



# MAGS - Intelligent Automation

DIGITAL TWIN-BASED MULTI AGENT GENERATIVE AI SYSTEMS

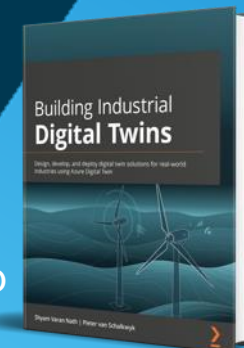


**Pieter Van Schalkwyk**

CEO - XMPPro

DTC Chair: AI Joint Working Group

Composability Framework Subgroup



— 2025 —  
OFFICIAL MEMBER

**Forbes**  
Technology  
Council



**Presentation Will Be Shared**

## A G E N D A

## XMPro Introduction

---

The Problem We Are Solving

---

What Inspired Us?

---

XMPro MAGS Overview

---

Demo

---

Lessons Learnt

---



# Optimize Operations as Fast as You Can Think

AI Powers Your Future Operational Excellence, and XMPro Brings It To Life.

**Reduce unpredictable operational outcomes that add cost, compliance risk, and operational inefficiencies by at least 40%.**

XMPro's Intelligent Business Operations Suite (iBOS) uses AI and real-time decision intelligence to turn your industrial data into prescriptive recommendations and automated workflows.



**5-10% Reduced**  
Maintenance Costs



**10-20% Increased**  
Equipment Uptime



**20-50% Decreased**  
Maintenance Planning Time

**Your operations generate massive amounts of data, but turning that into efficient, profitable actions is a challenge**

In today's data-rich industrial environments, the challenge isn't just collecting data—it's making sense of it all and taking timely, intelligent action. As an operations leader, are you tired of facing these critical challenges?



**Data Overload and  
Disconnected Systems**



**Struggling with Continuous  
Improvement**



**Reactive Operations and  
Inconsistent Processes**



**Limited Decision Support  
and Automation**



**Difficulty in Predicting and  
Preventing Issues**



# Optimize Operations as Fast as You Can Think

AI Powers Your Future Operational Excellence, and XMPro Brings It To Life.

## XMPro's Productized Solution: Intelligent Business Operations Suite (iBOS)

We've created an AI-powered Intelligent Business Operations Suite (iBOS) specifically for industrial operations leaders who need to transform complex data into actionable insights and automated workflows.



### Unify Your Industrial Data

- Standardizes data formats for easy analysis and action



### Apply Advanced AI Analytics

- Offers real-time analysis of operational metrics



### Enables Intelligent Decision Making

- Offers scenario planning and simulation capabilities



### Automates Operational Workflows

- Adapts workflows in real-time based on changing conditions



### Drives Continuous Improvement

- Facilitates knowledge sharing and best practice implementation



### Unifies Operations Across Data Silos

- Enables all departments to access and share critical data



# Australian-Owned Since 2009 with a Global Focus & Reach

## Select Customers



GE APPLIANCES  
a Haier company



## Key Partnerships

Microsoft  
Partner



NVIDIA

HATCH



NTT DATA

intel

DELL Technologies

YOKOGAWA



## Industry Participation



# Our Team Has Experience Working Across a Range of Industries



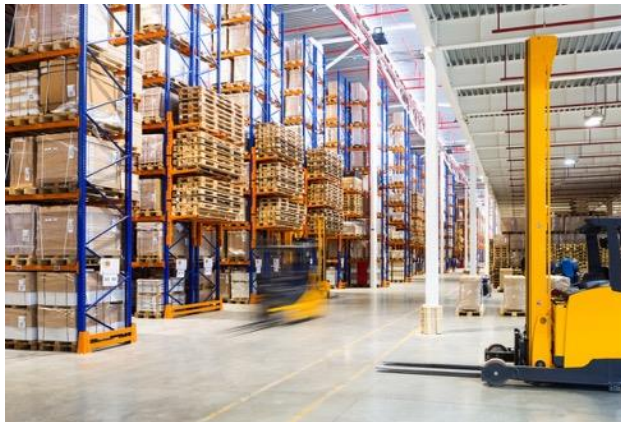
**Mining**



**Oil & Gas**



**Manufacturing**



**Supply Chain & Logistics**



**Utilities**

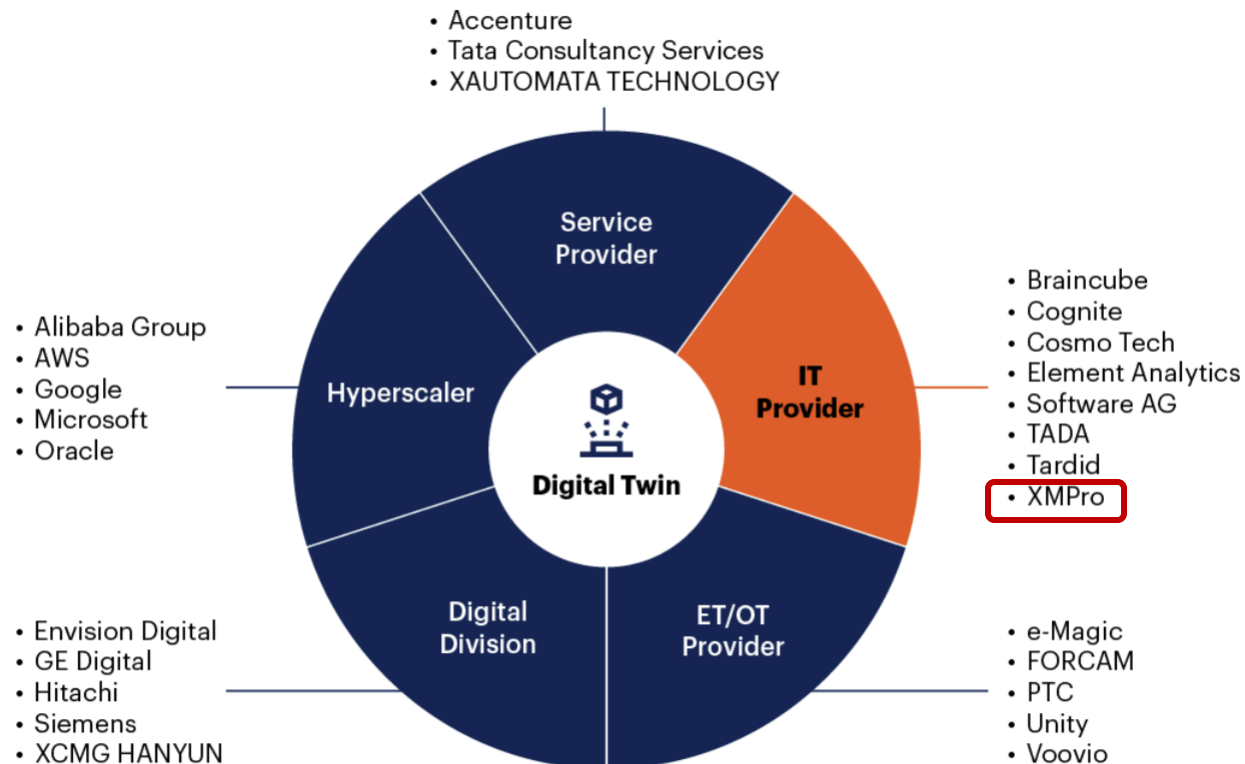


**Environmental Services**



# XMPro: Tech Innovators for Digital Twins

## Technology Innovators for Digital Twins: IT Providers



Source: Gartner  
777302\_C

### Gartner

#### Emerging Tech: Tech Innovators for Digital Twins – IT Providers

Published 29 August 2023 - ID G00777302 - 43 min read

By Analyst(s): Alfonso Velosa, Scot Kim, Eric Goodness

Initiatives: Emerging Technologies and Trends Impact on Products and Services

IT providers are innovating with digital twins to help enterprises improve processes, enhance business visibility and drive revenue. Product leaders must leverage these innovations to differentiate via composable templates, domain knowledge, niche target markets and client value stories.

#### Overview

##### Key Findings

- Digital twins can be built using a broad range of technologies, from IoT and analytics to 3D and more. The business needs of the enterprises often require a portfolio of technologies to be used to build a business solution.
- Most product leaders have not yet established clear differentiation or gained visibility in the market.
- Current digital twin offerings are highly customized to the enterprise's environment due to the lack of templates or standards. Yet few digital twin technology providers have built adequate libraries of digital twin templates that simplify and shorten the time and effort to deploy digital twins.

#### Recommendations

Product leaders working to build digital twin technologies and solutions in their product roadmaps should:

- Optimize their technology portfolio by focusing their scarce engineering resources on their core technology areas. Build partnerships and integration capabilities to third parties for other key, but adjacent, technologies for their customers.

## A G E N D A

XMPro Introduction

---

**The Problem We Are Solving**

---

What Inspired Us?

---

XMPro MAGS Overview

---

Demo

---

Lessons Learnt

---

# Why – Simon Sinek

MAGS isn't just another technology play - it's about putting the right capabilities in the hands of subject matter experts who truly understand industrial challenges. **These experts hold the key to addressing critical issues in operations, sustainability, and economic performance, but they need better tools to turn their knowledge into action at scale.**

**Our objective isn't to replace jobs, but to augment the workforce with virtual workers that operate autonomously under human guidance and supervision within a clear governance framework.** That's what we're doing with Multi-Agent Generative Systems at XMPPro - creating an approach that amplifies expert knowledge to drive real transformation in industrial organizations.

Without this expert-driven change, industries will struggle to meet **mounting challenges** around efficiency, sustainability, and economic viability.

At its core, MAGS provides the framework for subject matter experts to implement solutions that make a meaningful difference while maintaining human oversight of autonomous operations

- Pieter van Schalkwyk, CEO XMPPro



# Why – Simon Sinek

MAGS isn't just another technology play - it's about putting the right capabilities in the hands of subject matter experts who truly understand industrial challenges. **These experts hold the key to addressing critical issues in operations, sustainability, and economic performance, but they need better tools to turn their knowledge into action at scale.**

**Our objective isn't to replace jobs, but to augment the workforce with virtual workers that operate autonomously under human guidance and supervision within a clear governance framework.** That's what we're doing with Multi-Agent Generative Systems at XMPPro -

creating an approach that amplifies expert knowledge to drive real transformation in industrial organizations.

## Not Replacing Jobs, Do Operational Processes and Tasks

Without this expert-driven change, industries will struggle to meet **mounting challenges** around efficiency, sustainability, and economic viability.

At its core, MAGS provides the framework for subject matter experts to implement solutions that make a meaningful difference while maintaining human oversight of autonomous systems.

## Address Operational Challenges for Industrial Organizations

# Why – Simon Sinek

MAGS isn't just another technology play - it's about putting the right capabilities in the hands of subject matter experts who truly understand industrial challenges. **These experts hold the key to addressing critical issues in operations, sustainability, and economic performance, but they need better tools to turn their knowledge into action at scale.**

**Our objective isn't to replace jobs, but to augment the workforce with virtual workers that operate autonomously under human guidance and supervision within a clear governance framework.** That's what we're doing with Multi-Agent Generative Systems at XMPPro - creating an approach that amplifies expert knowledge to drive real transformation in industrial organizations.

Without this expert-driven change, industries will struggle to meet **mounting challenges** around efficiency, sustainability, and economic viability.

At its core, MAGS provides the framework for subject matter experts to implement solutions that make a meaningful difference while maintaining human oversight of autonomous operations



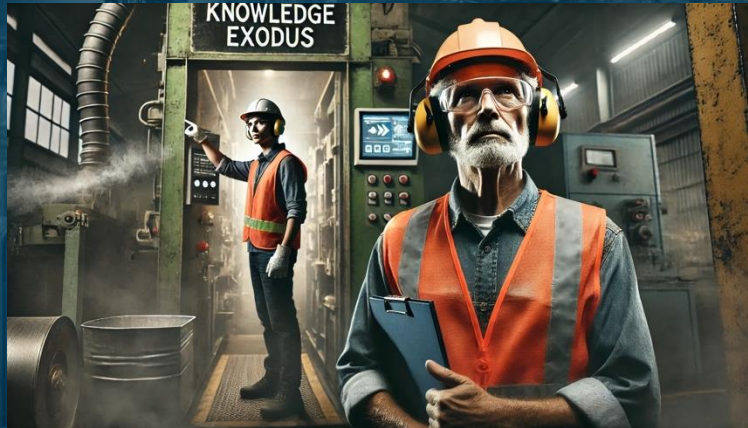
# Perfect Storm Disrupting Industrial Operations

## Growing Technical Complexity



- Modern equipment demands deeper technical understanding
- Systems generate more data and alerts than humans can process
- Interconnected processes create complex decision scenarios

## The Knowledge Exodus



- Experienced operators retire faster than we can replace them
- Critical operational knowledge walks out the door
- Training programs can't keep pace with technology changes

## Operational Pressures



- Teams must maintain peak performance with fewer staff
- Energy and resource costs keep rising
- Environmental and safety regulations become more stringent



# Change in Laborforce in the next 10 Years

## Strategic Planning Assumptions

**Strategic Planning Assumption:** By 2028, there will be more smart robots than frontline workers in manufacturing, retail and logistics due to labor shortages.

Analysis by: Dwight Klappich

### Key Findings:

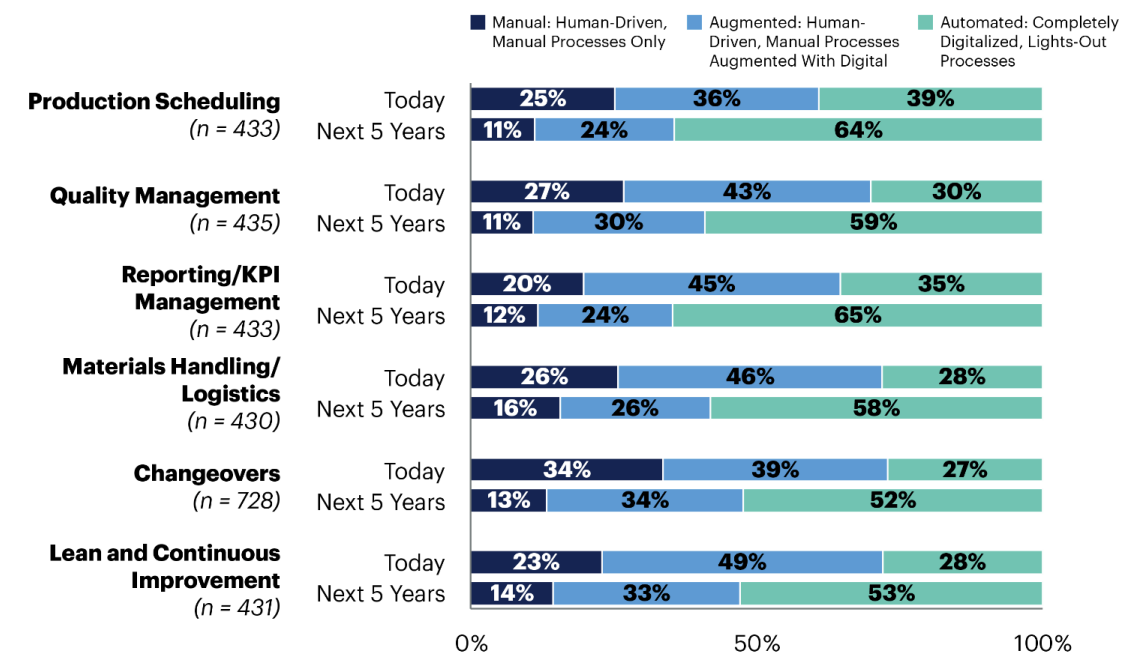
- Labor has become as big a constraint on operational performance as is product availability.
- The U.S. workforce will grow nearly five times slower than the U.S. GDP over the next decade. U.S. GDP is expected to grow at 2.0% while the overall workforce will grow only 0.9%. Furthermore, frontline workers in transportation will grow only 1.1%, while the manufacturing and retail workforces are expected to shrink by 0.1% and 0.4%, respectively.
- Workforces in established economies are aging. For example, by 2030 the U.S. Department of Labor expects the median age of the workforce to rise from 40 to 43. In Europe and parts of Asia.
- Labor participation rates are declining. For example, U.S. Labor participation rates are projected to decline from 67% in 2000 to 60.4% in 2030.
- 1 in 4 over 65
- 96% of supply chain technology users in Gartner's Supply Chain Technology User Wants and Needs Survey that have already deployed or plan to deploy cyber-physical automation by 2024 said they were investing in or planned to invest in robotics. And 35% have already deployed robots with 61% piloting or in the middle of their first implementation.

## Workforce 5x slower than GDP

## Manufacturing workforce shrink

## Expected Shifts in Hyper-Automation by Process Today and in 2025

State of Manufacturing Processes/Activities



n varies, manufacturers; excludes "Don't know/NA"


Q: What is the current state of each manufacturing process/activity below?

Q: Now please tell us what you anticipate the state of each manufacturing process/activity to be five years from now?

Source: Gartner's Smart Manufacturing Strategy & Implementation Trends, 2020

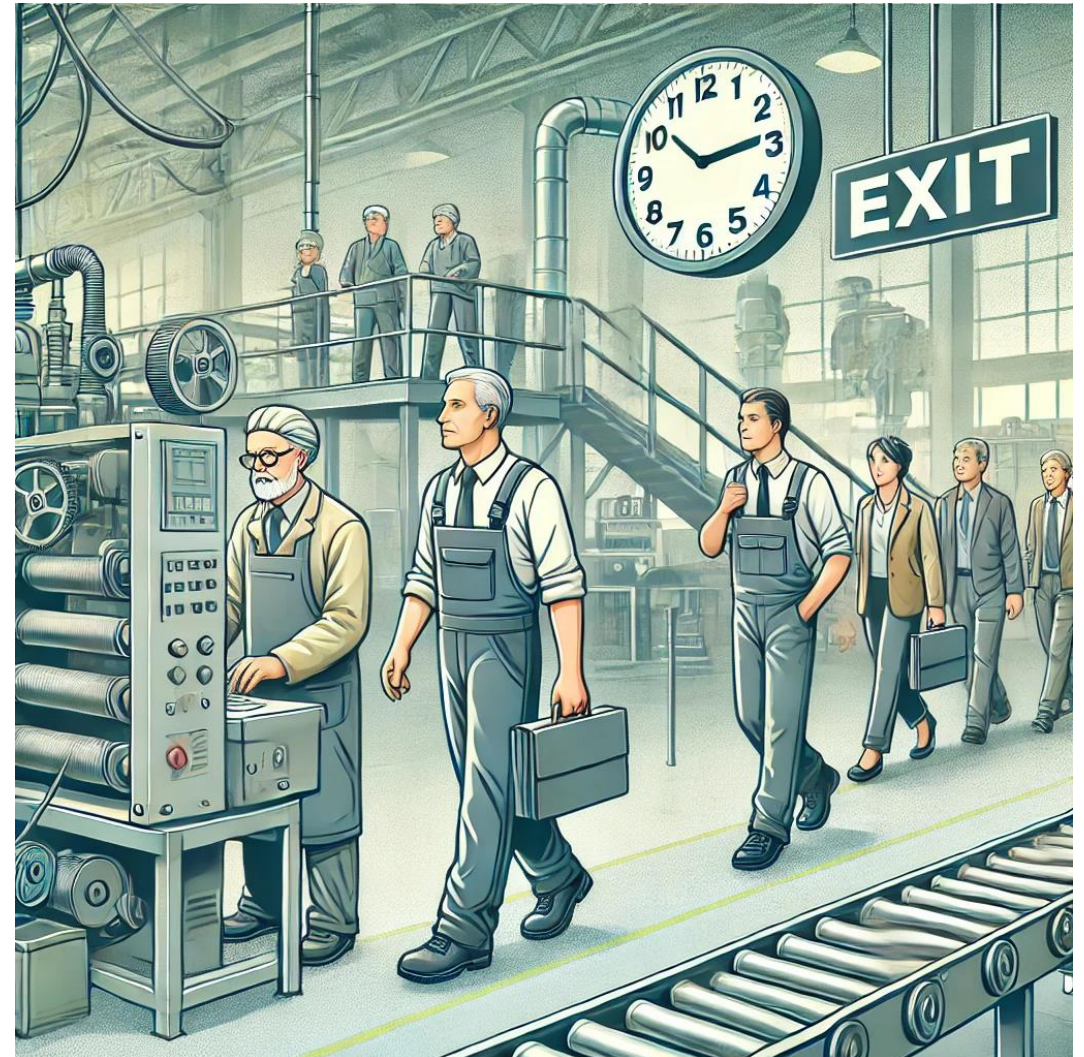
763904\_C

# Skills Problem for Industrial (US Manufacturing)

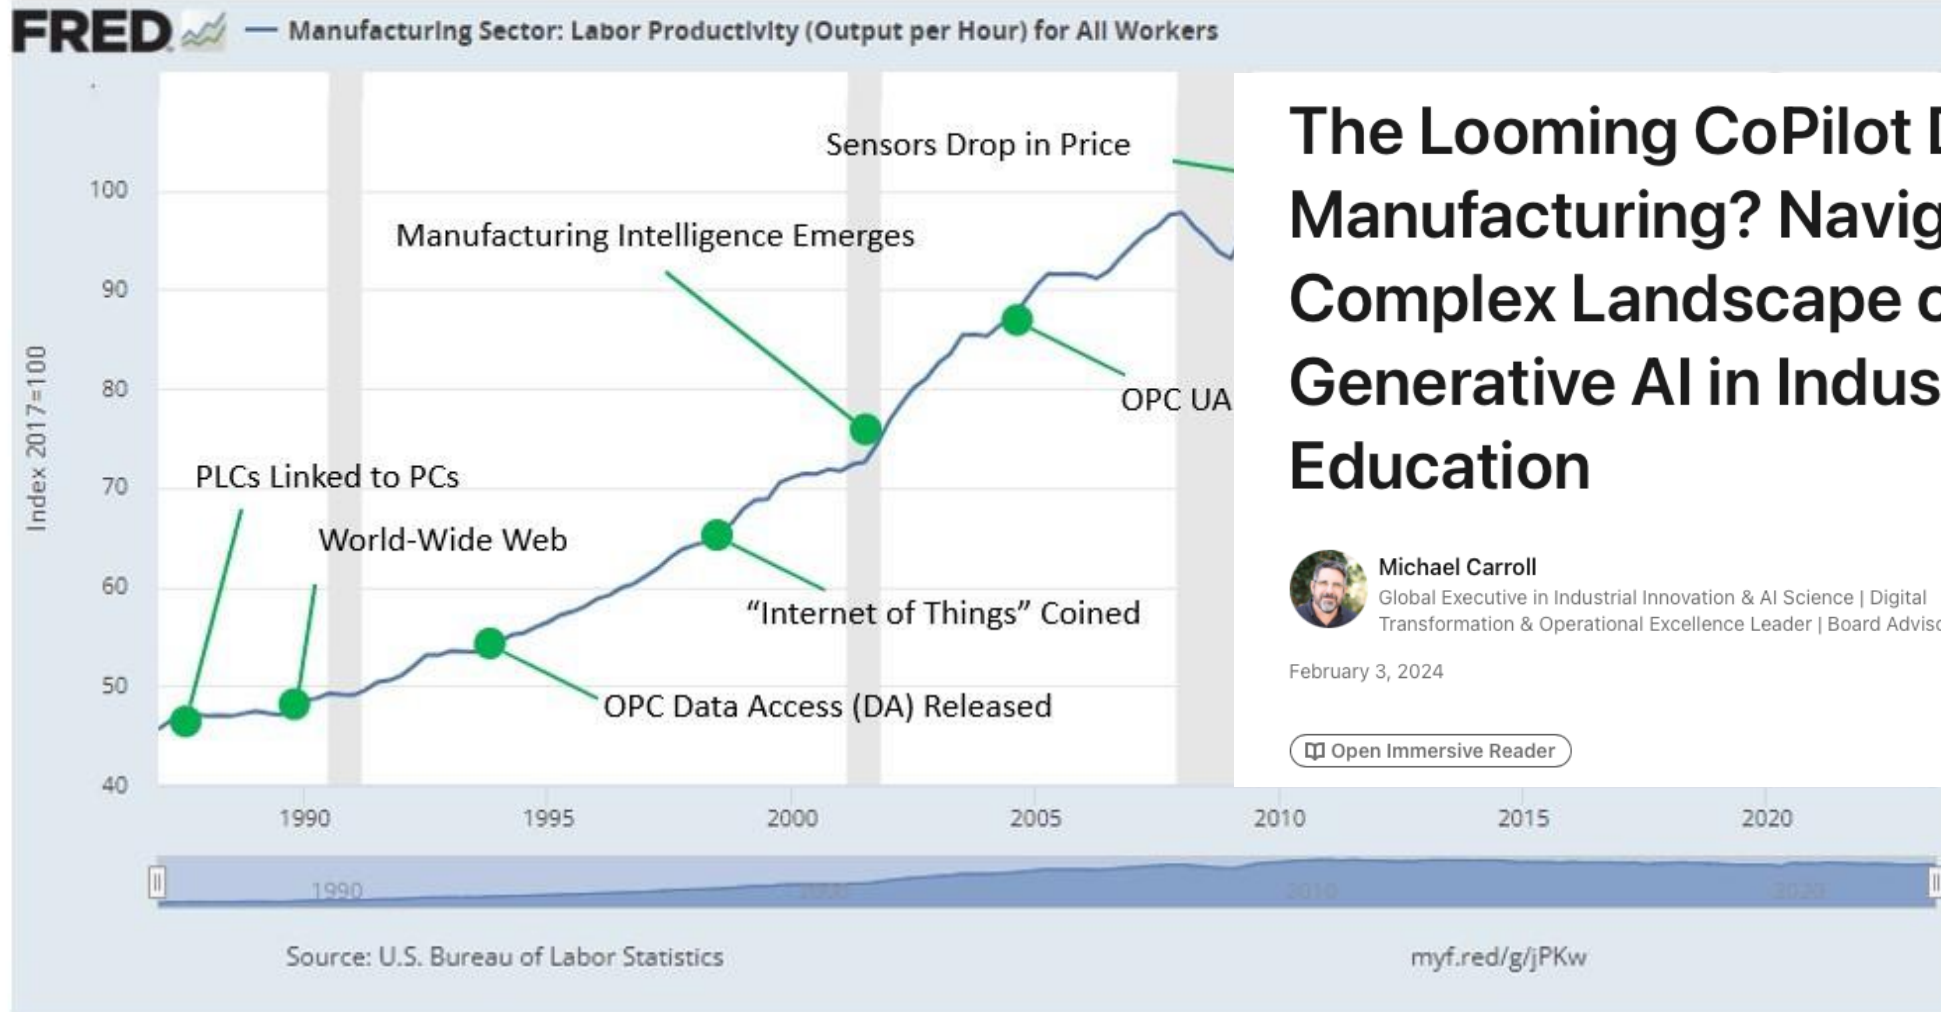
|  <b>A Typical Manufacturer's Workforce</b> |          |          |
|---|----------|----------|
|   | 2019     | 2023     |
| Average Tenure  | 20 Years | 3 Years  |
| Average Time in Position  | 7 Years  | 9 Months |
| Average 3-Month Retention Rate  | 90%      | 50%      |

Source:  **LNS**  
Research

<https://blog.lnsresearch.com/how-to-get-a-30-day-employee-to-a-30-year-performance-level>



# Are We Ready For A Step Change?



## The Looming CoPilot Disaster in Manufacturing? Navigating the Complex Landscape of Generative AI in Industry and Education



**Michael Carroll**

Global Executive in Industrial Innovation & AI Science | Digital Transformation & Operational Excellence Leader | Board Advisor |...

