist tasks	Travel booking assistance
Maker-Checker	Structured production loop	Quality control agents
Goal Setting	Plan steps to reach a goal	Business process automation
Debate	Compare options and decide	Product design review
Learn & Adapt	Refine through experience	Trading strategy refinement

Multi-Agent Systems

Multiple agents work together to achieve a goal. Collaboration of specialised agents who each have a role and expertise.

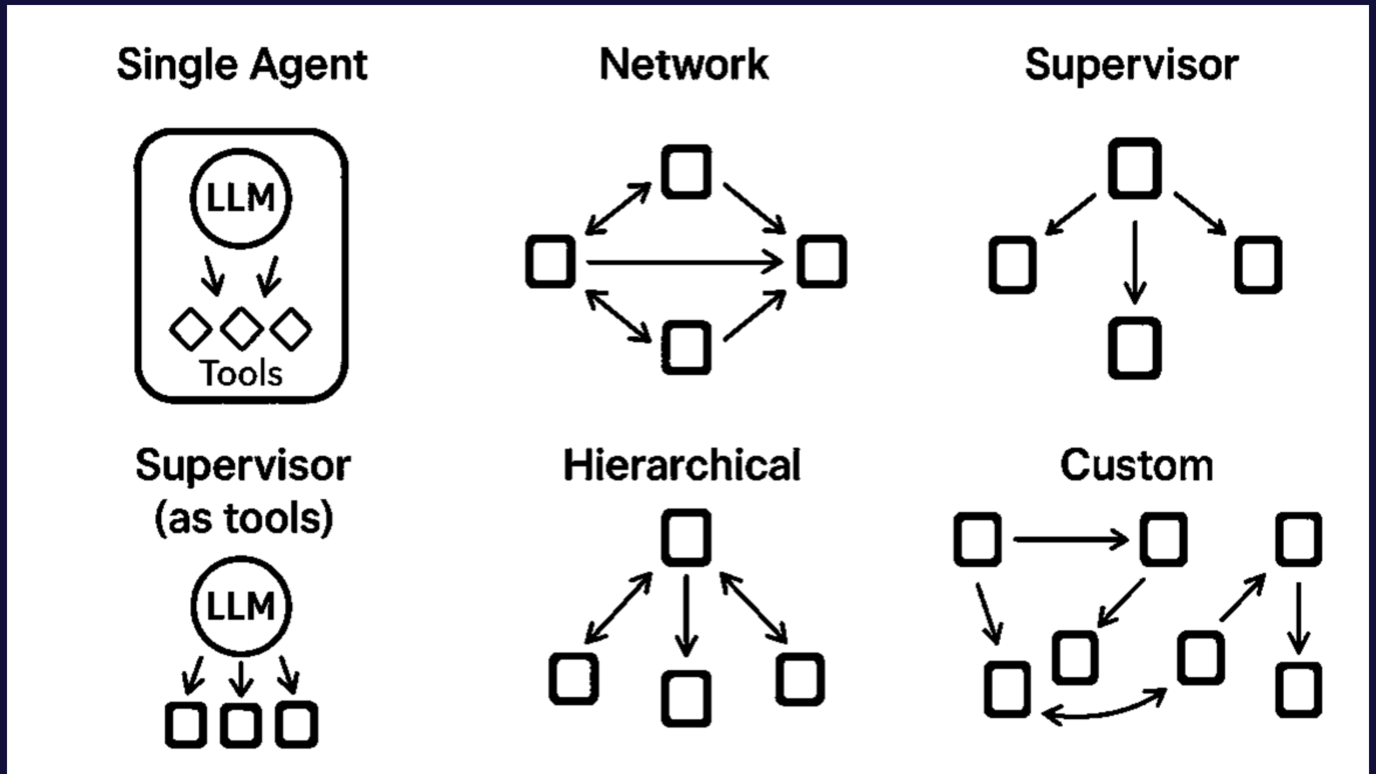

Source: The Learning Space

Common Agent Patterns

01

Tool / Function Calling

Agent invokes external functions or APIs as tools to extend its capabilities

Calls a SQL query tool to fetch sales data, then summarizes results

02

Model Context Protocol

Standardized protocol for sharing context between models and data sources

Connects to a live DB via MCP server to answer with fresh data

03

Connectors

Pre-built integrations linking agents to SaaS platforms with minimal code

Connector reads/writes SharePoint lists or sends Teams messages

04

Agent-to-Agent (A2A)

Multiple agents collaborate by exchanging messages and delegating subtasks

Planning agent assigns research and writing to specialist agents

05

REST API

Agent sends HTTP requests to web services via standard endpoints

Calls a weather API to provide real-time forecasts in chat

06

Automation Workflows

Agent triggers or embeds in orchestration flows like Power Automate

Request triggers a flow that runs an agent, then emails output

07

Computer Use

Agent interacts with desktop or browser UI, clicking and navigating like a human

Opens a legacy ERP, fills a purchase order form, and submits it

CHAPTER 3

Where do agents fit?

