


ANALYST SESSIONS SUMMARY

# STKI Summit 2026

## Internet of Agents

### Executive Analyst Synthesis

Consolidated insights across all analyst sessions


01 CIO Macro Trends

02 The Trust Layer

03 Agentic Operating Model

04 Future of Work & ARP

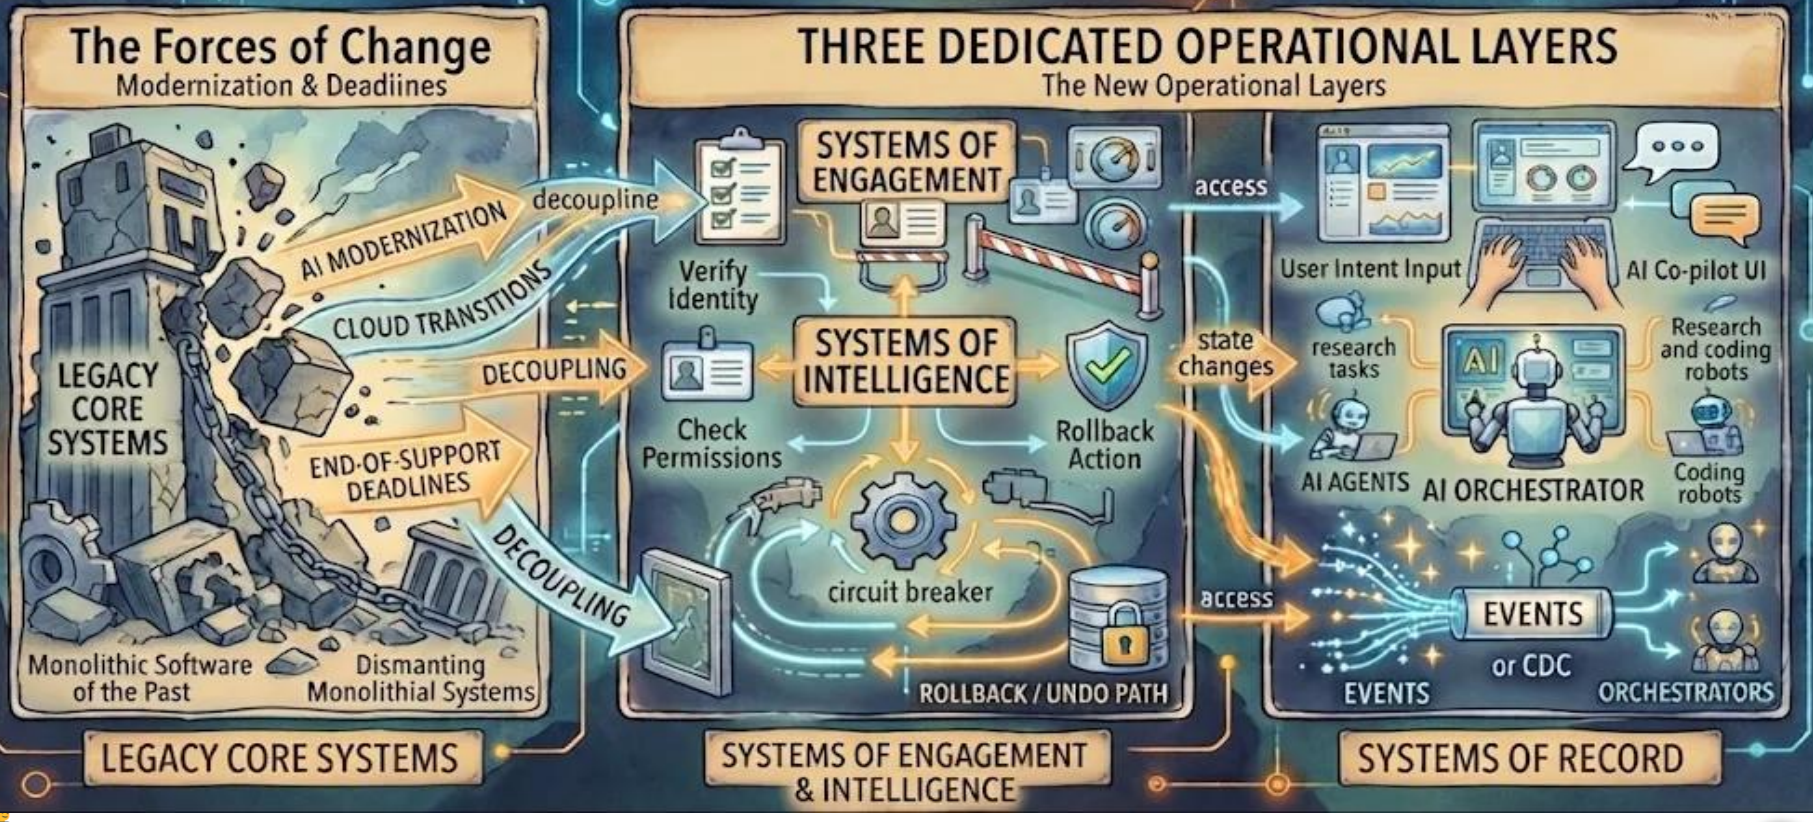
05 MCP & A2A Foundations



Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Enterprise technology is undergoing a once-in-a-generation upgrade cycle.

The monolithic software applications of the past are decoupling into three dedicated operational layers.




STKI.INFO Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# 2027–2028: Horizon 1: the clear road (first 100 mt)

## Implement

### Agentic workflows

Identify & implement workflows outside your old legacy systems (ERP, CRM, Core)

### Transactional systems evolve

Prepare transactional systems for headless agentic access as systems of records.

### Data Strategy & Implementation

Implement a context-driven data strategy that will support agentic workflows





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# The first 100 meters: priorities for 2027–2028

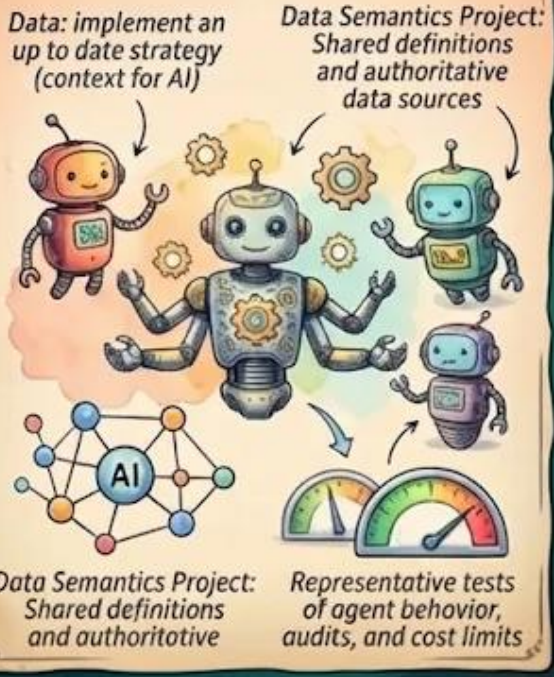
## Systems of Record (SoR)

*Selective modernization of transaction core*




## Systems of Intelligence (SoI)

*Build trusted enterprise context*




## Systems of Engagement (SoE)

*Redesign engagement*





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# 2029–2030: Horizon 2: the next stretch (100-300 mt)
# plan & start implementing the architecture

## Start decoupling

Separate engagement, intelligence and record systems.

## Coordination

Develop agentic workflows that will work across systems.

## Agent marketplace

Prepare the enterprise's agent marketplace





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

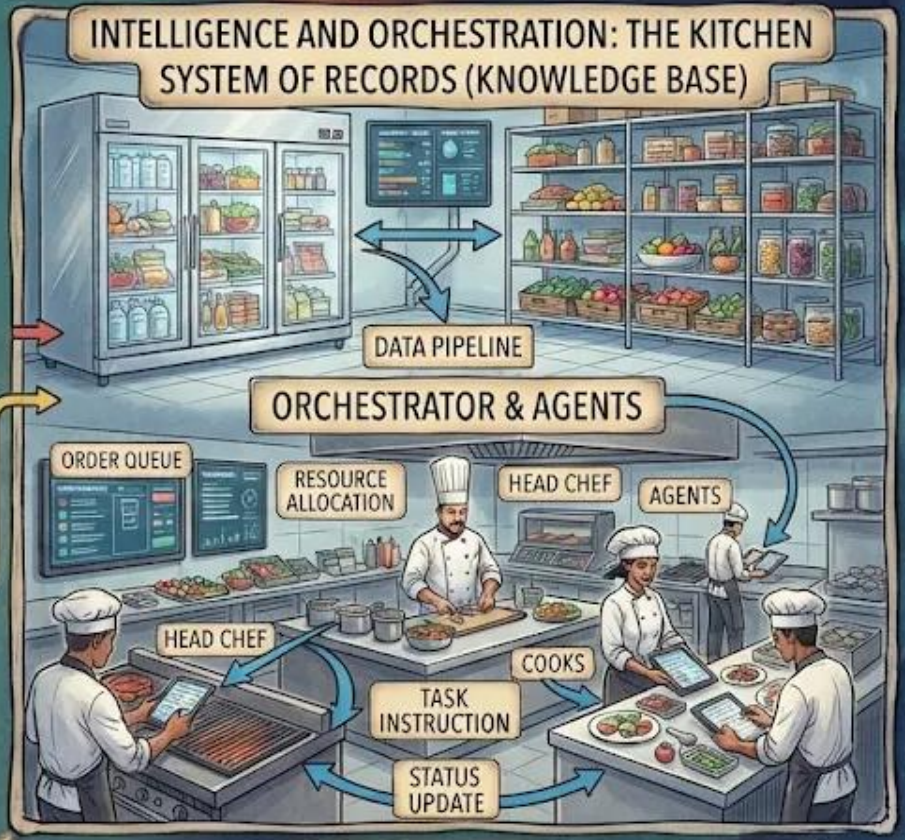




Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Architectural Metaphor: The High-Level Restaurant as an Enterprise System




INTEGRATED ENTERPRISE: **WAITER** (Engagement) → **CHEF** (Orchestrator/Intelligence) → **AGENTS** (Cooks) → **FRIDGE/PANTRY** (Records)

# The Decoupled Enterprise

identity, authorization, and audit span every layer —
a target architecture for interaction, reasoning, and transactional state

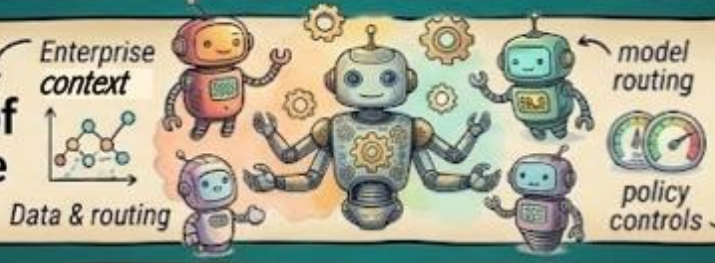
Identity, Authorization & Auditability

## SoE: Systems of Engagement




**Express Intent**
(Task-specific interfaces, approvals, and exceptions)

## SoI: Systems of Intelligence




**Coordinate Work**
(Enterprise context, model routing, and policy controls)

## SoR: Systems of Record




**Protect State**
(Authoritative records and validated transactions)


Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# The Headless Shift: Execution Beyond the GUI

identity, authorization, and audit span every layer —
less reliance on manual navigation, agents control where it matters

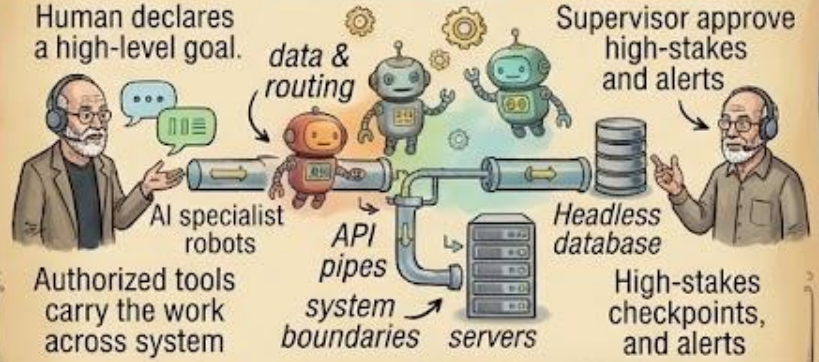
Systematic Transition & Automation

## Human-Operated Workflow
(Legacy GUI)




## Agent-Operated Workflow
(Headless Execution)

**1. Express Intent** → **2. Execute via APIs** → **3. Review Exceptions**



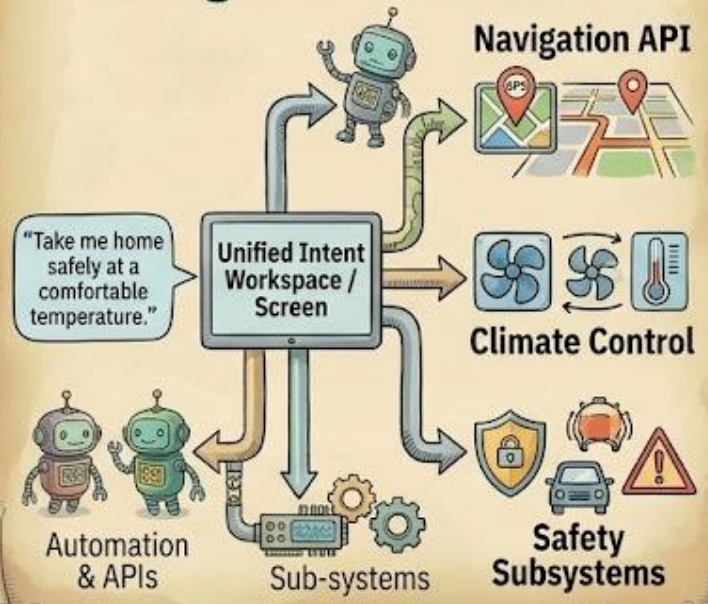

**People retain oversight of ambiguity and consequential decisions.**


Copyright@STKI, 2026 Do not remove source or attribution from any slide, graph or portion of graph

# Example of a “unified intent workspace” across systems

Systematic Transition & Automation

## One intent screen coordinates “several” workflows: navigation and comfort



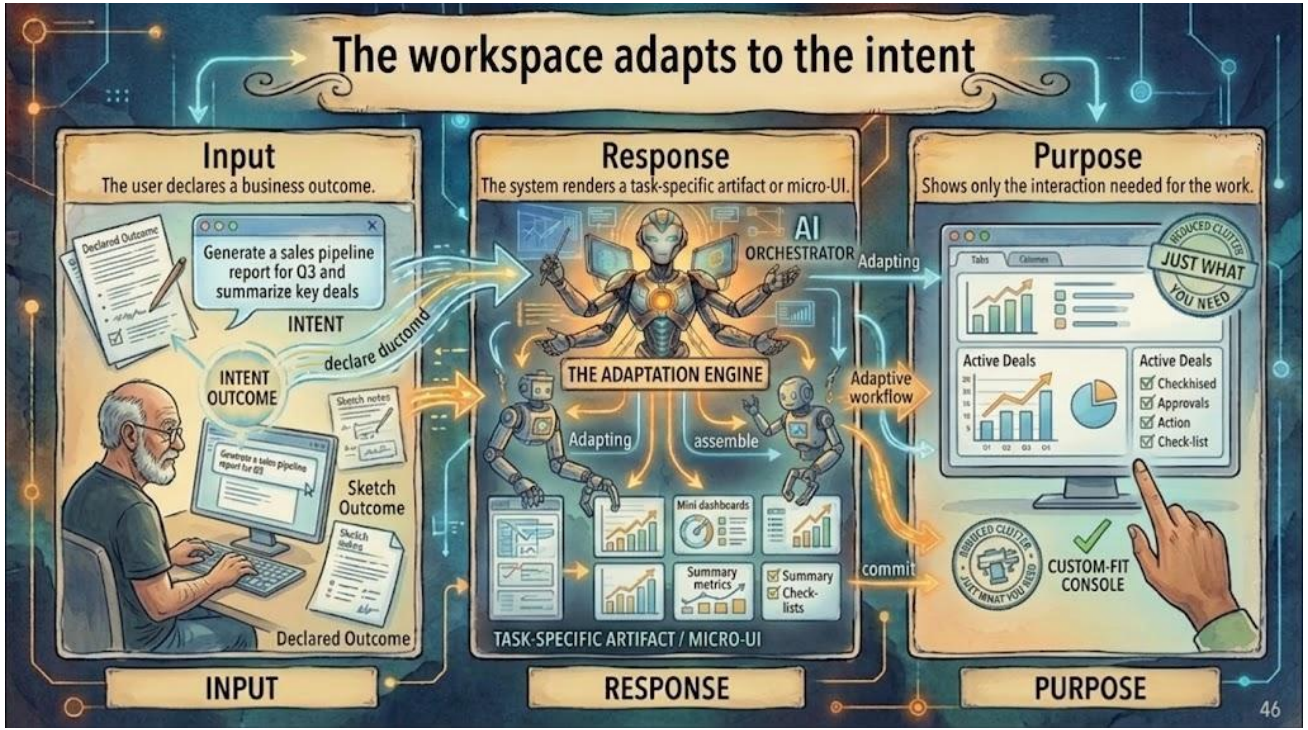





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Interfaces are *temporary* and built for a task