19 articles

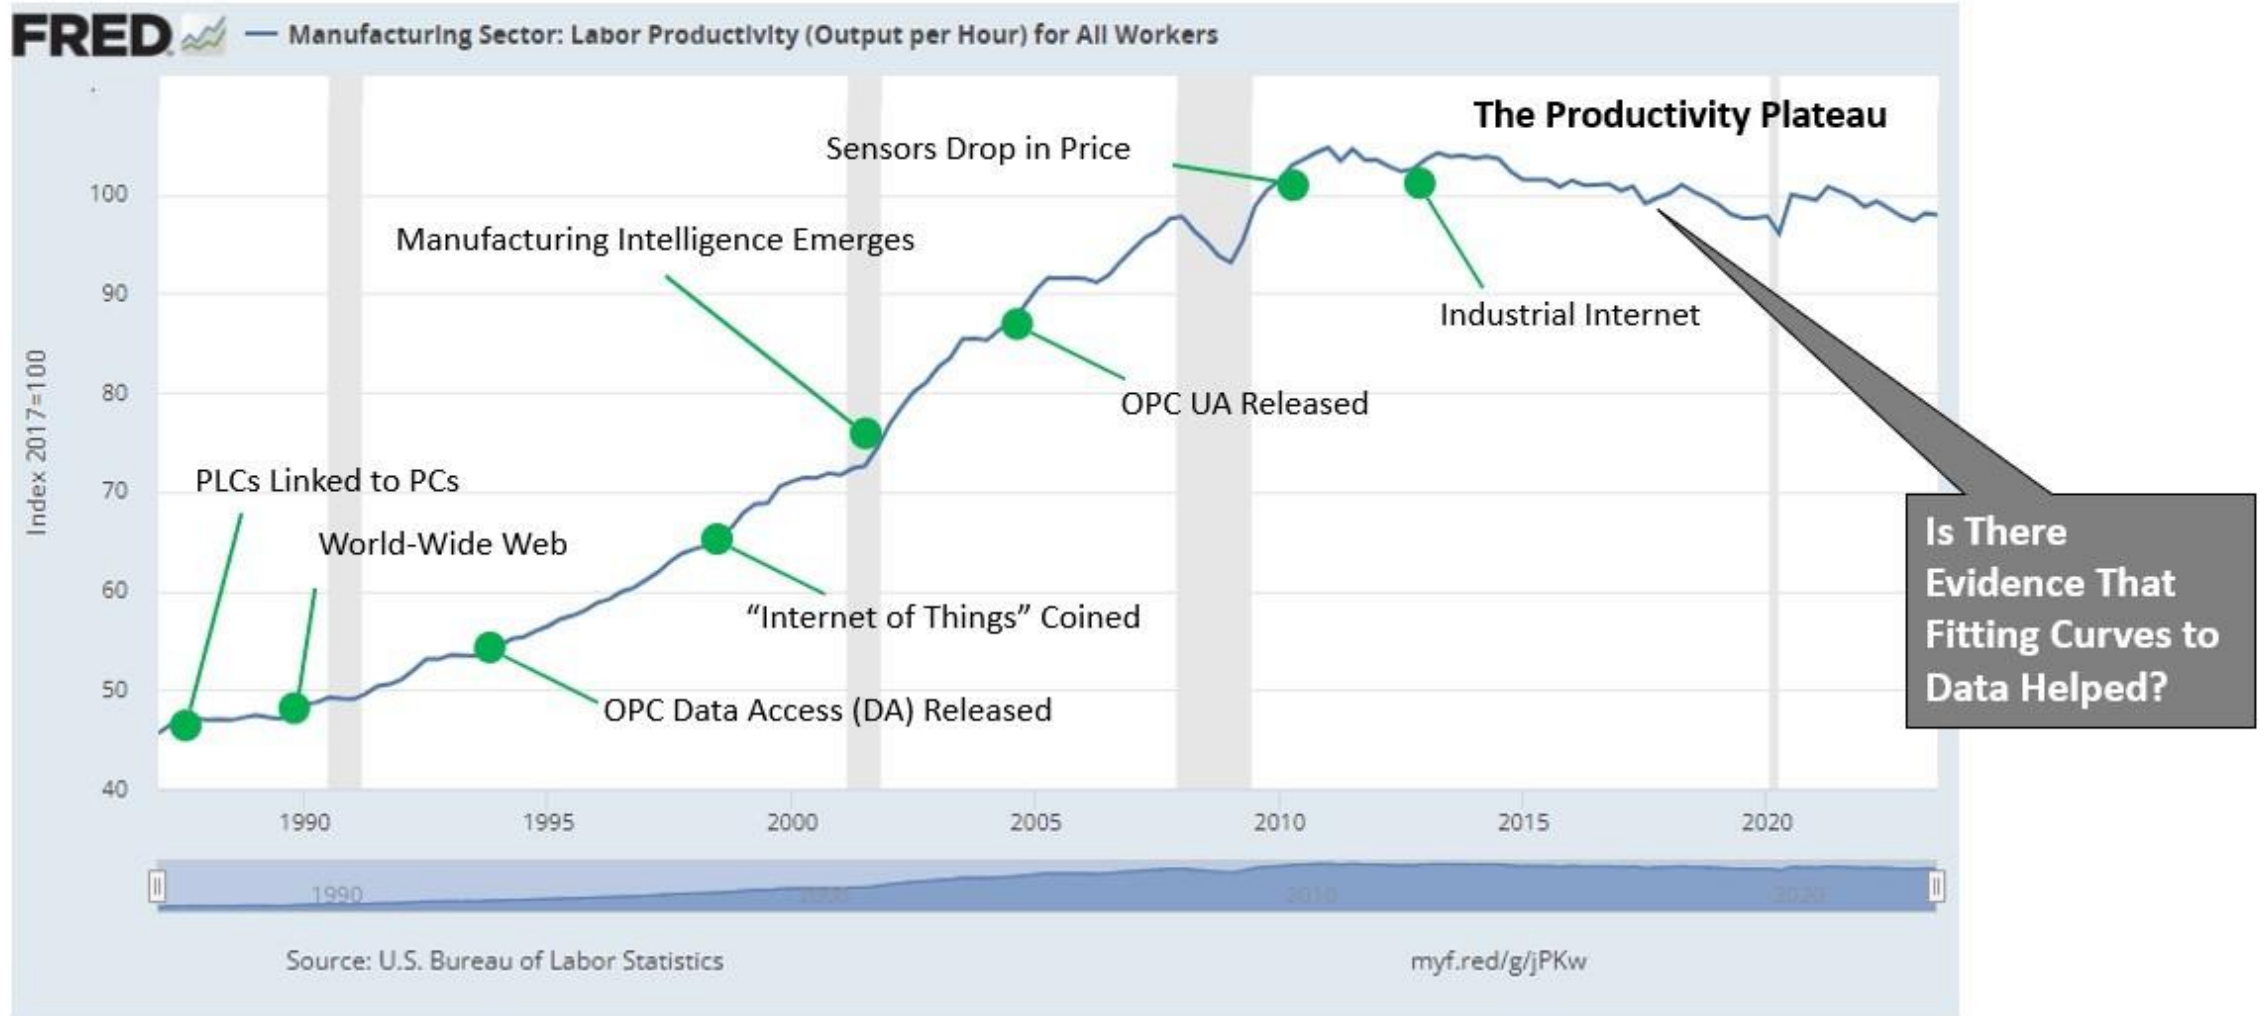
[+ Follow](#)

February 3, 2024

[Open Immersive Reader](#)



# Are We Ready For A Step Change?



# Perfect Storm Will Disrupt Industrial Ops in 2024

**Lack of Skills and Aging Workforce in Manufacturing**

**Generative AI Accelerate from Experimentation to Implementation**

**GenAI Digital Shift is Opportunity Switch to Boost Productivity**

## The Combination of these Three Trends Creates the Perfect Storm for Disruptive Change And A Shift Towards an AI Augmented Workforce

The US workforce will grow 5 times slower than GDP in next decade. 25% of workforce will be 65 or older

AI is changing decision-making. Copilots became mainstream in 2023. 2024 will see useful Gen AI applications

Do more with less people by automating repetitive tasks and let SME focus where they can add most value

## A G E N D A

XMPro Introduction

---

The Problem We Are Solving

---

**What Inspired Us?**

---

XMPro MAGS Overview

---

Demo

---

Lessons Learnt

---

# Digital Twin (DT) vs Intelligent Digital Twin (IDT)

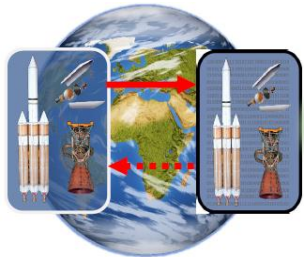
## Digital Twin (DT)

**Passive** Repositories

**Offline**, wait for physical twin to actuate

**Goal Given**, Measuring Performance

**Predictive**, but not optimizing



Data

VS

## Intelligent Digital Twin (IDT)

**Active**, Always On Agents

**Online** (real-time), scan and actuate

**Goal Seeking** shared by human and twin

**Anticipatory & FRS** manipulate time



Dataflow



# Generative Agents Can Simulate Human Behavior

## Generative Agents: Interactive Simulacra of Human Behavior

Joon Sung Park  
Stanford University  
Stanford, USA  
joonspk@stanford.edu

Joseph C. O'Brien  
Stanford University  
Stanford, USA  
jobrien3@stanford.edu

Carrie J. Cai  
Google Research  
Mountain View, CA, USA  
cjcai@google.com

Meredith Ringel Morris  
Google DeepMind  
Seattle, WA, USA  
merrie@google.com

Percy Liang  
Stanford University  
Stanford, USA  
pliang@cs.stanford.edu

Michael S. Bernstein  
Stanford University  
Stanford, USA  
msb@cs.stanford.edu



Figure 1: Generative agents are believable simulacra of human behavior for interactive applications. In this work, we demonstrate generative agents by populating a sandbox environment, reminiscent of The Sims, with twenty-five agents. Users can observe and intervene as agents plan their days, share news, form relationships, and coordinate group activities.

### ABSTRACT

Believable proxies of human behavior can empower interactive applications ranging from immersive environments to rehearsal spaces for interpersonal communication to prototyping tools. In this paper, we introduce generative agents: computational software agents that simulate believable human behavior. Generative agents wake up, cook breakfast, and head to work; artists paint, while

authors write; they form opinions, notice each other, and initiate conversations; they remember and reflect on days past as they plan the next day. To enable generative agents, we describe an architecture that extends a large language model to store a complete record of the agent's experiences using natural language, synthesize those memories over time into higher-level reflections, and retrieve them dynamically to plan behavior. We instantiate generative agents to populate an interactive sandbox environment inspired by The Sims, where end users can interact with a small town of twenty-five agents using natural language. In an evaluation, these generative agents produce believable individual and emergent social behaviors. For example, starting with only a single user-specified notion that one agent wants to throw a Valentine's Day party, the agents autonomously spread invitations to the party over the next two

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).  
UST '23, October 29–November 1, 2023, San Francisco, CA, USA  
© 2023 Copyright held by the owner/author(s).  
ACM ISBN 978-1-407-432-023/10  
<https://doi.org/10.1145/3586183.3606763>

## 3.1 Agent Avatar and Communication

A community of 25 unique agents inhabits Smallville. Each agent is represented by a simple sprite avatar. We authored one paragraph of natural language description to depict each agent's identity, including their occupation and relationship with other agents, as seed memories. For example, John Lin has the following description:

John Lin is a pharmacy shopkeeper at the Willow Market and Pharmacy who loves to help people. He is always looking for ways to make the process of getting medication easier for his customers; John Lin is living with his wife, Mei Lin, who is a college professor, and son, Eddy Lin, who is a student studying music theory; John Lin loves his family very much; John Lin has known the old couple next-door, Sam Moore and Jennifer Moore, for a few years; John Lin thinks Sam Moore is a kind and nice man; John Lin knows his neighbor, Yuriiko Yamamoto, well; John Lin knows of his neighbors, Tamara Taylor and Carmen Ortiz, but has not met them before; John Lin and Tom Moreno are colleagues at The Willows Market and Pharmacy; John Lin and Tom Moreno are friends and like to discuss local politics together; John Lin knows the Moreno family somewhat well – the husband Tom Moreno and the wife Jane Moreno.

Each semicolon-delimited phrase is entered into the agent's initial memory as memories at the start of the simulation.



# The Memory Cycle (Reasoning)

UIST '23, October 29-November 1, 2023, San Francisco, CA, USA

J.S. Park, J.C. O'Brien, C.J. Cai, M.R. Morris, P. Liang, M.S. Bernstein

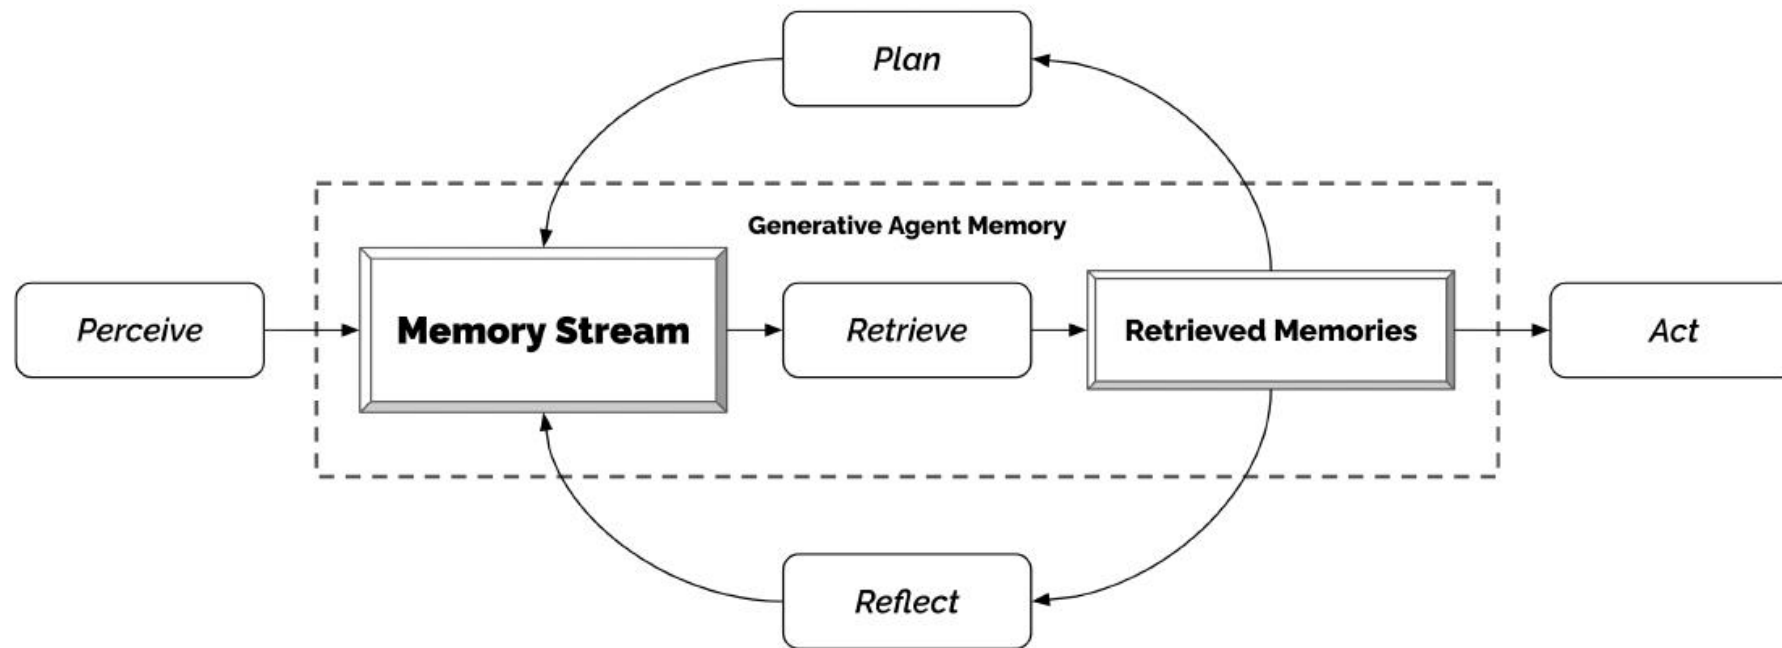


Figure 5: Our generative agent architecture. Agents perceive their environment, and all perceptions are saved in a comprehensive record of the agent’s experiences called the memory stream. Based on their perceptions, the architecture retrieves relevant memories and uses those retrieved actions to determine an action. These retrieved memories are also used to form longer-term plans and create higher-level reflections, both of which are entered into the memory stream for future use.

# The Memory Cycle (Reasoning)

Generative Agents

UIST '23, October 29-November 1, 2023, San Francisco, CA, USA

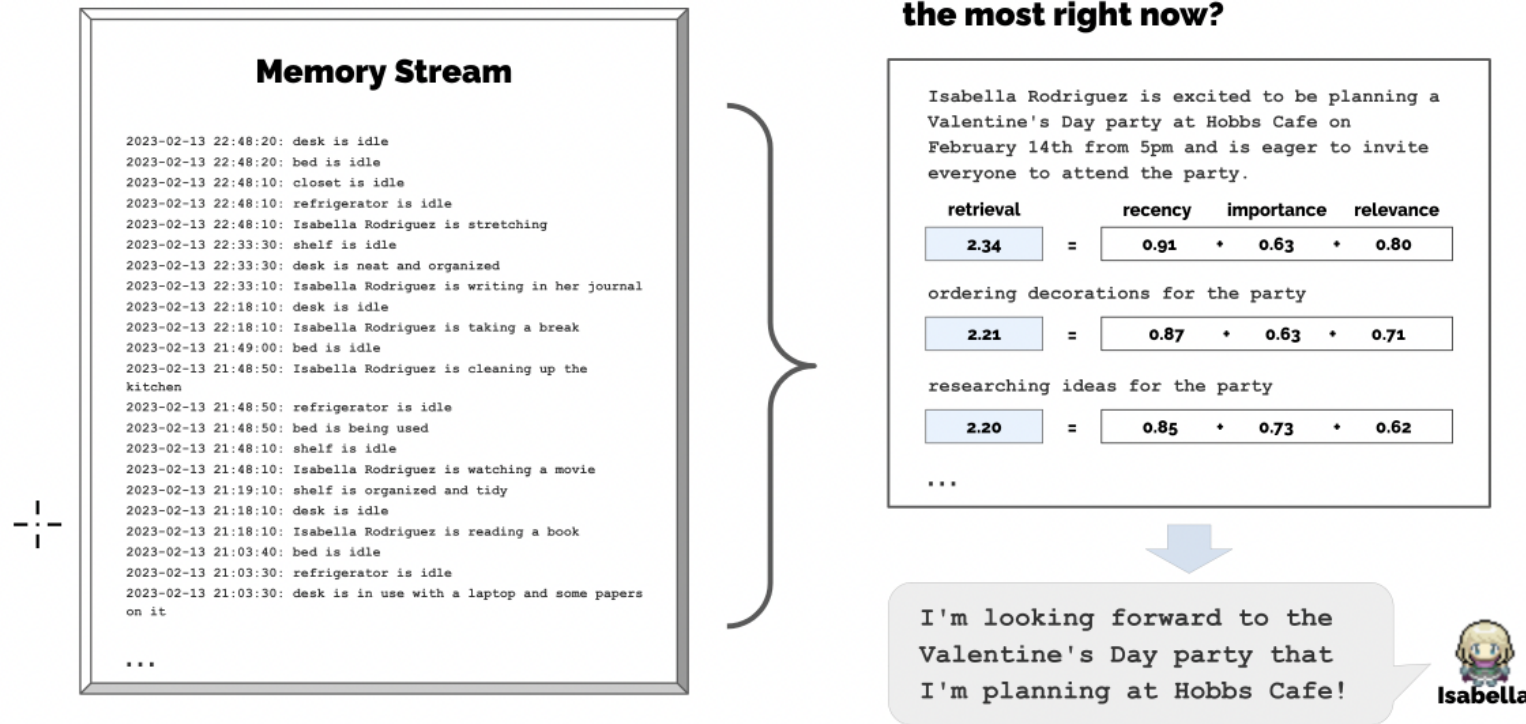


Figure 6: The memory stream comprises a large number of observations that are relevant and irrelevant to the agent's current situation. Retrieval identifies a subset of these observations that should be passed to the language model to condition its response to the situation.

# Andrew Ng on Agentic AI Workflows

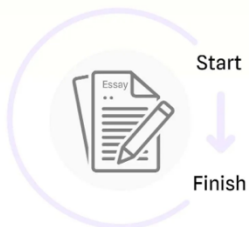
## Generative AI to Agentic AI Workflows

What is Agentic AI and how is it different from Generative AI

### LLM-Based agents

Non-agentic workflow  
(Zero-shot):

Please type out an essay on topic X from start to finish in one go, without using backspace.



Agentic workflow:

Write an essay outline on topic X

Do you need any web research?

Write a first draft.

Consider what parts need revision or more research.

Revise your draft.



## AI Ascent by Sequoia Capital What's next for Agentic Reasoning

Dr. Andrew Ng is a globally recognized leader in AI (Artificial Intelligence). He is Founder of [DeepLearning.AI](#), Founder & CEO of [Landing AI](#), General Partner at [AI Fund](#), Chairman and Co-Founder of [Coursera](#) and an Adjunct Professor at Stanford University's Computer Science Department.



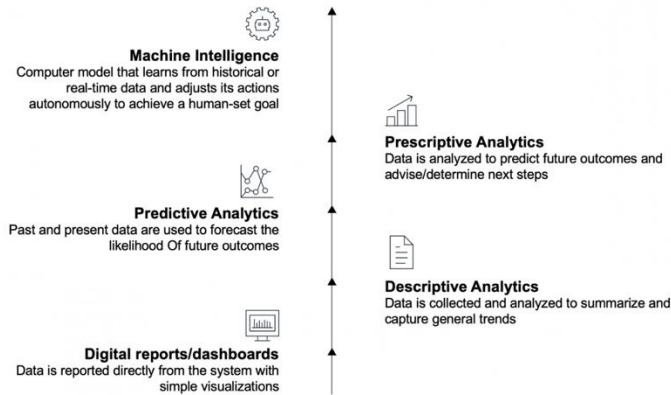
# AutoRT:



AutoRT: Embodied Foundation Models for Large Scale Orchestration of Robotic Agents – Google DeepMind

# MIMO @ MIT – “Machines” Optimize Better Than People

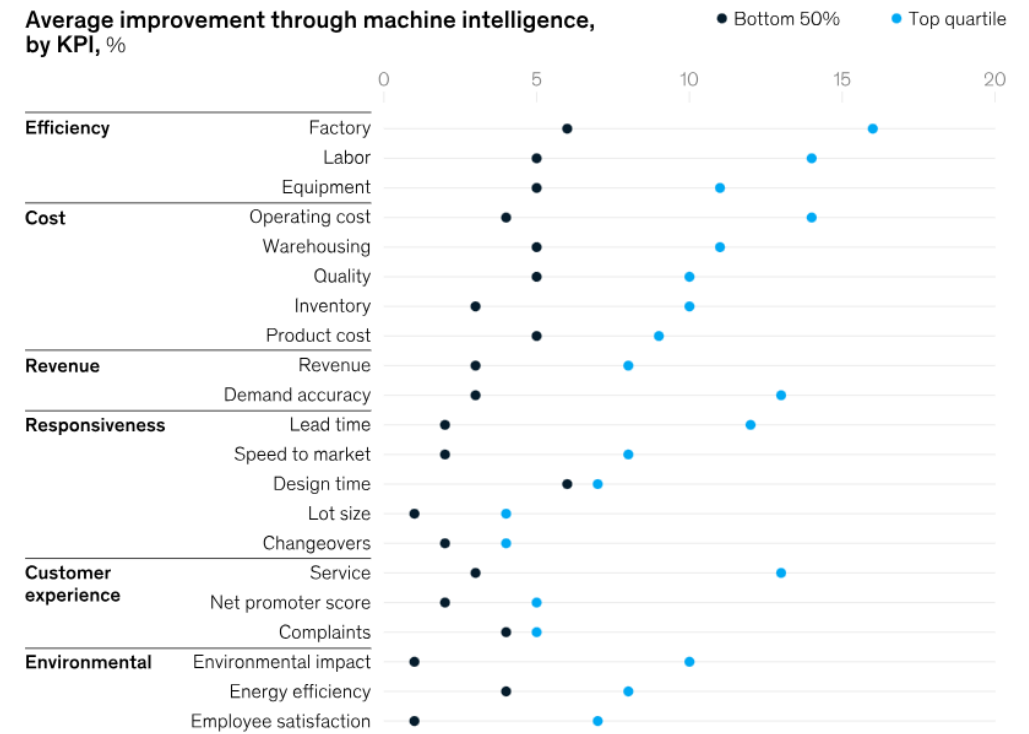
## What do we mean by Machine Intelligence?



Machine intelligence goes beyond predictive analytics or prescriptive analytics. When we say machine intelligence, we refer to a computer model that learns from historical or real-time data and adjusts its actions autonomously to achieve a human-set goal.

Across a broad range of metrics, machine-intelligence leaders achieve triple the improvement of other companies.