THE FRONTIER FIRM

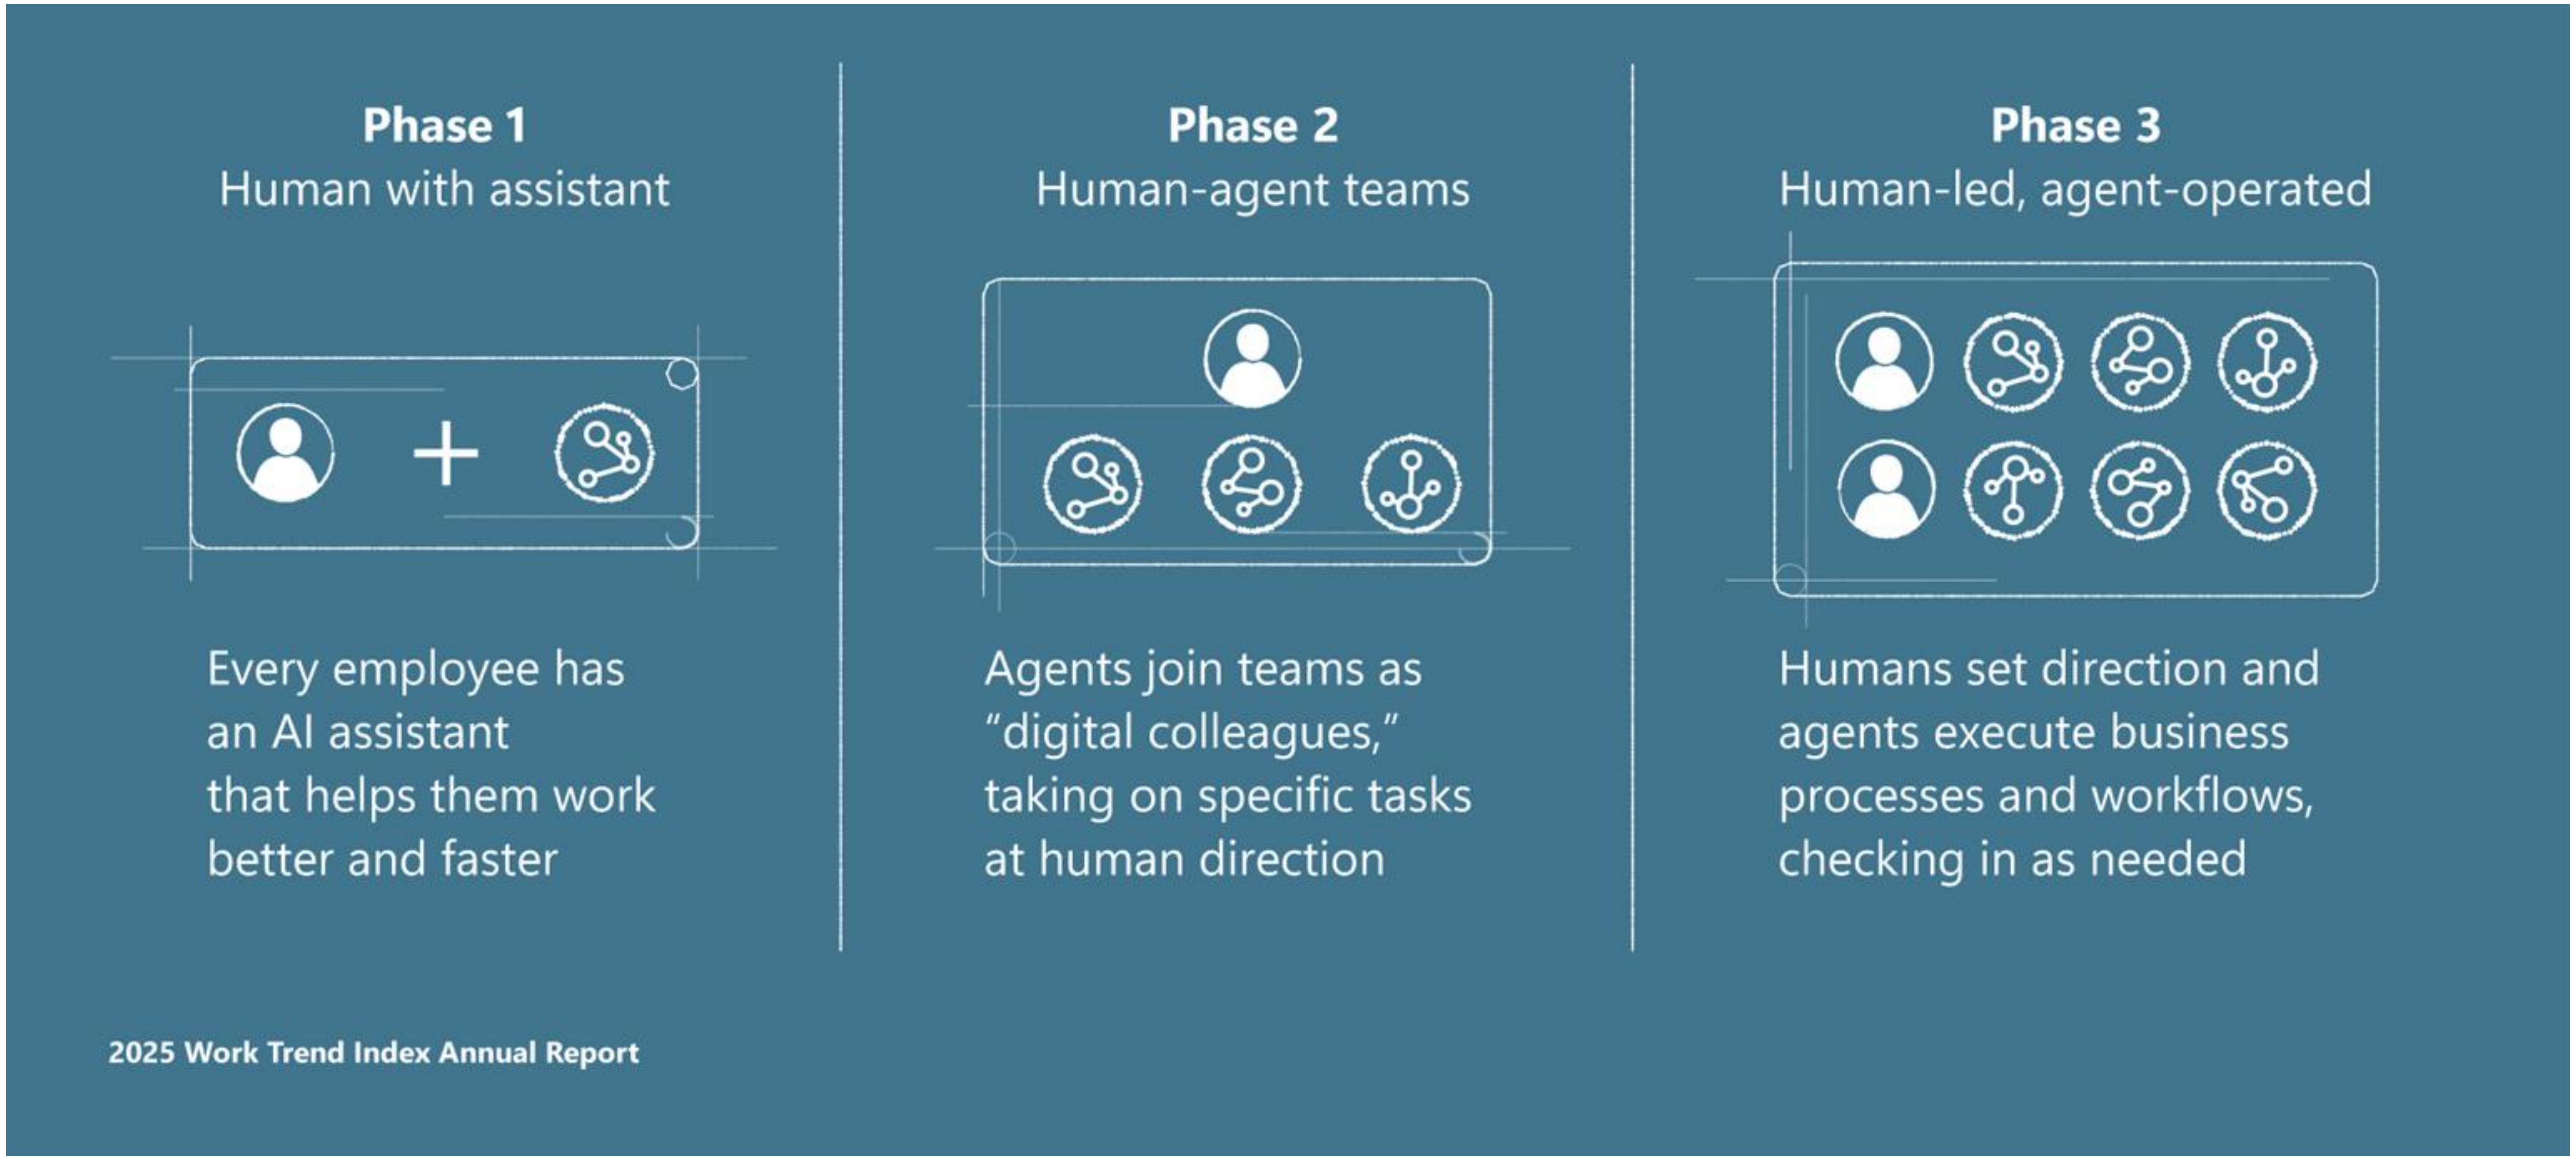

Phase 1
Human with assistant

Every employee has an AI assistant that helps them work better and faster

Phase 2
Human-agent teams

Agents join teams as “digital colleagues,” taking on specific tasks at human direction

Phase 3
Human-led, agent-operated

Humans set direction and agents execute business processes and workflows, checking in as needed

2025 Work Trend Index Annual Report

Where do agents fit?

CHANGE OF CUSTOMER DETAILS SOLUTION ARCHITECTURE

Microsoft Azure · Power Platform · Fabric

EXPERIENCE
Staff & Customer Interfaces

Power App	Digital Capture	Power Apps
Branch Scanner / Inbox	Paper intake workflow	Power Automate
Branch Staff Agent	Step-by-step teller assist	Copilot Studio
Operations Agent	Back-office queue & review	Copilot Studio
Power App	Doc Processing UI	Power Apps

ORCHESTRATION
Workflow & Case Management

Multi-Agent Orchestrator	Routes tasks, manages state	Copilot Studio
Automated Workflows	Trigger on events, SLA track	Power Automate
Case Management	Change request lifecycle	Dataverse
Notification Hub	SMS, email & in-app confirms	Azure Comms Svcs

INTELLIGENCE
AI & Rules Engines

Document Intelligence	Extract form & ID fields	Azure AI Doc Intel
Identity Verification	Face match, liveness check	Azure AI Vision
AI Decisioning	Fraud scoring, deduplication	Azure ML / Prompt Flow
Fabric Data Agents	CIF lookups, data enrichment	Microsoft Fabric
Policy / Rules	Compliance & field validation	Logic Apps / Rules

INTEGRATION
APIs, Events & RPA

Banking API Gateway	CIF, CRM, Core Banking, Cards	Azure API Mgmt
Event Propagation	Fan-out change events real-time	Service Bus / Event Grid
Fallback RPA / CUA	Legacy endpoints, screen-scrape	Power Automate Desktop / Computer Use
MCP Connectors	Data, Process & System tools	Azure Functions
Execution Services	pAPIs, sAPIs, zAPIs	Azure Functions

DATA
Storage, Analytics & Compliance

SQL Validation	Operational rules & pre-commit	Azure SQL DB
Fabric Lakehouses	OneLake · raw, curated, gold	Microsoft Fabric
Master Data Views	Golden customer record, MDM	Fabric / Purview
Audit & Event Log	Immutable compliance change log	Azure Monitor
Semantic Models	SLA & volume reporting	Power BI / Fabric
Source Systems	CIF, CRM, Core Banking, Cards	Azure Data Factory

Where do agents fit?