## The user declares an outcome; the system renders just the dynamic-UI needed and shows only what the work requires.





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Systems of Intelligence: the reasoning fabric

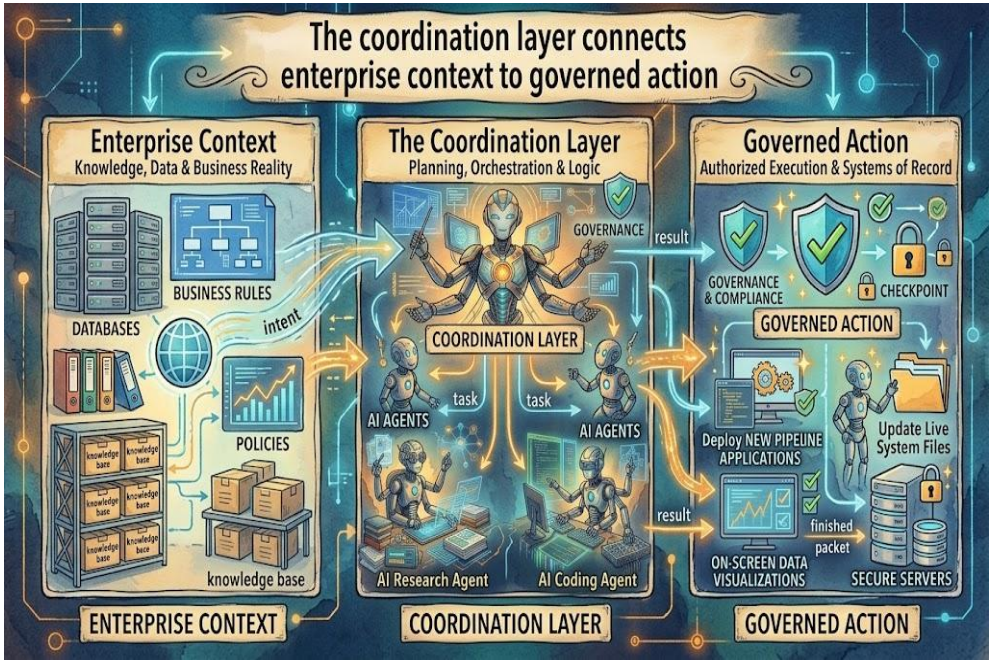
*The coordinated layer of agents connects enterprise context to governed action*

## Enterprise context

Connects business definitions, relevant documents, and live system states.

## Model routing

Selects models by capability, risk, latency, and cost.



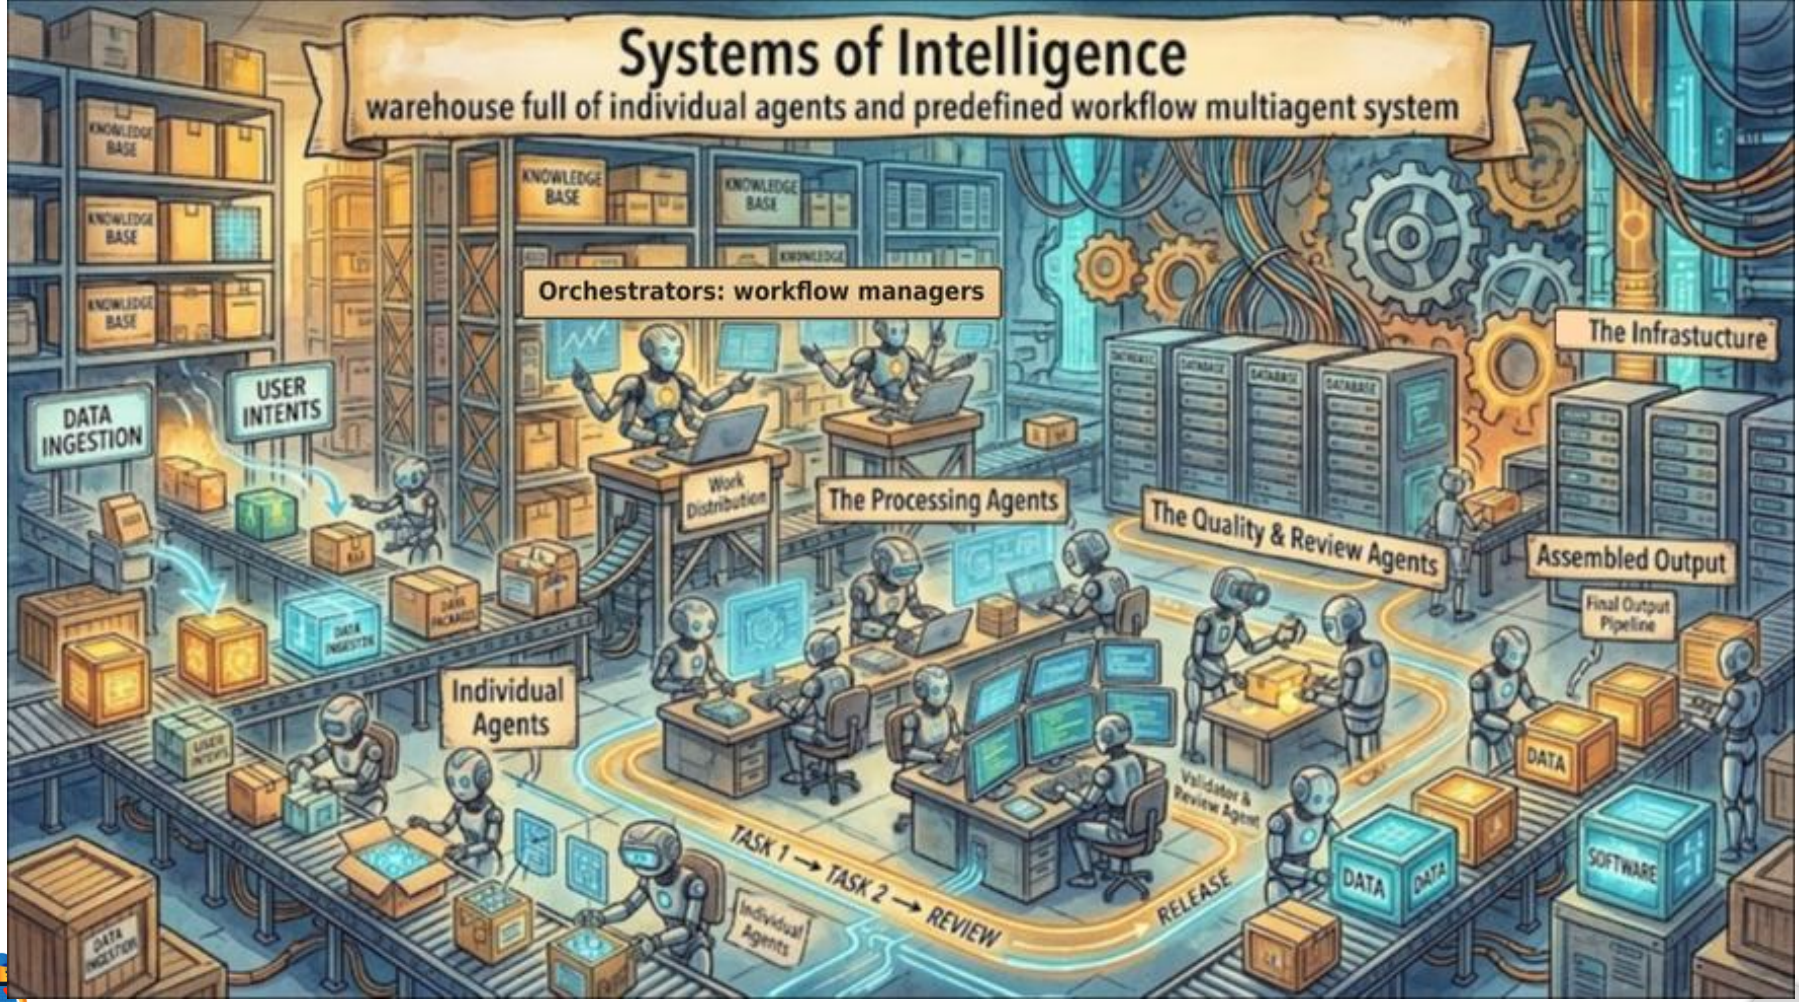

## Task orchestration

Breaks goals into steps. Coordinates specialist agents and tools.

## Runtime governance

Applies authorization checks, approval gates, spending limits, and audit logging.


Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

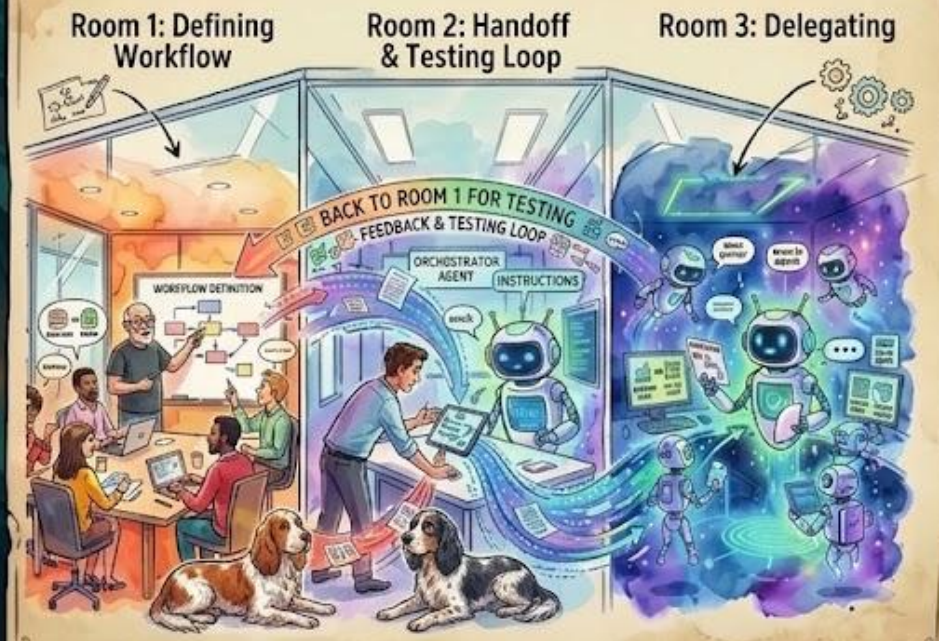




Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# SYSTEMS OF INTELLIGENCE: ENTERPRISES MUST DECIDE BUY OR BUILD

## BUILD: CONFIGURE WORKFLOW MULTIAGENT SYSTEMS IN-HOUSE



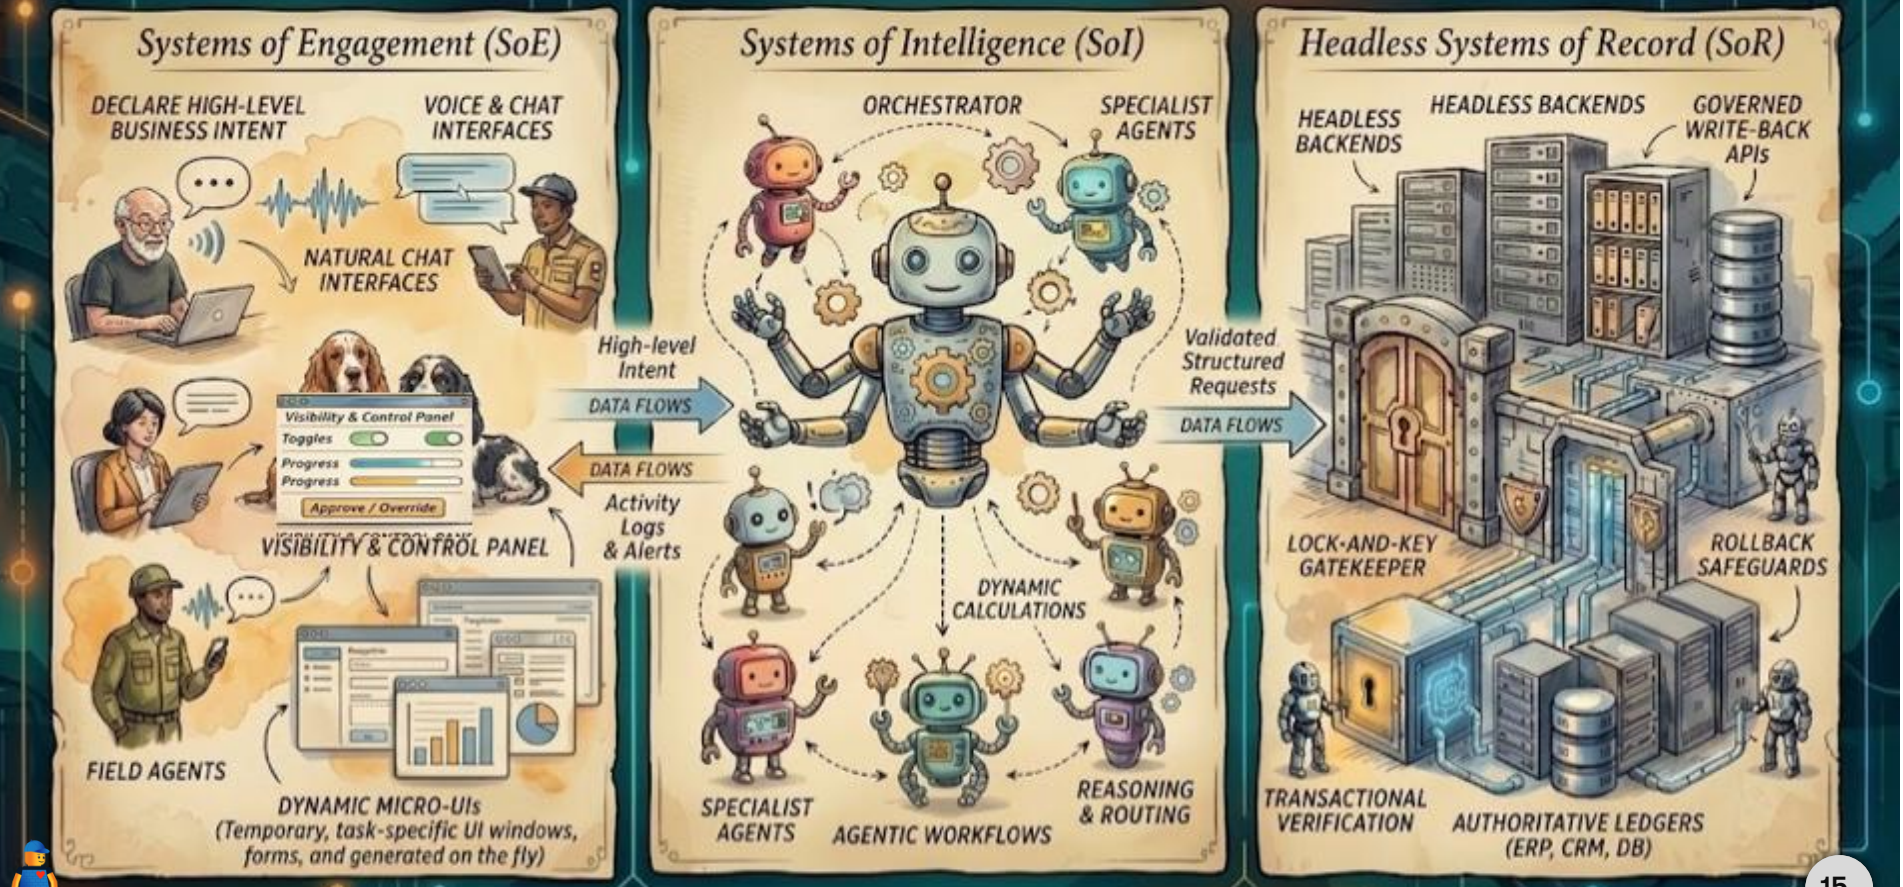

## BUY: BUY AND SELL INDIVIDUAL AGENTS (AGENT SUPERMARKET)



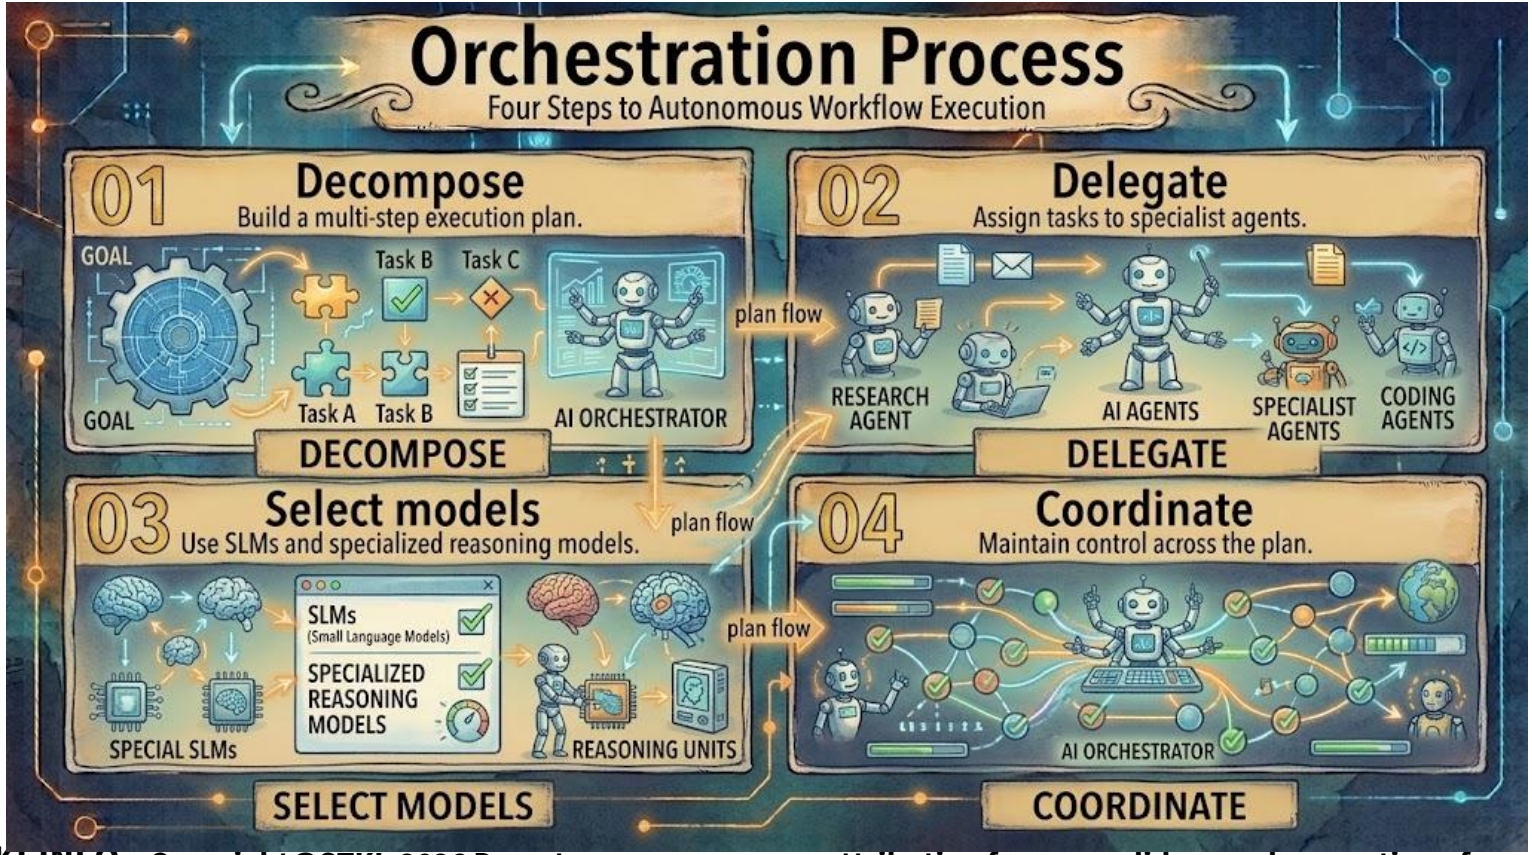



STKI.INFO Copyright@STKI 2026 Do not remove source or attribution from any slide, graph or portion of graph

# Transforming Legacy Core into Headless Systems of Record



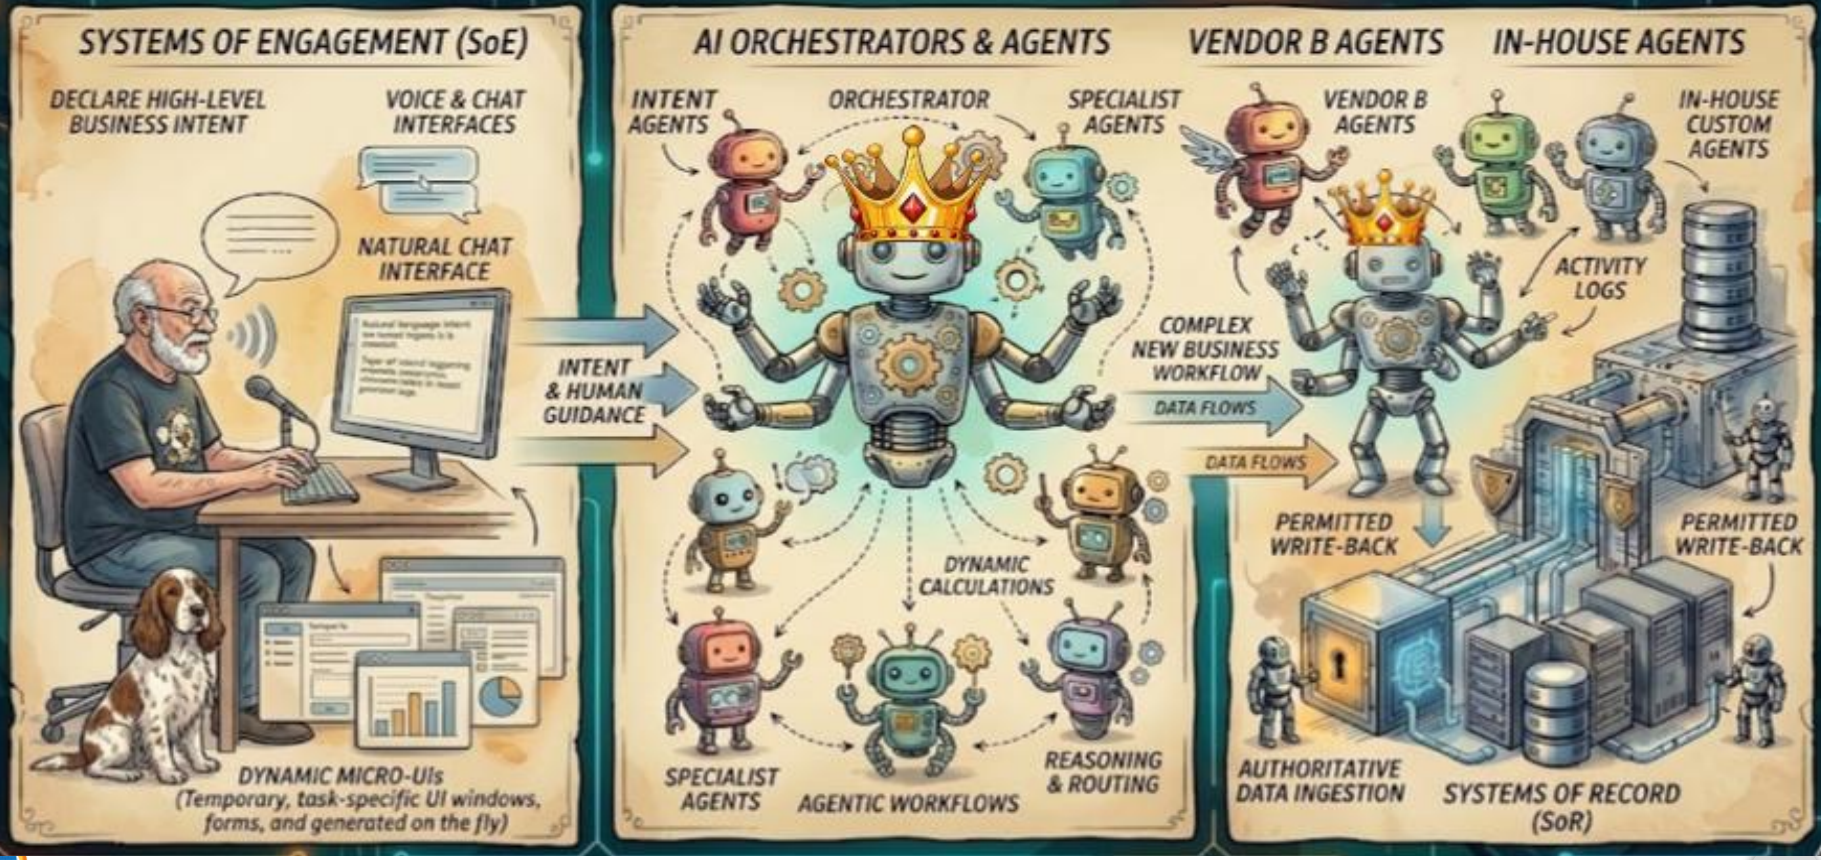
# The ”orchestrator” breaks a goal into work





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Systems of Intelligence (SoI): Multi-Orchestrator & Multi-Vendor Agent Fabric





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# *agents are probabilistic, records are deterministic.*

## A model can produce a *plausible but wrong answer*, while core records must follow explicit rules.





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# 2031 and beyond: Horizon 3: (over 300 mt)

## Position without pretending to know

***long-term project choices need assumptions that can change.***

### Finish the uncoupling project

Systems of engagement, intelligence and records work together based on the new ***uncoupled architecture***

### What remains uncertain

Technological advancements & maturity, adoption speed, and economic models.

### What still helps

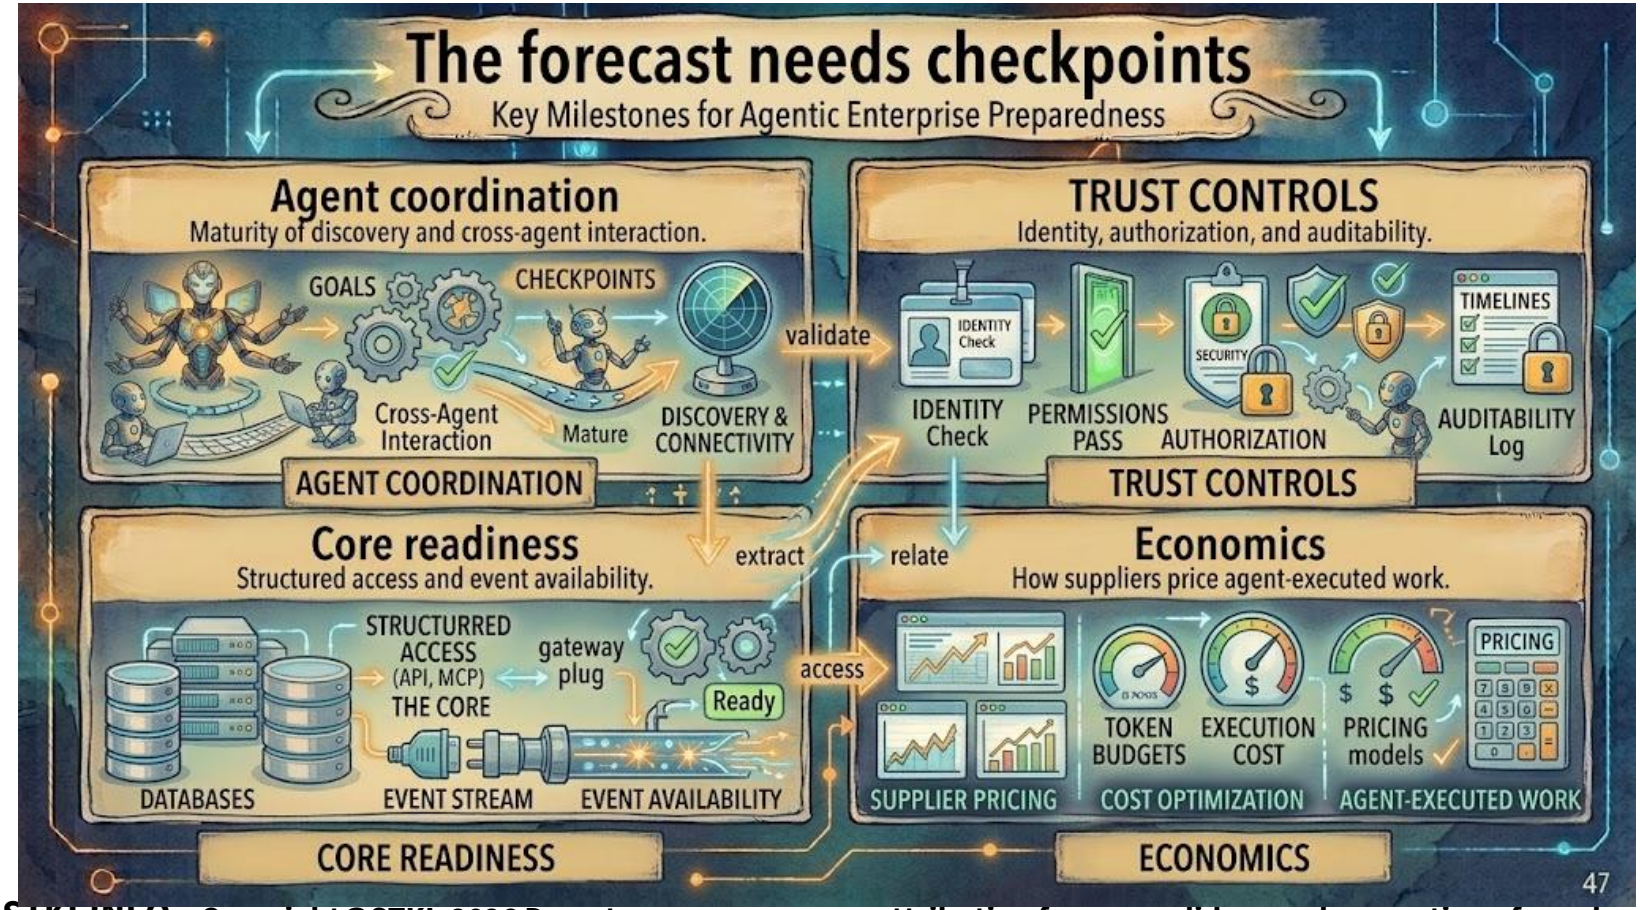
A clear ***architecture direction*** and explicit technology hypotheses.