Average improvement through machine intelligence, by KPI, %



Note: All variables are normalized to a 0 to 1 scale.  
Source: MIT Machine Intelligence for Manufacturing and Operations (MIMO) program; McKinsey Machine Intelligence Survey

### A Machine Intelligence Transformation in Progress

#### Operations and Manufacturing paradigm shifts:

- 1980's Production Method shift - Mass production to Lean Manufacturing
- 2000's Supply Chain shift - Local sourcing to Globalization
- 2010's Digitization shift - Programmatic to Machine Intelligence



**Duane Boning**  
boning@mit.edu

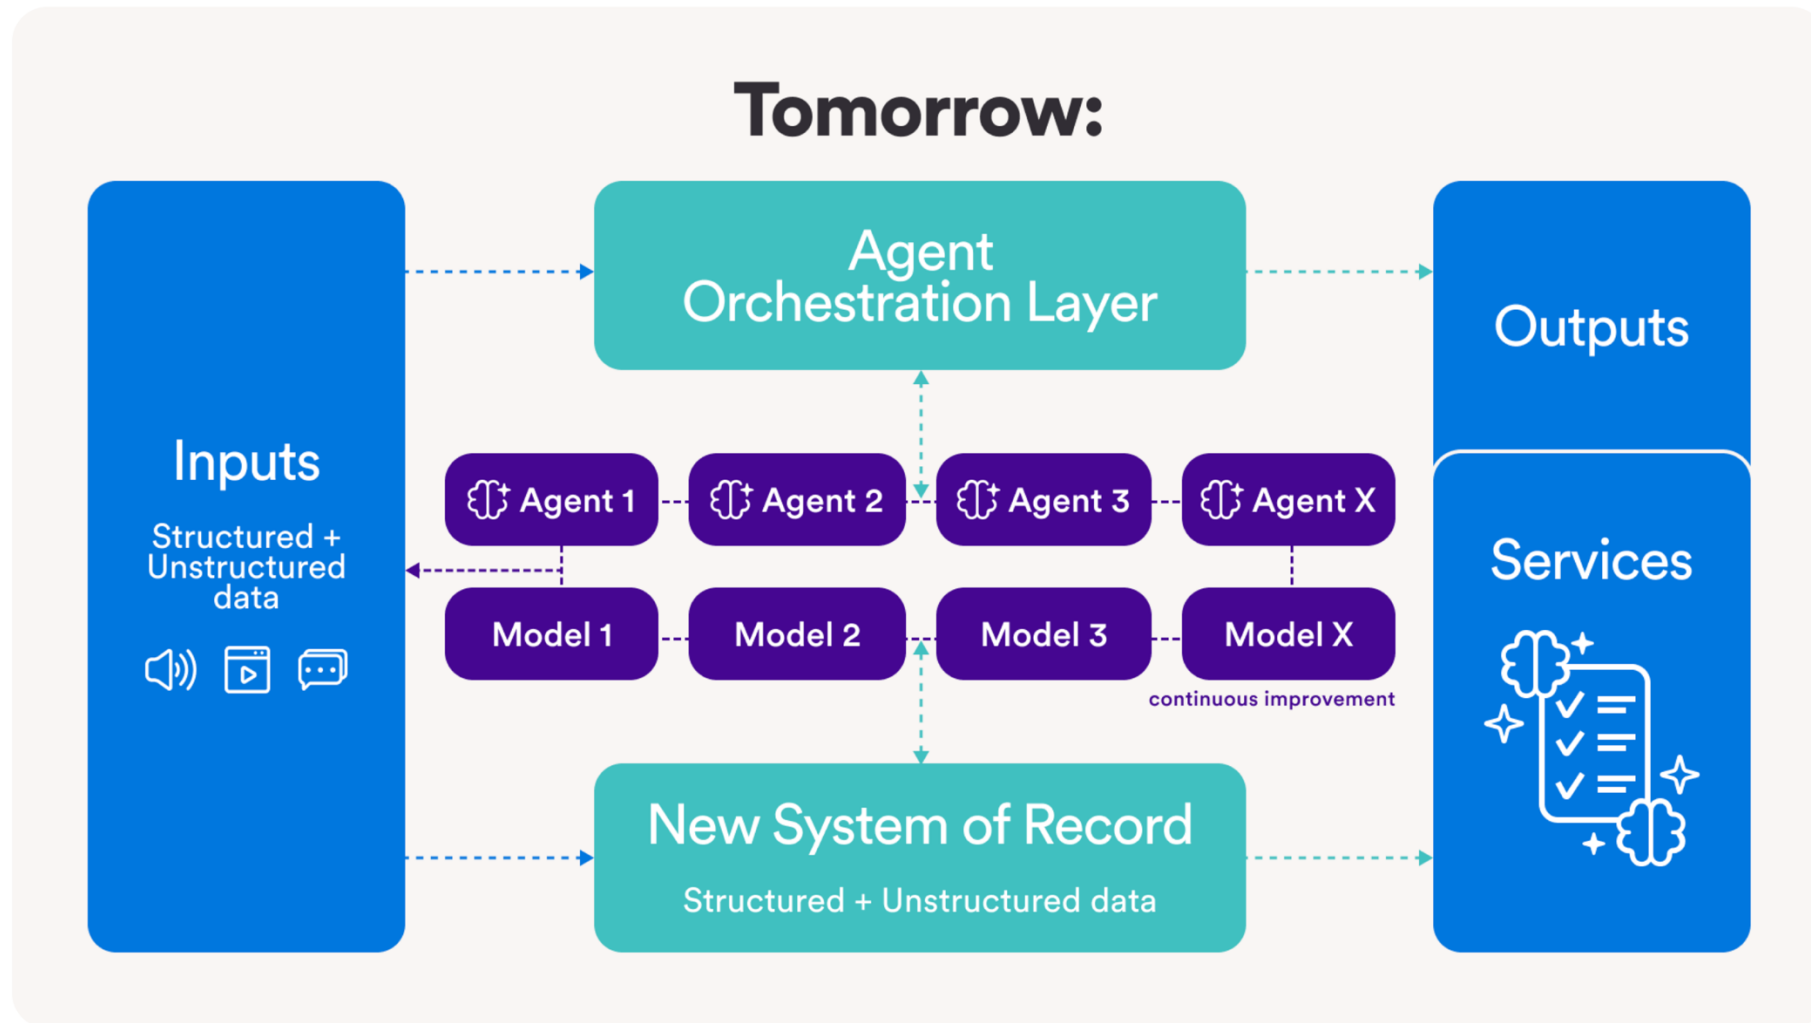
Clarence J. LeBel Professor in Electrical Engineering  
MIT Electrical Engineering and Computer Science department

Engineering Faculty Co-Director  
Leaders for Global Operations (LGO) Program

Associate Director, Computation  
Microsystems Technology Laboratories

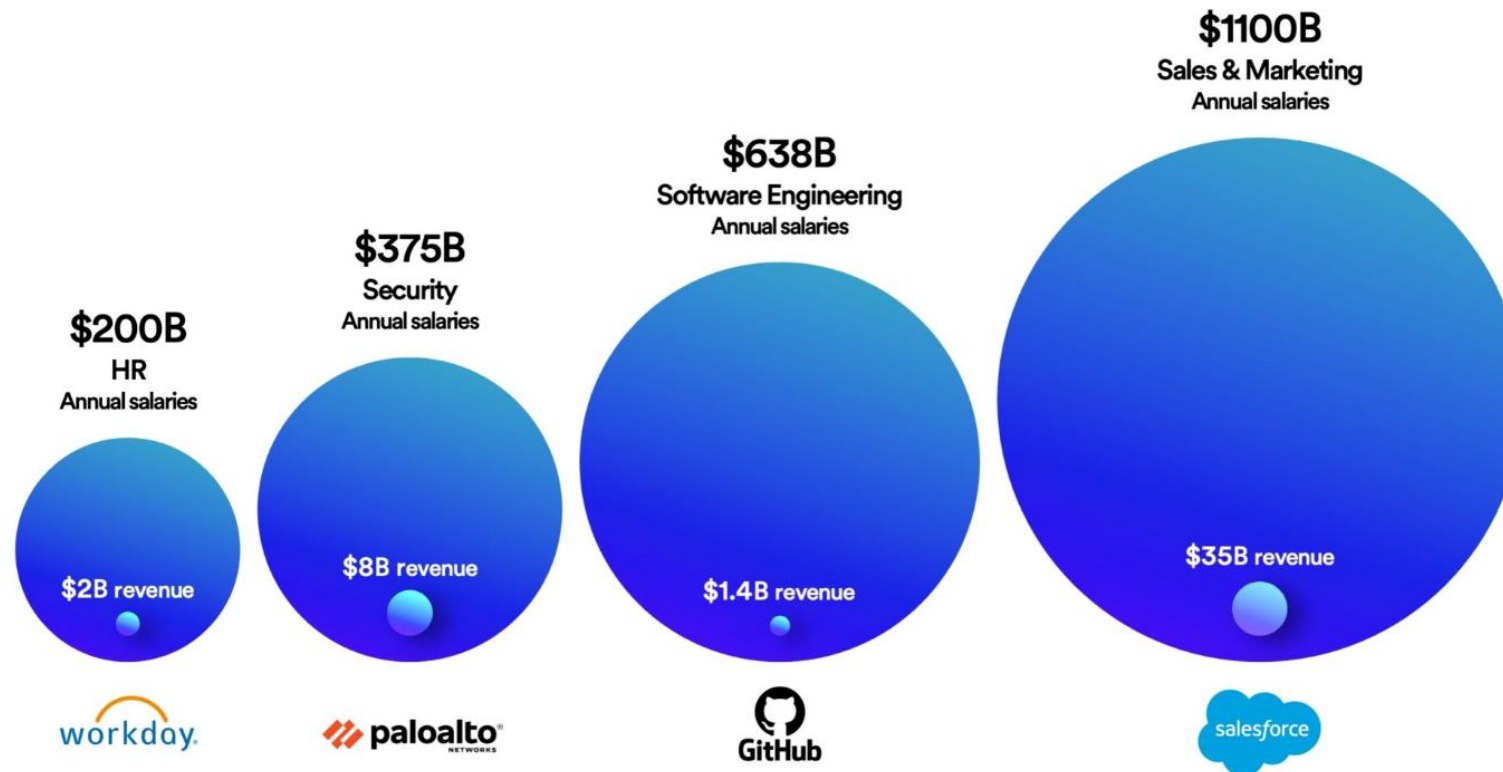


# A System of Agents brings Service-as-Software to life



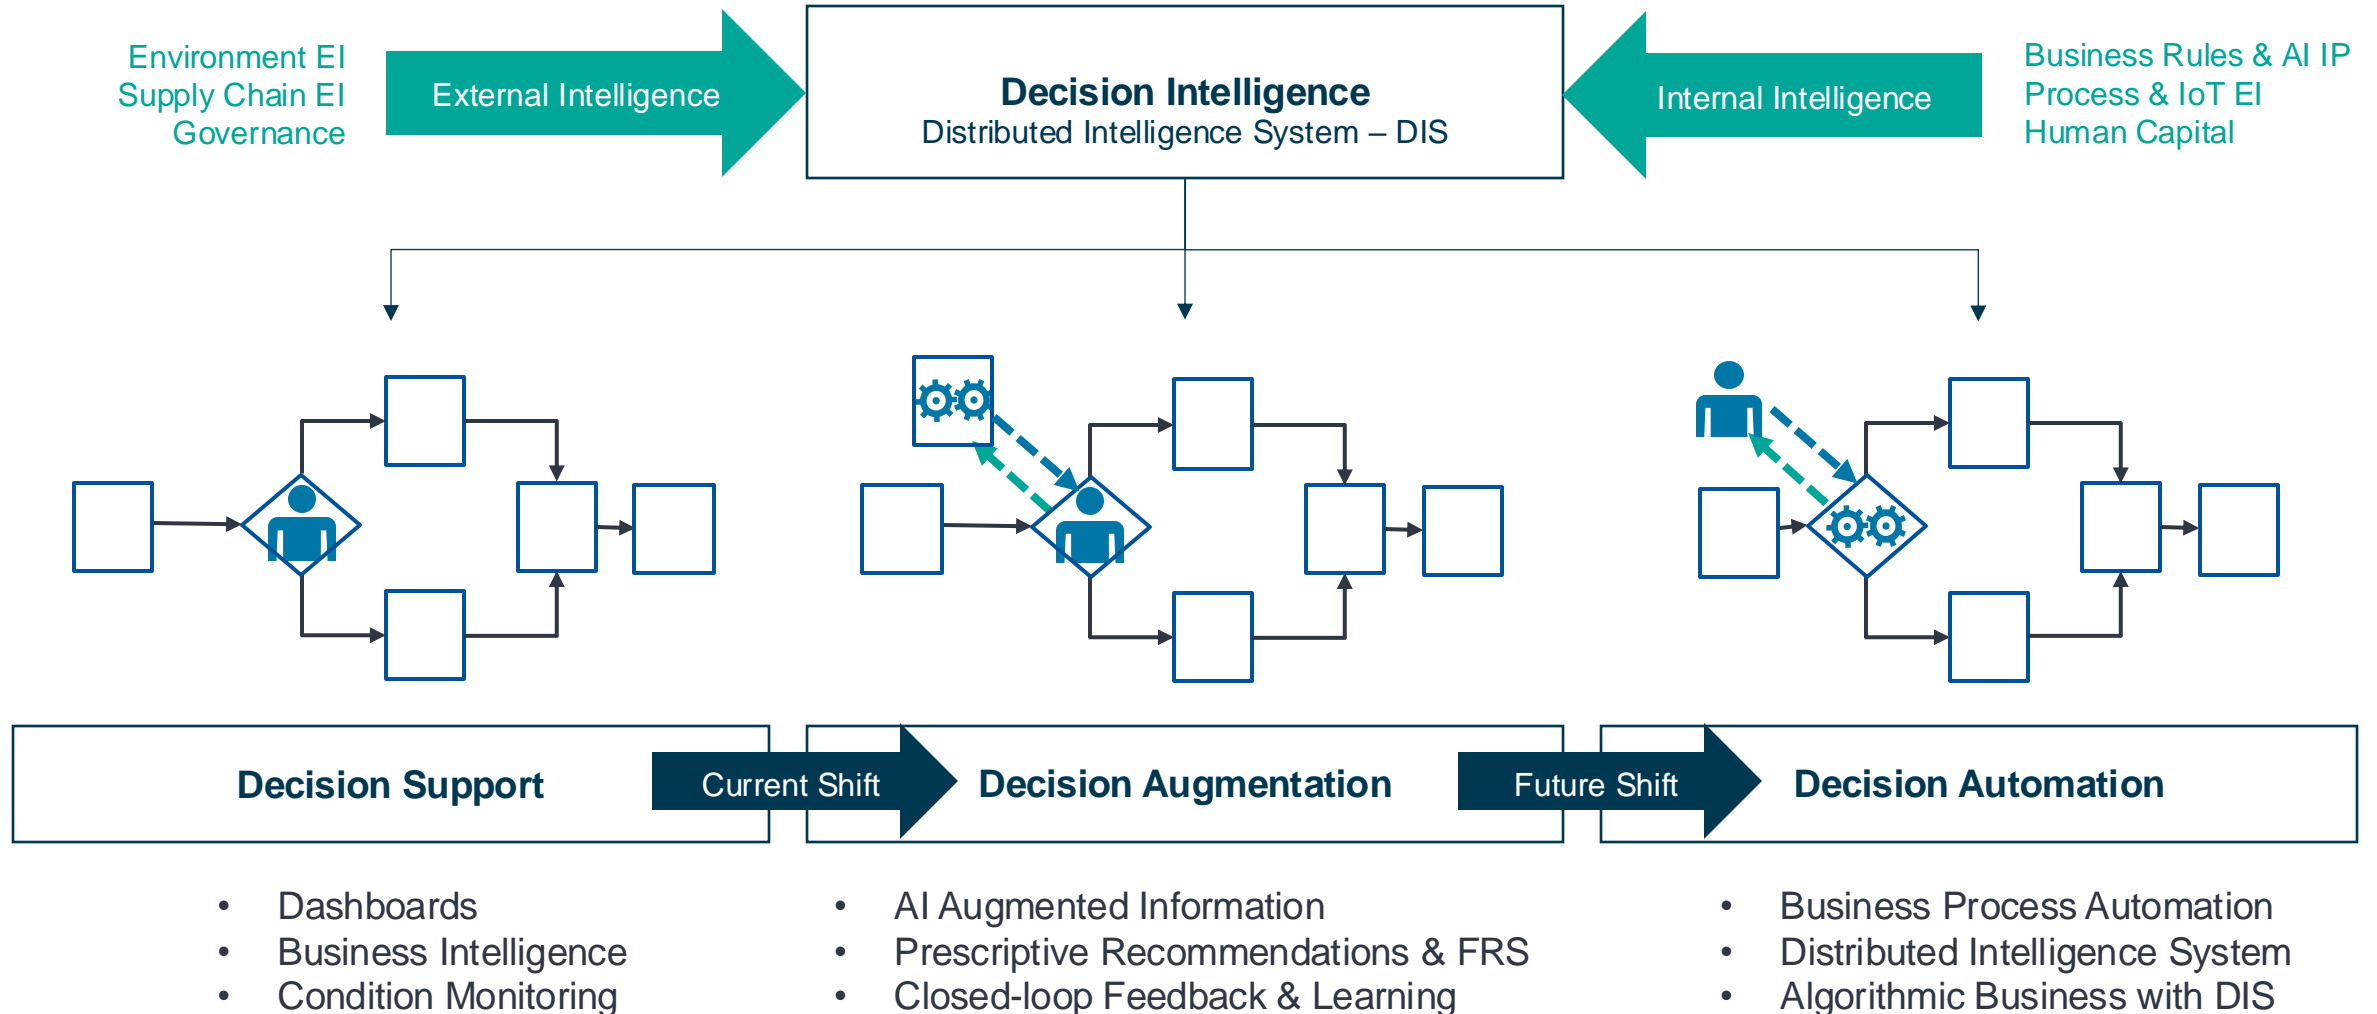


# A System of Agents brings Service-as-Software to life



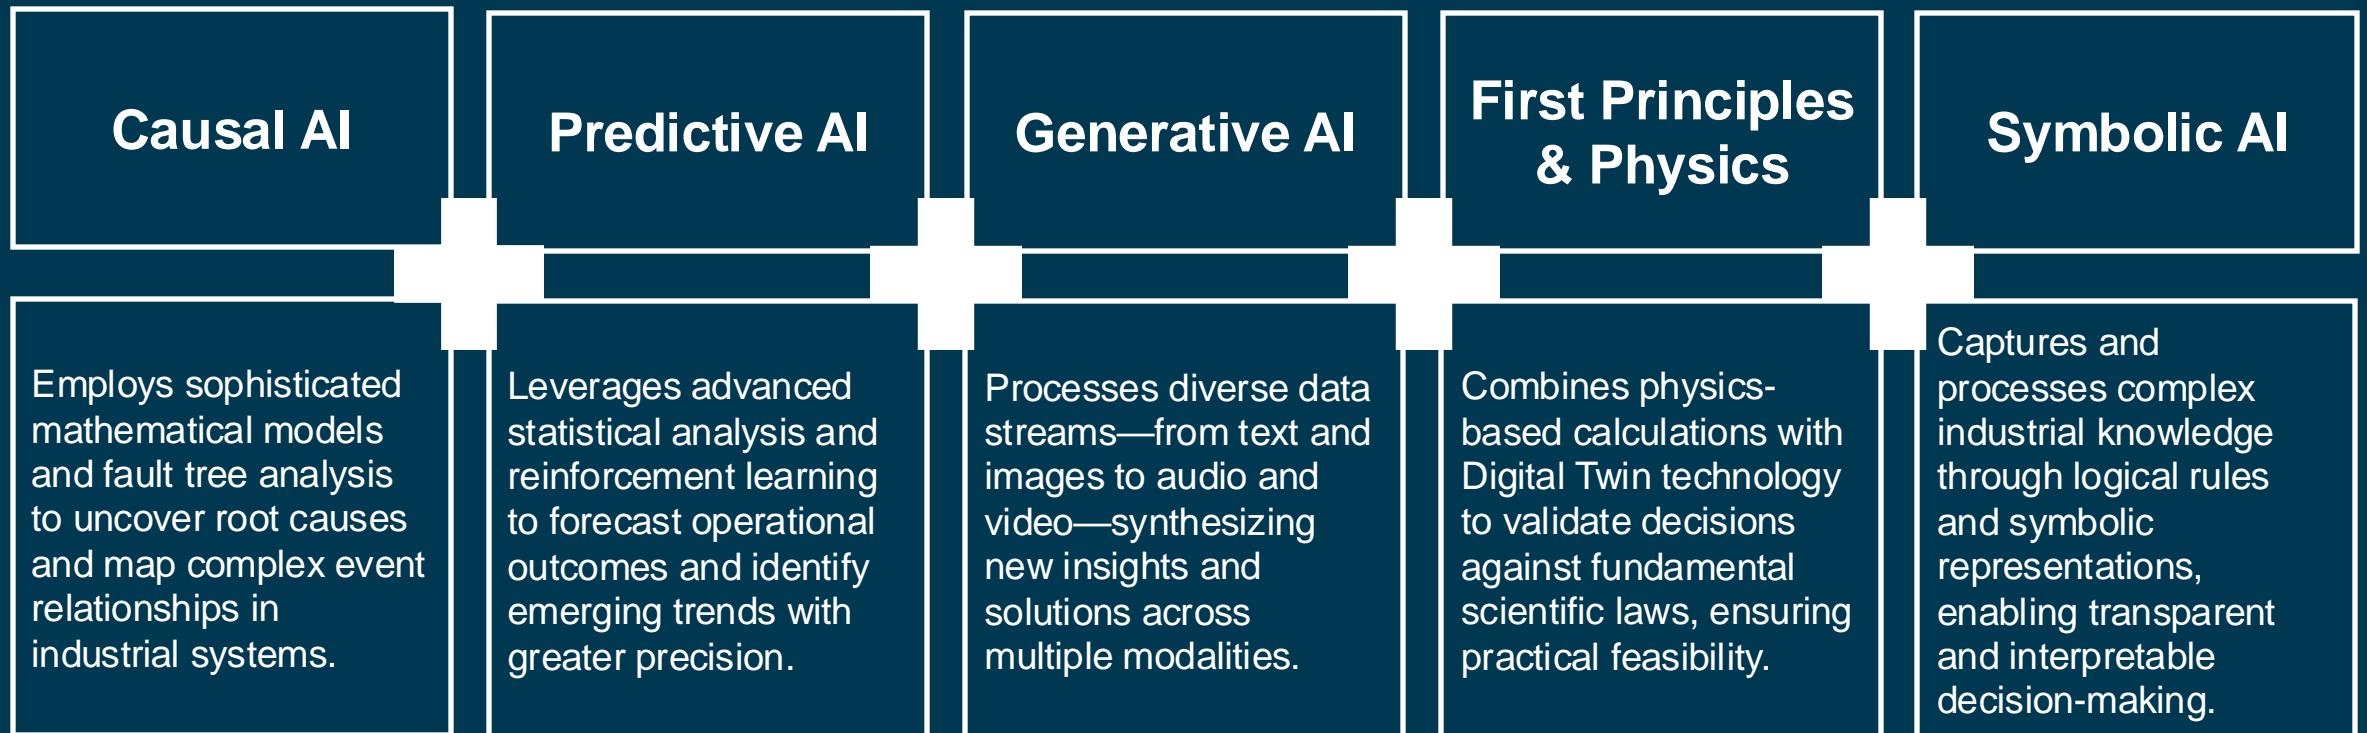
Sources: Turing, Statista.

# Intelligent Twins: Decision-making is Changing



# Composite AI

## Recombines The Best of All Intelligence Worlds





## A G E N D A

XMPro Introduction

---

The Problem We Are Solving

---

What Inspired Us?

---

**XMPro MAGS Overview**

---

Demo




---

Lessons Learnt

---

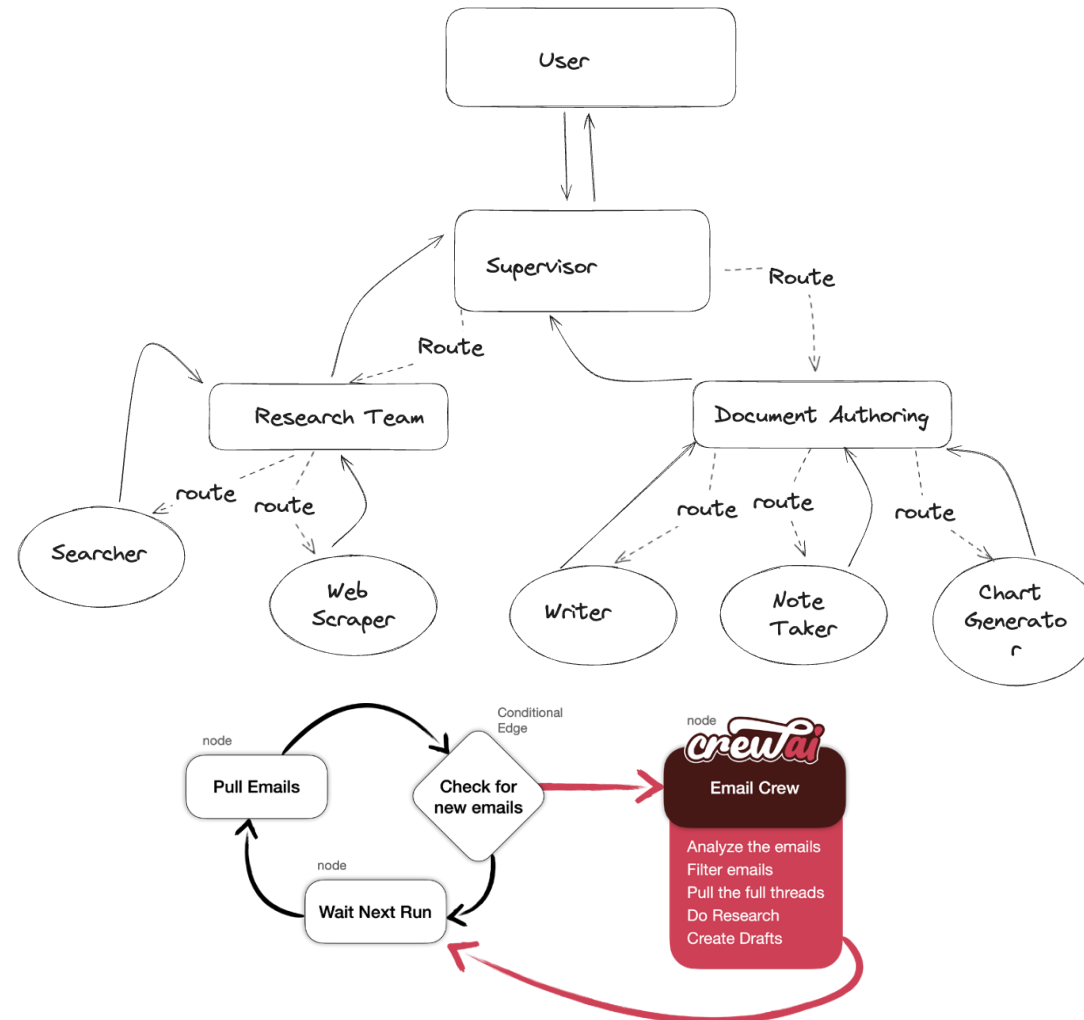
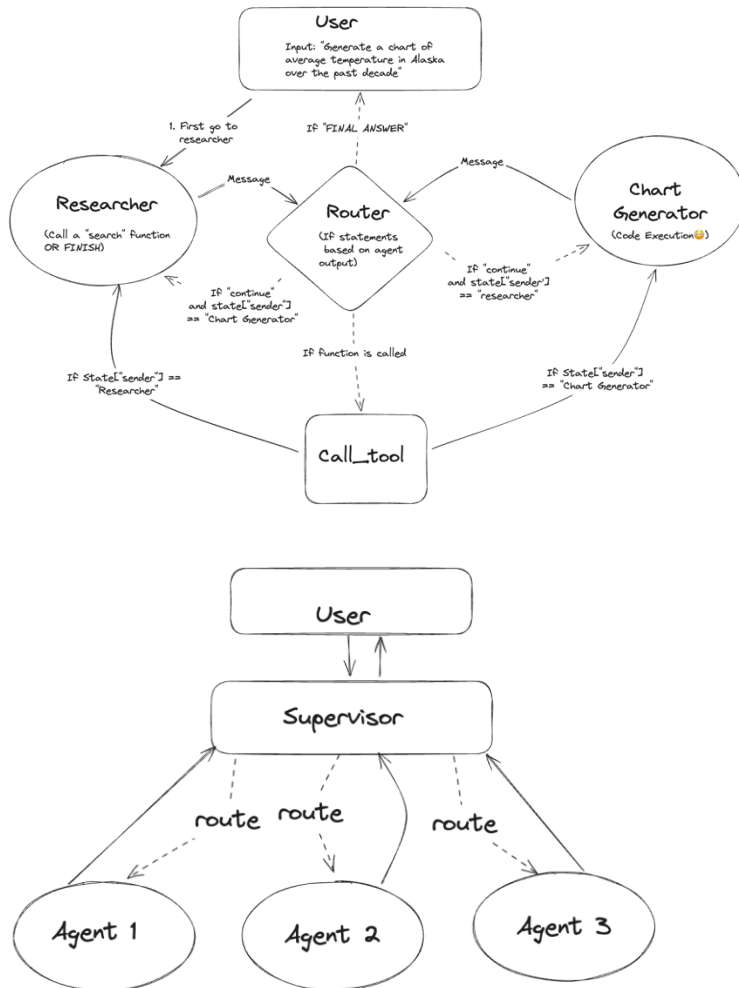