CHANGE OF CUSTOMER DETAILS SOLUTION ARCHITECTURE

Microsoft Azure · Power Platform · Fabric

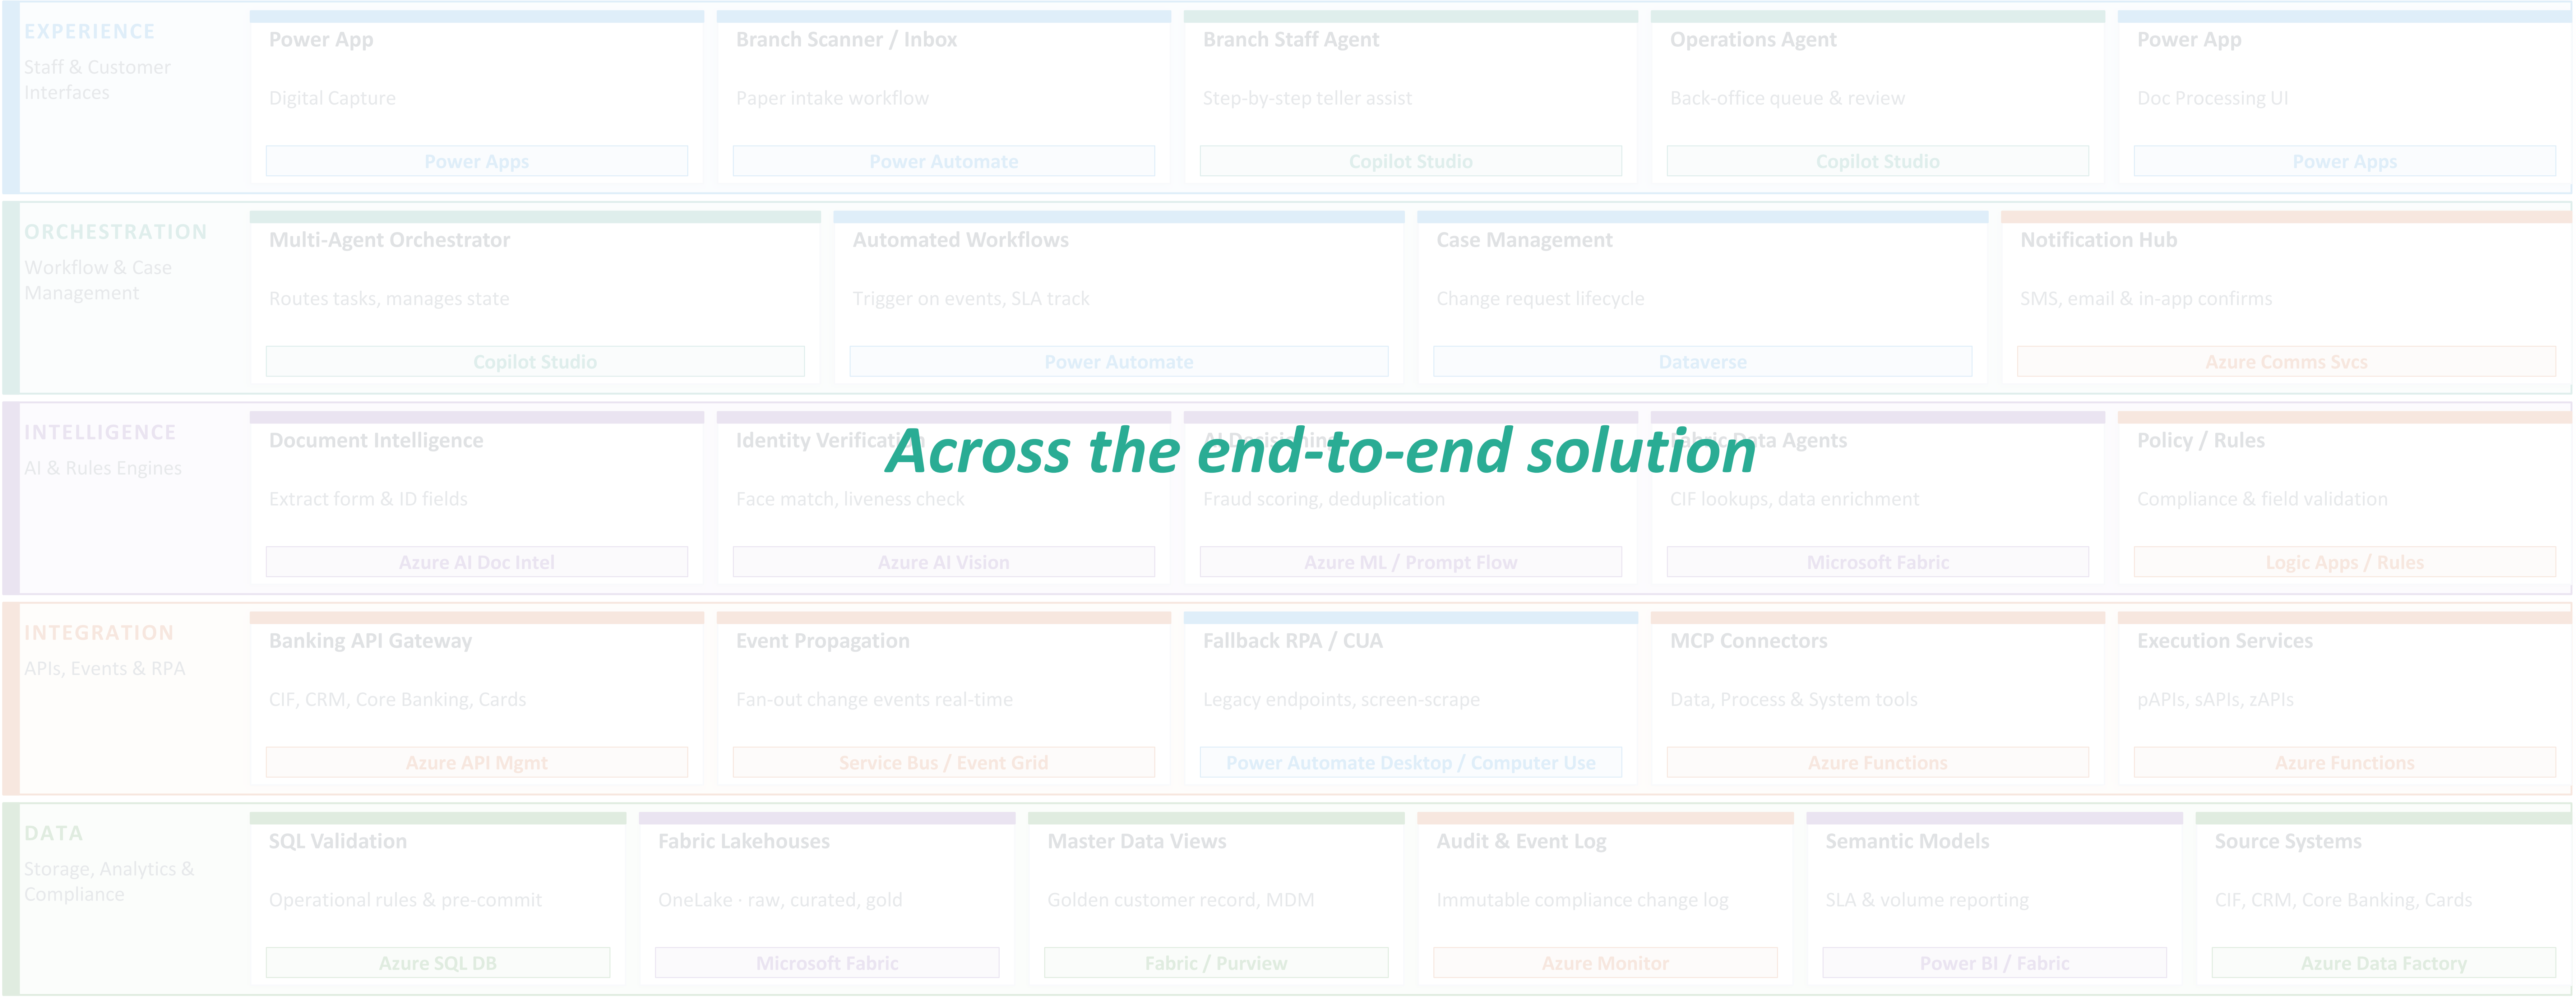

Across the end-to-end solution

EXPERIENCE
Staff & Customer Interfaces

- Power App: Digital Capture (Power Apps)
- Branch Scanner / Inbox: Paper intake workflow (Power Automate)
- Branch Staff Agent: Step-by-step teller assist (Copilot Studio)
- Operations Agent: Back-office queue & review (Copilot Studio)
- Power App: Doc Processing UI (Power Apps)

ORCHESTRATION
Workflow & Case Management

- Multi-Agent Orchestrator: Routes tasks, manages state (Copilot Studio)
- Automated Workflows: Trigger on events, SLA track (Power Automate)
- Case Management: Change request lifecycle (Dataverse)
- Notification Hub: SMS, email & in-app confirms (Azure Comms Svcs)

INTELLIGENCE
AI & Rules Engines

- Document Intelligence: Extract form & ID fields (Azure AI Doc Intel)
- Identity Verification: Face match, liveness check (Azure AI Vision)
- AI Decisioning: Fraud scoring, deduplication (Azure ML / Prompt Flow)
- Fabric Data Agents: CIF lookups, data enrichment (Microsoft Fabric)
- Policy / Rules: Compliance & field validation (Logic Apps / Rules)

INTEGRATION
APIs, Events & RPA

- Banking API Gateway: CIF, CRM, Core Banking, Cards (Azure API Mgmt)
- Event Propagation: Fan-out change events real-time (Service Bus / Event Grid)
- Fallback RPA / CUA: Legacy endpoints, screen-scrape (Power Automate Desktop / Computer Use)
- MCP Connectors: Data, Process & System tools (Azure Functions)
- Execution Services: pAPIs, sAPIs, zAPIs (Azure Functions)

DATA
Storage, Analytics & Compliance

- SQL Validation: Operational rules & pre-commit (Azure SQL DB)
- Fabric Lakehouses: OneLake · raw, curated, gold (Microsoft Fabric)
- Master Data Views: Golden customer record, MDM (Fabric / Purview)
- Audit & Event Log: Immutable compliance change log (Azure Monitor)
- Semantic Models: SLA & volume reporting (Power BI / Fabric)
- Source Systems: CIF, CRM, Core Banking, Cards (Azure Data Factory)

Where do agents fit?