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

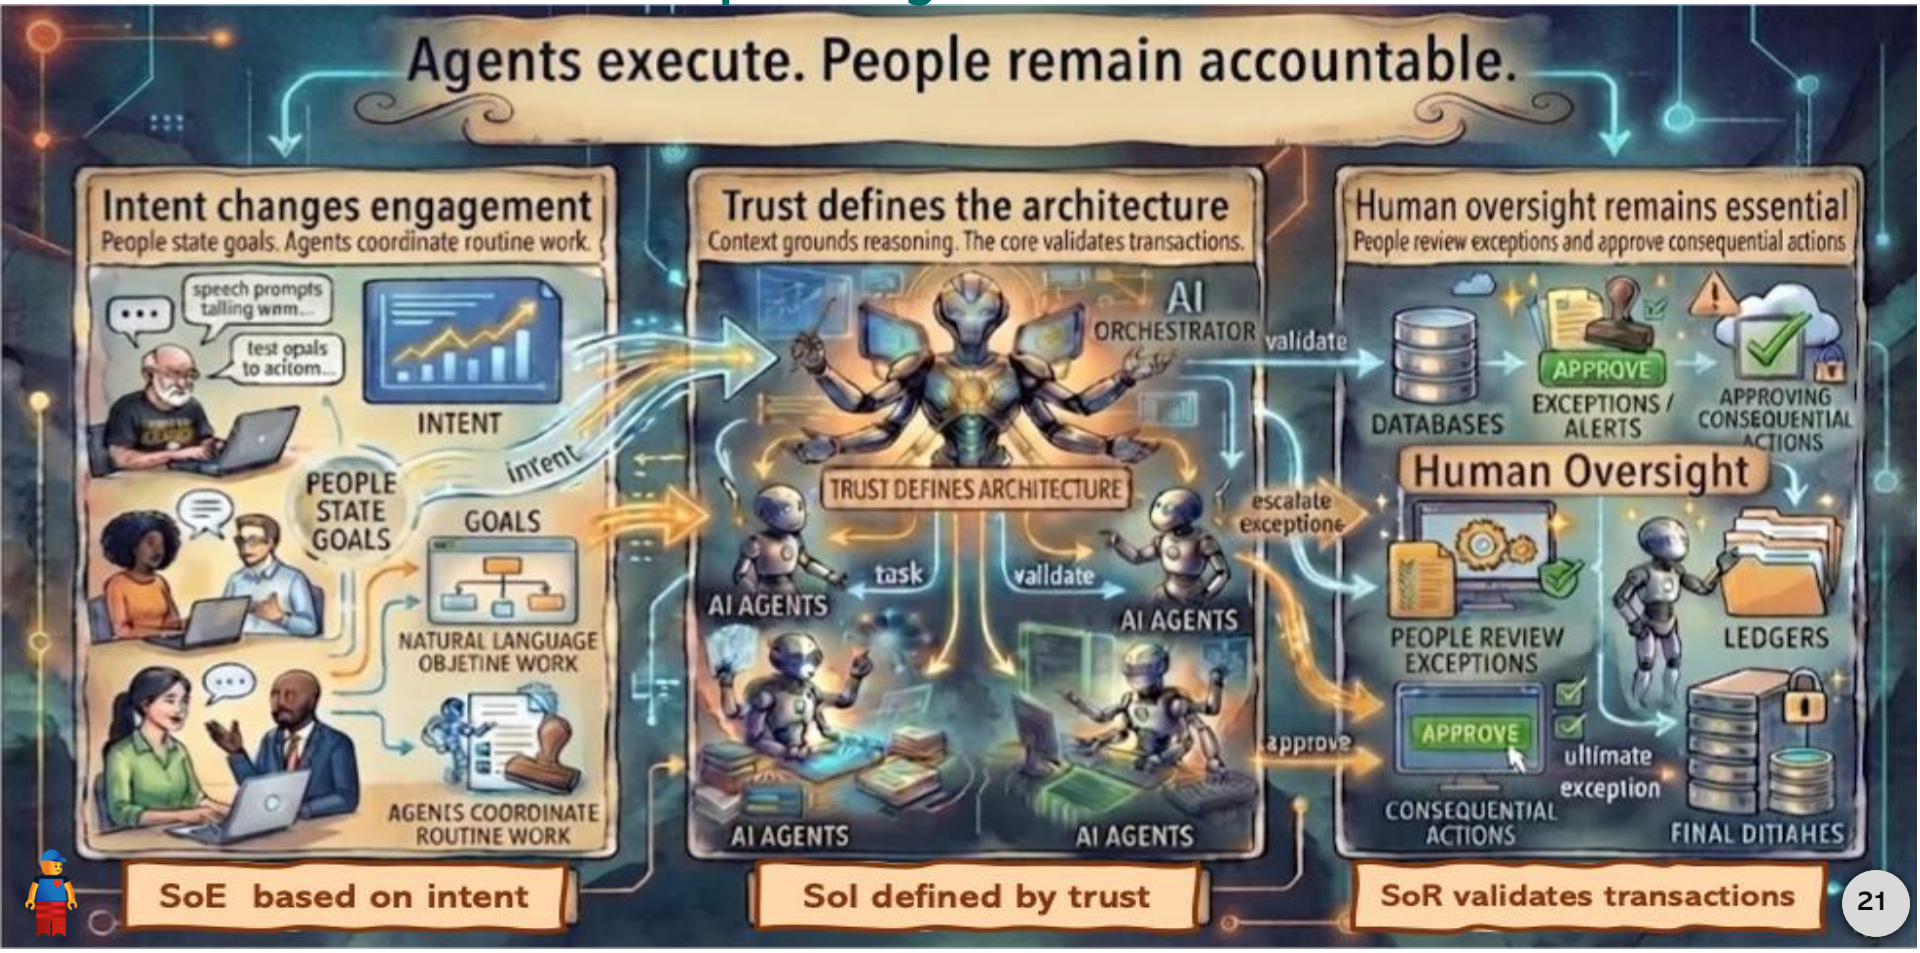
# After 300 mt (2031 and beyond) Signals worth watching





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# CONCLUSION:
# 2027–2031 is a planning architectural scenario



# SHORT-TERM PLAN 2027-2028:

## 2027–2028: the clear road (first 100 mt)

### Implement

**Agentic workflows**

Identify & implement workflows outside your old legacy systems (ERP, CRM, Core)

**Core modernization**

Prepare transactional systems for headless agentic access as systems of records.

**Data Strategy & Implementation**

Implement a context-driven data strategy that will support agentic workflows






Build a **context** layer that supports agentic **workflows**
**Talk to your data** as a 1st step towards agentic analytics
**Data management** will become **knowledge management**


Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# MID-TERM PLAN 2029-3030:

The context layer becomes a living, breathing language and rules platform for agents, BI, analytics & workflows
The next phase is building **agentic analytics workflows**
That will take data through insights all the way to **action**



## 2029–2030: the next stretch (100-300 mt)

## plan & start implementing the architecture

### Start decoupling the old architecture

Separate engagement, intelligence and record systems.

### Coordination

Develop agentic workflows that will work across systems.

### Agent marketplace

Prepare the enterprise's agent marketplace





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# LONG-TERM
# 2031 and beyond:

**The historical divide between Analytical & Operational data** will dissolve.
Data **=** analytical + operational + unstructured and tribal knowledge



## 2031 and beyond (over 300 mt) : strategic hypotheses

### Position without pretending to know

***long-term project choices need assumptions that can change.***

**Finish the uncoupling project**

Systems of engagement, intelligence and records work together based on the new ***uncoupled architecture***

**What remains uncertain**

Technology maturity, adoption speed, and economic models.

**What still helps**

A clear ***architecture direction*** and explicit technology hypotheses.




STKI.INFO

Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

13


Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

CDAOs need to **enable** to organization's decoupled architecture

But they also need to **build** one for analytics purposes




Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Existing (soon to be legacy) data architecture:



ML model

Data scientist






Reports

Data analyst



Data analyst





Dashboards

**SoE**
System of engagement:
**ML platform**

**SoE**
System of engagement:
**BI Tool B**

**SoE**
System of engagement:
**BI Tool A**

Tightly coupled

**SoI**
System of Intelligence:
**Feature Store**

↔ Separated ↔

**SoI**
System of Intelligence:
**Semantic layer B**

**SoI**
System of Intelligence:
**Semantic layer A**

**SoR**
System of record:
**Data Lake**

↔ Separated ↔

**SoR**
System of record:
**Data warehouse**

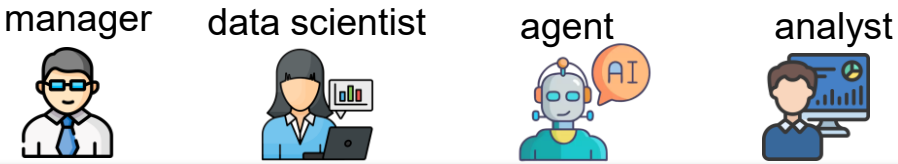


Data platform



STKI.INFO Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Modern decoupled data architecture







Reports



Dashboards

Interface

## SoE Systems of Engagement

Intent, personalization, UX



Talk to your data



Analytics chatbot



Orchestration

Meaning

## Sol Systems of Intelligence

Context & agent orchestration




Context



Ontology

Data Catalog

Semantic Layer



Verified records

Data

## SoR Systems of Record

Data repository, data products, strict data policies



Data platform



STKI.INFO Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph



**CDO 1.0:**
Data governance

Data governance
Compliance
Risk mitigation
Data quality



**CDO 2.0:**
Growth Partner

Dashboards
Predictions
Analytical models
Data products
Data monetization

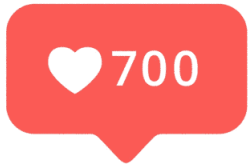


**CDO 3.0:**
Workflow trust governance

Data definitions
Governed context layer
By:
SoR modernization (data platform)
SoI (semantics and context)


Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# The market darling: Context




**“The New Critical Infrastructure for AI!”** - Gartner

"**Own the context, rent the rest**” – McKinsey

**“The next moat is not just data; it is proprietary context.”** – Accenture

**“Context engineering is the defining skill for AI”** – LangChain CEO

**“Context being the single biggest reason for AI failures”** – MIT

**“Context layer is THE best investment organizations should make”** - STKI



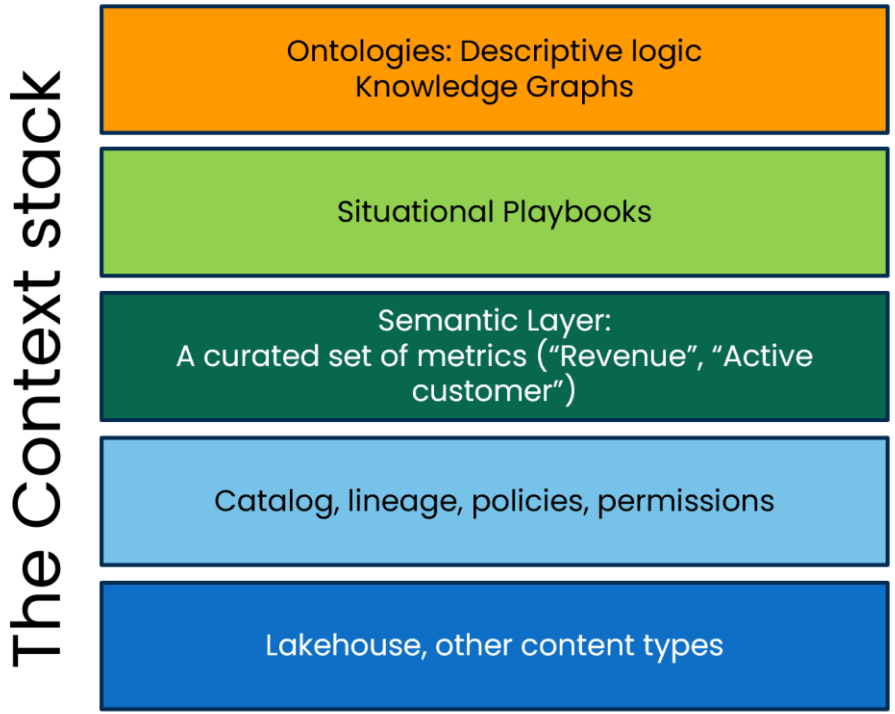

Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# What is a context layer?

Not a solution but a **stack** of solutions

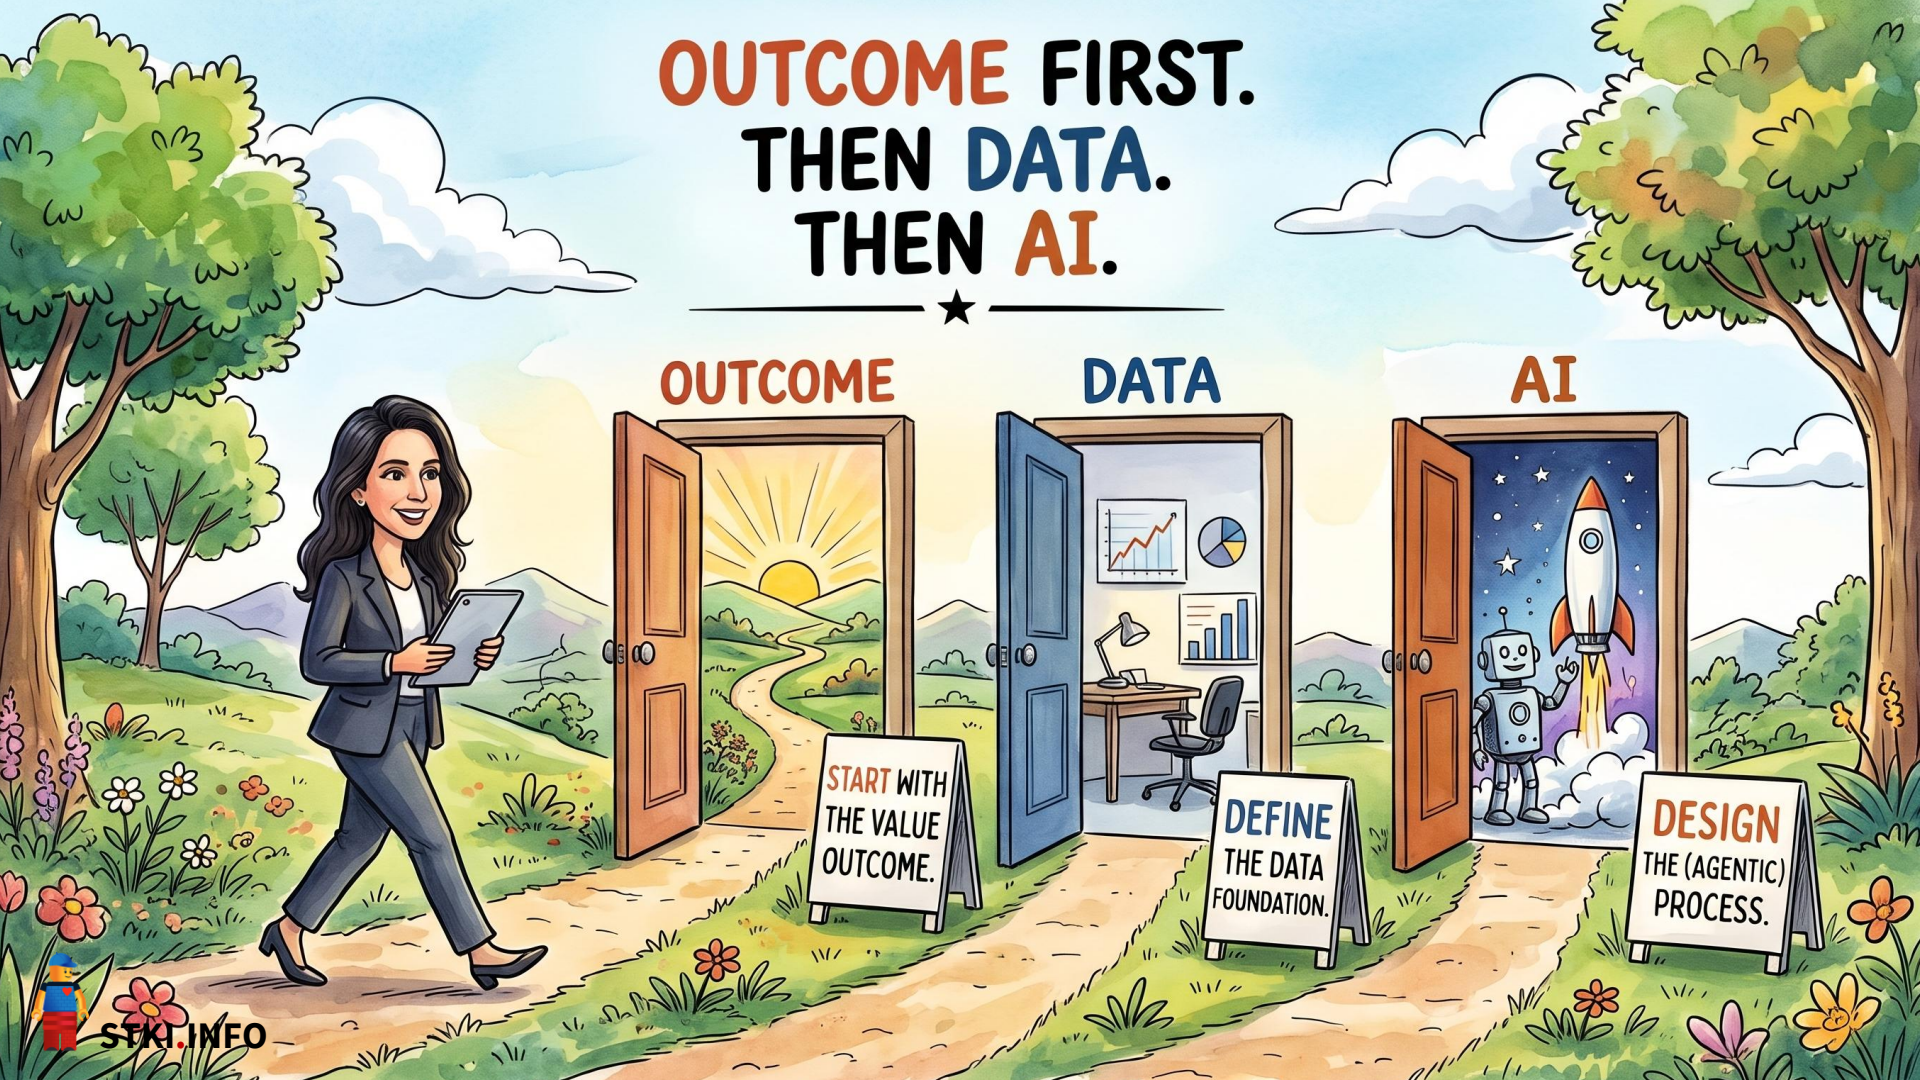
Includes:

- Identity resolution
- Canonical entities
- Specific instructions to dissect tribal knowledge
- Governance guidance
- How to handle ambiguity
- Situational/scenario-based rules
- Control costs (bonus!)