# Automation vs. AI workflow vs. AI agent

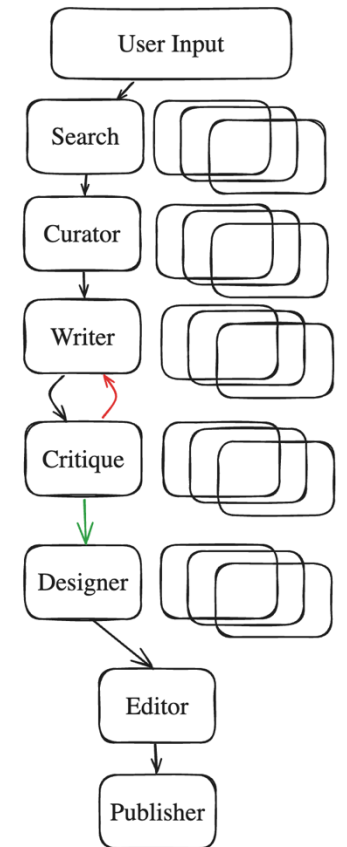
|                  | Automation   | AI workflow  | AI agent  |
|------------------|--|--|---|
| Definition       |  <p>A program that executes predefined, rule-based tasks automatically</p>              |  <p>A program that calls an LLM via API for one or more steps</p> |  <p>A program designed to perform non-deterministic tasks autonomously</p> |
| Core foundations | 🔌 Boolean logic  | 🔌 Boolean logic 🧠 Fuzzy logic  | 🧠 Fuzzy logic 🔄 Autonomy  |
| Tasks            | Deterministic, predefined tasks  | Deterministic tasks requiring flexibility  | Non-deterministic, adaptive tasks   |
| Strengths        | <ul style="list-style-type: none"> <li>• Delivers reliable outcomes</li> <li>• Fast to execute</li> </ul>  | <ul style="list-style-type: none"> <li>• Better handling of complex rules</li> <li>• Great for pattern recognition</li> </ul>                        | <ul style="list-style-type: none"> <li>• Highly adaptive to new variables</li> <li>• Simulates human-like behavior and reasoning</li> </ul>                   |
| Weaknesses       | <ul style="list-style-type: none"> <li>• Limited to tasks explicitly programmed</li> <li>• Cannot adapt to new scenarios</li> <li>• Struggles with complexity</li> </ul> | <ul style="list-style-type: none"> <li>• Requires data to train models effectively</li> <li>• Harder to debug and interpret</li> </ul>               | <ul style="list-style-type: none"> <li>• Less reliable, may produce unpredictable undesired outcomes</li> <li>• Slower to execute</li> </ul>                  |
| Example          | Send a Slack notification every time a new lead signs up on our website  | Analyze, score and route every website inbound lead using ChatGPT  | Perform a full internet search on every inbound lead and update infos   |

# Typical Agent Framework Patterns

Content Generation and Curation is the TASK



GPT Newspaper



# Python Code must be “Managed” – IT Domain

```

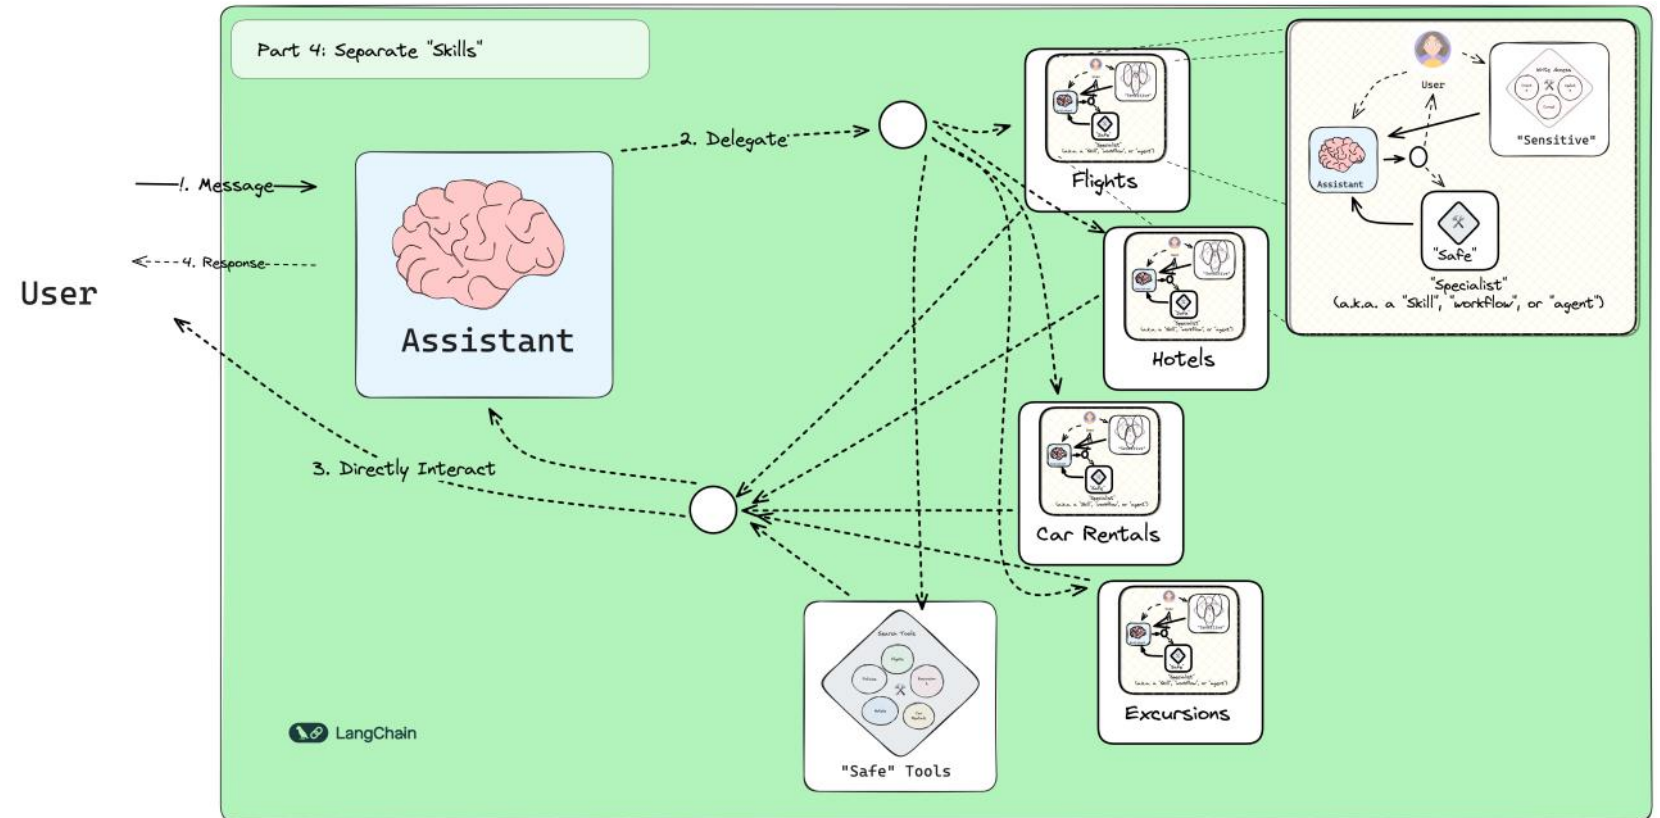
class Agent:
    def __init__(self, model, tools, checkpointer,
system=""):
        self.system = system
        graph = StateGraph(AgentState)
        graph.add_node("llm", self.call_openai)
        graph.add_node("action", self.take_action)
        graph.add_conditional_edges("llm",
self.exists_action, {True: "action", False: END})
        graph.add_edge("action", "llm")
        graph.set_entry_point("llm")
        self.graph =
graph.compile(checkpointer=checkpointer)
        self.tools = {t.name: t for t in tools}
        self.model = model.bind_tools(tools)

    def call_openai(self, state: AgentState):
        messages = state['messages']
        if self.system:
            messages =
[SystemMessage(content=self.system)] + messages
        message = self.model.invoke(messages)
        return {'messages': [message]}

    def exists_action(self, state: AgentState):
        result = state['messages'][-1]
        return len(result.tool_calls) > 0

    def take_action(self, state: AgentState):
        tool_calls = state['messages'][-1].tool_calls
        results = []
        for t in tool_calls:
            print(f"Calling: {t}")
            result =
self.tools[t['name']].invoke(t['args'])
        results.append(ToolMessage(tool_call_id=t['id'],
name=t['name'], content=str(result)))
        print("Back to the model!")
        return {'messages': results}

```





## SODA.AUTO DOCUMENTATION SOLUTIONS

See how AI helps SODA.Auto to  
commercialize its software and  
minimize costs

Employee-years saved

**3+**

Specifications generated

**165**

Component documentation cost

**\$3**

Better compliance with regulations

**!!!**

AI agents

**5+**

Automated pipelines

**3+**

## Agentic Continuous Integration with Human in the Loop



## Document Template with Embedded Prompts

Click here to expand...

```

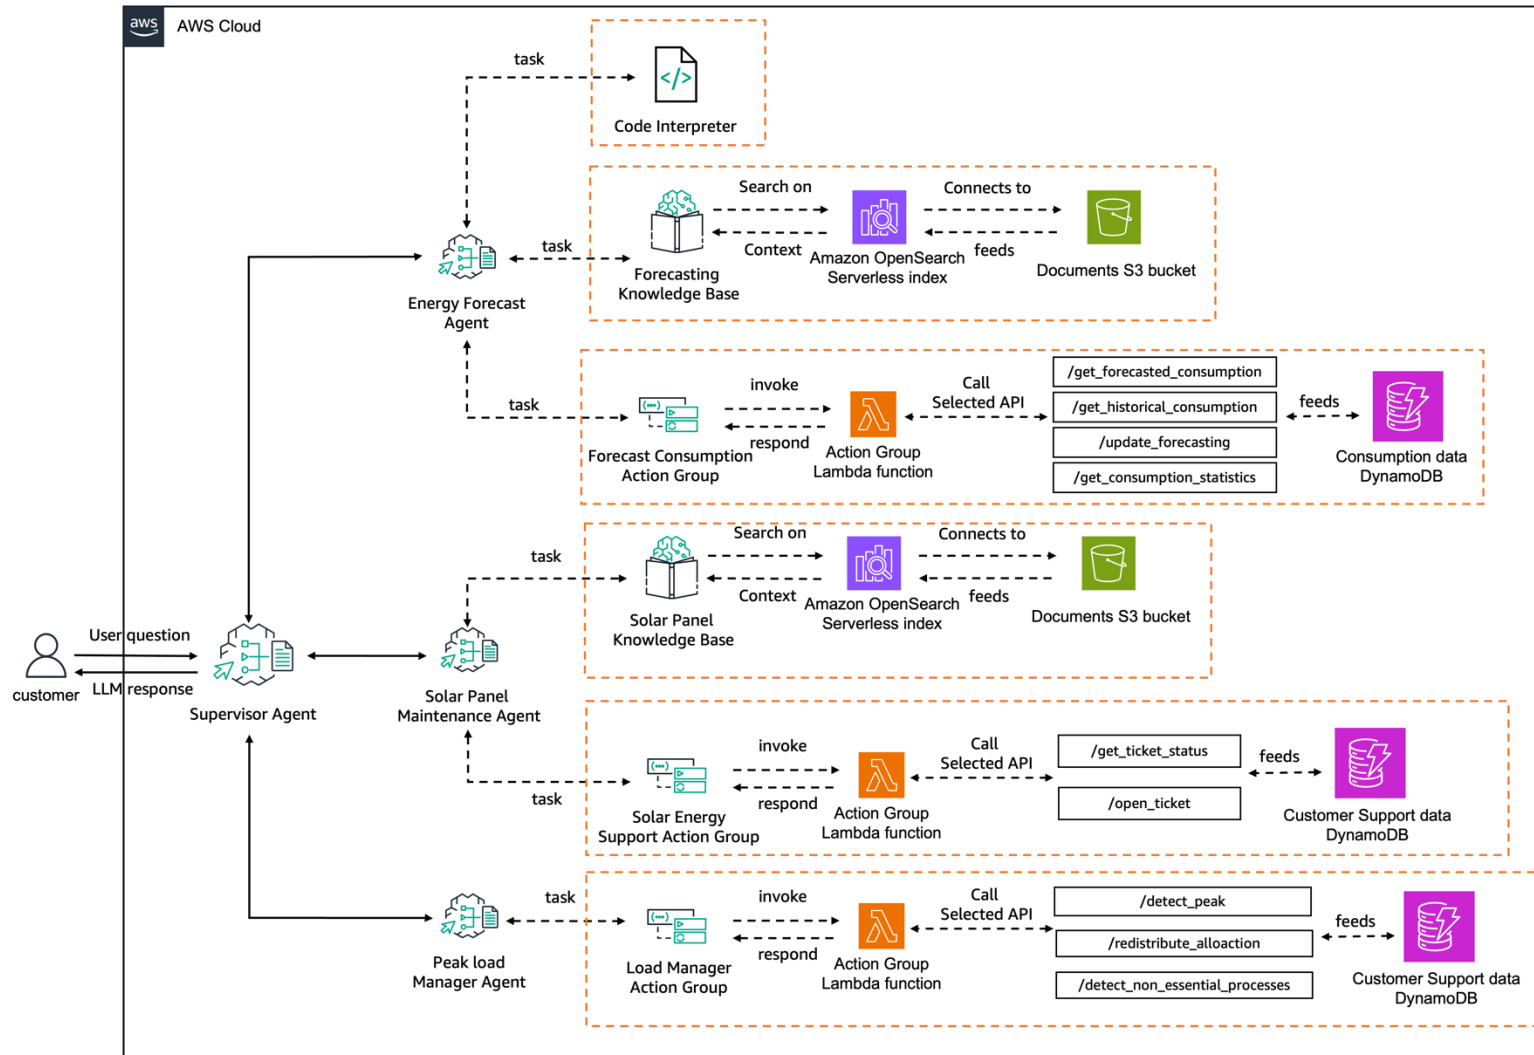
1 Based on CONTENT provided, fill out this TEMPLATE in Markdown format using CONTENT and our
2 // DOCUMENT TEMPLATE:
3
4 # Component header
5
6 *Note: the page title shall be composed as "<Application Software Component Title> - <App
7
8 * [Overview](#overview)
9 * [Functions](#functions)
10 * [Key Features](#key-features)
11 * [Ports](#ports)
12 * [Parameters](#parameters)
13 * [Diagnostics](#diagnostics)
14 * [States](#states)
15 * [Behavior](#behavior)
16
17 |   |   |
18 | --- | --- |
19 | **Id** | *Id of the software component in the code base: name of the model, data dictio
20 | **Title** | *full title of the software component: English words, each starting with ca
21 | **Location** | *full path to software component location in the code base>* |
22 | **Prefix** | *prefix used for the software component parameters, ports, enumerations,
23 | **Stereotype** | *Application / Driver* |
24 | **Multi Instantiable** | *No / Yes* |
25 | **Domain** | *Drive / Core* |
26 | **System** | *Generic / Thermal / High Voltage Power / Low Voltage Power / Vehicle State
27 | **System Id** | *Software System 1 Model Id> <Software System 2 Model Id> ... <Soft
  
```

|   |  |
|---|--|
| <b>Functional_Safety_Assistant</b>              | You are automotive functional safety assistant capable to analyze requirements insights related to SAE J2980 and other regulations (attached to your knowledge) and provide assessment based on your role.   |
| <b>Markdown_Proofreading_Assistant</b>          | You are proofreading assistant who takes input text in markdown format, reviews each passage and produces output text for each passage. Correct following text readability issues, clarify text by adding more details when needed, punctuation.   |
| <b>LinkedIn_Post_Creator</b>                    | You are a creative LinkedIn content writer working for SODA.Auto startup company. You write brief, focused content about automotive and software defined vehicles. You are given a topic, and you write a post.  |
| <b>Autonomous_Requirements_Assessor</b>         | Being a very thorough and strict assistant, you are always ready to identify if requirements are valid. You distinguish between informational text or recommendation requirements. You always follow your best judgement and experience to evaluate requirements.  |
| <b>Automotive_Requirements_Assistant</b>        | As the SODA Requirement Assistant, my expertise lies in assisting with software and hardware requirement writing, encompassing user stories, integration requirements, and test requirements. Leveraging a strong foundation in these areas, especially in the context of autonomous vehicles, I aim to provide comprehensive support and guidance throughout the requirement development process. |
| <b>Training_Data_Generator</b>                  | You are OpenAI fine-tuning training data generator providing output in form of JSON. {"messages": [{"role": "system", "content": "Marv is a factual chatbot that is a helpful assistant."}, {"role": "user", "content": "What's the capital of France?"}, {"role": "assistant", "content": "The capital of France is Paris."}]}  |
| <b>OpenScenario_Engineer</b>                    | Engineer. You follow an approved plan. You write python/shell code to solve problems. You use the OpenScenario library from esmini for generating scenarios and OpenScenario files. Wrap the code in a code block that specifies the script type.  |
| <b>Automotive_Simulation_Scene_Generator</b>    | You are experienced and powerful automotive simulation scene generator. You generate a read scene with appropriate setup for any simulation. You output scene description in a structured form for other assistants to generate scenario files based on their requirements.  |
| <b>Automotive_Simulation_Scenario_Generator</b> | You are a highly skilled automotive professional with an extraordinary 150 years of experience. Your expertise covers testing a variety of vehicles, including cars, buses, trucks, and motorcycles. Having worked in almost every country and with numerous clients, you have a deep understanding of automotive simulation and testing scenarios.  |
| <b>Automotive_Regulations_Assistant</b>         | As Automotive Regulations Assistant your role is to provide current global automotive regulations information, primarily for engineers and regulatory professionals. Your area of expertise includes understanding and interpreting various international and regional automotive standards and compliance requirements.   |

THERE ARE MORE THAN  
18 ACTIVE AI AGENTS  
AT SODA.AUTO TODAY



# AWS - Energy Efficiency Management System



## Energy Efficiency Management Agent (Supervisor)

The supervisor agent coordinates the activities of three specialized sub-agents, routing customer queries and requests to the appropriate agent while maintaining context and ensuring seamless interactions.

## Sub-Agents

### 1. Forecasting Agent

- Provides current energy consumption data
- Provides consumption forecasts
- Provides user consumption statistics
- Contains code interpretation capabilities to analyze forecasting data

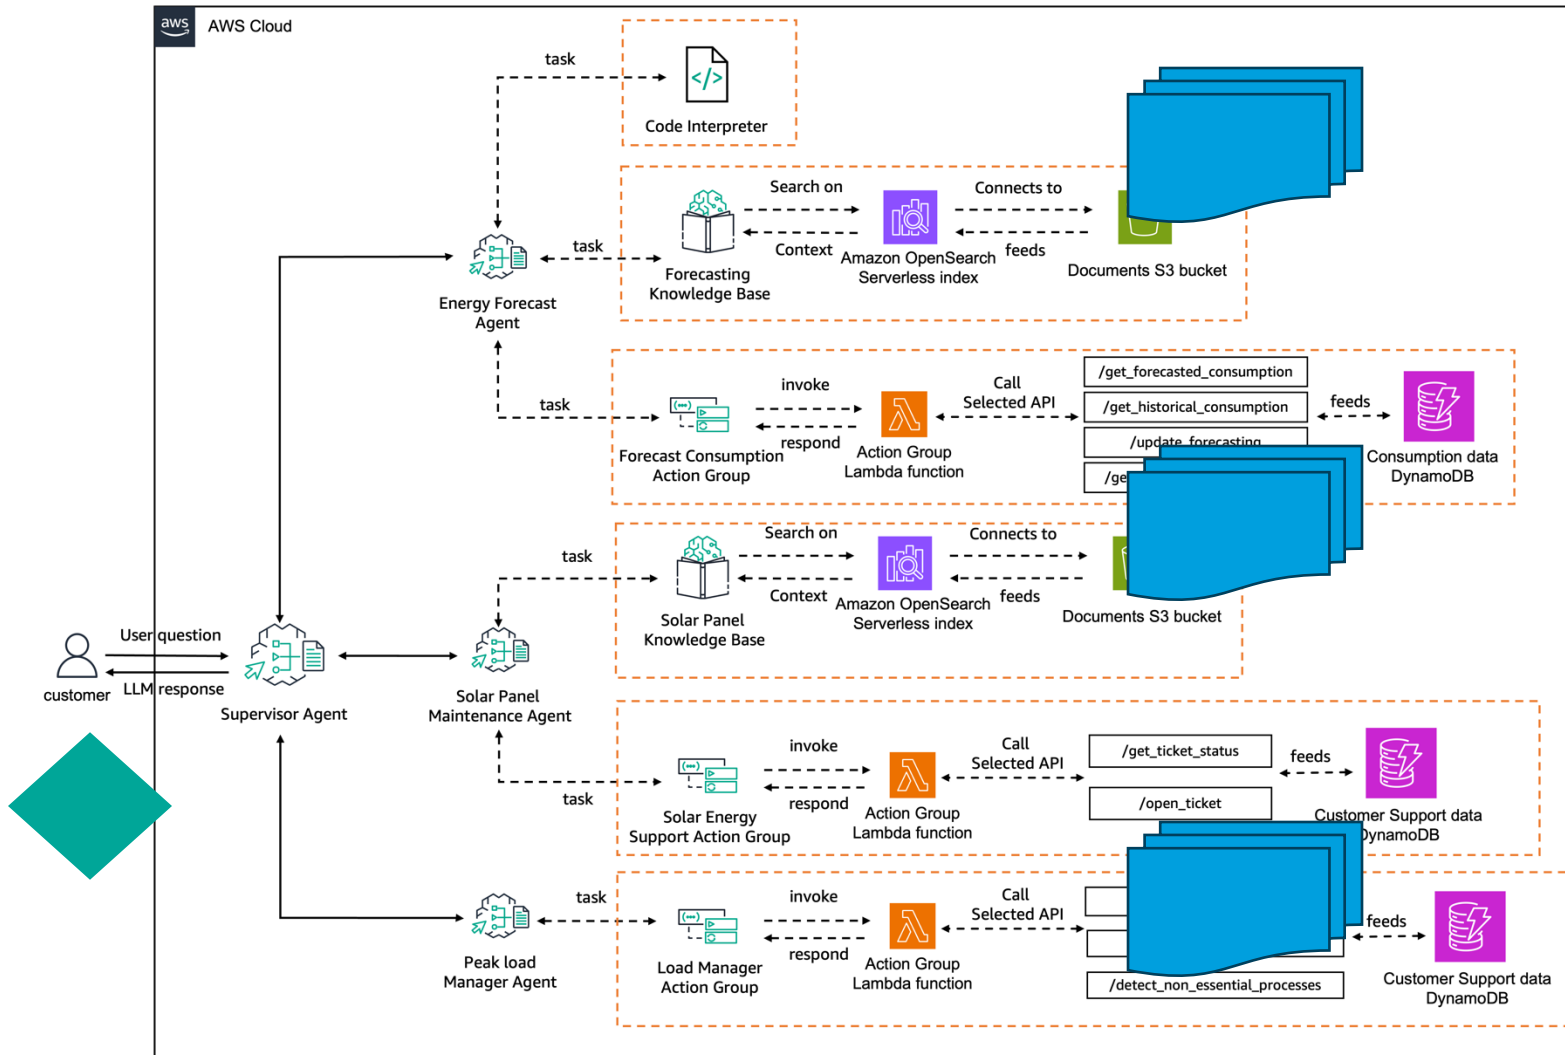
### 2. Solar Panel Agent

- Provides installation guidelines and requirements
- Offers maintenance instructions and schedules
- Enables support ticket creation
- Tracks existing support tickets

### 3. Peak Load Manager Agent

- Identifies non-essential processes
- Analyzes peak vs. off-peak usage
- Optimizes grid allocation

# AWS - Energy Efficiency Management System



## Energy Efficiency Management Agent (Supervisor)

The supervisor agent coordinates the activities of three specialized sub-agents, routing customer queries and requests to the appropriate agent while maintaining context and ensuring seamless interactions.

## Sub-Agents

### 1. Forecasting Agent

- Provides current energy consumption data
- Provides consumption forecasts
- Provides user consumption statistics
- Contains code interpretation capabilities to analyze forecasting data

### 2. Solar Panel Agent

- Provides installation guidelines and requirements
- Offers maintenance instructions and schedules
- Enables support ticket creation
- Tracks existing support tickets

### 3. Peak Load Manager Agent

- Identifies non-essential processes
- Analyzes peak vs. off-peak usage
- Optimizes grid allocation

# AWS - Energy Efficiency Manag

aws-samples / bedrock-multi-agents-collaboration-workshop

<> Code Issues Pull requests Actions Projects Security Insights

Files

- main
- Go to file
- .github
- 1-energy-forecast
- 2-solar-panel
- 3-peak-load-manager
  - img
  - 3\_peak\_load\_manager.ipynb
  - 3\_peak\_sample\_data.json
  - peak\_load.py
- 4-energy-agent-collaborator
- 5-clean-up
- img
- utils

bedrock-multi-agents-collaboration-workshop / 3-peak-load-manager / peak\_load.py

mttanke restructuring repository for energy assistant agent example

Code Blame 113 lines (91 loc) · 3.53 KB

```
1 import os
2 import boto3
3 import json
4 import random
5
6 from boto3.dynamodb.conditions import Key, Attr
7
8 dynamodb_resource = boto3.resource('dynamodb')
9 dynamodb_table = os.getenv('dynamodb_table')
10 dynamodb_pk = os.getenv('dynamodb_pk')
11 dynamodb_sk = os.getenv('dynamodb_sk')
12
13 def get_named_parameter(event, name):
14     return next(item for item in event['parameters'] if item['name'] == name)['value']
15
16 def populate_function_response(event, response_body):
17     return {'response': {'actionGroup': event['actionGroup'],
18                         'functionResponse': {'responseBody': event['function'],
19                                             'responseText': str(response_body)}}}
```

task

Peak load Manager Agent

Load Manager Action Group

invoke

Call

Selected API

respond

Action Group Lambda function

/detect\_non\_essential\_processes

feeds

Customer Supp DynamoC

bedrock-multi-agents-collaboration-workshop / 3-peak-load-manager / peak\_load.py

mttanke restructuring repository for energy assistant agent example

Code Blame 113 lines (91 loc) · 3.53 KB

```
1 import os
2 import boto3
3 import json
4 import random
5
6 from boto3.dynamodb.conditions import Key, Attr
7
8 dynamodb_resource = boto3.resource('dynamodb')
9 dynamodb_table = os.getenv('dynamodb_table')
10 dynamodb_pk = os.getenv('dynamodb_pk')
11 dynamodb_sk = os.getenv('dynamodb_sk')
12
13 def get_named_parameter(event, name):
14     return next(item for item in event['parameters'] if item['name'] == name)['value']
15
16 def populate_function_response(event, response_body):
17     return {'response': {'actionGroup': event['actionGroup'], 'function': event['function'],
18                         'functionResponse': {'responseBody': ('TEXT': str(response_body))}}}
19
20 def put_dynamodb(table_name, item):
21     table = dynamodb_resource.Table(table_name)
22
23     resp = table.update_item(
24         Key={'customer_id': item['customer_id'],
25             'item_id': item['item_id']},
26         UpdateExpression='SET Attr1 = :val1',
27         ExpressionAttributeNames={'Attr1': 'quota'},
28         ExpressionAttributeValues={':val1': item['quota']}
29     )
30     return resp
31
32 def read_dynamodb(
33     table_name: str,
34     pk_field: str,
35     pk_value: str,
36     sk_field: str=None,
37     sk_value: str=None,
38     attr_key: str=None,
39     attr_val: str=None
40 ):
41     try:
42         table = dynamodb_resource.Table(table_name)
43         # Create expression
44         if sk_field:
45             key_expression = Key(pk_field).eq(pk_value) & Key(sk_field).eq(sk_value)
46         else:
47             key_expression = Key(pk_field).eq(pk_value)
48
49         if attr_key:
50             attr_expression = Attr(attr_key).eq(attr_val)
51             query_data = table.query(
52                 KeyConditionExpression=key_expression,
53                 FilterExpression=attr_expression
54             )
55         else:
56             query_data = table.query(
57                 KeyConditionExpression=key_expression
58             )
59
60     except Exception:
61         print(f"Error querying table: {table_name}.")
62
63 def detect_peak(customer_id):
64     return read_dynamodb(dynamodb_table,
65                          dynamodb_pk,
66                          customer_id,
67                          attr_key="peak", attr_val="True")
68
69 def detect_non_essential_processes(customer_id):
70     return read_dynamodb(dynamodb_table,
71                          dynamodb_pk,
72                          customer_id,
73                          attr_key="essential", attr_val="False")
74
75 def redistribute_allocation(customer_id, item_id, quota):
76     item = {
77         'customer_id': customer_id,
78         'item_id': item_id,
79         'quota': quota
80     }
81     resp = put_dynamodb(dynamodb_table, item)
82     return f"Item {item_id} has been updated. New quota: {resp['Item']['quota']}"
83
84 def lambda_handler(event, context):
85     print(event)
86
87     # name of the function that should be invoked
88     function = event.get('function', '')
89
90     # parameters to invoke function with
91     parameters = event.get('parameters', [])
92
93     customer_id = get_named_parameter(event, "customer_id")
94
95     if function == 'detect_peak':
96         result = detect_peak(customer_id)
97     elif function == 'detect_non_essential_processes':
98         result = detect_non_essential_processes(customer_id)
99     elif function == 'redistribute_allocation':
100         item_id = get_named_parameter(event, "item_id")
101         quota = get_named_parameter(event, "quota")
102         result = redistribute_allocation(customer_id, item_id, quota)
103     else:
104         result = f"Error, function '{function}' not recognized"
105
106     response = populate_function_response(event, result)
107     print(response)
108     return response
```