CHANGE OF CUSTOMER DETAILS AGENTIC SOLUTION ARCHITECTURE

Microsoft Azure · Power Platform · Fabric

Illustrative example only

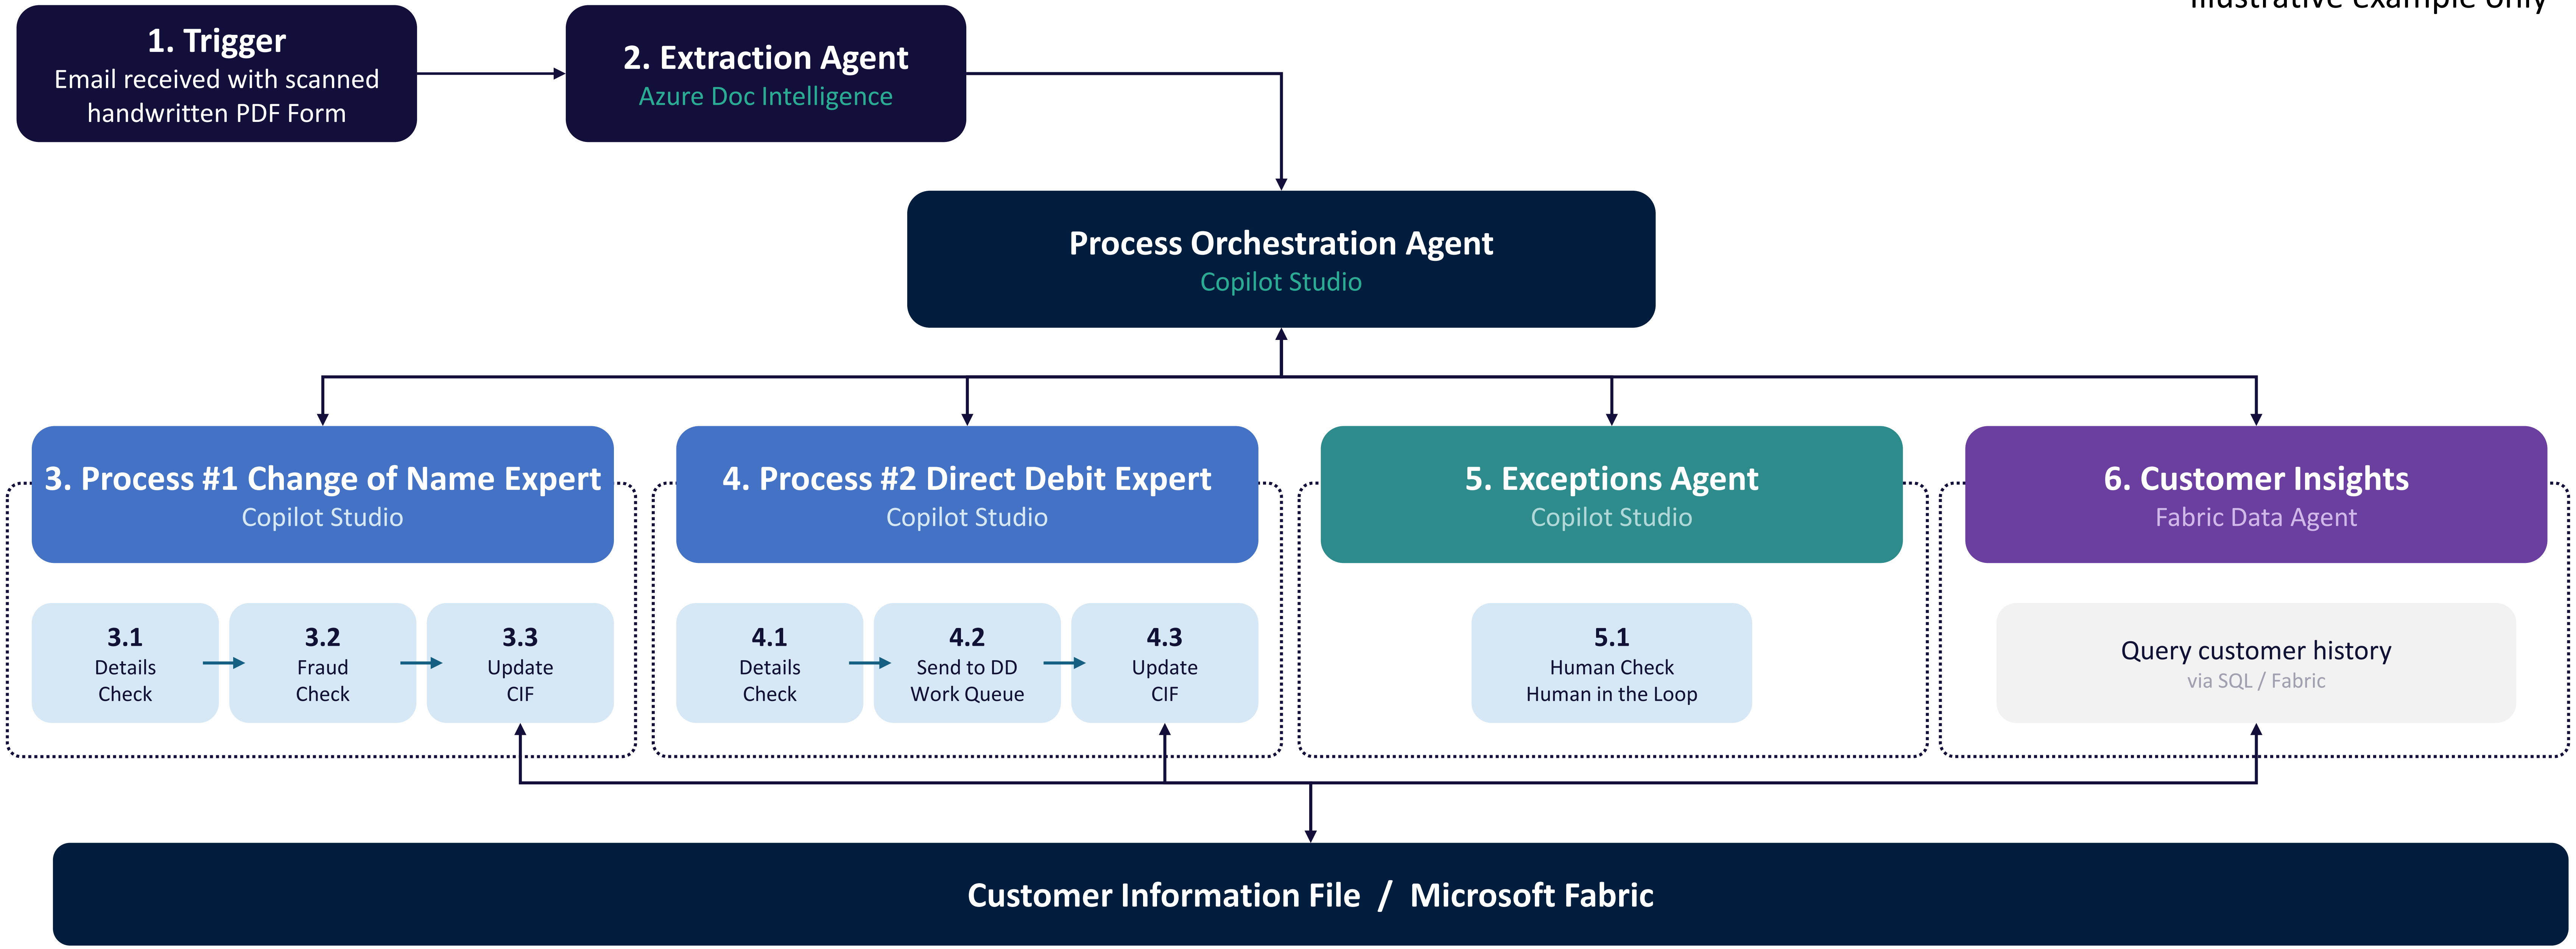

1. Trigger
Email received with scanned handwritten PDF Form

2. Extraction Agent
Azure Doc Intelligence

Process Orchestration Agent
Copilot Studio

3. Process #1 Change of Name Expert
Copilot Studio

- **3.1** Details Check
- **3.2** Fraud Check
- **3.3** Update CIF

4. Process #2 Direct Debit Expert
Copilot Studio

- **4.1** Details Check
- **4.2** Send to DD Work Queue
- **4.3** Update CIF

5. Exceptions Agent
Copilot Studio

- **5.1** Human Check Human in the Loop

6. Customer Insights
Fabric Data Agent

Query customer history
via SQL / Fabric

Customer Information File / Microsoft Fabric

DEMO

Solution Demo

CHANGE OF CUSTOMER AGENTIC SOLUTION

CHAPTER 4

Microsoft Fabric AI

1

DATA AGENTS

AI agents that converse with your OneLake data, enforce governance, and surface actionable insights

Agent queries customer history across lakehouse tables to validate address change requests

2

COPILOT IN FABRIC NOTEBOOKS

In-cell AI generates code, SQL completions, and data transformations across all Fabric workloads

Copilot writes the PySpark to detect anomalous address changes across customer records

3

AI FUNCTIONS

LLM-powered summarization, classification, and text generation on OneLake data in a single line of code

Classify change requests as routine, suspicious, or high-risk using AI functions over lakehouse data

CHAPTER 4

Power BI Copilot

1

NATURAL LANGUAGE INSIGHTS

Standalone Copilot finds and analyses any report or model. Ask questions, get instant visuals and summaries

"Show address changes by region this quarter" instantly generates a filtered dashboard

2

AI REPORT CREATION

Copilot creates entire report pages, selects best visuals, and auto-generates DAX measures with descriptions

Build a customer details audit report in seconds, with anomaly detection visuals auto-selected

3

EMBEDDED EVERYWHERE

Copilot in SharePoint, Teams, mobile apps, and org apps. Insights at every decision point

Service agent in Teams asks Copilot about a customer's change history mid-call via embedded report

SQL Server 2025 AI

1

COPILOT IN SSMS

Natural language to T-SQL in SSMS 21. Write, explain, fix, and optimise queries via chat

"Show all customers who changed address more than twice this year" generates the T-SQL instantly

2

VECTOR SEARCH

Native VECTOR data type and VECTOR_SEARCH() for semantic search directly in T-SQL. No external DB needed

Semantic search finds similar past address fraud cases by meaning, not just keyword matching

3

FABRIC MIRRORING

Zero-ETL mirroring streams SQL Server data to OneLake in near real time for AI analytics without impacting OLTP

Customer address changes in SQL Server are mirrored to Fabric for real-time anomaly analysis

SQL is the quiet hero

1

DETERMINISTIC ANCHOR

AI generates probabilistic outputs. SQL delivers exact, repeatable results every time. When precision matters, SQL is the final arbiter

An agent calculates a customer's exact account balance via SQL rather than estimating it from context

2

VALIDATION LAYER

Audit the SQL, not the prompt. Stored procedures and views create an explainable, version-controlled logic layer

A compliance team reviews the stored procedure an agent used, not the unpredictable prompt that triggered it

3

CONTROL PLANE

Schema engineering replaces prompt engineering. Well-designed tables, constraints, and relationships define what agents can and cannot do

Foreign keys prevent an agent from creating an order for a non-existent customer, no prompt guard needed

SQL is the quiet hero

4

STRUCTURED MEMORY

Agents need persistent, queryable memory. SQL tables give agents structured recall across sessions

An agent recalls a customer's full order history from SQL before recommending next steps

5

TRANSACTIONAL SAFETY

ACID guarantees mean agent actions either fully commit or fully roll back. No partial writes, no data corruption

An agent processing a refund either updates the order, adjusts inventory, and logs the action, or none of it happens

6

GOVERNANCE AT SCALE

Row-level security, audit logs, and role-based access built into SQL give agents guardrails enterprises already trust

Agents can only access rows they are authorised for, with every query logged for compliance

CHAPTER 5

Trust in AI Agents

The "**Agentic Era**" is here, marking a shift from chatbots that only converse to agents that take actions, invoking APIs, accessing databases, and overseeing their own long-term memory.

However, increased autonomy introduces new risks. How can we ensure that agents stay secure, compliant, and behave as expected?

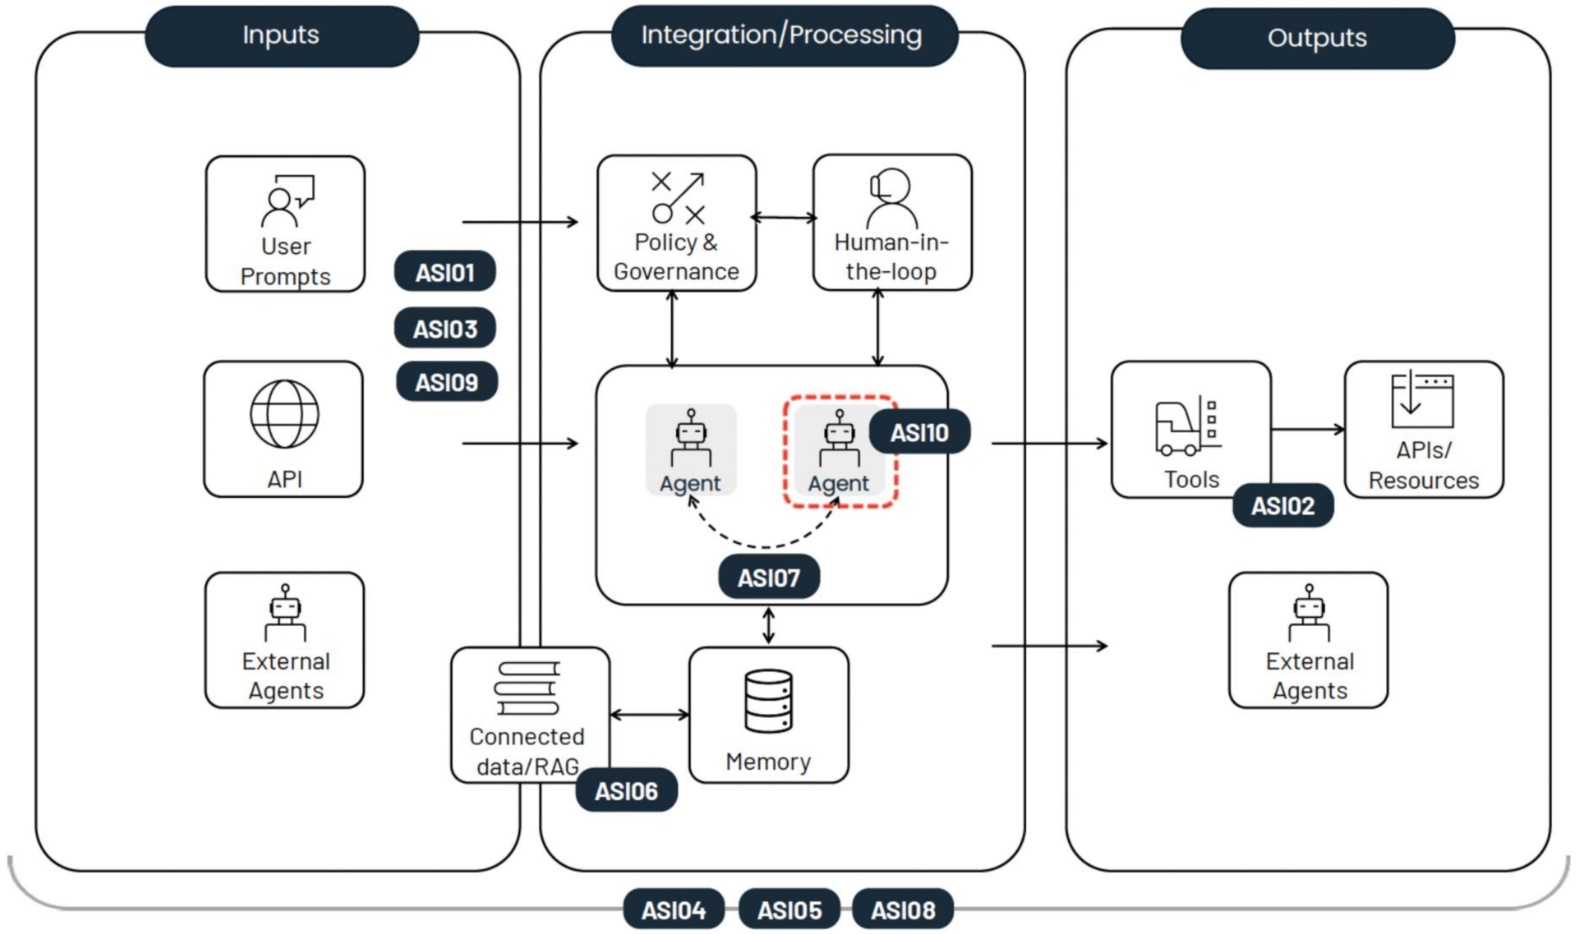

ASI01: Agent Goal Hijack

ASI02: Tool Misuse & Exploitation

ASI03: Identity & Privilege Abuse

ASI04: Agentic Supply Chain Vulnerabilities

ASI05: Unexpected Code Execution (RCE)