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph


OUTCOME FIRST.
THEN DATA.
THEN AI.
OUTCOME
DATA
AI
START WITH THE VALUE OUTCOME.
DEFINE THE DATA FOUNDATION.
DESIGN THE (AGENTIC) PROCESS.
STKI.INFO

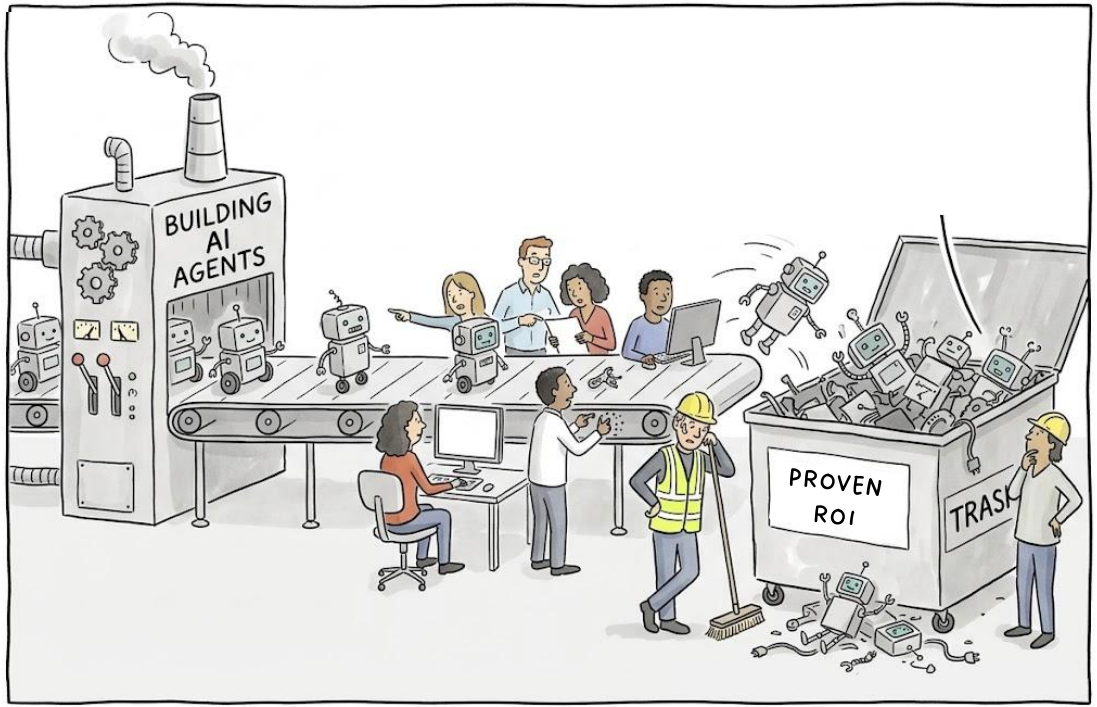
# AI as a Tool



- AI is a smart personal assistant that speeds up specific tasks (coding, drafting)
- It doesn't manage, decide, or hand off work
- Accountability is strictly 100% human


Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Everyone is building agents. Few prove




- Individuals report clear productivity jumps (a 20-30%+ faster dev!)
- End-to-end workflows remain completely unchanged (overall SDLC is not getting any shorter)

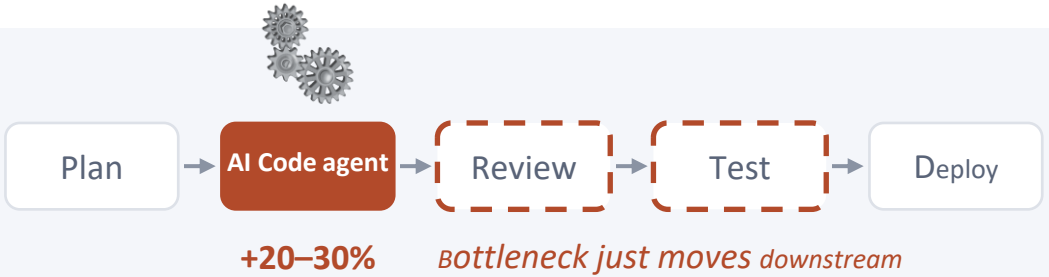

Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Why one agent is not a system

A standalone agent speeds up a task, not a process
But coding was never the whole process. The agent pushes more code into testing and security faster than those stages can absorb it

## Agent on an island

1 agent, 1 task

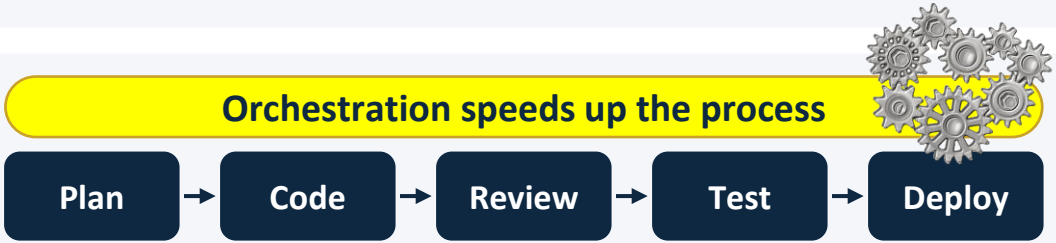



Cycle time

**Unchanged**

## Orchestrated system

Many agents, 1 flow




Multiple agents work together across all stages, eliminating wait times for approvals and creating a seamless flow between stations

Cycle time

**Shrinks**


Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

**CORE**

Interface

Business logic

Database



**ERP**

Interface

Business logic

Database



**CRM**

Interface

Business logic

Database

**The human**

copies, re-types and bridges gaps

**The human**

finds the right app, navigates the menus, fixes validation errors

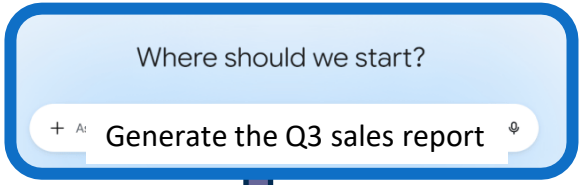
**Humans as an integration layer:** The old stack optimized the database, not the worker. People had to adapt to the system




Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

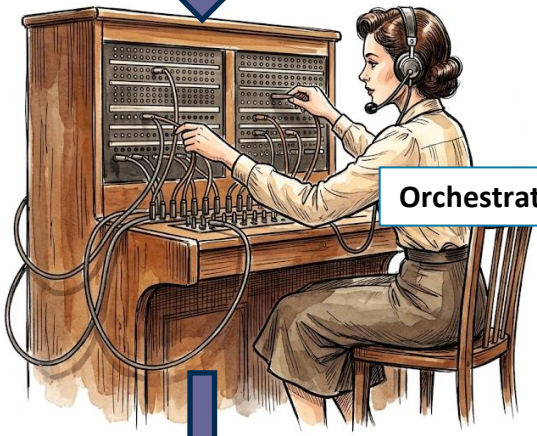
## Systems of Engagement

Capture Human Intent + Personalization




## Systems of Intelligence

Orchestrator understands the intent and **builds the workflow** using the right agents **under policy**




agents pull the correct data from the SoR

**NO GARBAGE DATA**

## Systems of Records

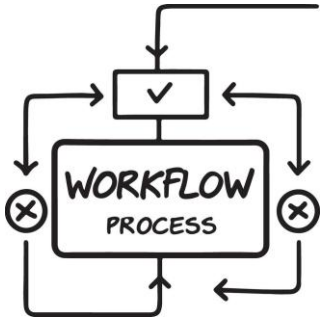
Protected rigid DATA schemas (ERP, CRM, Core)
Source of Truth





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# AI as a Workflow

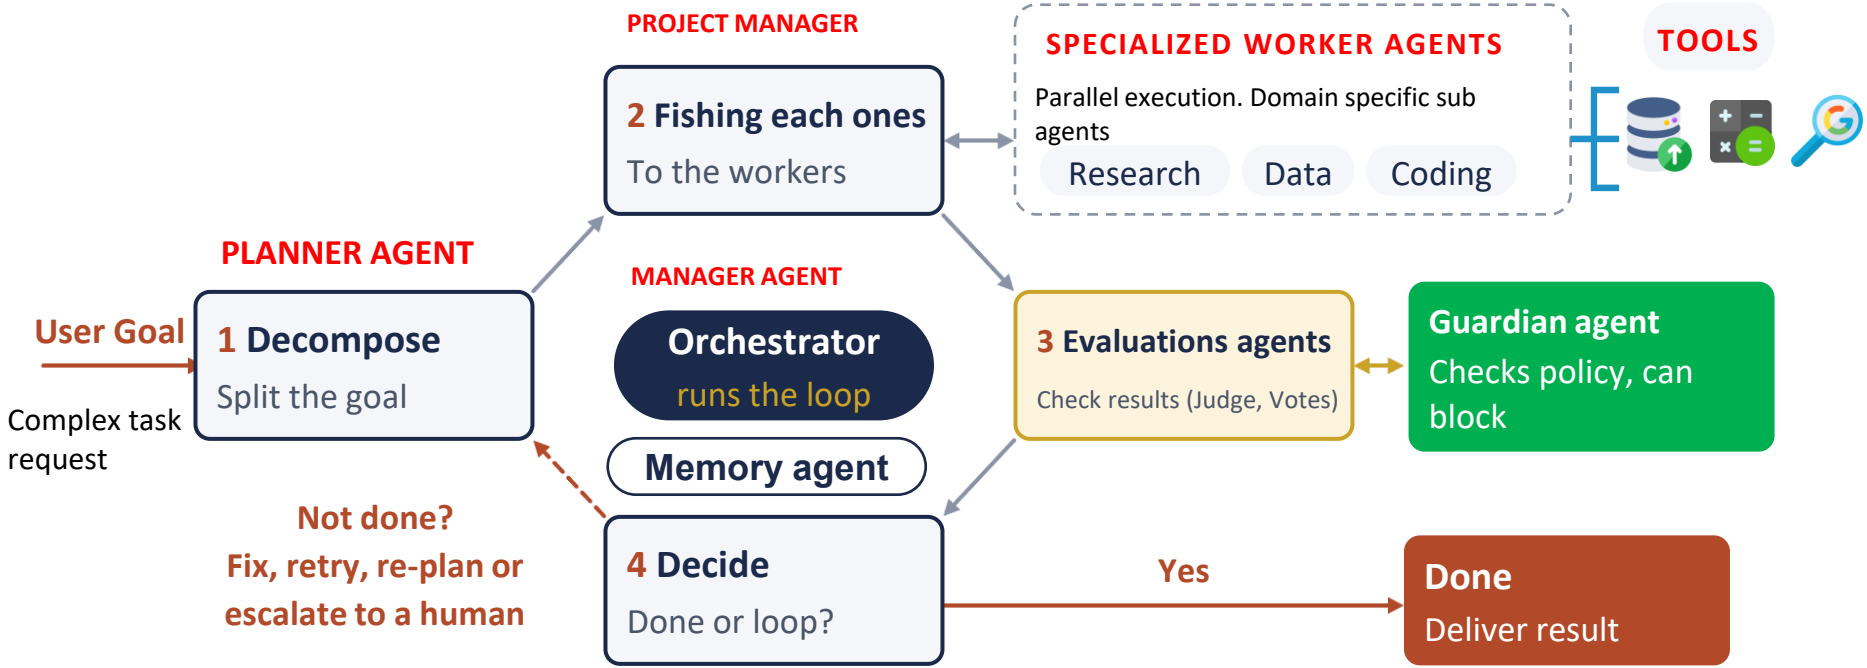



AI stops being just a helper on the side, and becomes the engine running the whole process end-to-end


Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Systems of Intelligence

## Multi-agent System and Orchestration Powering Each Workflow




**When AI becomes the engine,**
**it's time to reimagine the workflow from scratch**


Copyrig

For 30 years,
the software told us how to work

- Off-the-shelf ERP and CRM came with “vanilla” processes
- We adapted the organization to the system
- “Best practice” = the vendor’s practice
- RPA didn’t fix this either. It automated processes “as-is”



The lost art of designing processes from scratch


Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph



## Standard Automation/ RPA

Static and deterministic workflows based on rigid, predictable rules

**RPA is here to stay:**
When process is known in advance (and it's cheaper)



## AI Agent

Reasons and acts on a single task, handles variations, but still needs a human to steer

**One agent, one task:**
Speeds up a step, not the whole process



## Orchestrated Agentic Workflows

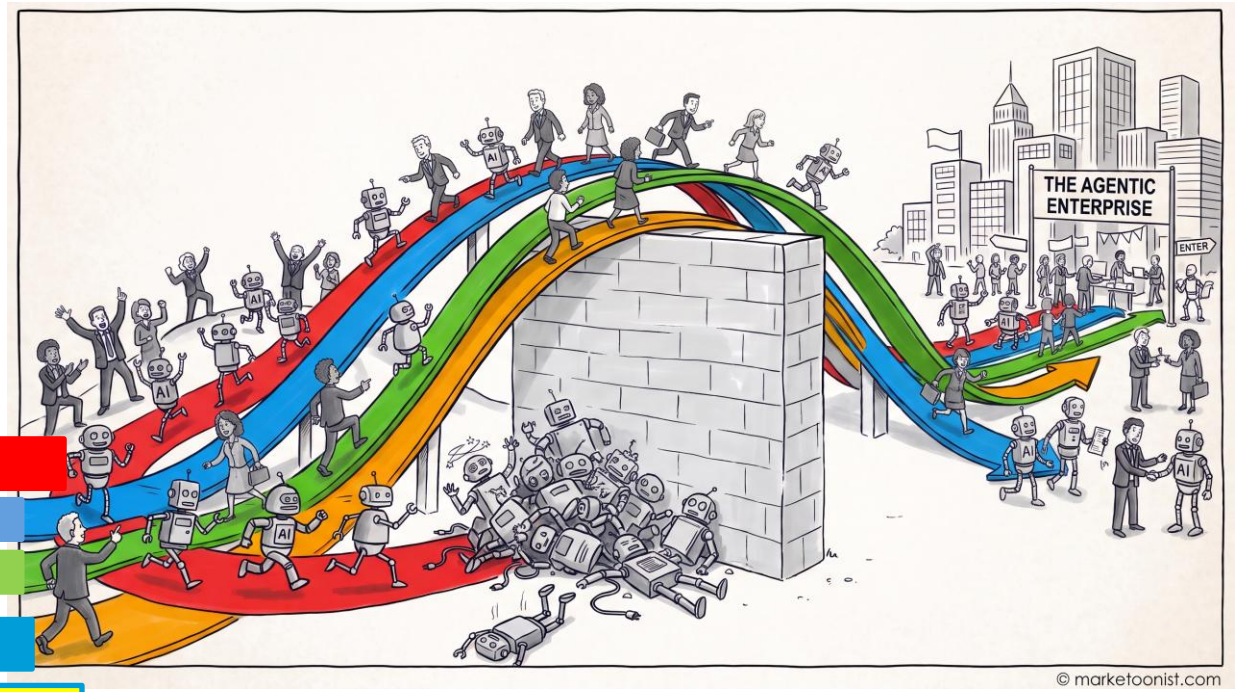
Stochastic (probabilistic) workflows that dynamically adapt to new data, exceptions, and changing conditions over time