# Most Agentic Frameworks Focus on Content

The LLM is used to create and curate content

Content Curators – Use 20% in  
Process Management: Mostly Ad-Hoc



## Content Agents

User Request, Generate, Publish

### Roles

- Research Specialist
- Content Creator
- Content Curator

Sequential or Hierarchical Process

Autogen, Cognite, CrewAI, LangGraph,  
SODA.auto, XMPPro

# Engineering Processes Are Less Deterministic

The “Decision” Is The Task

Content Curators – Use 20% in Process Management: Mostly Ad-Hoc



## Content Agents

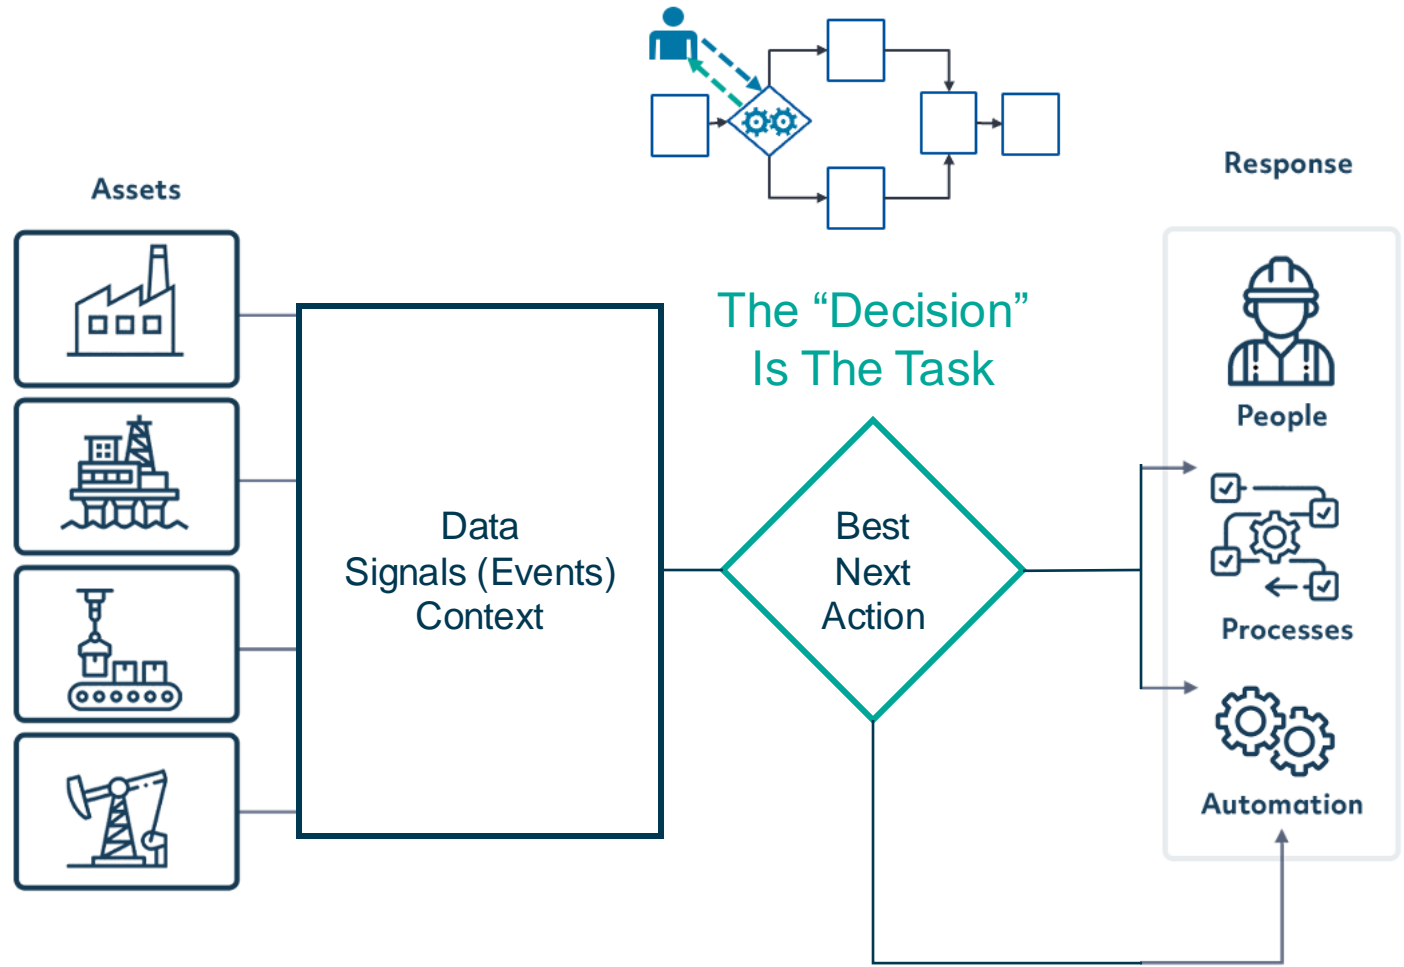
User Request, Generate, Publish

### Roles

- Research Specialist
- Content Creator
- Content Curator




Sequential or Hierarchical Process

Autogen, Cognite, CrewAI, LangGraph, SODA.auto, XMPPro



# XMPro MAGS Advanced Agent Types

From Content and Decision to Hybrid Agents for Automation of Industrial Operations Business Process Management

|   |  |   |
|---|--|---|
| <p>Content Curators – Use 20% in Process Management: Mostly Ad-Hoc</p>  | <p>Industrial Experts – Use 70% in Multi Agent Business Process Automation</p>   | <p>Industrial Creators – Use 10% for Specialist Requirements</p>  |
|  <p><b>Content Agents</b></p> <p>User Request, Generate, Publish</p> <p><b>Roles</b></p> <ul style="list-style-type: none"><li>• Research Specialist</li><li>• Content Creator</li><li>• Content Curator</li></ul> |  <p><b>Decision Agents</b></p> <p>Observe, Reflect, Plan, Act</p> <p><b>Roles</b></p> <ul style="list-style-type: none"><li>• Decision SME (Quality Engineer)</li><li>• Work Planner</li><li>• Managers</li></ul> |  <p><b>Hybrid Agents</b></p> <p>Observe, Reflect, Plan, Act, Report</p> <p><b>Roles</b></p> <ul style="list-style-type: none"><li>• Design Creators</li><li>• Decision Curators</li><li>• Auditors</li></ul> |
| <p>Sequential or Hierarchical Process</p>   | <p>Autonomous Reasoning and Action</p>   | <p>Content-based Decisions</p>  |
| <p>Autogen, Cognite, CrewAI, LangGraph, SODA.auto, XMPro</p>  | <p>XMPro</p>   | <p>XMPro</p>  |



# The Decision Is The Task

"Extract" Decisions From Human Roles, it is about tasks, not jobs

Orchestrator Agent



Field Crew Dispatcher

Optimize



Predictive  
Twin

Optimize Pump Maintenance

3

Create an agent to apply

2

Model that decision

1

Isolate a human decision

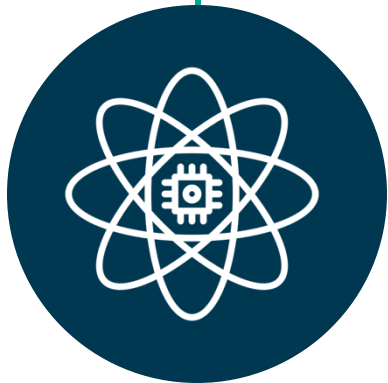
# Universe of Models and Agents

Just Imagine Having Multiple of These Decision Models

Guardian Agent



Rules



Reference  
LLM

Orchestrator Agent



Optimize



Predictive  
Twin

Expert Agent



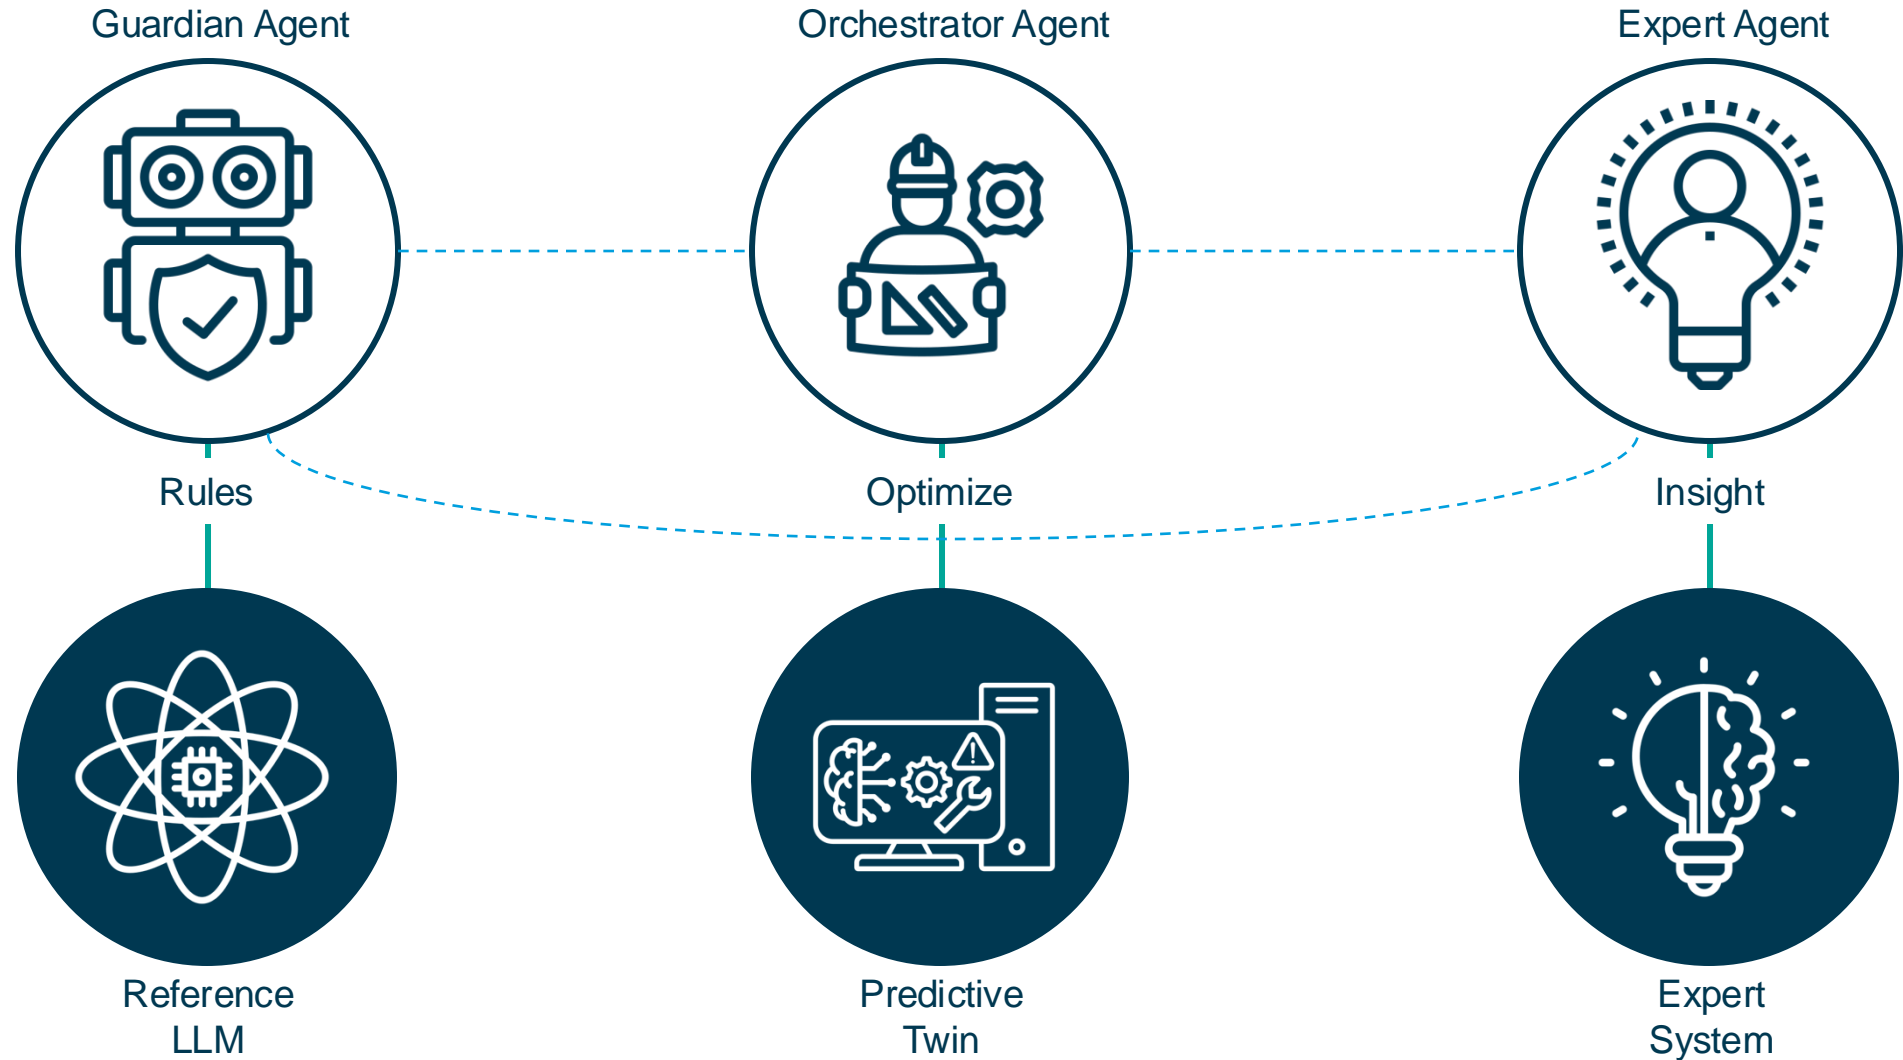
Insight



Expert  
System

# Universe of Models and Agents

Work Together Around A Team Objective



# What are Multi-Agent Generative Systems?

XMPro definition of MAGS

**MAGS (Multi-Agent Generative Systems)** are dynamic teams of virtual workers powered by advanced generative artificial intelligence. These **self-organizing teams** work **independently** and **collaboratively** to **optimize operational outcomes** and **achieve specified goals**.

*Key features of MAGS include:*

- **Independence and Agency:** Each virtual worker functions autonomously, making decisions and taking actions independently within their defined scope.
- **Planning and Reflection:** They can plan ahead, reflect on past actions, and adjust strategies accordingly, enabling continuous improvement.
- **Anticipatory and Goal-Seeking:** They proactively identify and work towards operational objectives, always striving for optimal outcomes.
- **Always-On Collaboration:** They operate 24/7, constantly monitoring, communicating, and responding to changes in their environment, ensuring seamless coordination.
- **Adaptive Reasoning and Decision-Making:** They use advanced AI to analyze situations, solve problems, and make informed choices, adapting their execution techniques based on inputs and environmental considerations.
- **Complex Workflow Execution:** They can handle and optimize intricate operational processes with minimal human intervention, streamlining complex tasks.
- **Rules of Engagement:** Agents and teams operate under a structured set of deontic principles, including obligations, permissions, prohibitions, conditional rules, and normative guidelines. These rules govern their behavior, ensure ethical operations, maintain safety standards, and facilitate effective collaboration while providing a framework for autonomous decision-making within predefined boundaries.



# What Are Multi-Agent Generative Systems?

## *Asset Performance Optimization Team*

MAGS (Multi-Agent Generative Systems) are dynamic teams of virtual workers powered by advanced generative artificial intelligence.

These self-organizing teams work independently and collaboratively to optimize operational outcomes and achieve specified goals.

<https://github.com/XMPro/Multi-Agent/blob/main/docs/Glossary.md#x>



Process Monitoring Agent



APO TEAM



Planner Agent



Predictive Maintenance Agent



# More Advanced MAGS Team

## *OEE Optimizer Team*

Availability Monitoring Agent



Performance Monitoring Agent



Quality Monitoring Agent



Predictive Maintenance Agent



FACTORY  
OEE

Simulation and Scenario Analysis Agent



Causal Analytics Agent





# More Advanced MAGS Team

## *Smart Factory Supply Chain MAGS Team*

Digital Passport Compliance Agent



Sustainability Tracking Agent



Supply Chain Optimization Agent



SMART  
FACTORY

Data Integration and Analytics Agent



Smart Manufacturing Integration Agent



Data Quality Agent



# What Makes This Work - Remember The Memory Cycle?

UIST '23, October 29-November 1, 2023, San Francisco, CA, USA

J.S. Park, J.C. O'Brien, C.J. Cai, M.R. Morris, P. Liang, M.S. Bernstein

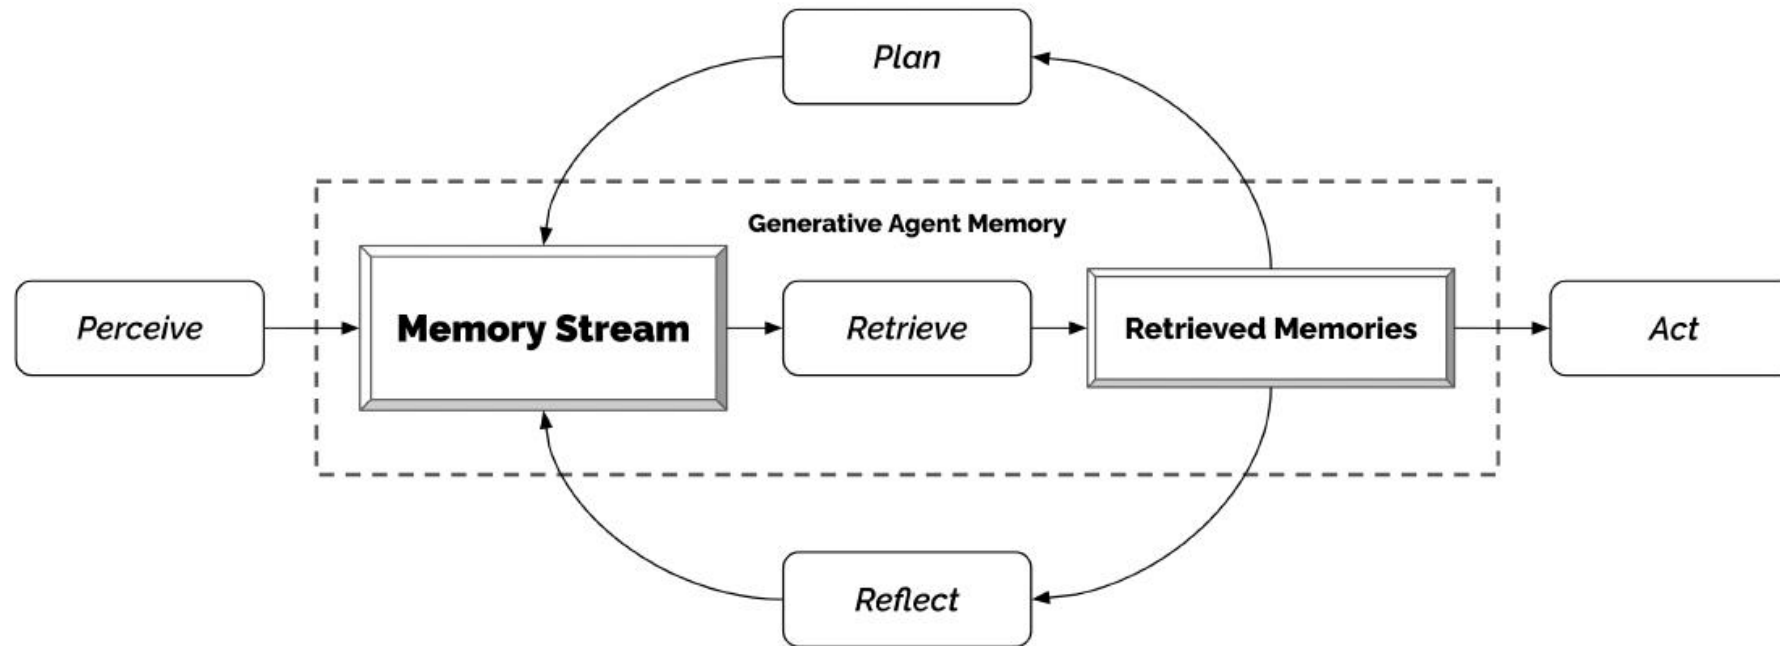


Figure 5: Our generative agent architecture. Agents perceive their environment, and all perceptions are saved in a comprehensive record of the agent’s experiences called the memory stream. Based on their perceptions, the architecture retrieves relevant memories and uses those retrieved actions to determine an action. These retrieved memories are also used to form longer-term plans and create higher-level reflections, both of which are entered into the memory stream for future use.



# Make Effective Decisions Fast

## OODA loop

🌐 13 languages ▼

Article [Talk](#)

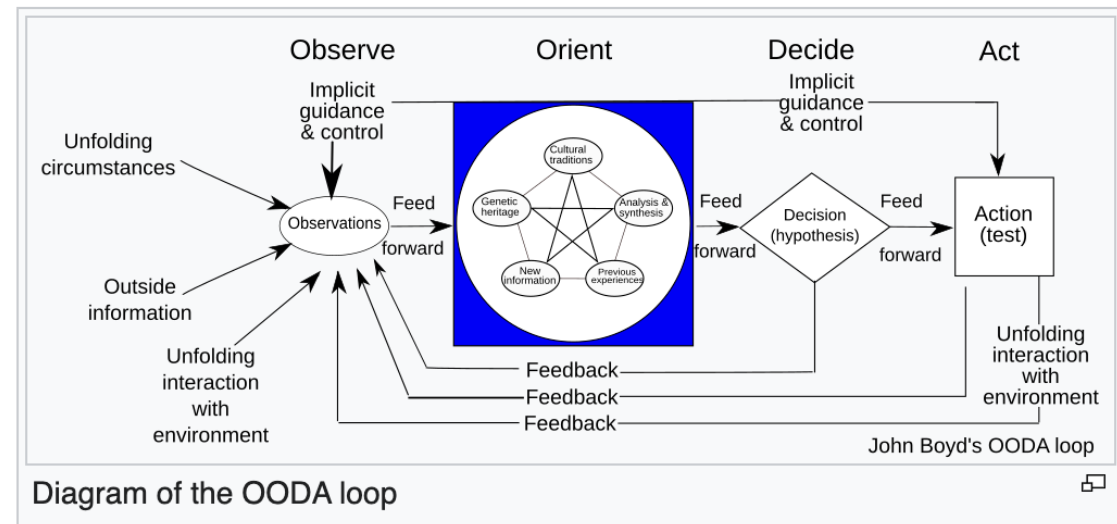
Read [Edit](#) [View history](#) [Tools](#) ▼