ASI06: Memory & Context Poisoning

ASI07: Insecure Inter-Agent Communication

ASI08: Cascading Failures

ASI09: Human-Agent Trust Exploitation

ASI10: Rogue Agents

Source: OWASP Top 10 for Agentic Applications 2026

CHAPTER 5

Trust in AI Agents

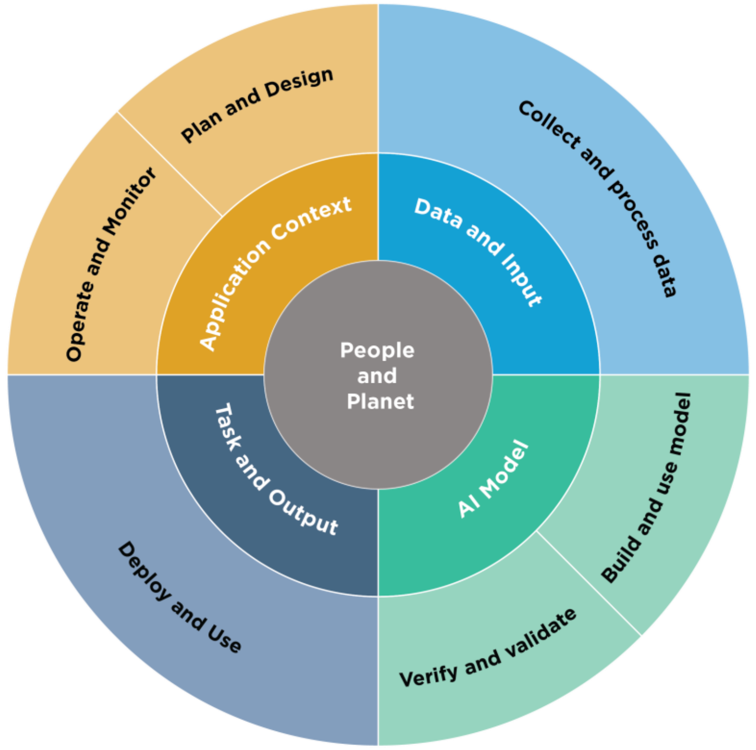

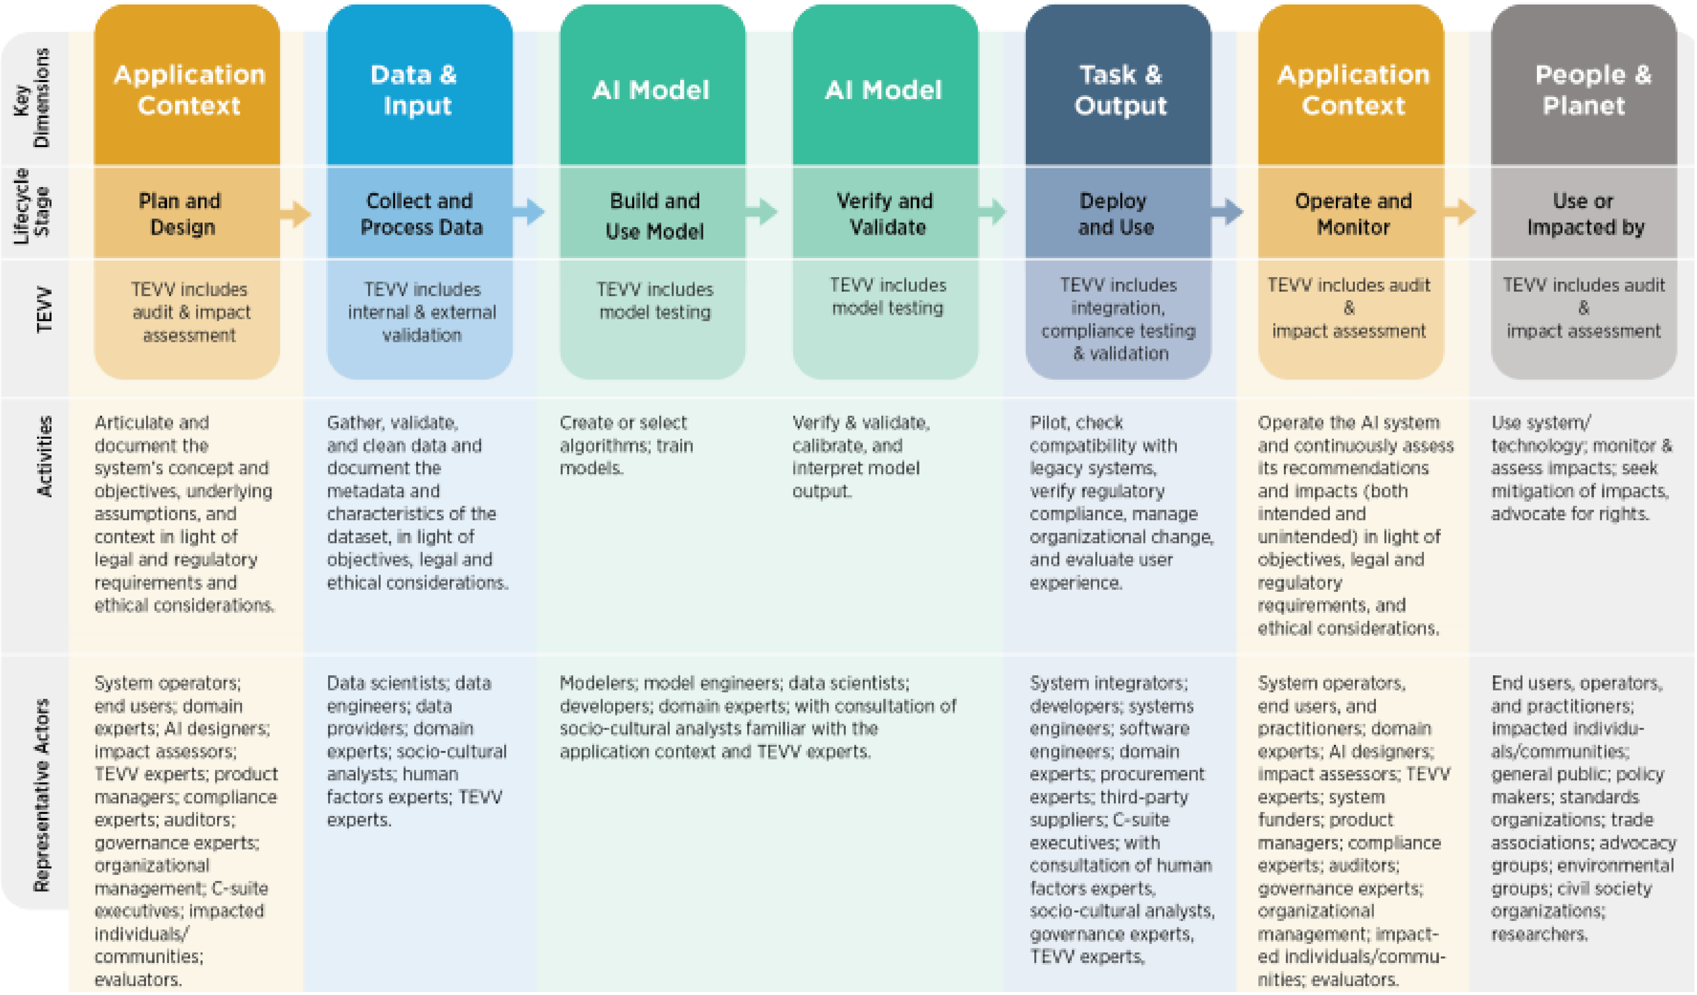

Key Dimensions	Application Context	Data & Input	AI Model	AI Model	Task & Output	Application Context	People & Planet
Lifecycle Stage	Plan and Design	Collect and Process Data	Build and Use Model	Verify and Validate	Deploy and Use	Operate and Monitor	Use or Impacted by
TEVV	TEVV includes audit & impact assessment	TEVV includes internal & external validation	TEVV includes model testing	TEVV includes model testing	TEVV includes integration, compliance testing & validation	TEVV includes audit & impact assessment	TEVV includes audit & impact assessment
Activities	Articulate and document the system's concept and objectives, underlying assumptions, and context in light of legal and regulatory requirements and ethical considerations.	Gather, validate, and clean data and document the metadata and characteristics of the dataset, in light of objectives, legal and ethical considerations.	Create or select algorithms; train models.	Verify & validate, calibrate, and interpret model output.	Pilot, check compatibility with legacy systems, verify regulatory compliance, manage organizational change, and evaluate user experience.	Operate the AI system and continuously assess its recommendations and impacts (both intended and unintended) in light of objectives, legal and regulatory requirements, and ethical considerations.	Use system/ technology; monitor & assess impacts; seek mitigation of impacts, advocate for rights.
Representative Actors	System operators; end users; domain experts; AI designers; impact assessors; TEVV experts; product managers; compliance experts; auditors; governance experts; organizational management; C-suite executives; impacted individuals/ communities; evaluators.	Data scientists; data engineers; data providers; domain experts; socio-cultural analysts; human factors experts; TEVV experts.	Modelers; model engineers; data scientists; developers; domain experts; with consultation of socio-cultural analysts familiar with the application context and TEVV experts.		System integrators; developers; systems engineers; software engineers; domain experts; procurement experts; third-party suppliers; C-suite executives; with consultation of human factors experts, socio-cultural analysts, governance experts, TEVV experts,	System operators, end users, and practitioners; domain experts; AI designers; impact assessors; TEVV experts; system funders; product managers; compliance experts; auditors; governance experts; organizational management; impacted individuals/communities; evaluators.	End users, operators, and practitioners; impacted individuals/communities; general public; policy makers; standards organizations; trade associations; advocacy groups; environmental groups; civil society organizations; researchers.