**Agentic AI:** Probabilistic workflows (the process and outcome change based on statistics)


Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

The way over the wall, in 5 moves:

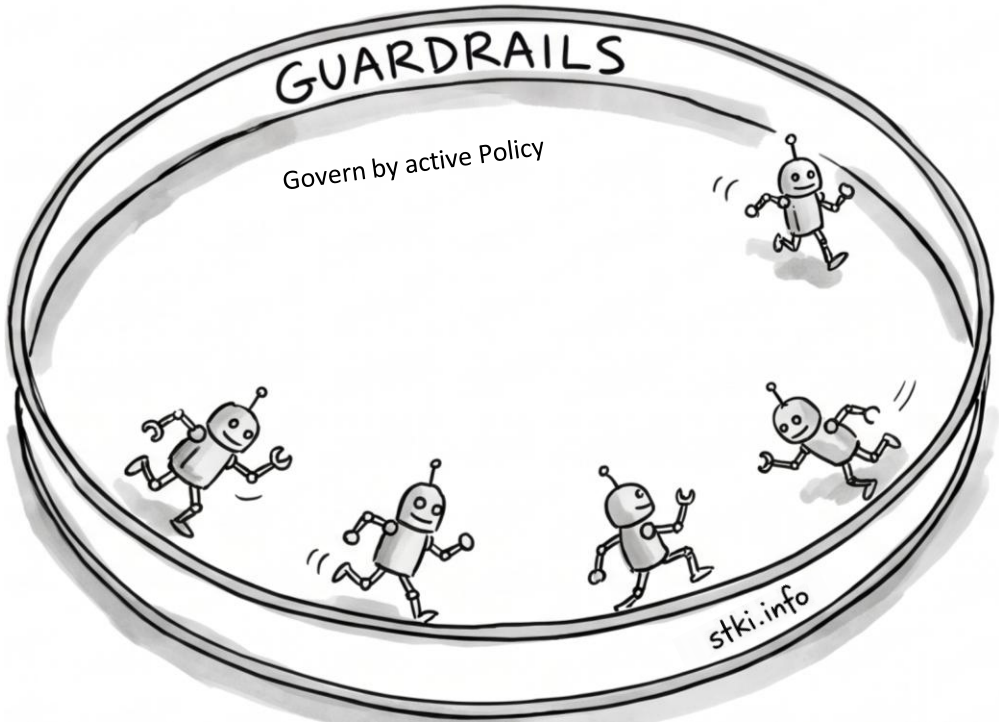
# Governance, Tokenomics, ROI, Processes and People




1. Take control: Governance
2. Fix economics: Tokenomics
3. Prove the Value: ROI
4. Redesign processes: WF
5. Bring the people: Change Mng

STKI.INFO Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Take control: Governance




- **Agent Registry**
  You cannot govern what you do not register
- **Active Policy**
  Real-time execution rules
- **Agent IAM**
  Define agent access exactly like a human worker
- **Audit Trail**
  Reasoning Loop traceability
- **Human-in-the-Loop**
  Escalation
- **AI FinOps**


Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

הטמעה

Break silos

SDLC

Less Hierarchy

Value-stream org

Biz KPIs

# The New Operating Model

Outcome VS Output

Ongoing budget

Product Management

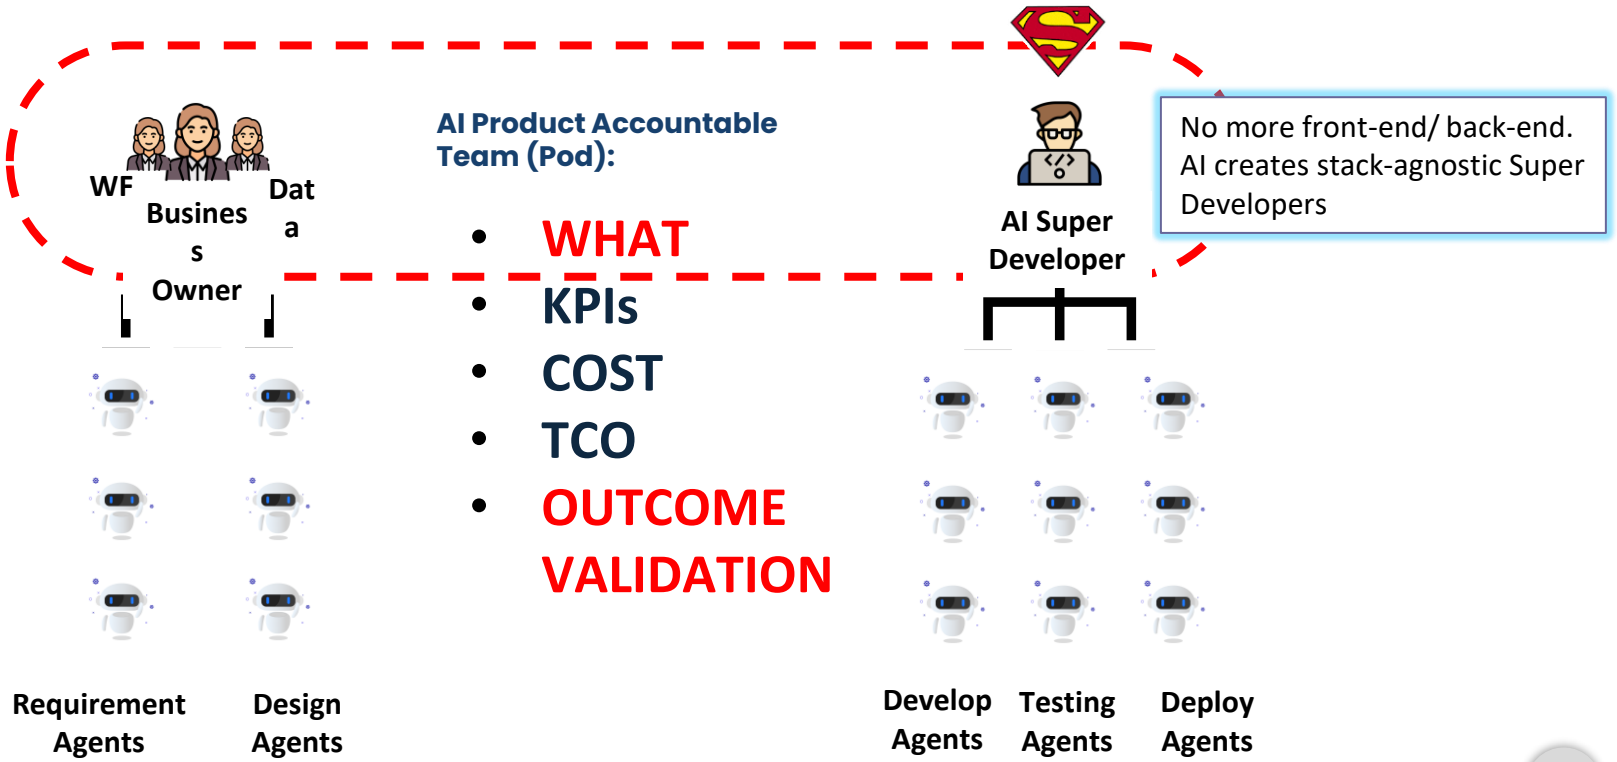
Business/IT co-ownership



STKI.INFO

*Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph*

# AI Software Development Life Cycle





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# STKI's Architectures for "Agentic Decoupled Enterprise"



We must be the ones at the wheel

**Fluid Human Intent**

Where should we start?

Ask Gemini Flash

**GOVERNANCE**

## Intelligent Orchestration

Dynamic workflows composed from specialized agents

## Stable Transactional Record

Protected rigid DATA schemas (ERP, CRM, Core)


Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

The rise of **AI Agents** is reshaping the workforce, but greater efficiency does not automatically create greater business value. While agents can automate tasks and accelerate execution, organizations still need human creativity, judgment, and innovation to transform productivity gains into meaningful outcomes.

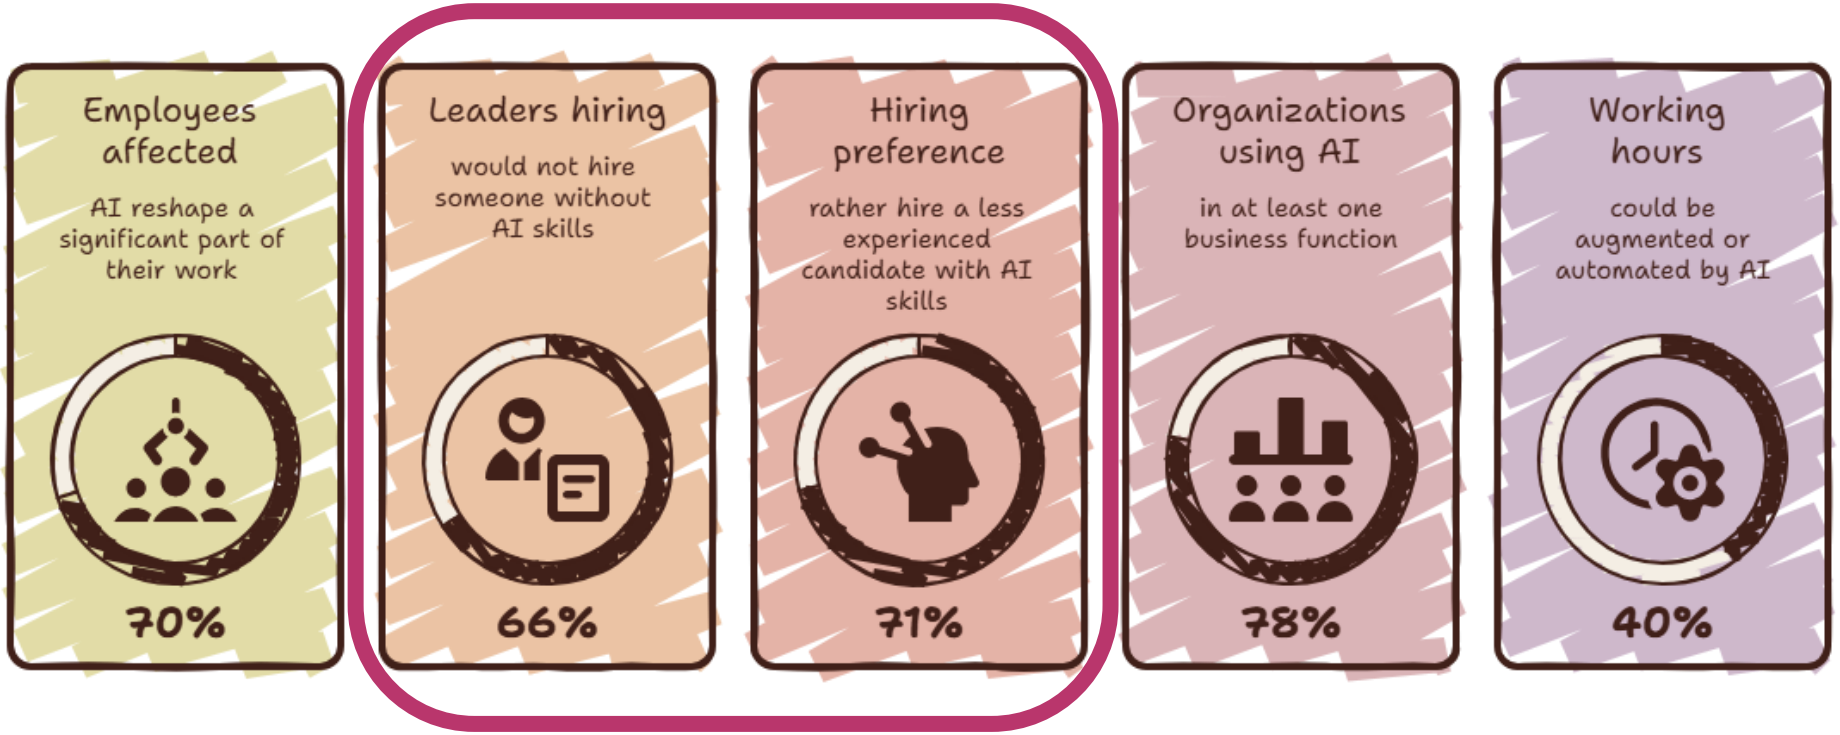
At the same time, enterprise software is evolving from application-centric experiences to **agent-driven, intent-based interactions**. In the emerging three-layer enterprise, **Systems of Record, Systems of Intelligence, and Systems of Engagement** are decoupled, allowing agents to understand user intent, orchestrate capabilities, and dynamically assemble the right experience instead of forcing users through predefined workflows.

In the **Agentic Enterprise**, employees express a business objective, and agents coordinate **processes, systems, and data on their behalf**. Success, however, depends on a robust intelligence layer built on trusted data and a unified enterprise context, because effective agents require high-quality organizational knowledge to make reliable decisions.




Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# The Great Workforce Shift



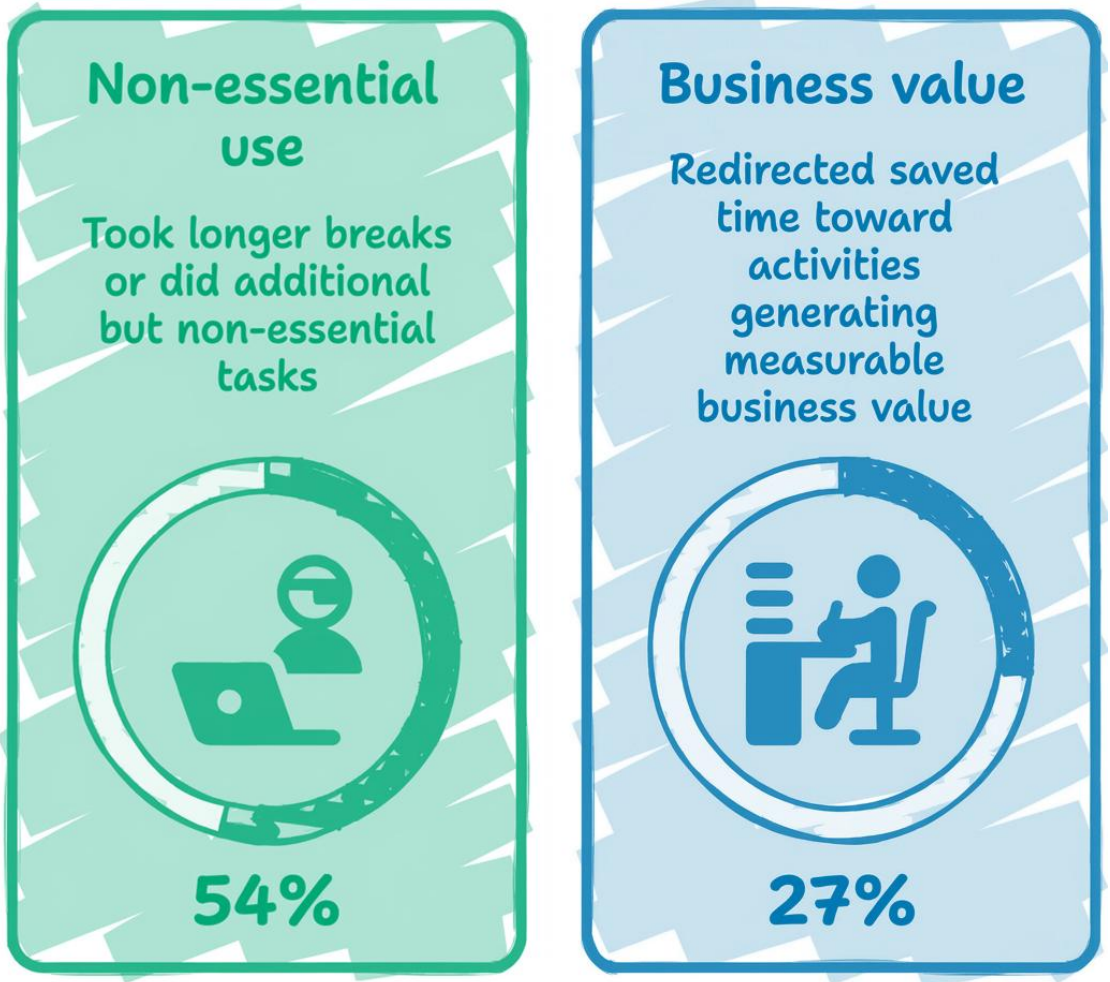

STKI.INFO Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

Sources: World Economic Forum, Microsoft, McKinsey

# It Saves Time, but Not Always Creating Value

## Productivity Theater:

**Illusion of progress: AI makes work look faster, busier, and more automated, while the underlying outcomes barely change.**