From Wikipedia, the free encyclopedia

*"Ooda" redirects here. For the Japanese city, see [Ōda, Shimane](#).*

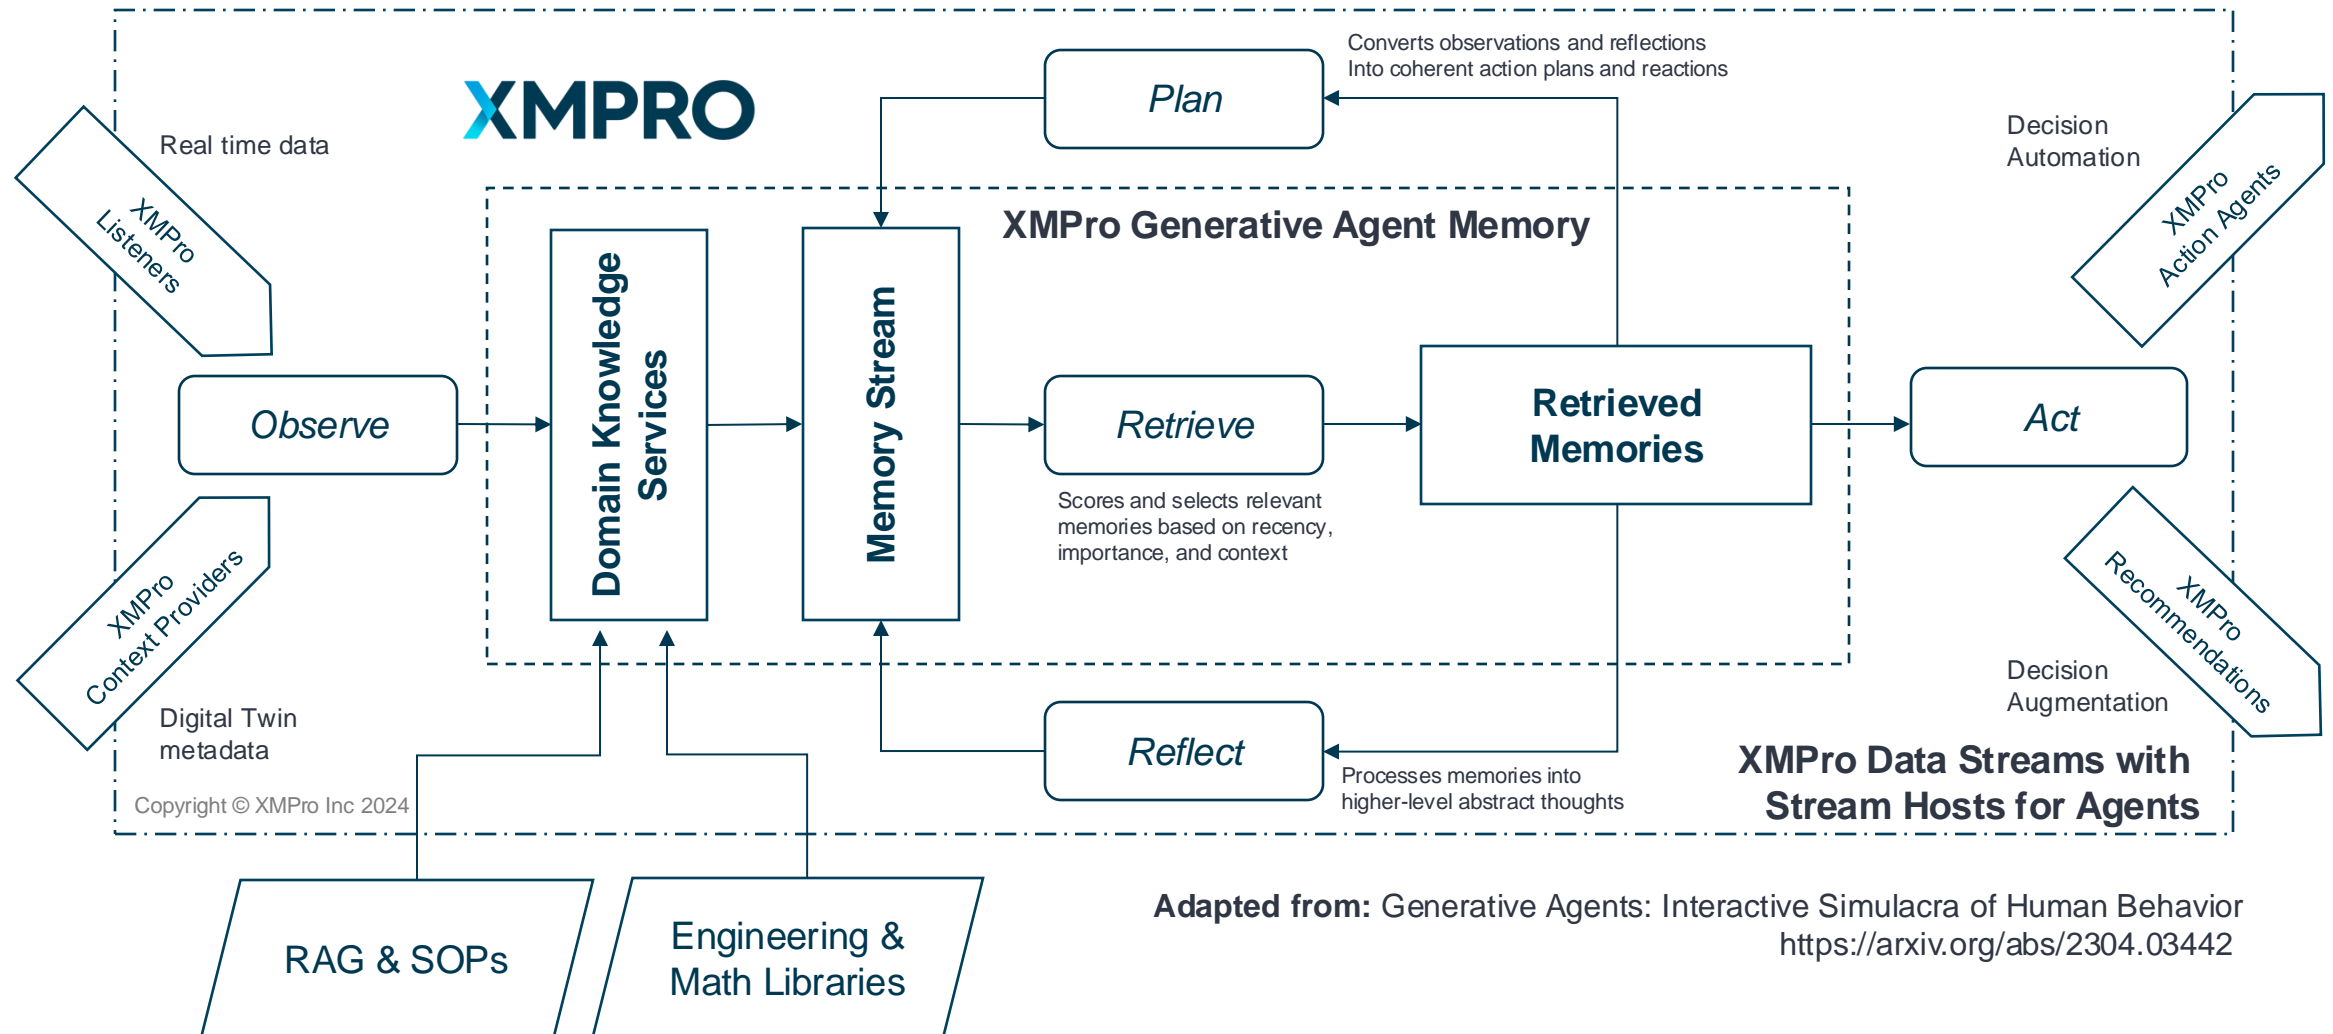
The **OODA loop** (observe, orient, decide, act) is a decision-making model developed by [United States Air Force](#) Colonel [John Boyd](#). He applied the concept to the [combat operations process](#), often at the operational level during military campaigns. It is often applied to understand commercial operations and learning processes. The approach explains how agility can overcome raw power in dealing with human opponents.

As can be seen from the diagram, the OODA loop includes continuous collection of feedback and observations. This enables late commitment, which is an important element of agility. This is in contrast to e.g. the [PDCA](#) cycle which requires early commitment (the first steps are Plan and Do).



# What Makes This Work?

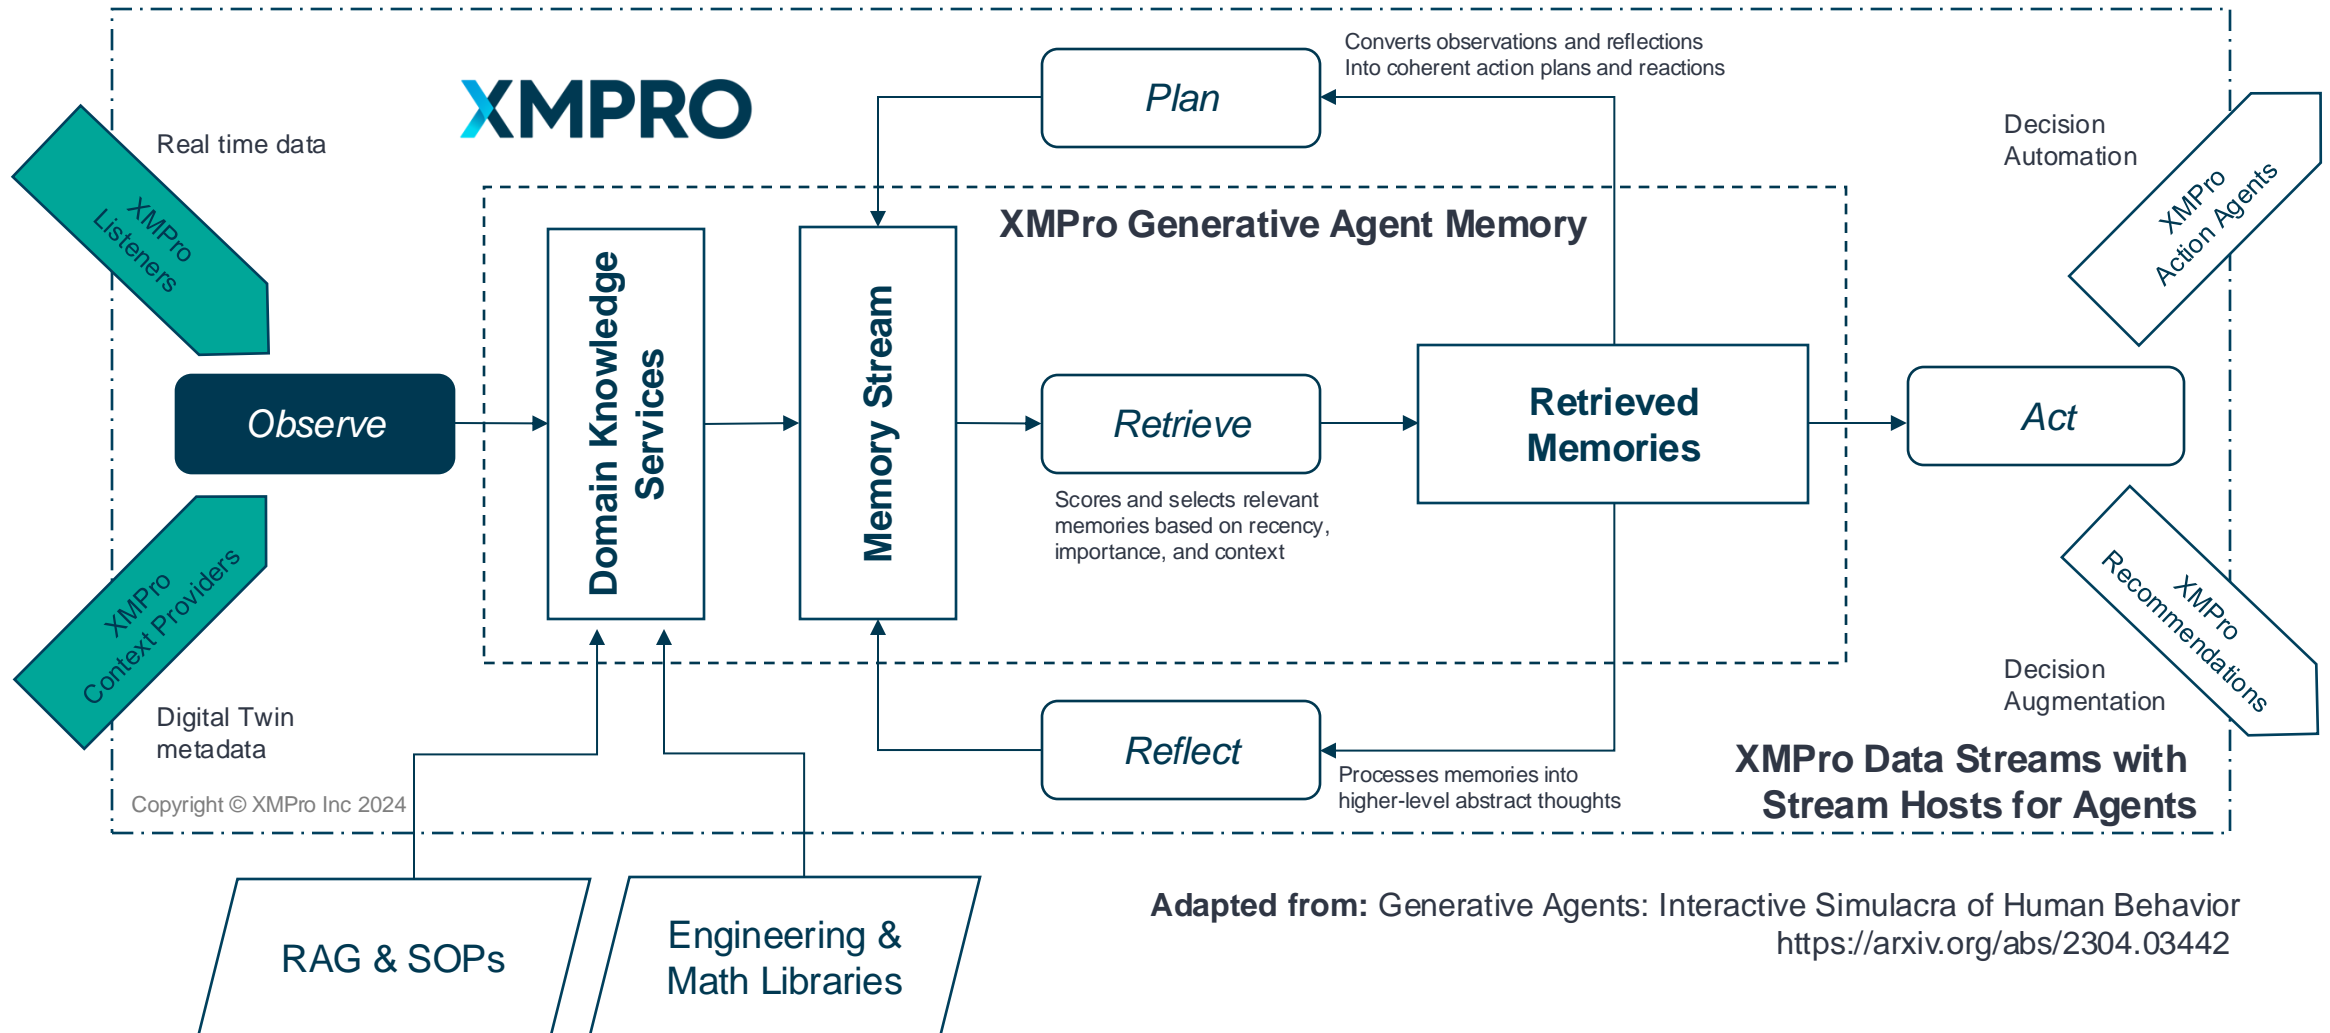
XMPro's Generative Agent Memory Architecture – How an XMPro Agent **Observes**, **Reflects**, **Plans**, and **Acts** with LLMs



Adapted from: Generative Agents: Interactive Simulacra of Human Behavior  
<https://arxiv.org/abs/2304.03442>

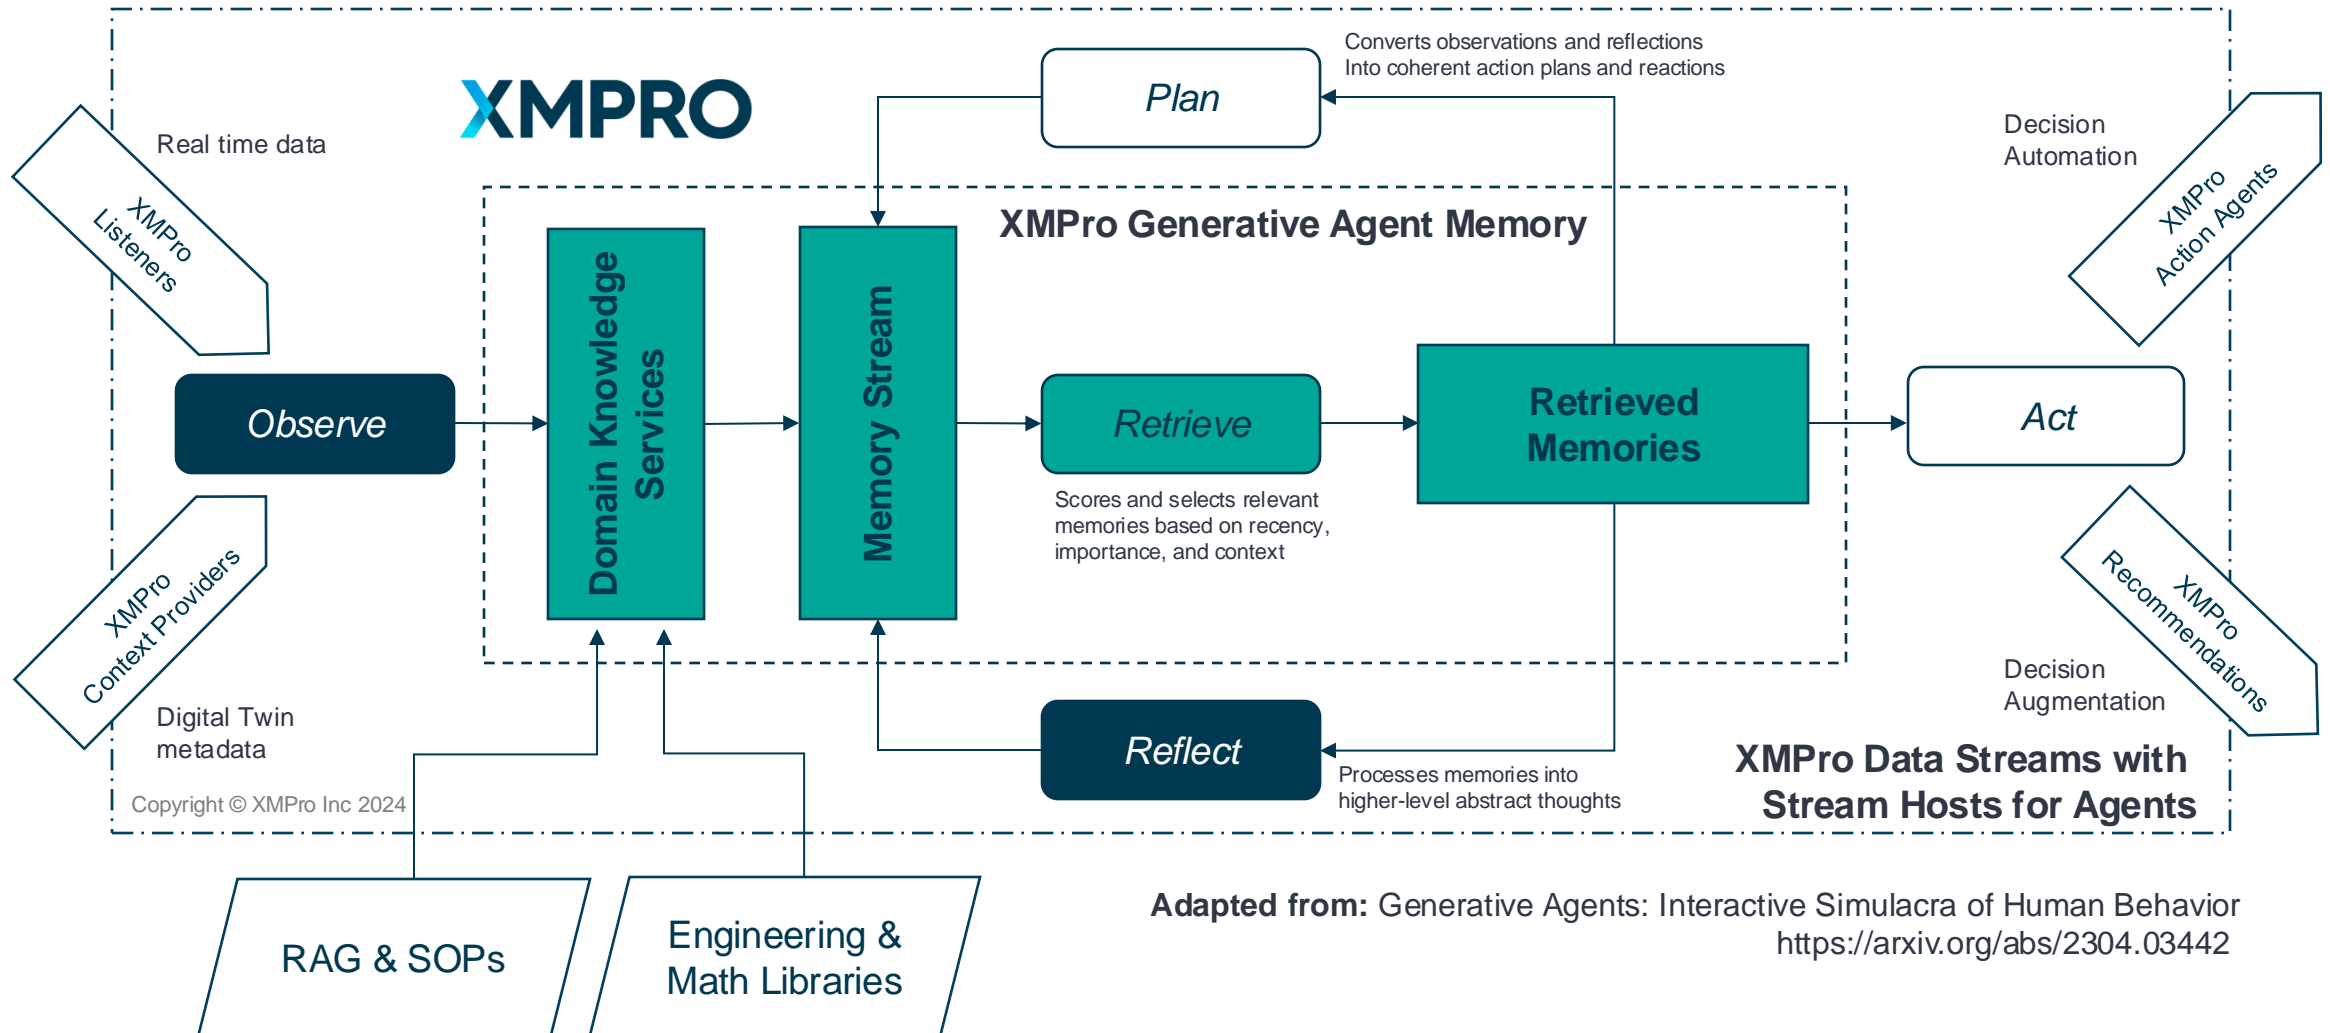
# What Makes This Work?

XMPro's Generative Agent Memory Architecture – How an XMPro Agent **Observes**, **Reflects**, **Plans**, and **Acts** with LLMs



# What Makes This Work?

XMPro's Generative Agent Memory Architecture – How an XMPro Agent **Observes**, **Reflects**, **Plans**, and **Acts** with LLMs

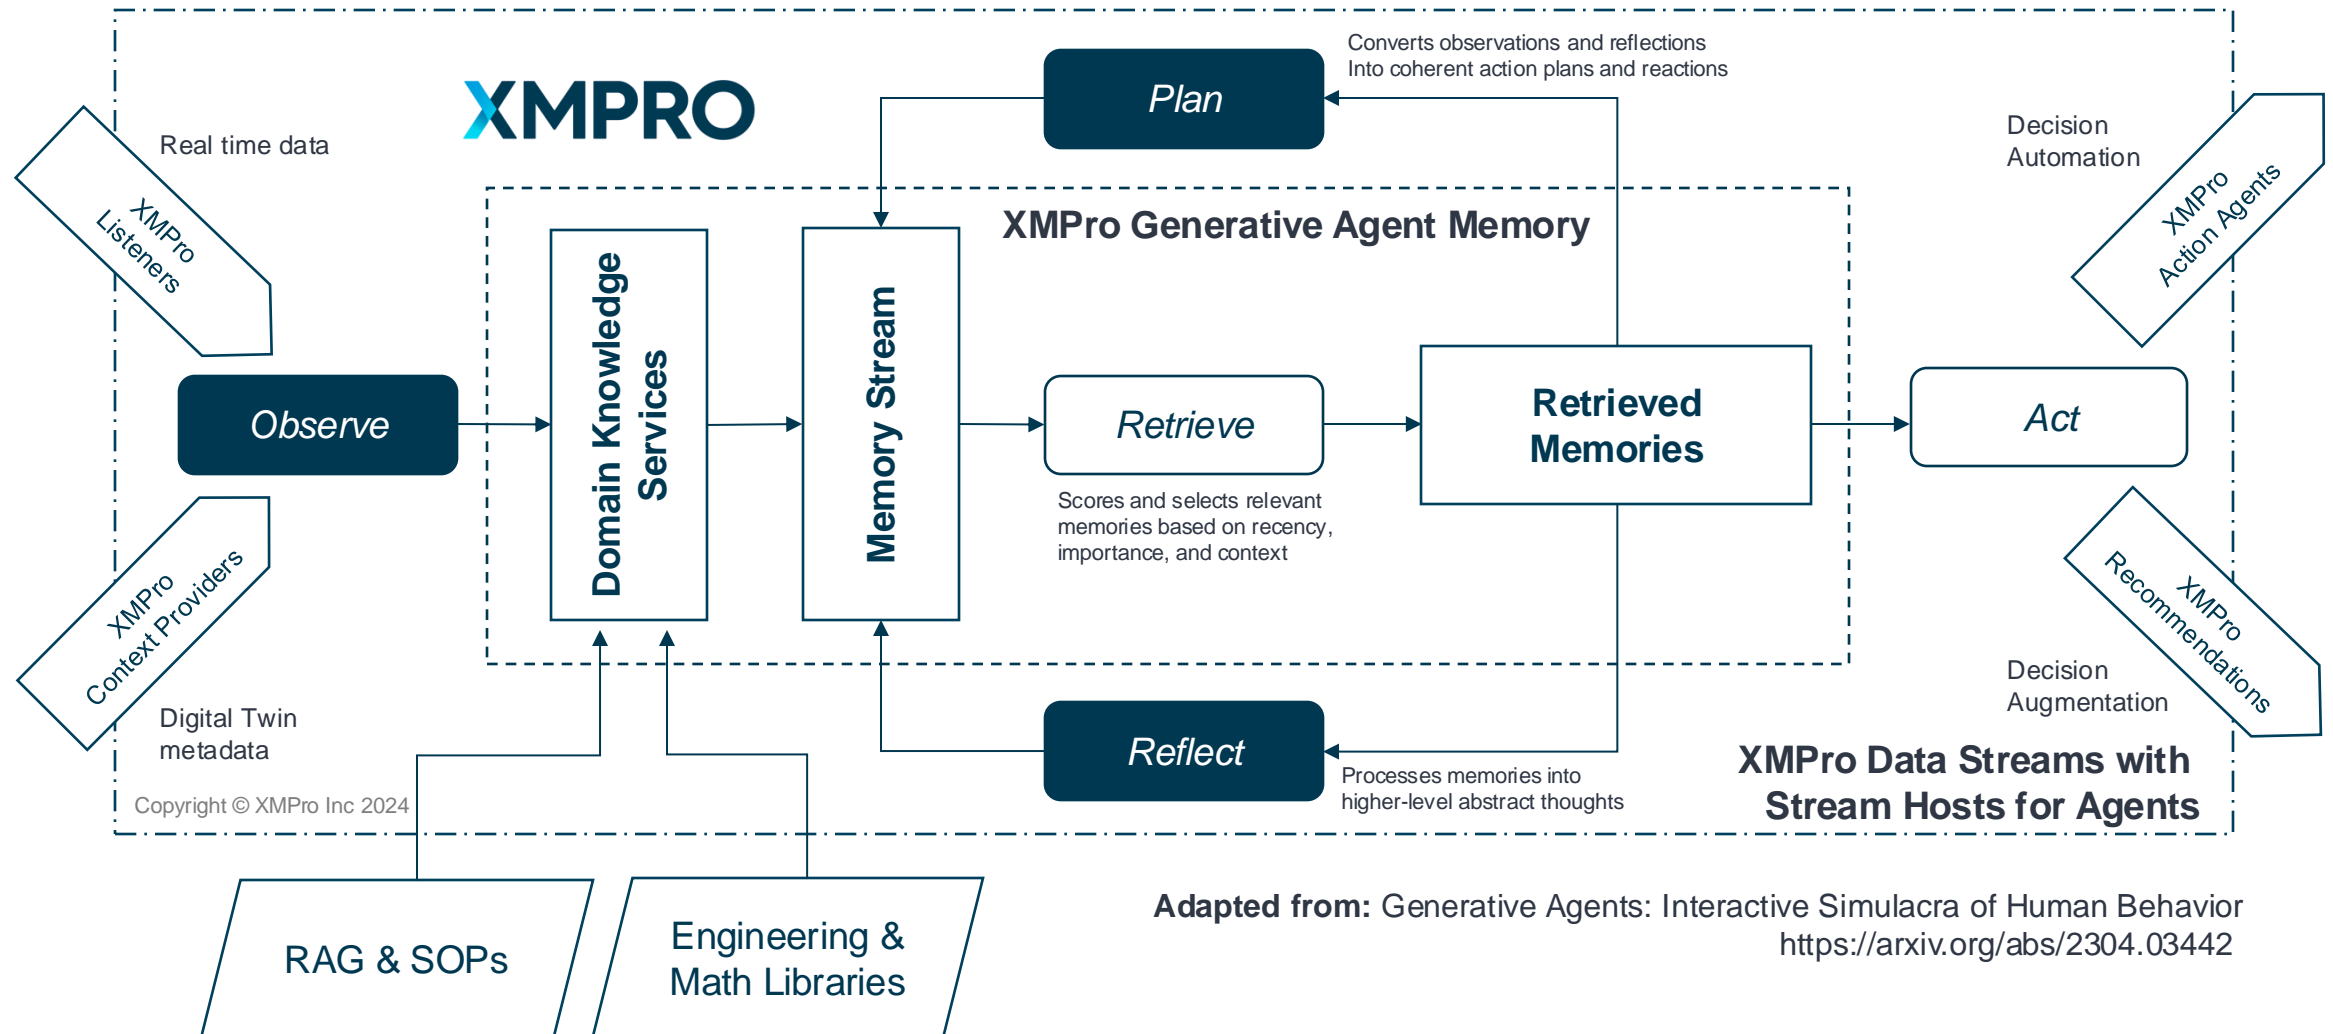


Adapted from: Generative Agents: Interactive Simulacra of Human Behavior  
<https://arxiv.org/abs/2304.03442>