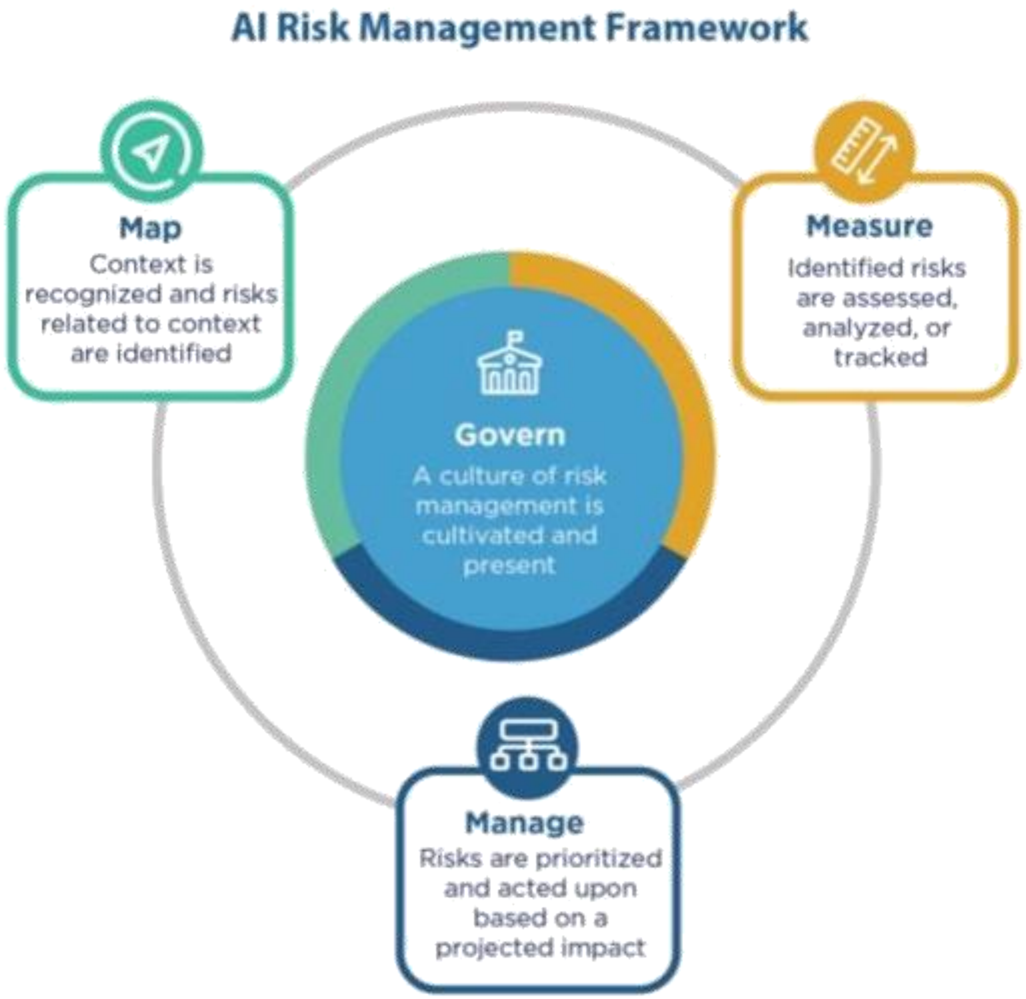

AI Risk Management Framework

Map
Context is recognized and risks related to context are identified

Measure
Identified risks are assessed, analyzed, or tracked

Govern
A culture of risk management is cultivated and present

Manage
Risks are prioritized and acted upon based on a projected impact

Source: NIST AI RMF 100-1

Trust in AI Agents

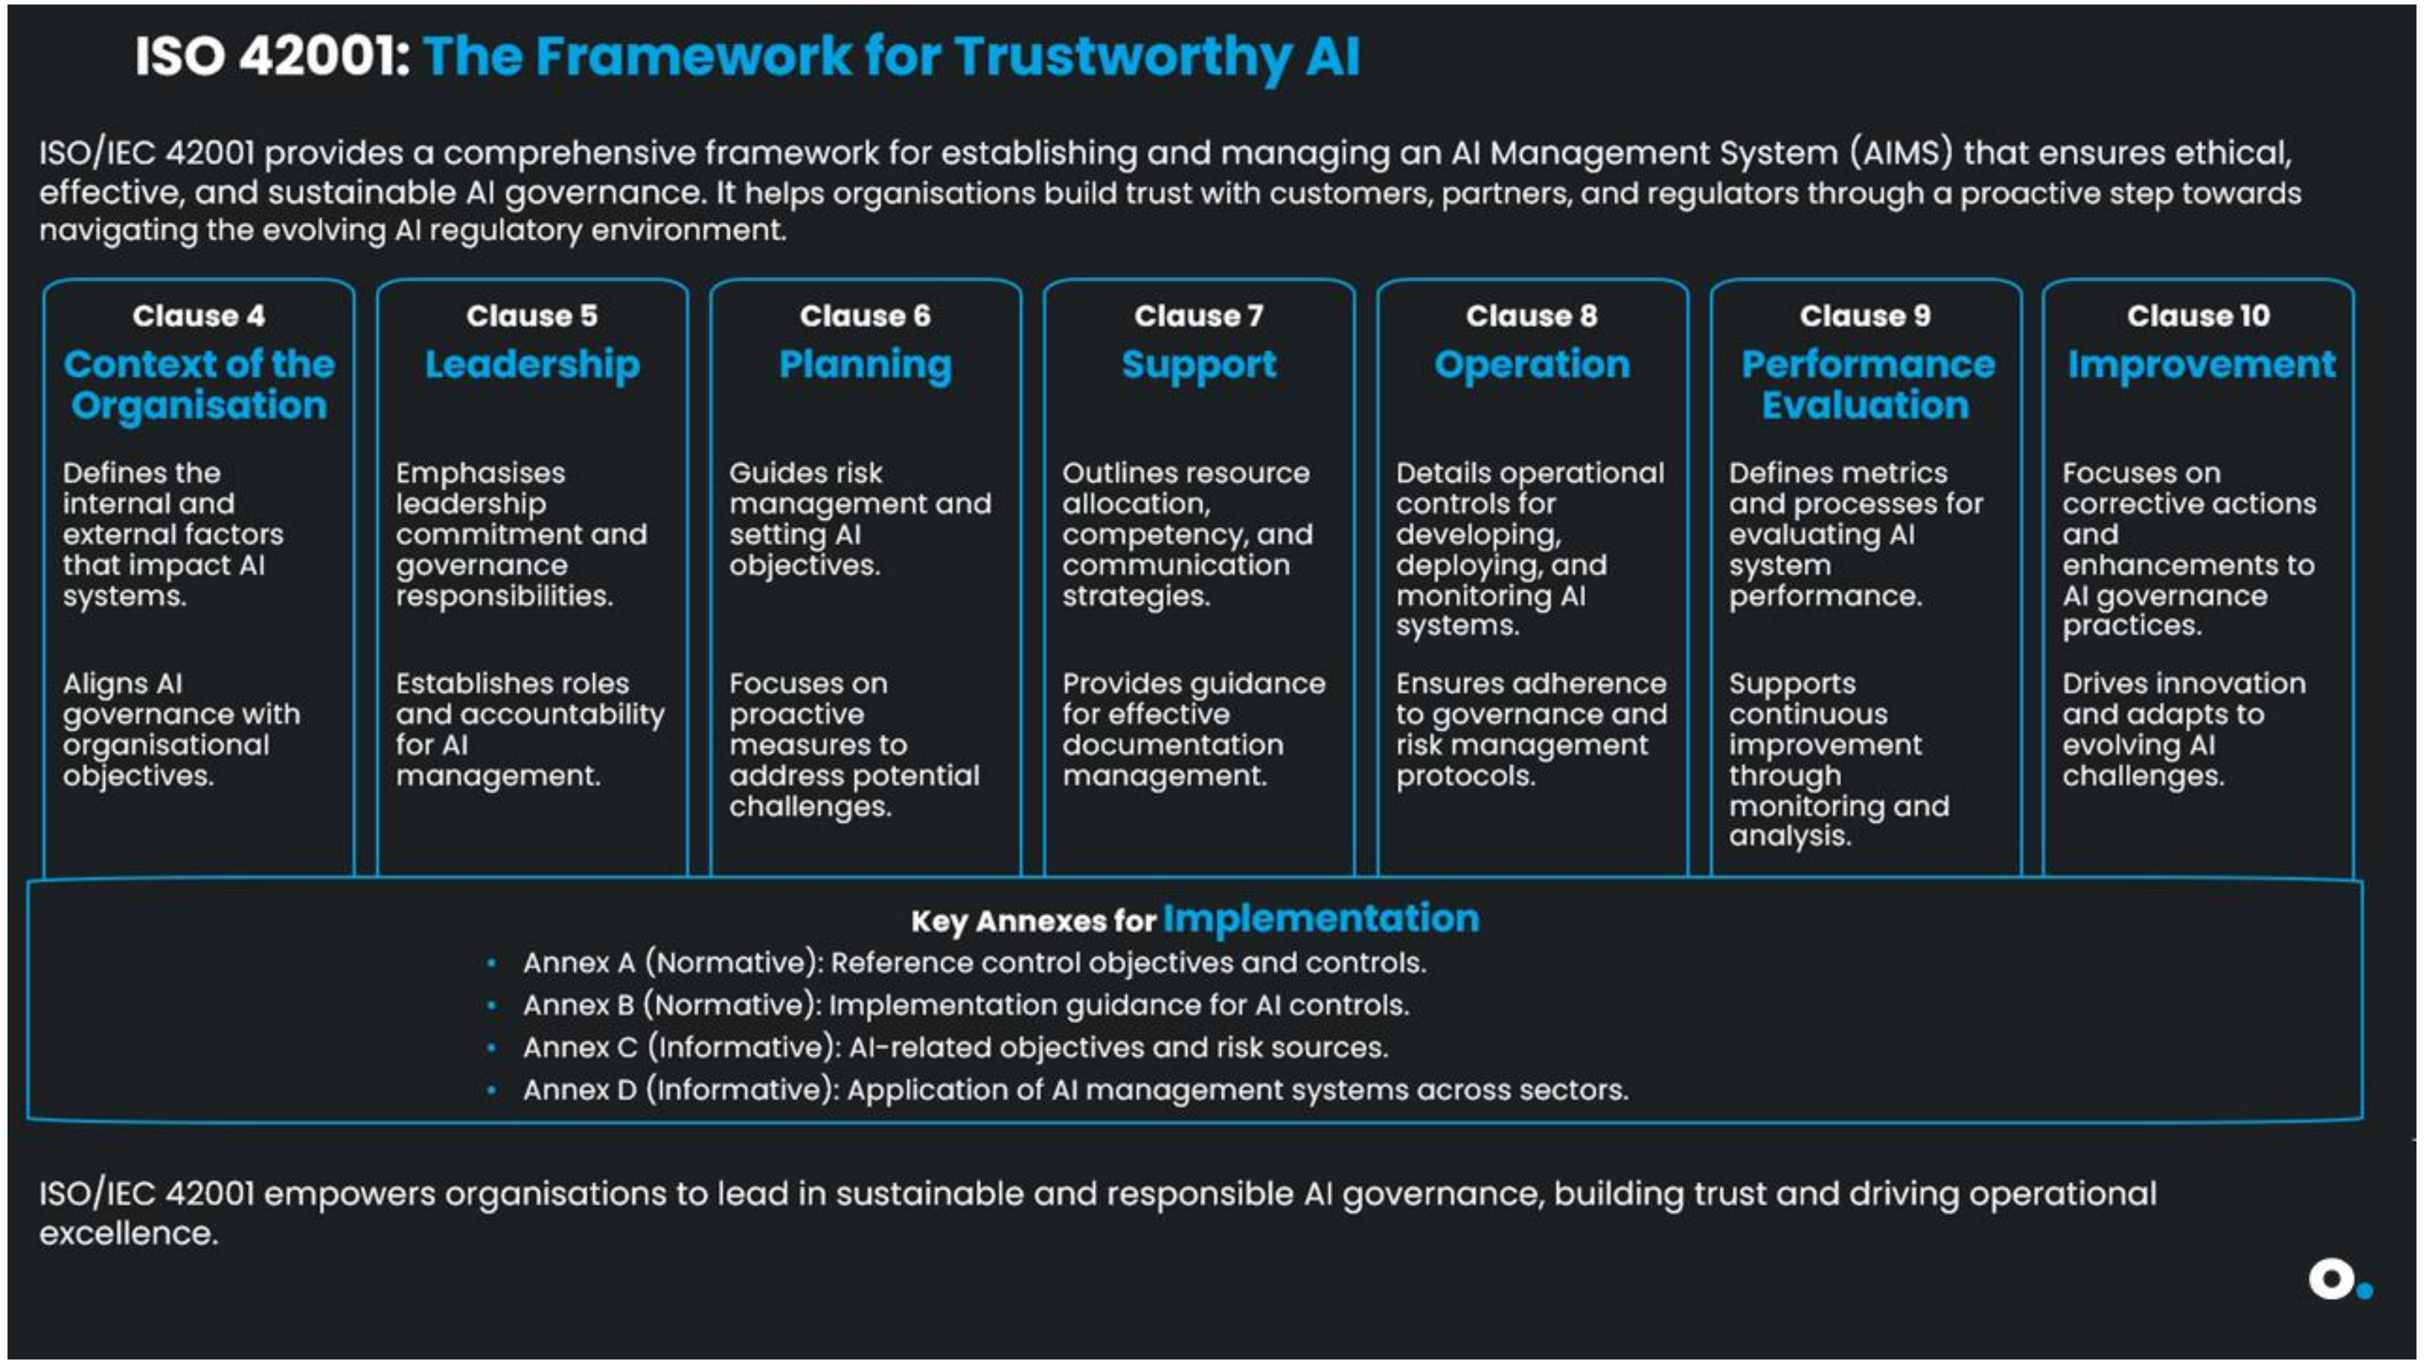

ISO 42001: The Framework for Trustworthy AI

ISO/IEC 42001 provides a comprehensive framework for establishing and managing an AI Management System (AIMS) that ensures ethical, effective, and sustainable AI governance. It helps organisations build trust with customers, partners, and regulators through a proactive step towards navigating the evolving AI regulatory environment.

Clause 4	Clause 5	Clause 6	Clause 7	Clause 8	Clause 9	Clause 10
Context of the Organisation	**Leadership**	**Planning**	**Support**	**Operation**	**Performance Evaluation**	**Improvement**
Defines the internal and external factors that impact AI systems.	Emphasises leadership commitment and governance responsibilities.	Guides risk management and setting AI objectives.	Outlines resource allocation, competency, and communication strategies.	Details operational controls for developing, deploying, and monitoring AI systems.	Defines metrics and processes for evaluating AI system performance.	Focuses on corrective actions and enhancements to AI governance practices.
Aligns AI governance with organisational objectives.	Establishes roles and accountability for AI management.	Focuses on proactive measures to address potential challenges.	Provides guidance for effective documentation management.	Ensures adherence to governance and risk management protocols.	Supports continuous improvement through monitoring and analysis.	Drives innovation and adapts to evolving AI challenges.

Key Annexes for Implementation

- Annex A (Normative): Reference control objectives and controls.
- Annex B (Normative): Implementation guidance for AI controls.
- Annex C (Informative): AI-related objectives and risk sources.
- Annex D (Informative): Application of AI management systems across sectors.

ISO/IEC 42001 empowers organisations to lead in sustainable and responsible AI governance, building trust and driving operational excellence.