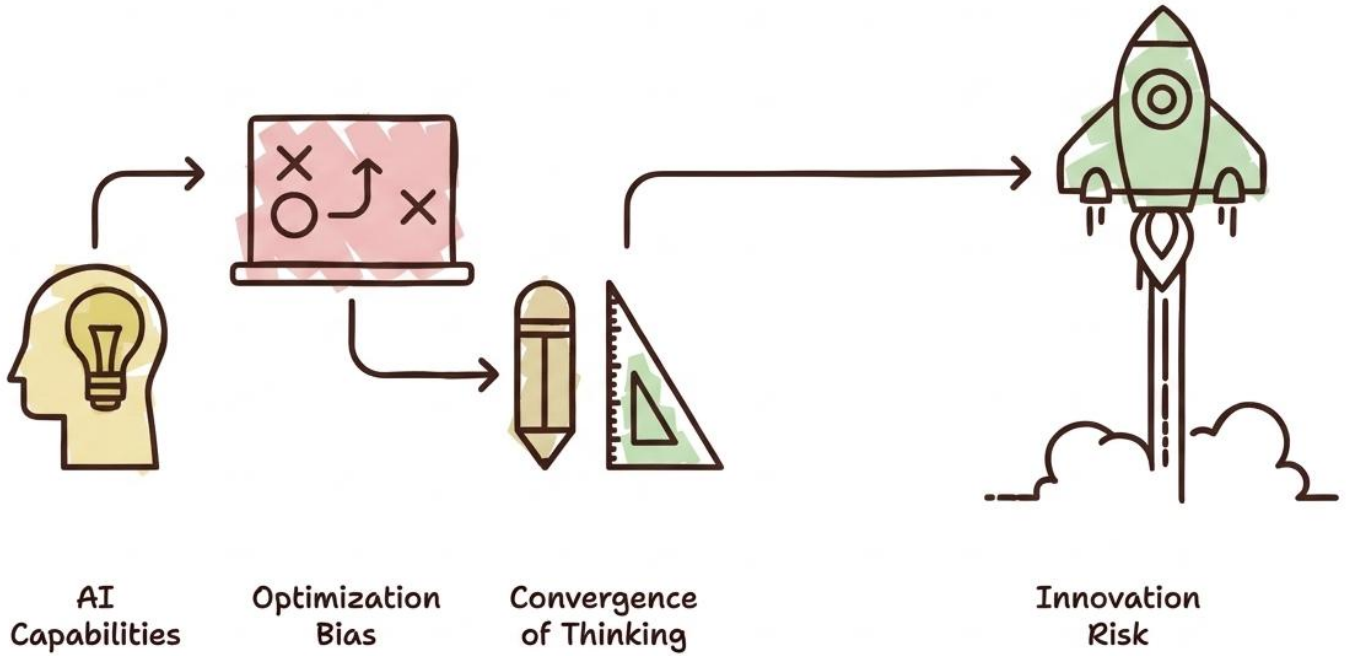


Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph
Source:BCG

# The creativity paradox

**AI optimizes for what is statistically likely.**

**Innovation often depends on what is statistically UNLIKELY.**




**As AI drives convergence toward proven answers, breakthrough thinking may become harder to achieve without human creativity, curiosity, and judgment.**


Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

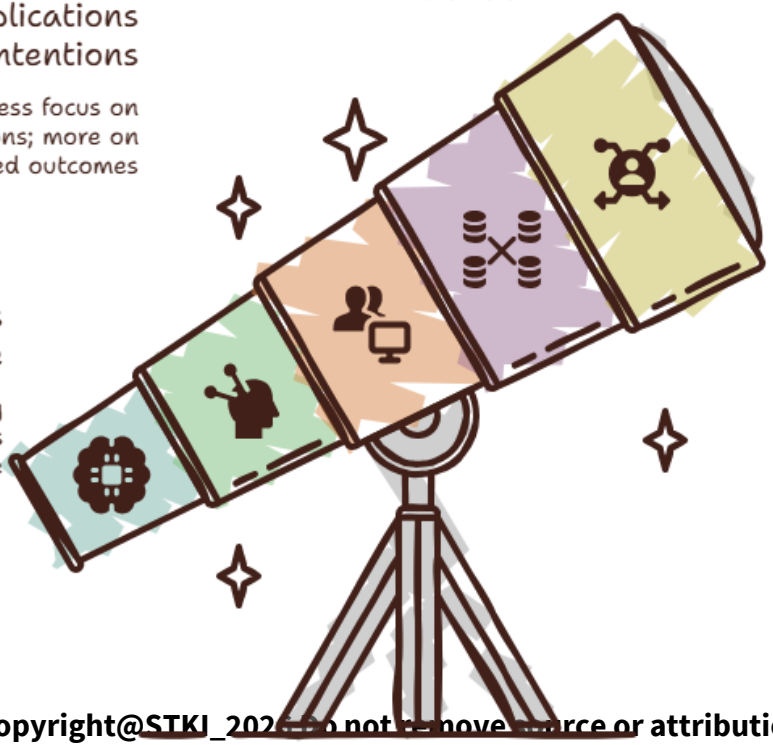
# Decoupling
# The Three-Layer Enterprise





Copyright@STKI_2026 Do not remove source of attribution from any slide, graph or portion of graph

# The Rise of the Three-Layer Enterprise



**Rise of Digital Agents**

Systems designed for humans and digital workers

**From Applications to Intentions**

Less focus on applications; more on desired outcomes

**Decoupling Experience from Execution**

Experience layer increasingly independent from systems of record

**AI Becomes Infrastructure**

A foundational capability embedded across enterprise software

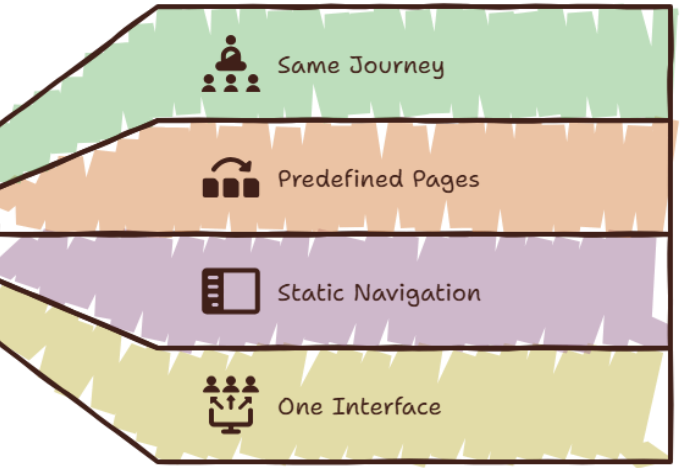
**Consumerization of Enterprise Expectations**

Intuitive, personalized, responsive as the best consumer applications

**Intent is the new interface. Applications are the new infrastructure.**

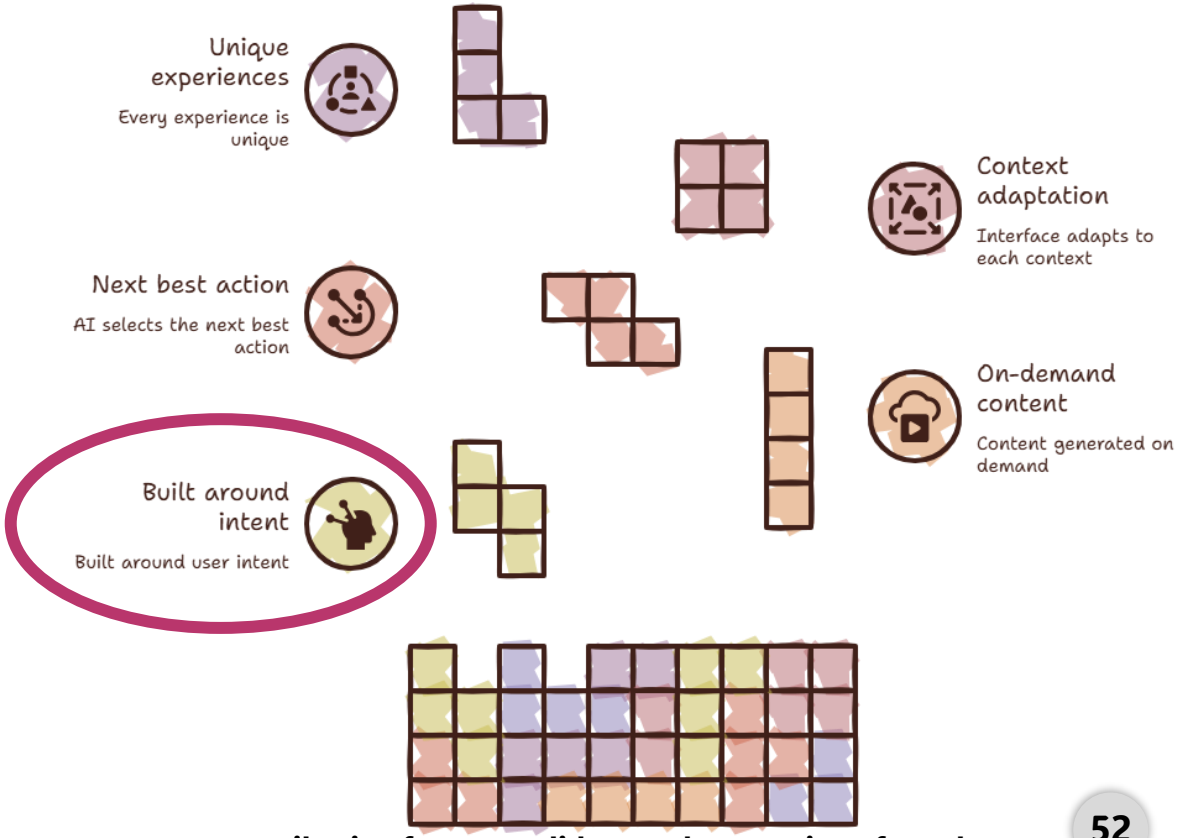

Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Traditional




The traditional model assumes every user follows the same path through the same application.

## System of Engagement: Dynamic Experiences





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# How an Agentic Enterprise Procures 50 Laptops

Purchase 50 laptops for the new sales team within approved budget.

## System of Engagement

Captures Intent: Request: Buy 50 laptops; User Interaction; Copilot; Business Chat; Conversational Experience.

1

## System of Intelligence

Determines What Should Happen: Understand Intent; Gather Context; Apply Policies; Check Budget; Select Supplier; Obtain Approvals; Orchestrate Workflow.



2

## System of Records

Executes Transactions: ERP; Procurement; Finance; Vendor Management. Actions: Create Purchase Requisition; Create Purchase Order; Commit Budget; Update Financial Records.

3

HITL





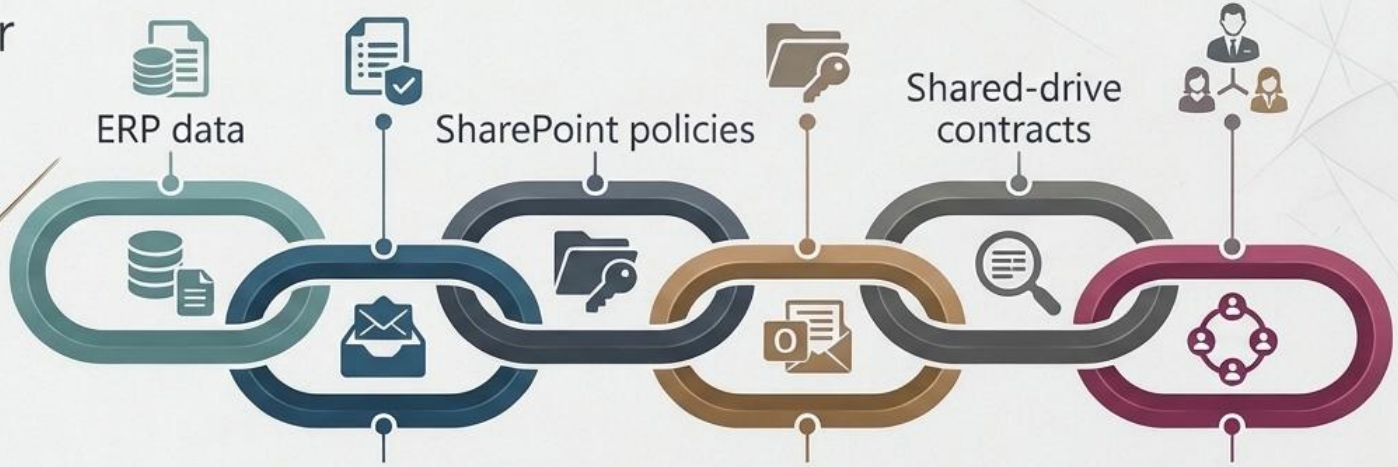
STKI.INFO Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of g

# Broken Systems of Record Create Broken Intelligence

The Orchestrator

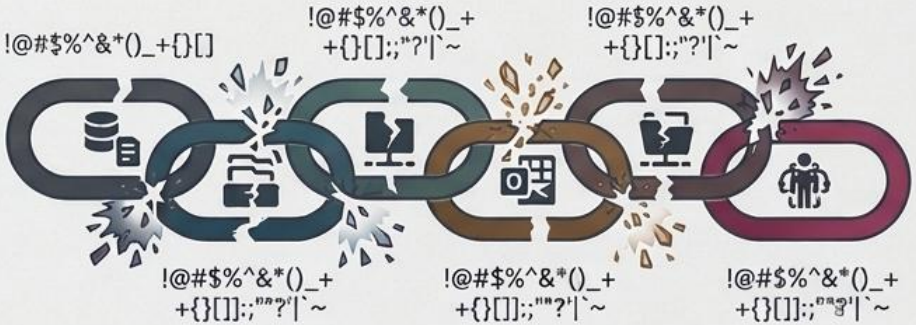


Human Orchestrator

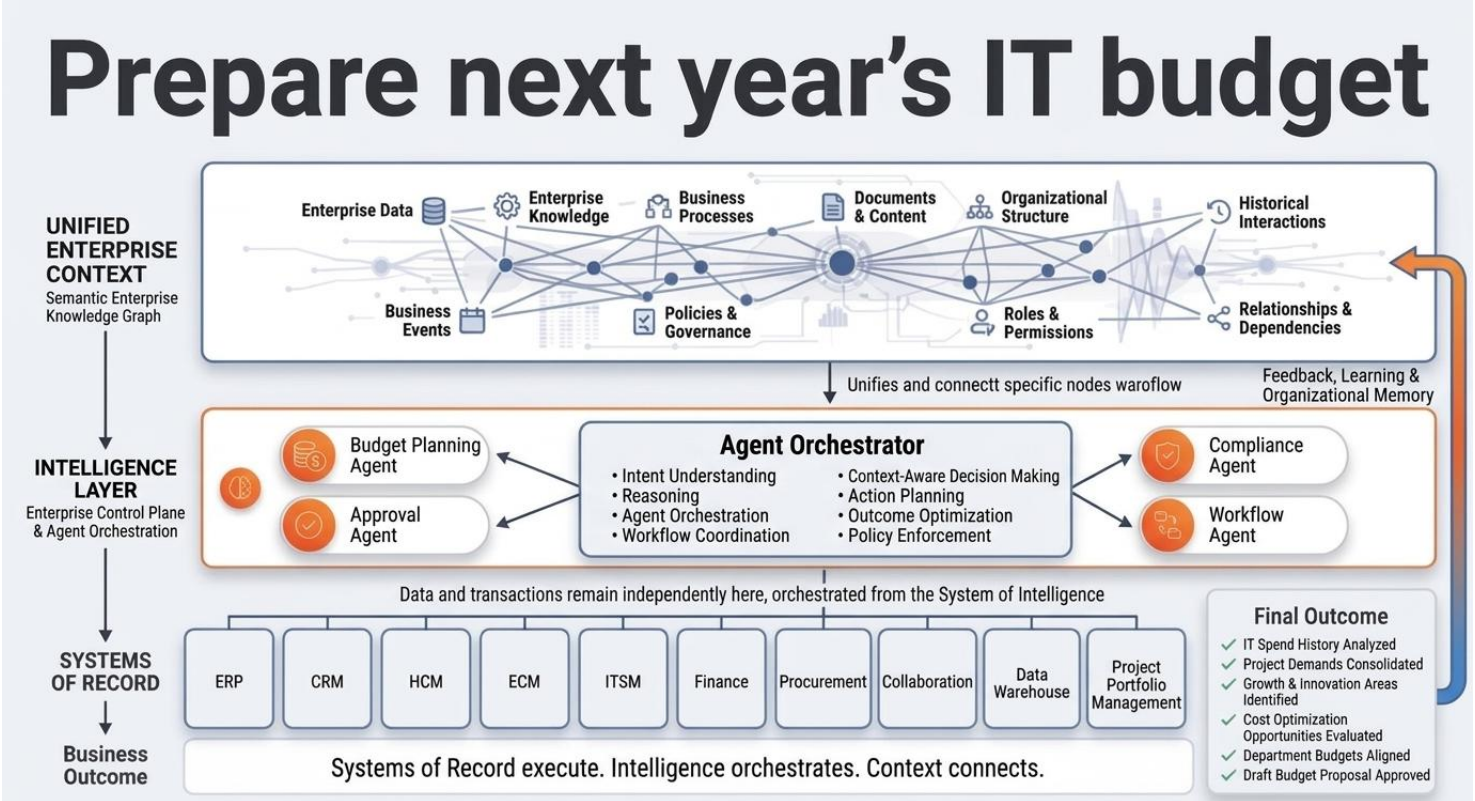



System of Intelligence





# Starts with Unified Enterprise




STKI.INFO Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# 2027–2028: Horizon 1: the clear road (first 100 mt)