# What Makes This Work?

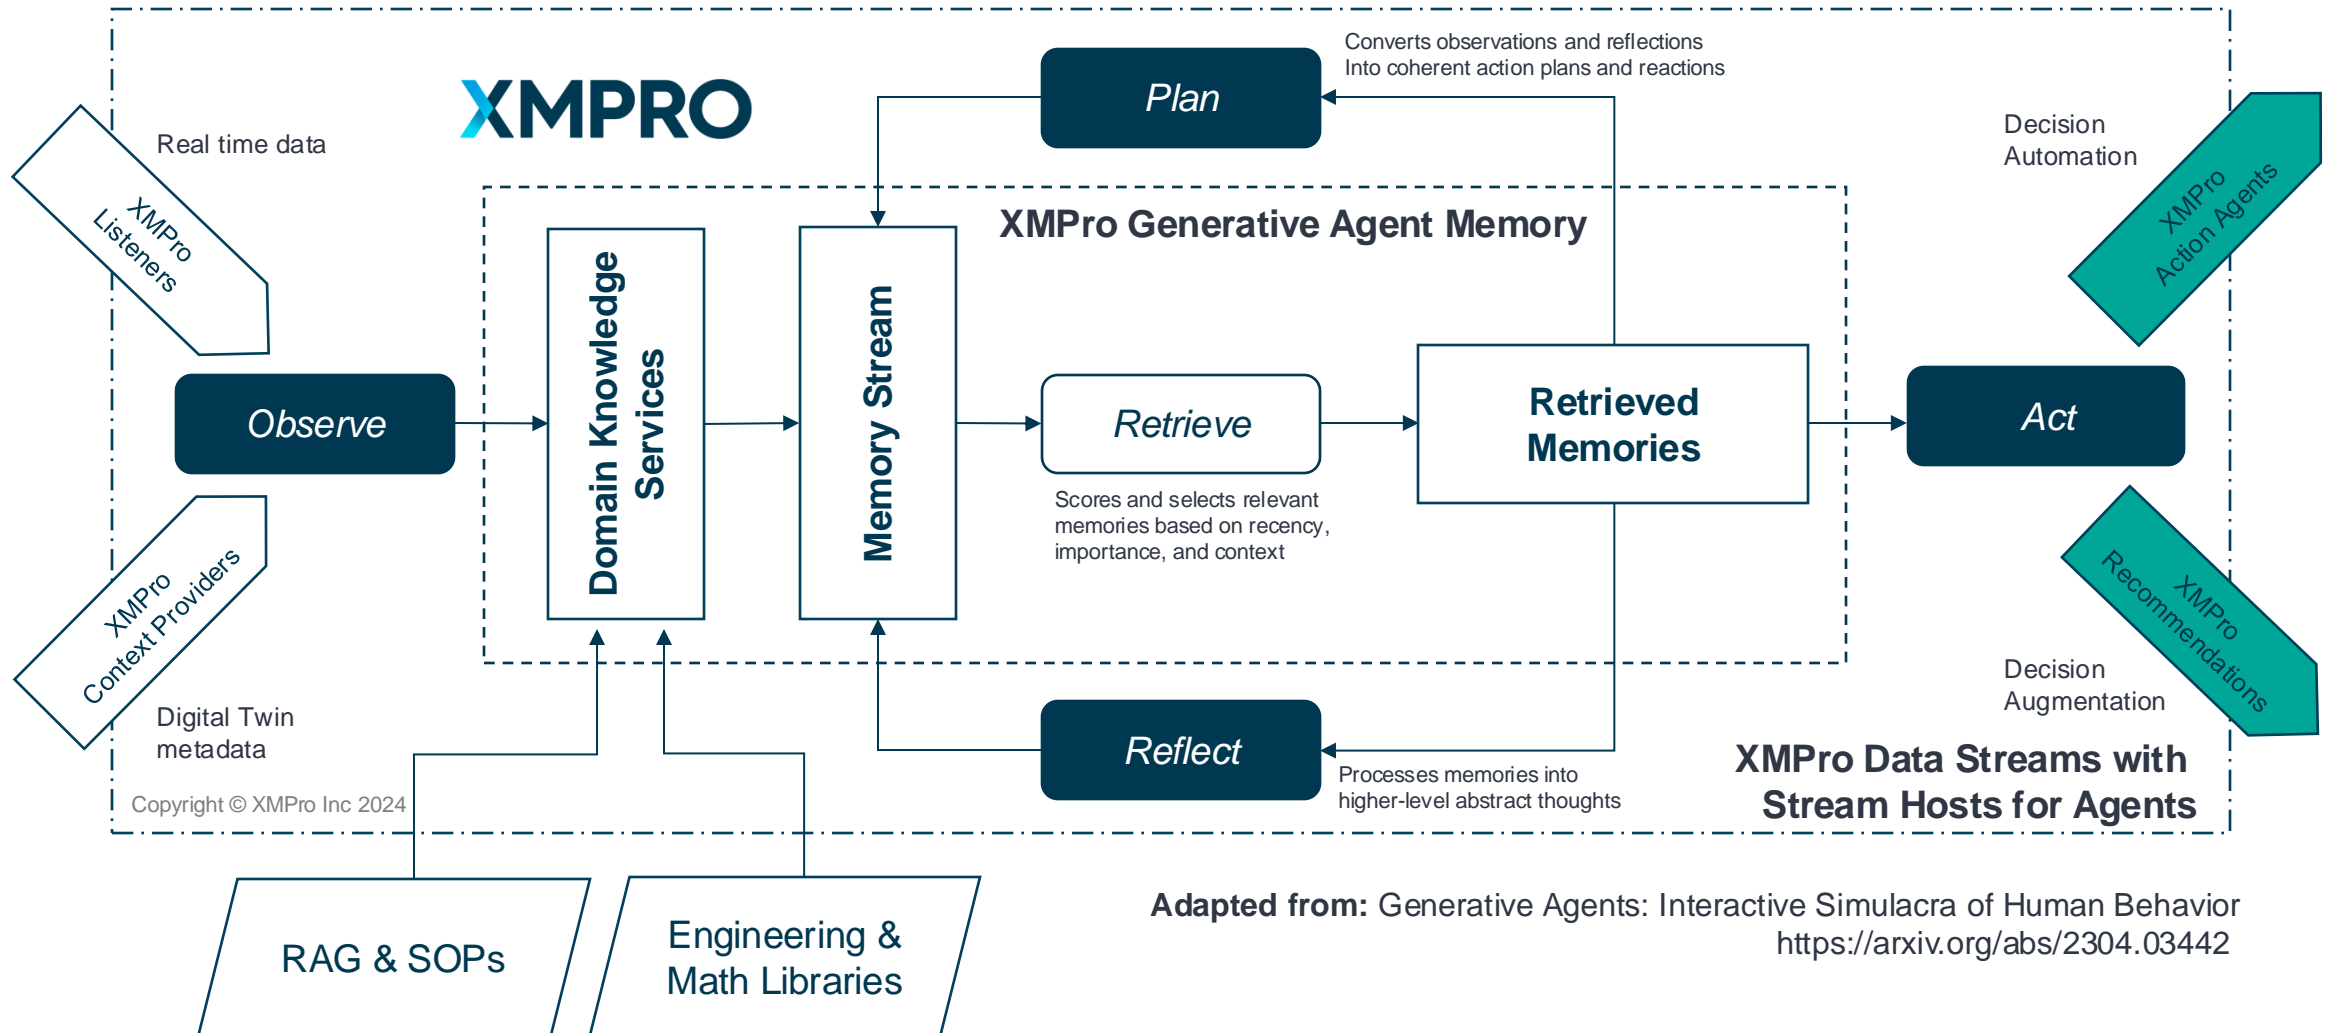
XMPro's Generative Agent Memory Architecture – How an XMPro Agent **Observes**, **Reflects**, **Plans**, and **Acts** with LLMs



Adapted from: Generative Agents: Interactive Simulacra of Human Behavior  
<https://arxiv.org/abs/2304.03442>

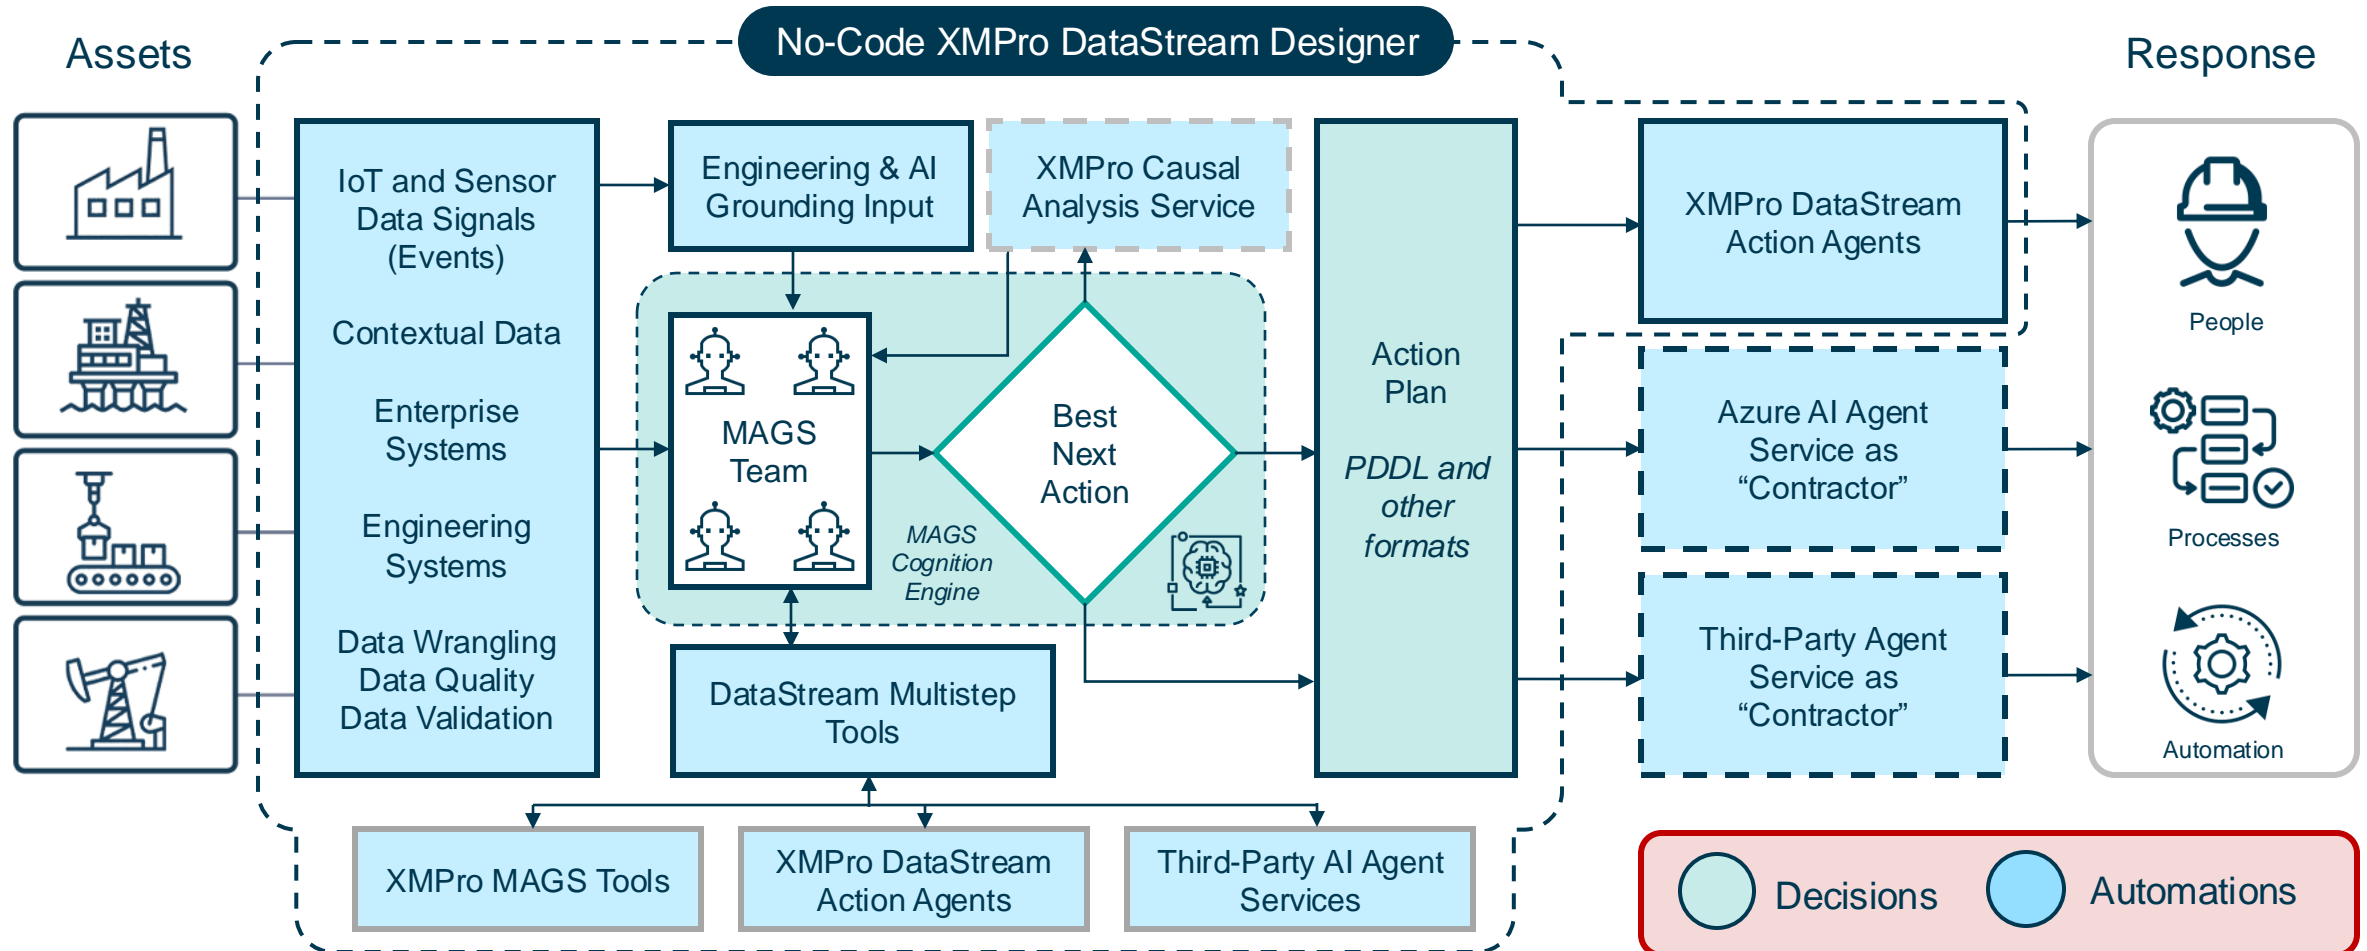
# What Makes This Work?

XMPro's Generative Agent Memory Architecture – How an XMPro Agent **Observes**, **Reflects**, **Plans**, and **Acts** with LLMs

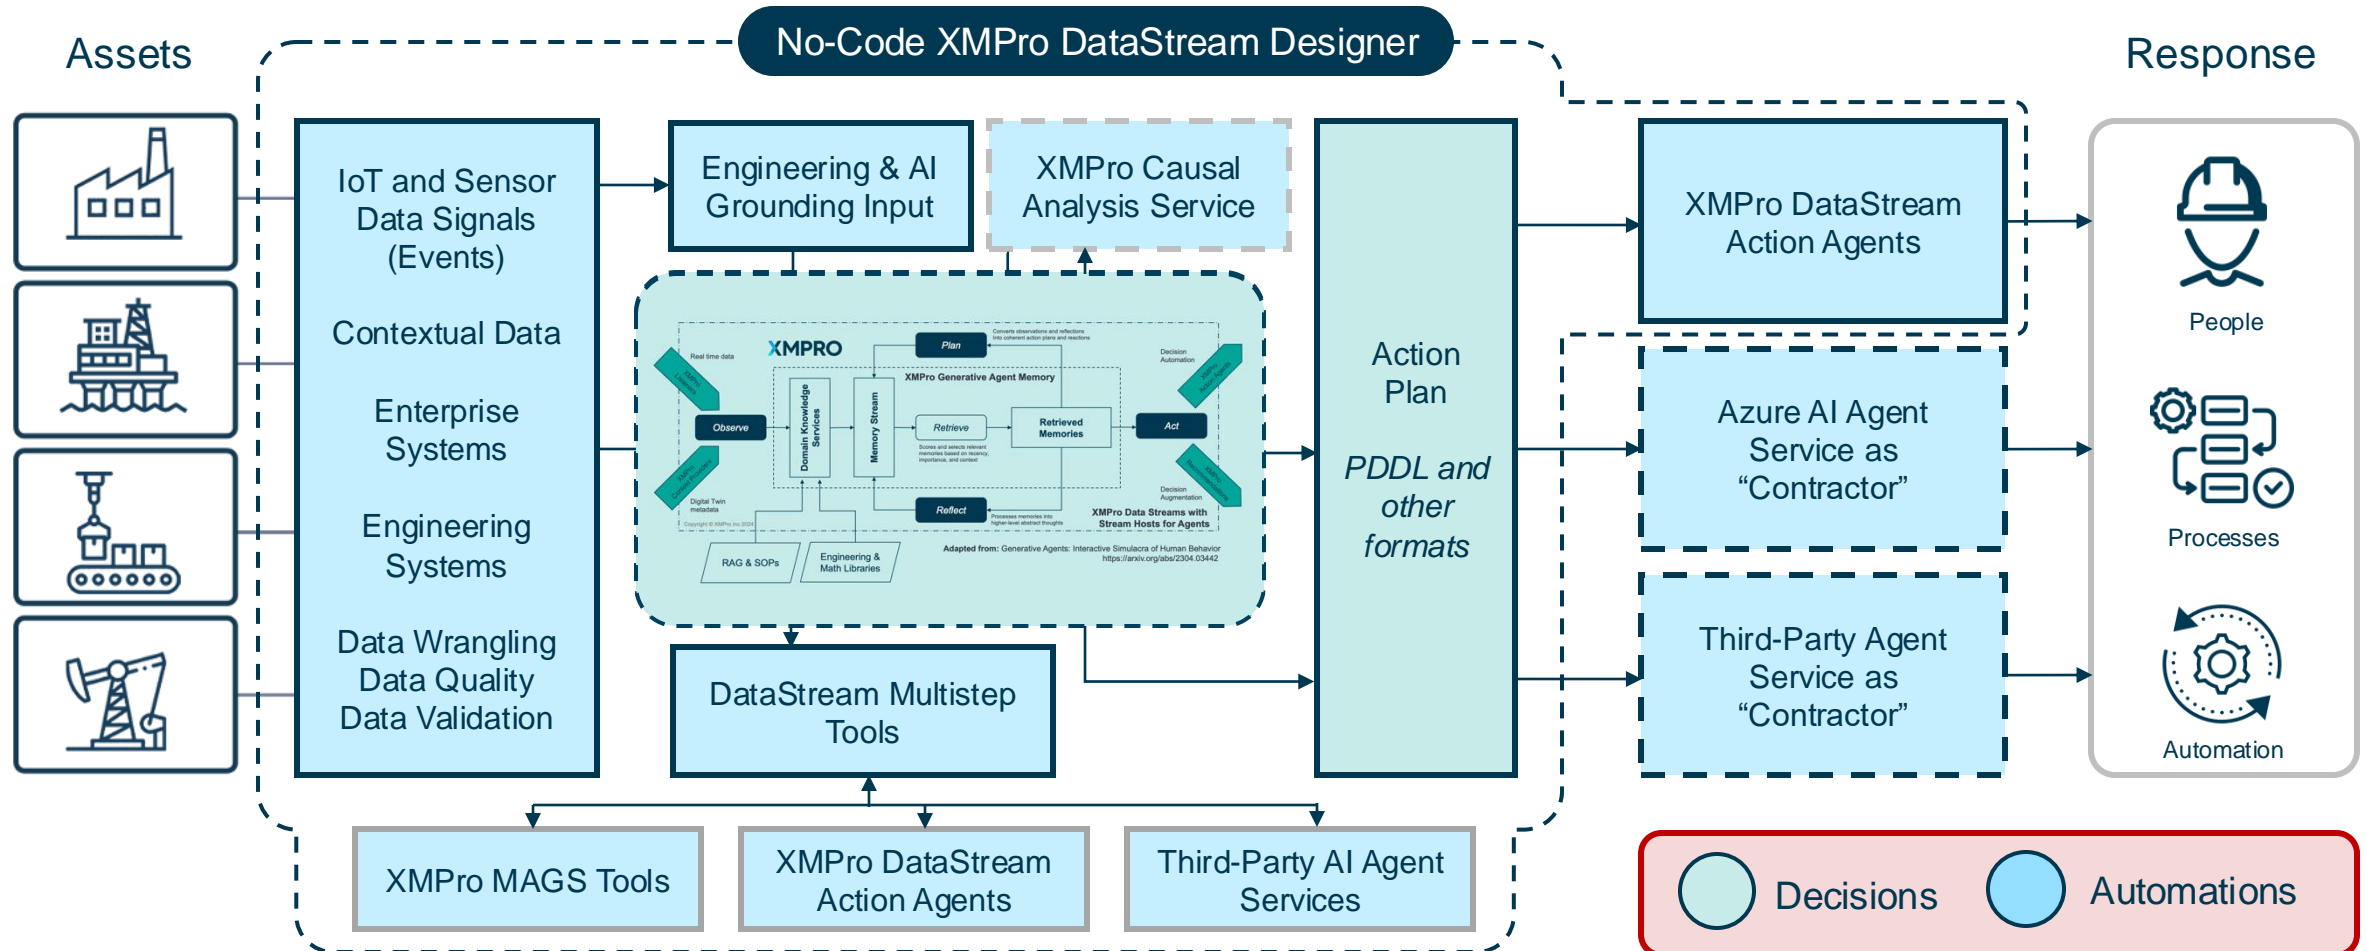


Adapted from: Generative Agents: Interactive Simulacra of Human Behavior  
<https://arxiv.org/abs/2304.03442>