Source: ISO/IEC 42001 Certification

Trust in AI Agents

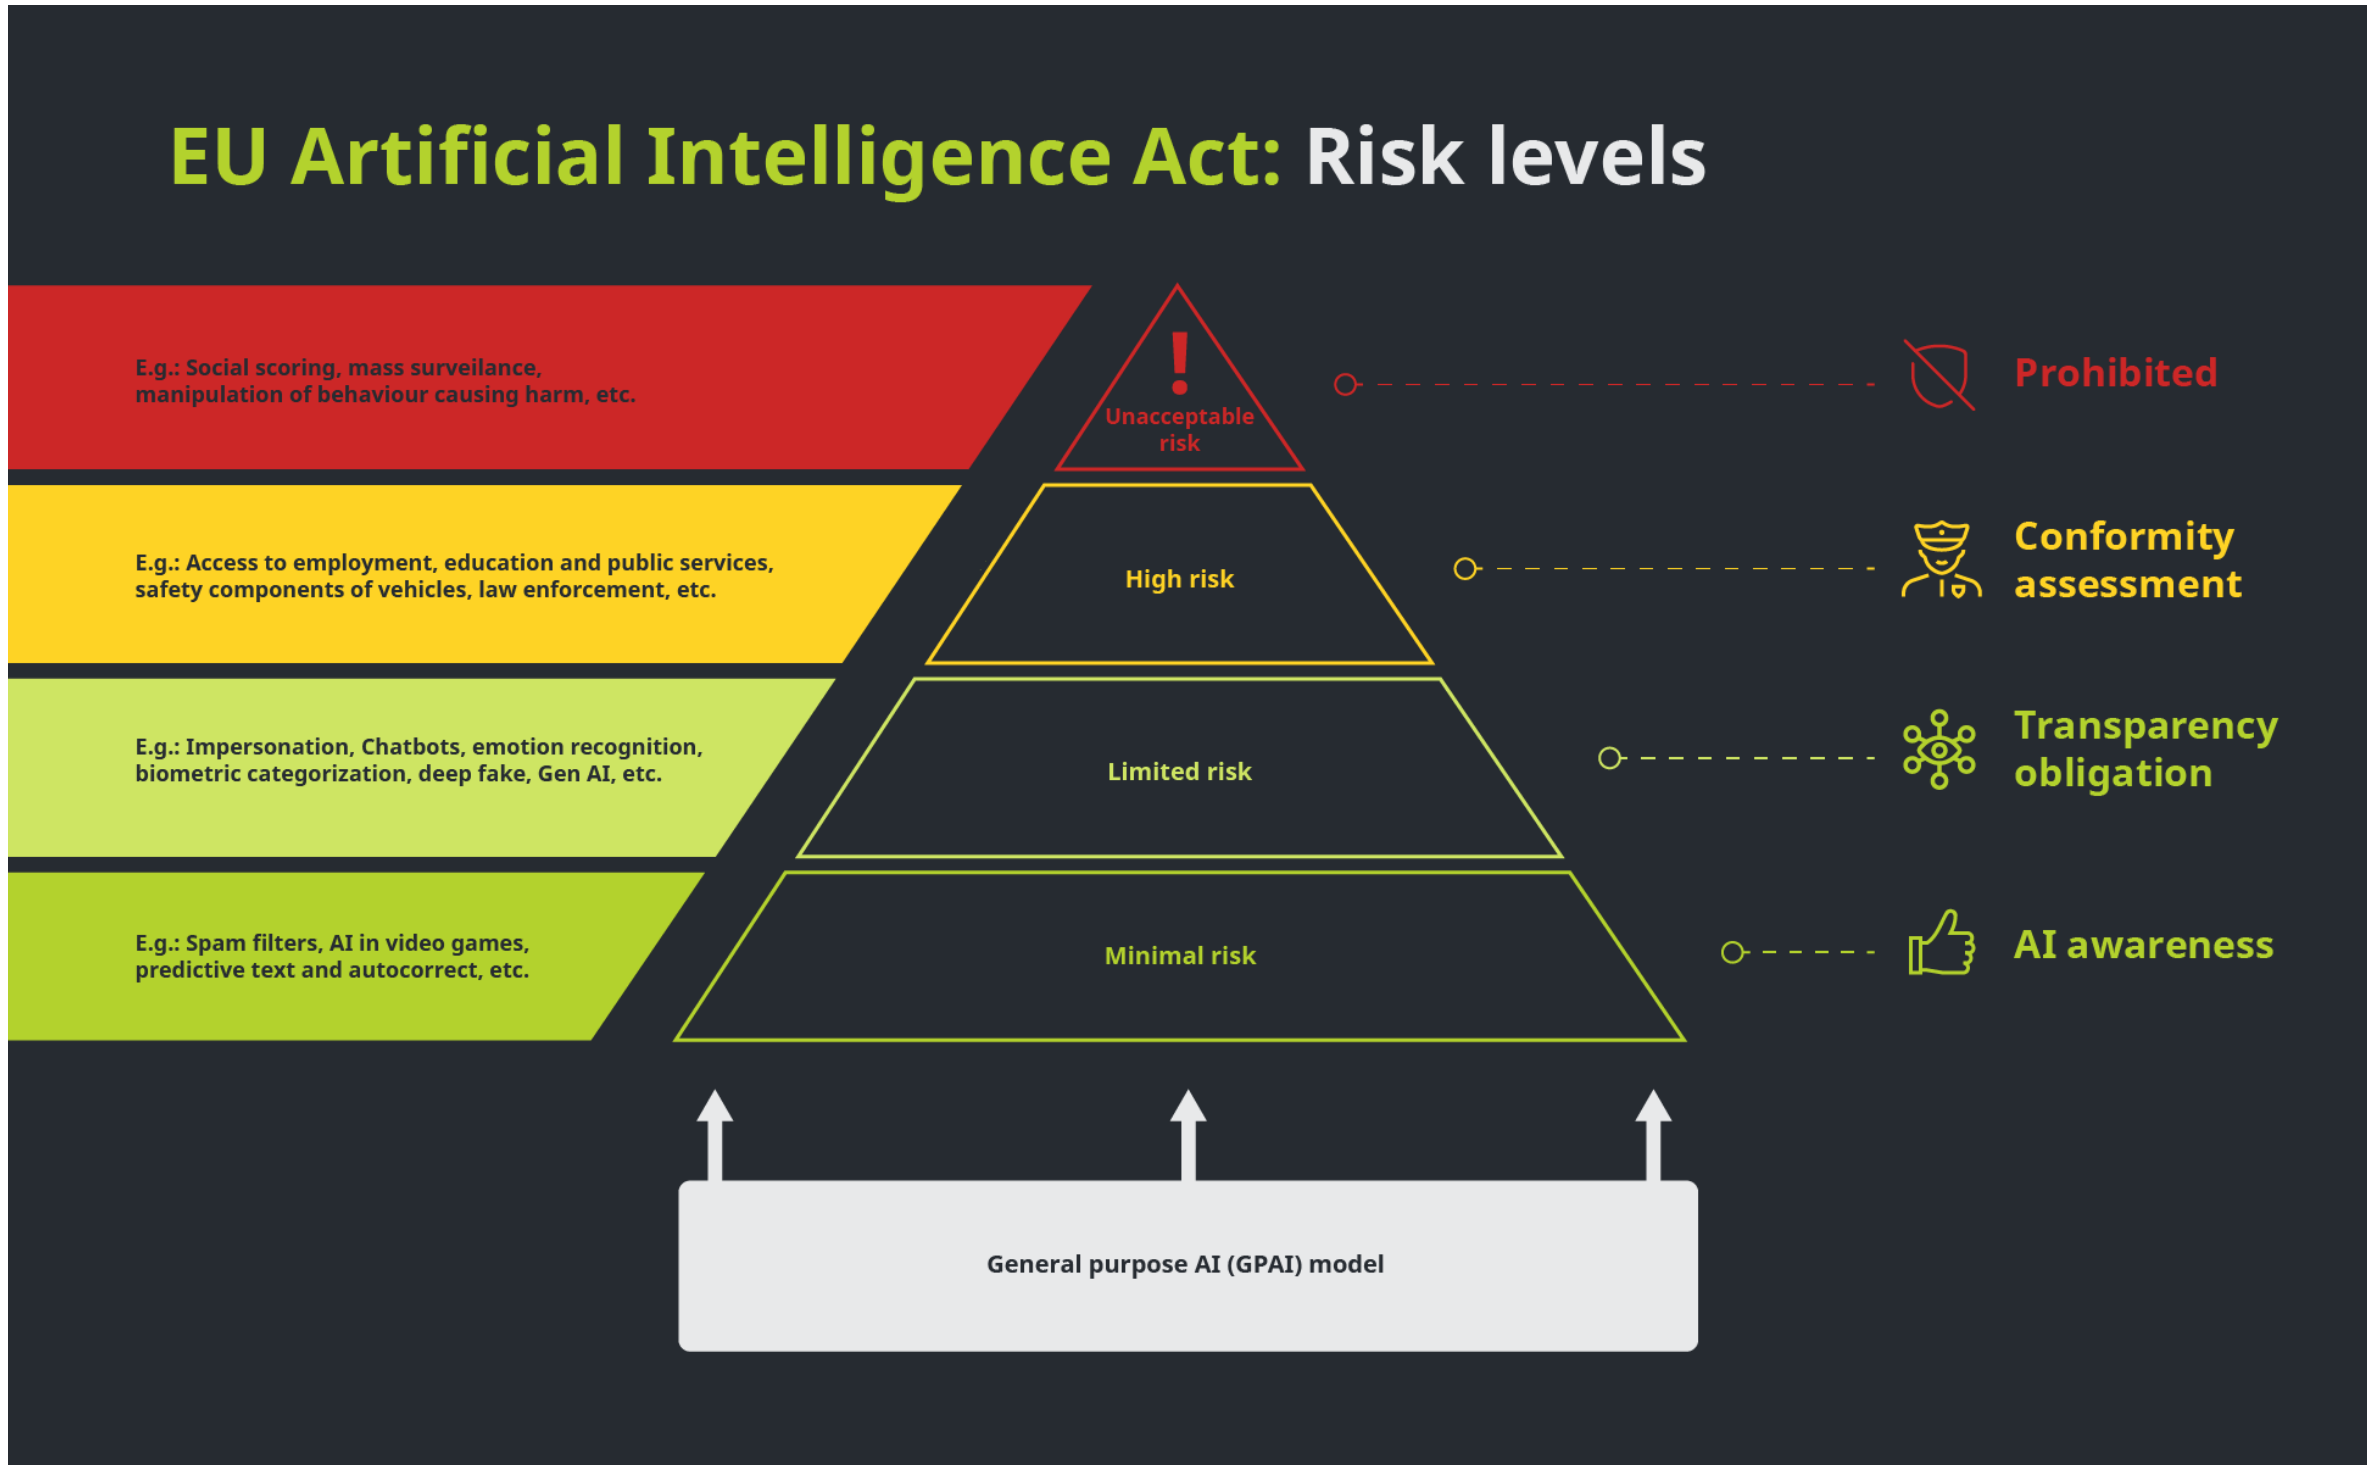

EU Artificial Intelligence Act: Risk levels

E.g.: Social scoring, mass surveilance, manipulation of behaviour causing harm, etc.

Unacceptable risk

Prohibited

E.g.: Access to employment, education and public services, safety components of vehicles, law enforcement, etc.

High risk

Conformity assessment

E.g.: Impersonation, Chatbots, emotion recognition, biometric categorization, deep fake, Gen AI, etc.

Limited risk

Transparency obligation

E.g.: Spam filters, AI in video games, predictive text and autocorrect, etc.

Minimal risk

AI awareness

General purpose AI (GPAI) model

Source: EU AI ACT Risk Levels

Trust in AI Agents

GOVERN

- Establish **Security by Design** culture across agent lifecycle
- Define **accountability chains**: who is liable for unauthorized API calls?
- Evidence-based governance: policies, attestations, tollgates
- Align with **NIST AI RMF 100-1**

MAP

- Identify full **attack surface** of each agent
- Catalog: LLM, system prompt, tools (APIs), RAG sources
- Map **agent-to-agent trust boundaries**
- Assess identity chain: **Entra Agent ID**, RBAC, service principals

MEASURE