## Development & Architecture

- Identify pivot developers that can act as “development orchestrators” building and using ai development agents (app architecture, design, UX, design tests, run tests, etc.)
- Define and build trust layer for AI code
- Define context engineering role
- Implement AI Gateway with basic token policy
- MCP, A2A and AI gateways form a new agent integration layer
- Use legacy understating tools





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# 2027–2028: Horizon 1: the clear road (first 100 mt)

## Infrastructure & Operations

- Experiment and become familiar with private LLM's (+SLM)
- Plan for expected Hardware Shortage changing conditions
- Lean new AI based DC components (RoCE memory, GDS storage, water cooling)
- Build RCA (root cause analysis) AI agent that is connected to all infra and observability systems
- Build first AI agents that replace some manual work ("write" agents) in infrastructure and end user support (ITSM)





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# 2027–2028: Horizon 1: the clear road (first 100 mt)

## Cyber Security

- Treat AI agents as a primary security domain
- Adopt zero trust principles (“perimeter protection” is not “the thing”)
- Deploy controls built for agentic systems
- Automate selected security operations with AI
- Implement NHI  (none human identity) capabilities





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# 2029–2030: Horizon 2: the next stretch (100-300 mt)
# plan & start implementing the architecture

## Development & Architecture

- Implement granular token policy
- Enterprise development process is completely based on AI achieving at least 3 times speed than today
- Build solid governance tools and processes for business users development
- Business users create first small scale business applications
- Integration is based on MCP A2A





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# 2029–2030: Horizon 2: the next stretch (100-300 mt)
# plan & start implementing the architecture

## Infrastructure and Operations

- AI agents take over much of day-to-day infrastructure operations
- Private AI inference spreads beyond defense to regulated sectors and general use
- AI agents resolve most routine service desk and ops incidents
- Endpoint management becomes largely autonomous





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# 2029–2030: Horizon 2: the next stretch (100-300 mt)
# plan & start implementing the architecture

## Cyber Security

- Execute mandatory post-quantum migration
- Move most Tier 1 SOC work to AI agents
- Most remediation actions are pre-build & autonomous





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# 2031 and beyond: Horizon 3: (over 300 mt)

## Position without pretending to know

***long-term project choices need assumptions that can change.***

### Development & Architecture

- Software is generated on demand from business intent
- Legacy modernization is done by AI





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# 2031 and beyond: Horizon 3: (over 300 mt)

## Position without pretending to know

***long-term project choices need assumptions that can change.***

### Infrastructure & Operations

- Fault-tolerant quantum computing arrives as a cloud service
- Most Infrastructure and Operations is done with AI
- Self-operating IT estates, with humans setting intent and guardrails
- The PC becomes a personal AI agent host rather than an app desktop





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# 2031 and beyond: Horizon 3: (over 300 mt)

## Position without pretending to know

***long-term project choices need assumptions that can change.***

### Cyber Security

- Monitor the quantum “Q-Day” risk window
- Design for autonomous, self-defending security
- Make crypto-agility a permanent architecture principle





Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Semantic layer for AI agents: operational + analytical, not analytical only



## STKI Recommends

Recommendation: Scope semantic data initiatives to include operational systems (ERP, CRM, core systems), not only the analytical estate (data warehouse, lakehouse, BI). AI agents don't just answer questions. They execute tasks in operational systems. To do that reliably, they need shared definitions of key business entities (customer, order, account, claim) that mean the same thing in both the operational and analytical worlds.

**Semantic Project**

**Only analytical data**

**Analytical + Operational data**




Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Markdown is the delivery format for LLMs

### Up to -80% tokens vs. HTML / PDF !!

Cloudflare: one blog post = 16,180 tokens in HTML, 3,150 in Markdown. Lower cost, more context headroom. (PPC Land)

### The native language of agents

AGENTS.md (Linux Foundation's AAIF, Dec 2025) and SKILL.md (Agent Skills open standard) are plain Markdown. (OpenAI, Strapi)

### Structure survives the pipeline

Headings, lists and tables keep their meaning for RAG chunking. Microsoft's open-source MarkItDown turns Office and PDF files into Markdown. (GitHub)

**AGENTS.md** tells an agent how to work on a specific project, and **SKILL.md** teaches any agent how to do a specific task

**What IT should do:** make Markdown the default for content delivered to LLMs. Convert document stores (MarkItDown), publish an llms.txt index, ship AGENTS.md / SKILL.md with internal systems, and keep it all versioned in Git.



**STKI Recommends** | Use MD documents (or other semi-structured formats like JSON) as much as possible




Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Markdown is the delivery format for LLMs

### Up to -80% tokens vs. HTML / PDF !!

Cloudflare: one blog post = 16,180 tokens in HTML, 3,150 in Markdown. Lower cost, more context headroom. (PPC Land)

### The native language of agents

AGENTS.md (Linux Foundation's AAIF, Dec 2025) and SKILL.md (Agent Skills open standard) are plain Markdown. (OpenAI, Strapi)

### Structure survives the pipeline

Headings, lists and tables keep their meaning for RAG chunking. Microsoft's open-source MarkItDown turns Office and PDF files into Markdown. (GitHub)

**AGENTS.md** tells an agent how to work on a specific project, and **SKILL.md** teaches any agent how to do a specific task

**What IT should do:** make Markdown the default for content delivered to LLMs. Convert document stores (MarkItDown), publish an llms.txt index, ship AGENTS.md / SKILL.md with internal systems, and keep it all versioned in Git.



**STKI Recommends** | Use MD documents (or other semi-structured formats like JSON) as much as possible




Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# Nine trust layers

## Outside the agent itself !!

**01 Supply chain & AI-BOM**
Every part of the agent is inventoried and approved.

**02 Identity & least privilege**
Its own identity and narrow, purpose-built tools only.

**03 Human-in-the-loop (external)**
A person approves irreversible actions. This is **done not by the agent but by the "trust" layer**

**04 Policy engine**
Deterministic rules: who, what, where, when, how much.

**05 Hard guardrails**
Limits live in code, not in the prompt.

**06 Verification & simulation**
Dry run, sandbox, digital twin, replay, second checker.

**07 Prompt injection & outside content**
Outside content is data, never orders ; packages to install are vetted first.

**08 Rollback & kill switch**
Any action can be stopped or undone.

**09 Audit & observability**
Every prompt, call and decision is logged.

Some are activated in the single action layer and some in a scenario layer




Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# What Eval tools do

**01**
### Build tests (+expected result)
Collect scenarios, including real cases pulled from production traces using different “personas” (new immigrant, retired, youth, family with kids, hacker, etc.)

**02**
### Run the agent on them
In batches, on every change, often wired into the CI/CD pipeline.

**03**
### Score the results
Rule checks, an LLM acting as judge, and human review.

**04**
### Compare versions
Catch regressions when the model, prompt or tools change.

**05**
### Trace and debug
See every step, tool call and decision in a failed run.

**06**
### Attack it on purpose
Red-team tests with prompt injections and unsafe requests.

**ALSO COVERED** **Guardrail validation · Consistency & hallucination checks · Comparison to golden responses**




Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

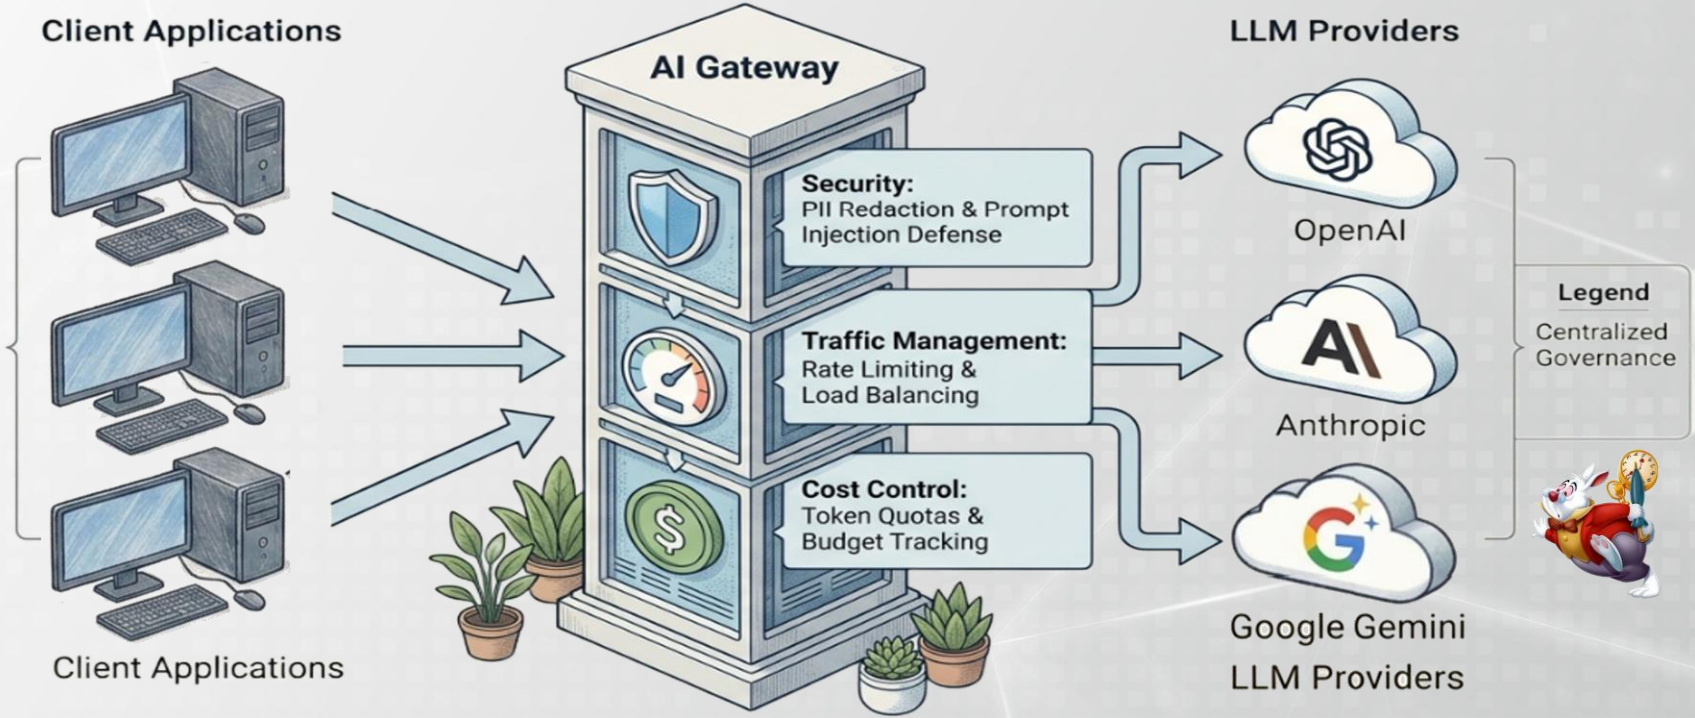
# Agents development = special case of AI Agent that is developing code for us



- AI is used (AI models are accessed for the development)
- Agent (AI) is installed on the PC and developing the code
- Code is developed

| | | | |
|---|---|---|---|
| Design | **Threat modeling**<br>Find design risks | **Secure design**<br>Define requirements | |
| Code | **SAST**<br>Scans source code | **Secrets scan**<br>Leaked credentials | **Code review**<br>Manual, by peers |
| Build | **SCA**<br>Libraries + SBOM | **IaC scan**<br>Cloud misconfigs | **Container scan**<br>Image CVEs |
| Test | **DAST**<br>Attacks running app | **IAST**<br>Agent inside app | **Pen testing**<br>Manual attack test |

# The AI Gateway is a must




**STKI Recommends** | Place AI gateway even if you do not need it today (it's a must)

# Potential AI token policy

| כללי חריגה | תקציב ברירת מחדל | מודלים מורשים | רמת מידע | סוג משימה מאושר | אוכלוסייה |
|---|---|---|---|---|---|
| מעבר למודל בינוני רק אם איכות נכשלת במדידה | עד Y ₪ למשתמש לחודש; עד N טוקנים לבקשה | מודל קטן/חסכוני, או מודל ארגוני מקומי | ציבורי או פנימי | תיוג, סיווג, חילוץ שדות, תקציר עד 500 מילים, ניסוח שגרתי | עובדי מטה |
| מחייב תיוג workflow ואיסור על PII ללא מסלול מאושר | תקציב גבוה יותר לפי יחידה | מודל בינוני מאושר בסביבת Enterprise | פנימי | ניתוח מסמכים ארוכים, סיכום רב־מסמכי, טיוטות רגישות | עובדי מטה |
| מודל מתקדם רק בפרויקט מאושר או לאחר כשל איכות | עד B ₪ למשתמש או לפרויקט | מודל קוד בינוני | קוד פנימי ללא סודות/נתוני לקוח | השלמת קוד, unit tests, תיעוד API, refactoring מוגבל | מפתחים |
| הצדקת צורך, תיעוד שימוש ובקרה תקופתית | עד Z ₪ לפרויקט אוworkflow | מודלים מתקדמיA | פנימי/סודי, רק בסביבה מאושרת | תכנון מערכת, debugging, רב־שלבי, threat modeling, סקירת ארכיטקטורה | מפתחים בכירים / ארכיטקטים |
| fallbackאוטומטי למודל זול; חסימה בחריגה | תקציב לכלtransaction, session ויום | רשימה מצומצמת וקבועה | לפי סיווגworkflow | פעולות מוגדרות מראש ומדידות, כמו RAG, סיווגticketsאו חילוץ נתונים | App / Agentבייצור |
| אין הזנת secrets, credentials או נתוני חקירה רגישים ללא אישור | תקציב workflow נפרד | סביבת אבטחה או מודל פרטי בלבד | נתוני אבטחה לפי אישור | triage,סיכום התראות, יצירת שאילתות, ניתוח ראשוני | SOC / אבטחה |



**STKI Recommends** | Create and enforce basic token policy




Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph

# New roles in the agent era

**Agent Owner**

using the agent, accountable for the agent, from what "correct" means to its business value

**Internal FDE**
forward deployed engineer

builds the agent with (for) the owner until it works reliably in production

**AI security engineer**

builds the protections that enforce those rules as instructed by the Governance Office & CISO

**AI evaluation engineer**

builds the tests and runs them before release and after every change

**Context engineer**

helps the FDE make the agent accurate, with the right data, instructions and tools

**AI Governance Office**

sets the rules for safety, ethics, compliance and risk levels

**More "personas" here**




Copyright@STKI_2026 Do not remove source or attribution from any slide, graph or portion of graph