- Conduct **Red Teaming** for agentic-specific threats
- Assess **Groundedness** scores and hallucination rates
- Proactive stress-testing before deployment
- Evaluate against **OWASP ASI01–ASI10**

MANAGE

- Deploy **guardrails** and continuous monitoring
- Prioritize: **Excessive Agency** (LLM08) **Goal Hijack** (ASI01)
- Integrate **Defender for Cloud** + Defender for AI
- Implement **kill switches** and circuit breakers

Source: Abhi Singh, Microsoft Defender for Cloud Blog 2026; Microsoft Architecting Trust Framework

CLOSE

Next Steps

When you go back to work on Monday:

1

Continue the AI Agents conversation

2

Learn by building low risk, high ROI agents

3

Promote trust and security by design

CLOSE

Thank you.

How was the session?

Complete Session Surveys in Whova for your chance to WIN PRIZES!

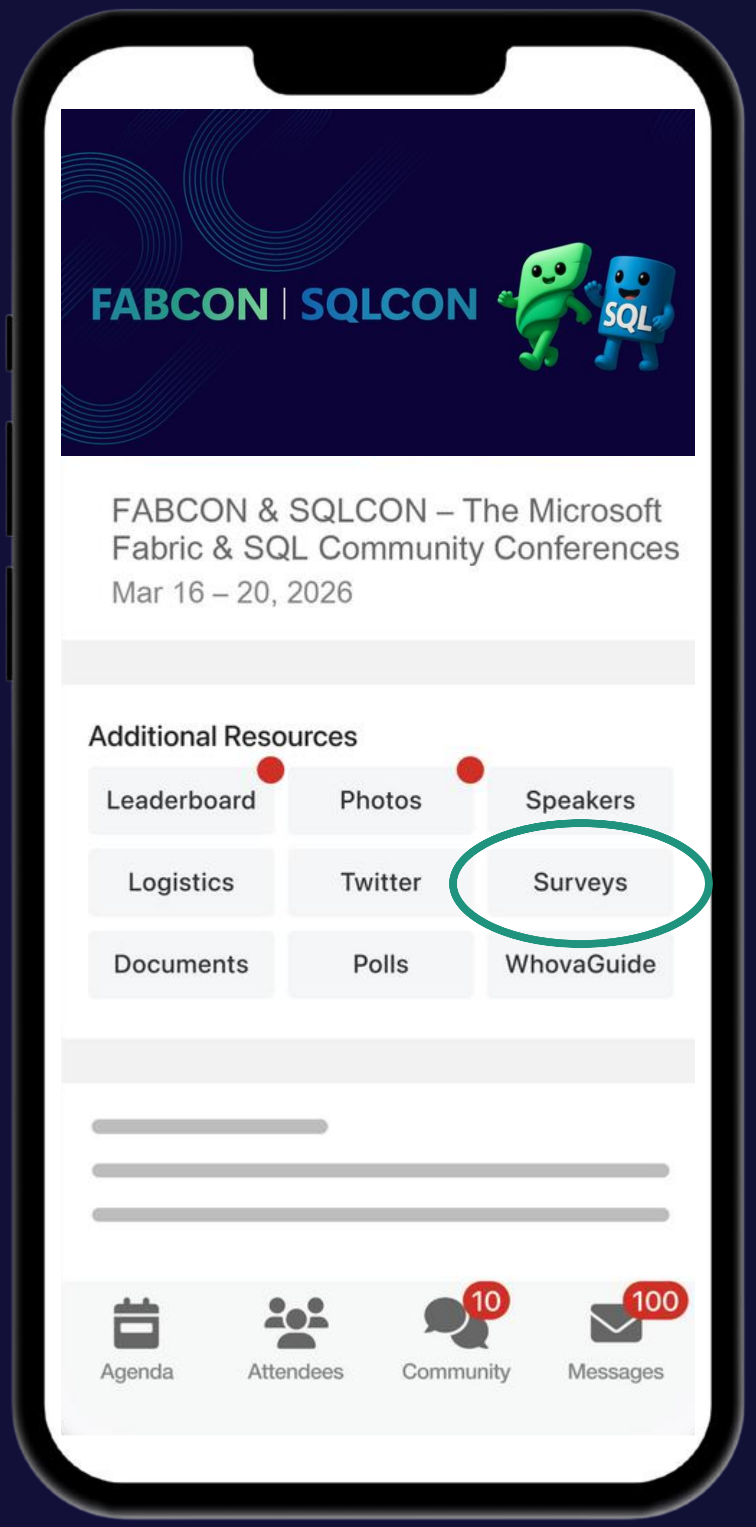

FABCON & SQLCON – The Microsoft Fabric & SQL Community Conferences
Mar 16 – 20, 2026

Get Two Fabric Certifications for FREE

Attendees of FABCON can take the Fabric Analytics Engineer or Fabric Data Engineer exam for free. Be part of the 2 fastest growing role-based certifications in Microsoft history.

Request your voucher by March 23, 2026.

https://aka.ms/fabcon/cert100