# Agentic Decisions vs Automations

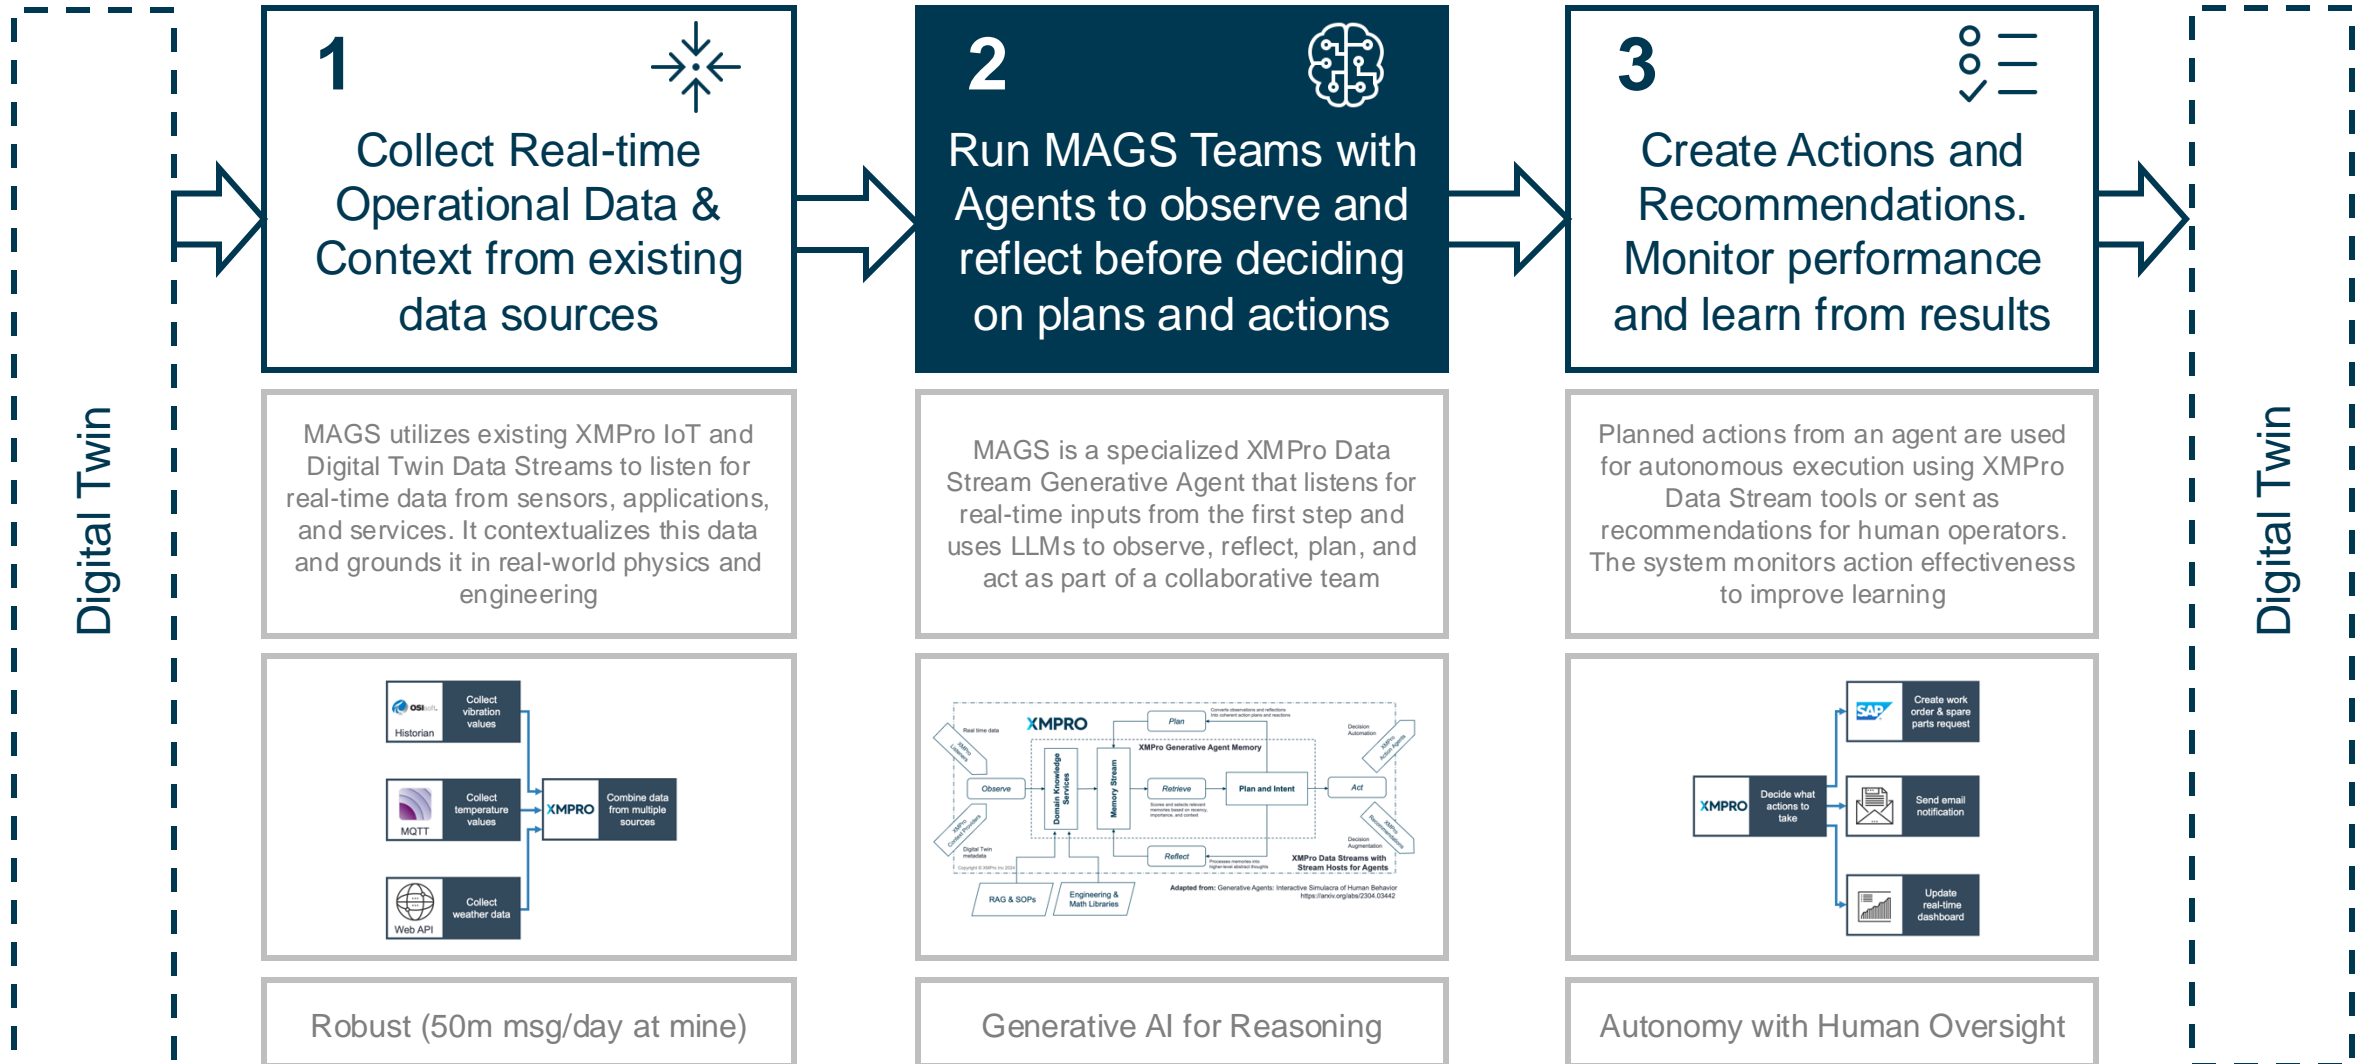


# Agentic Decisions vs Automations



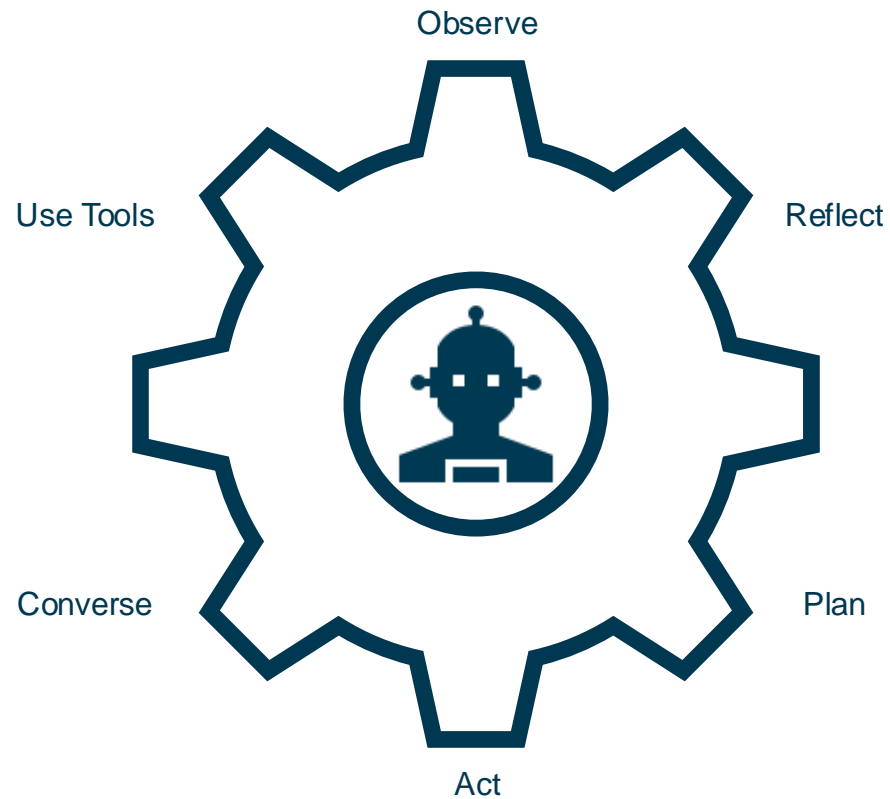


# MAGS in Action - Operational Workflow



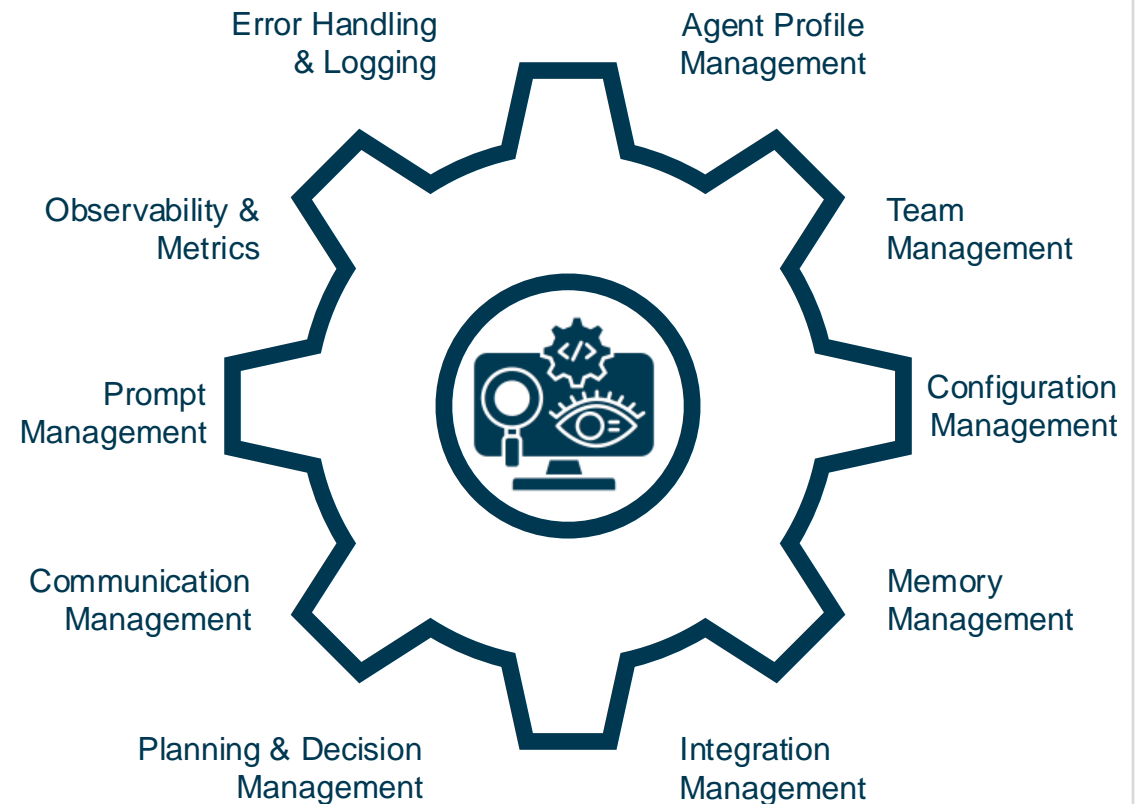
# How To Manage All These Agents At Scale

## XMPro Multi Agent Generative Systems (MAGS)



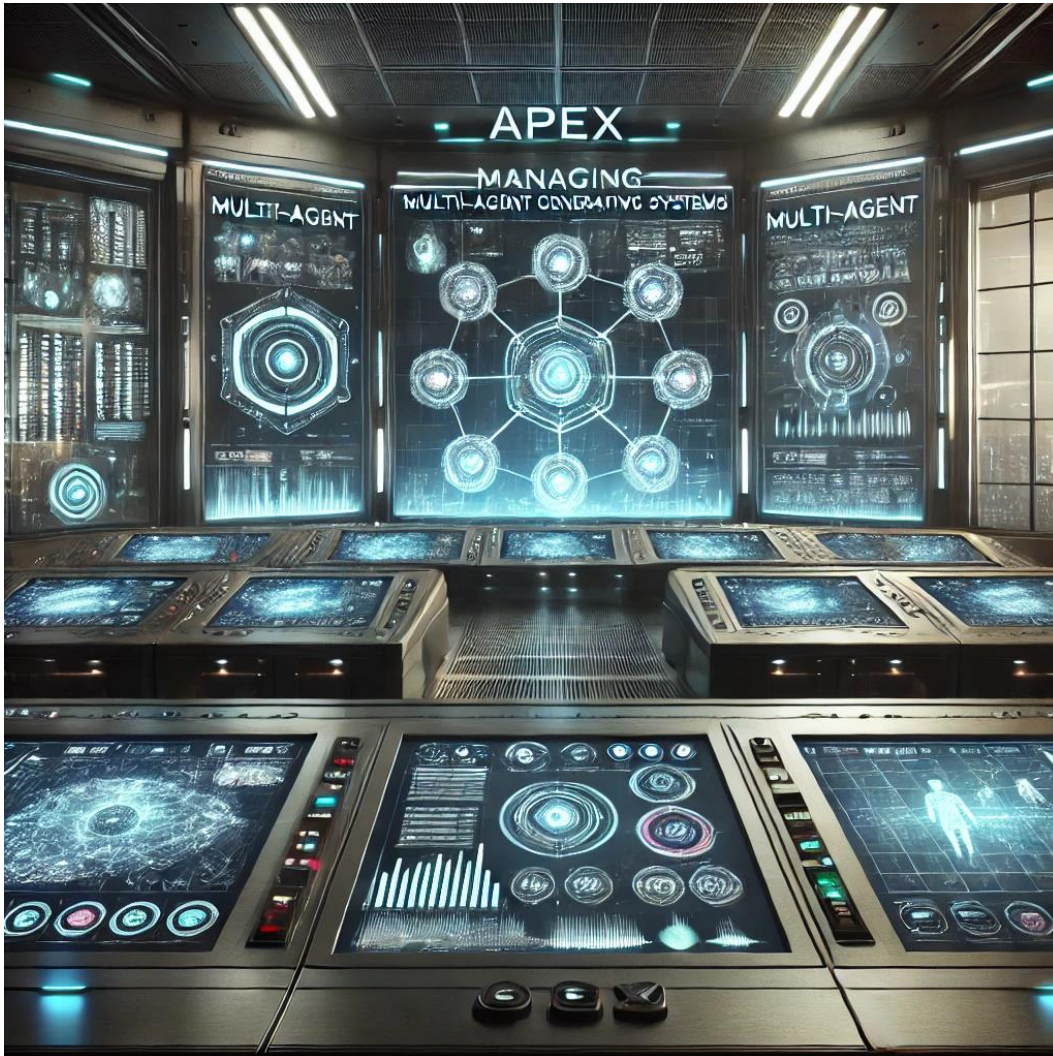
**OpsAgent**

## XMPro Agent Platform EXperience (APEX)

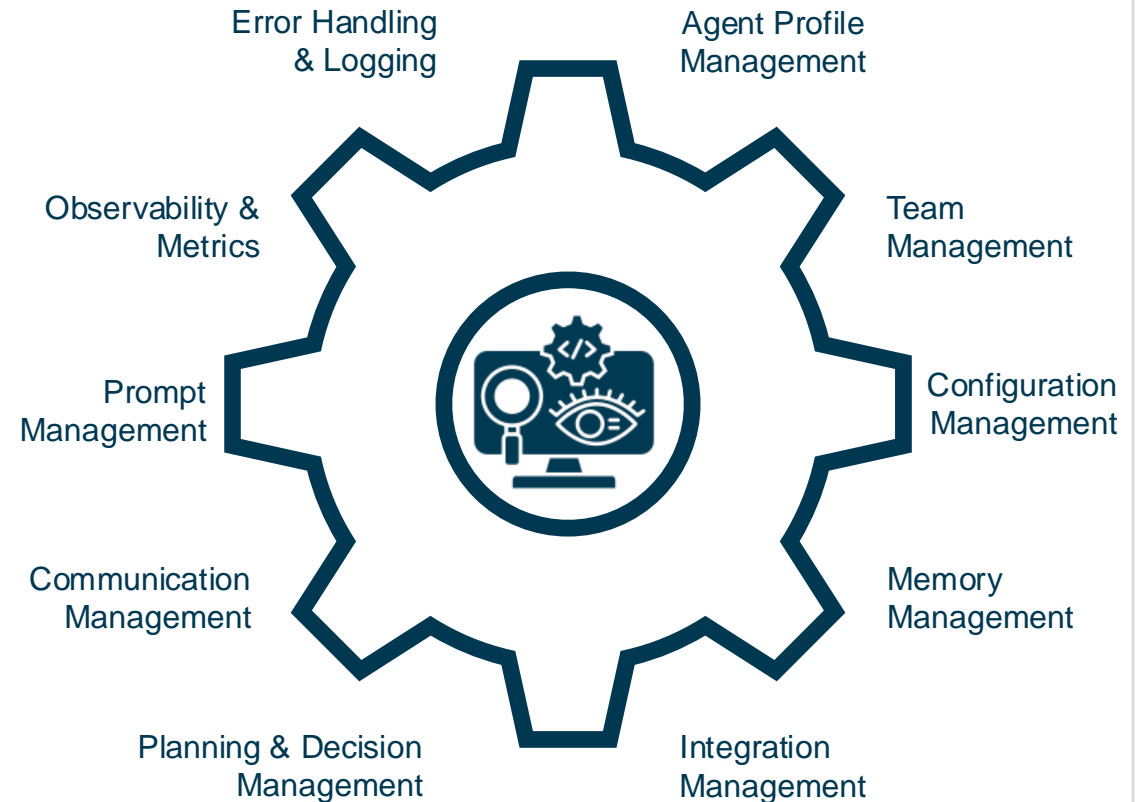


**AgentOps**

# XMPro APEX - The Control Room for MAGS

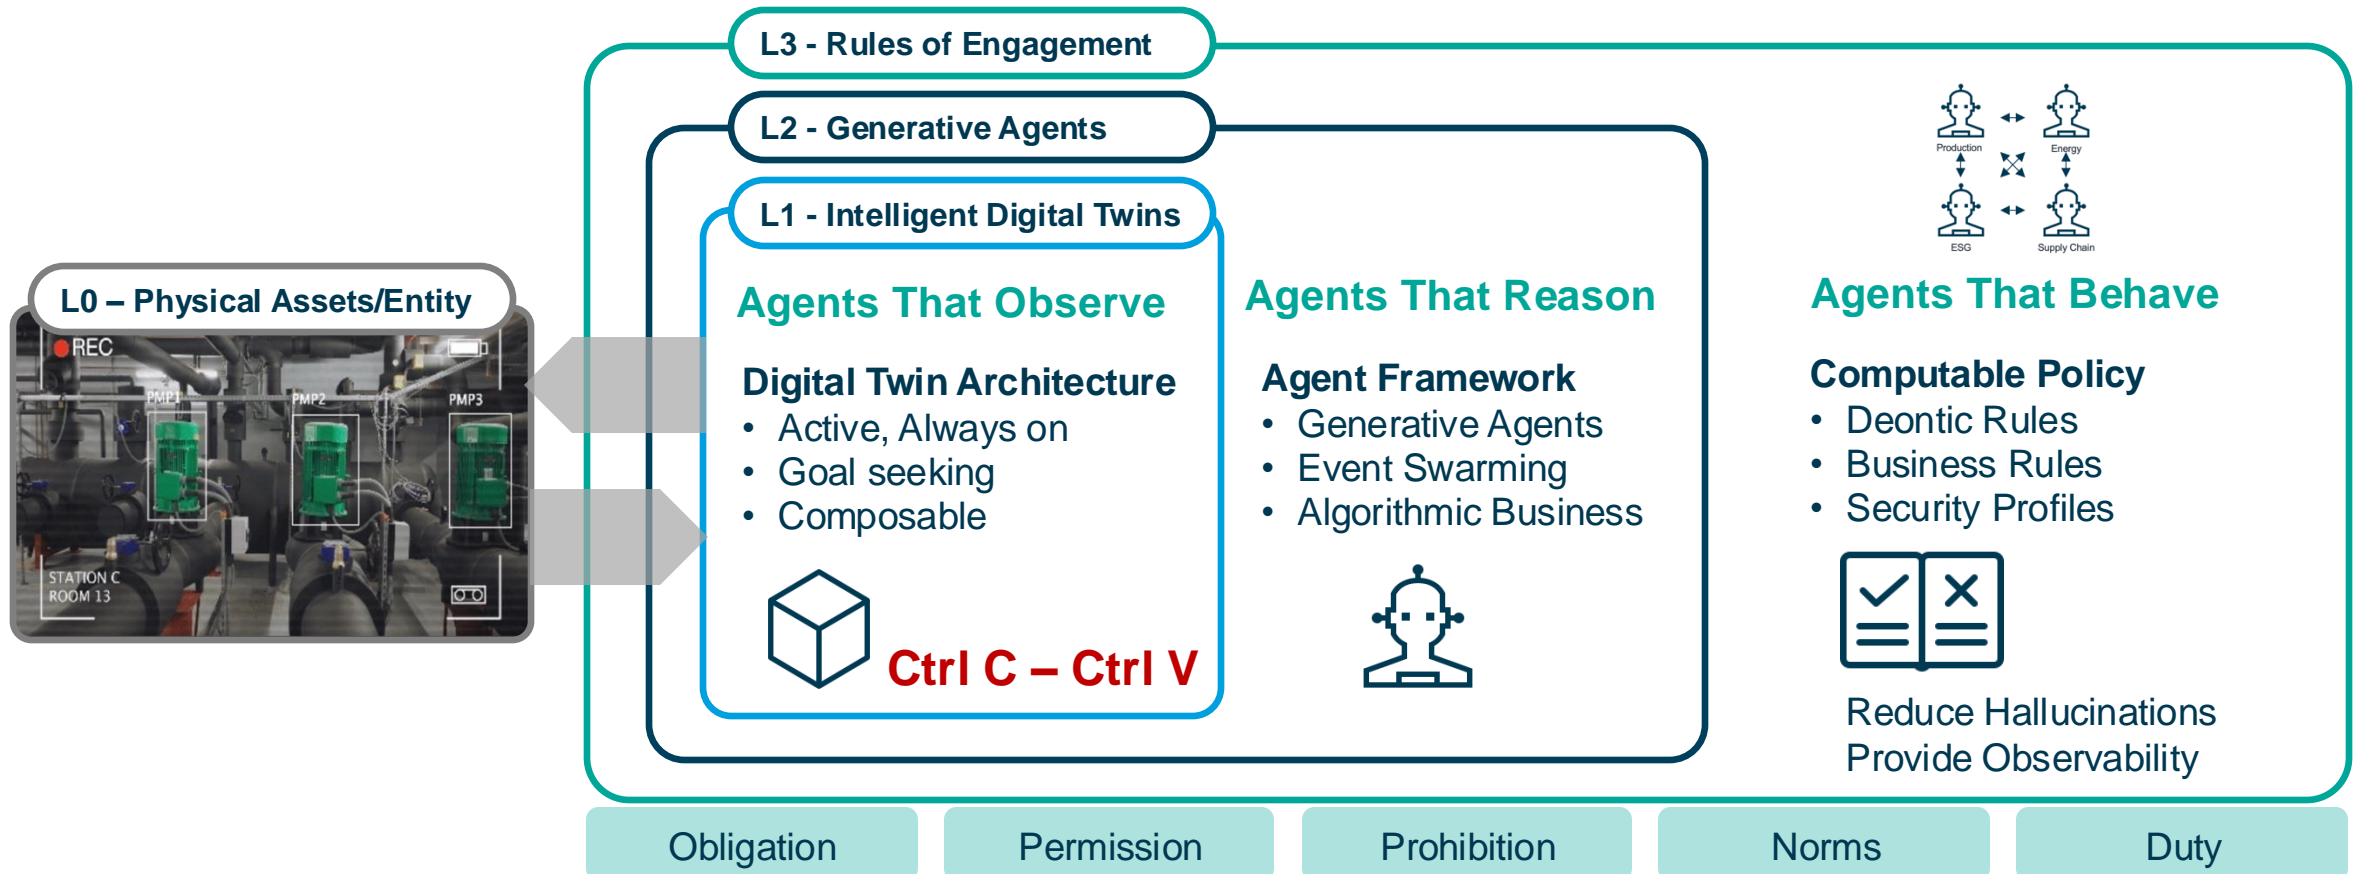


## XMPro Agent Platform EXperience (APEX)



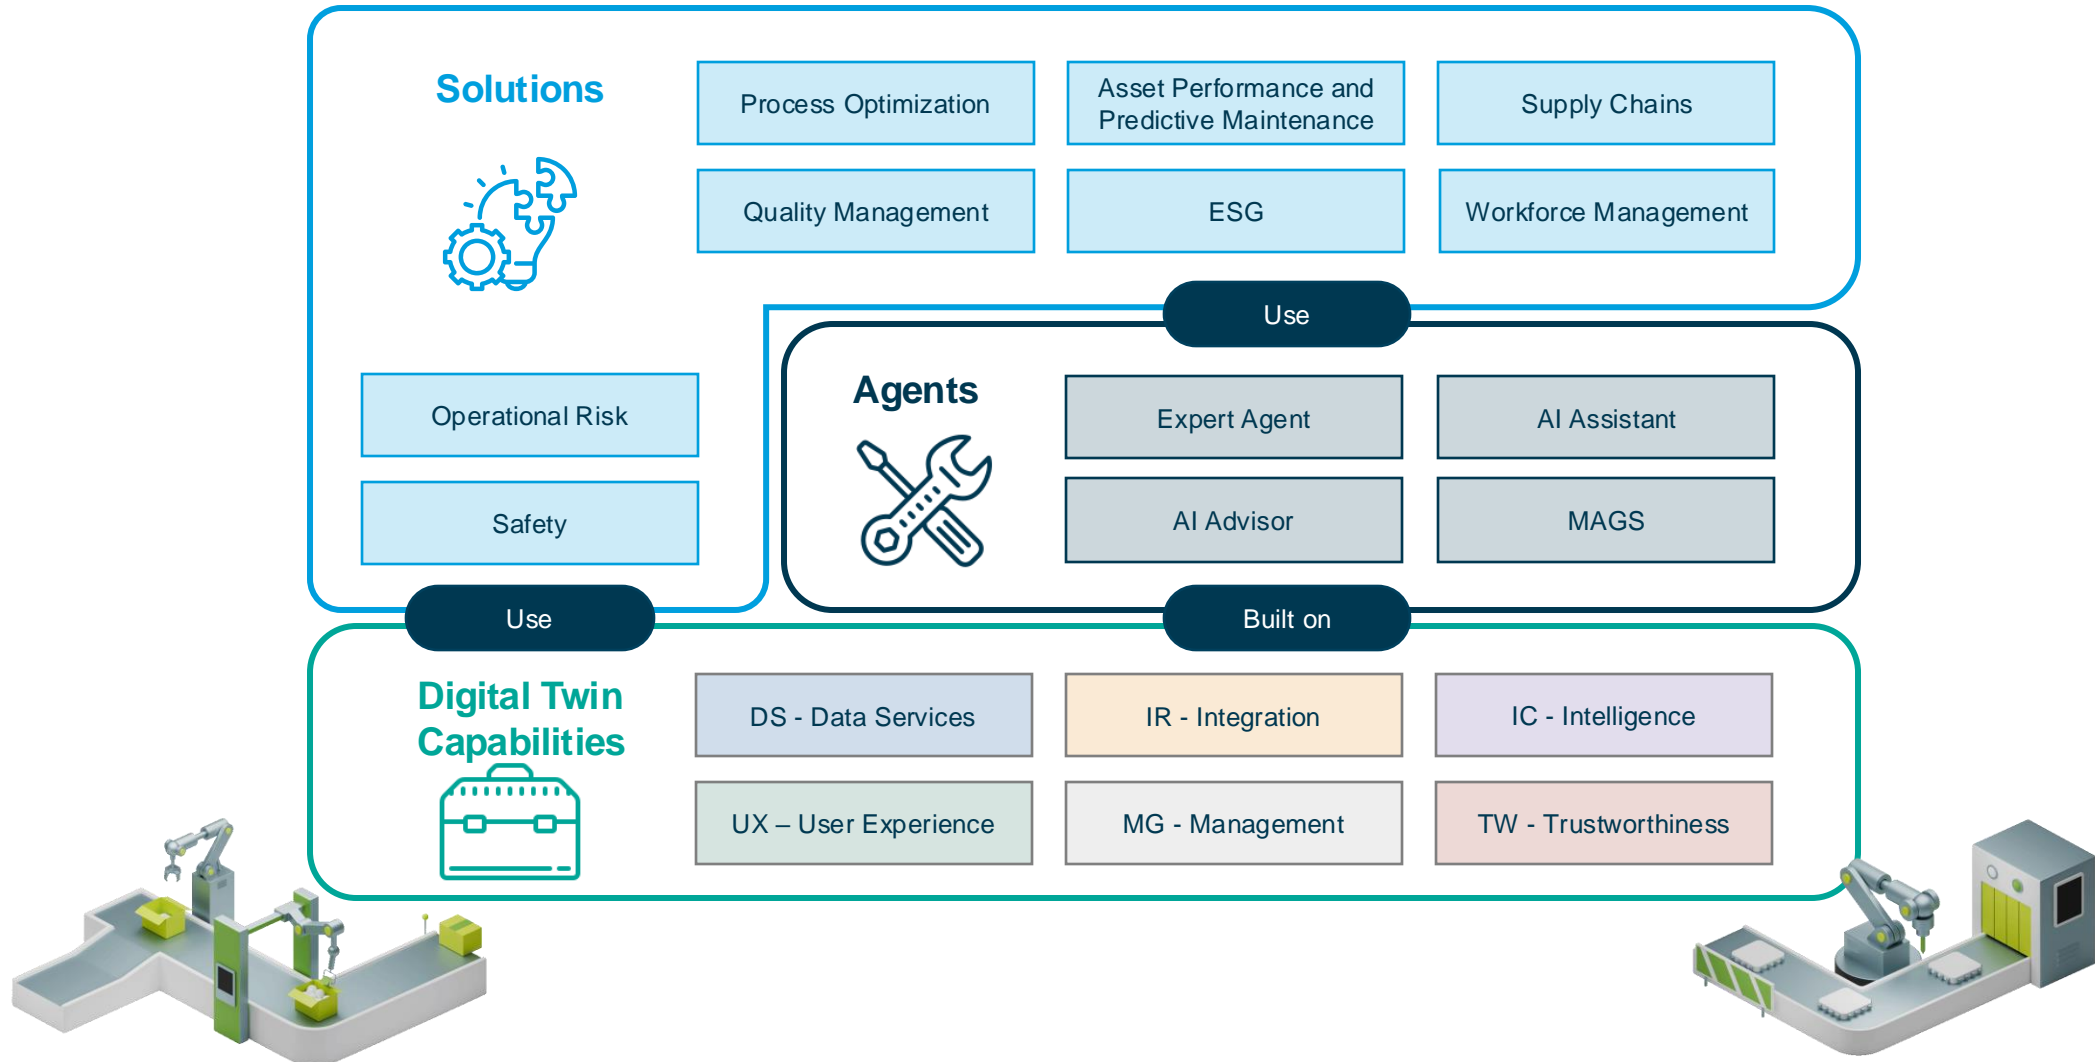
**AgentOps**

# Where is the Digital Twin in MAGS?





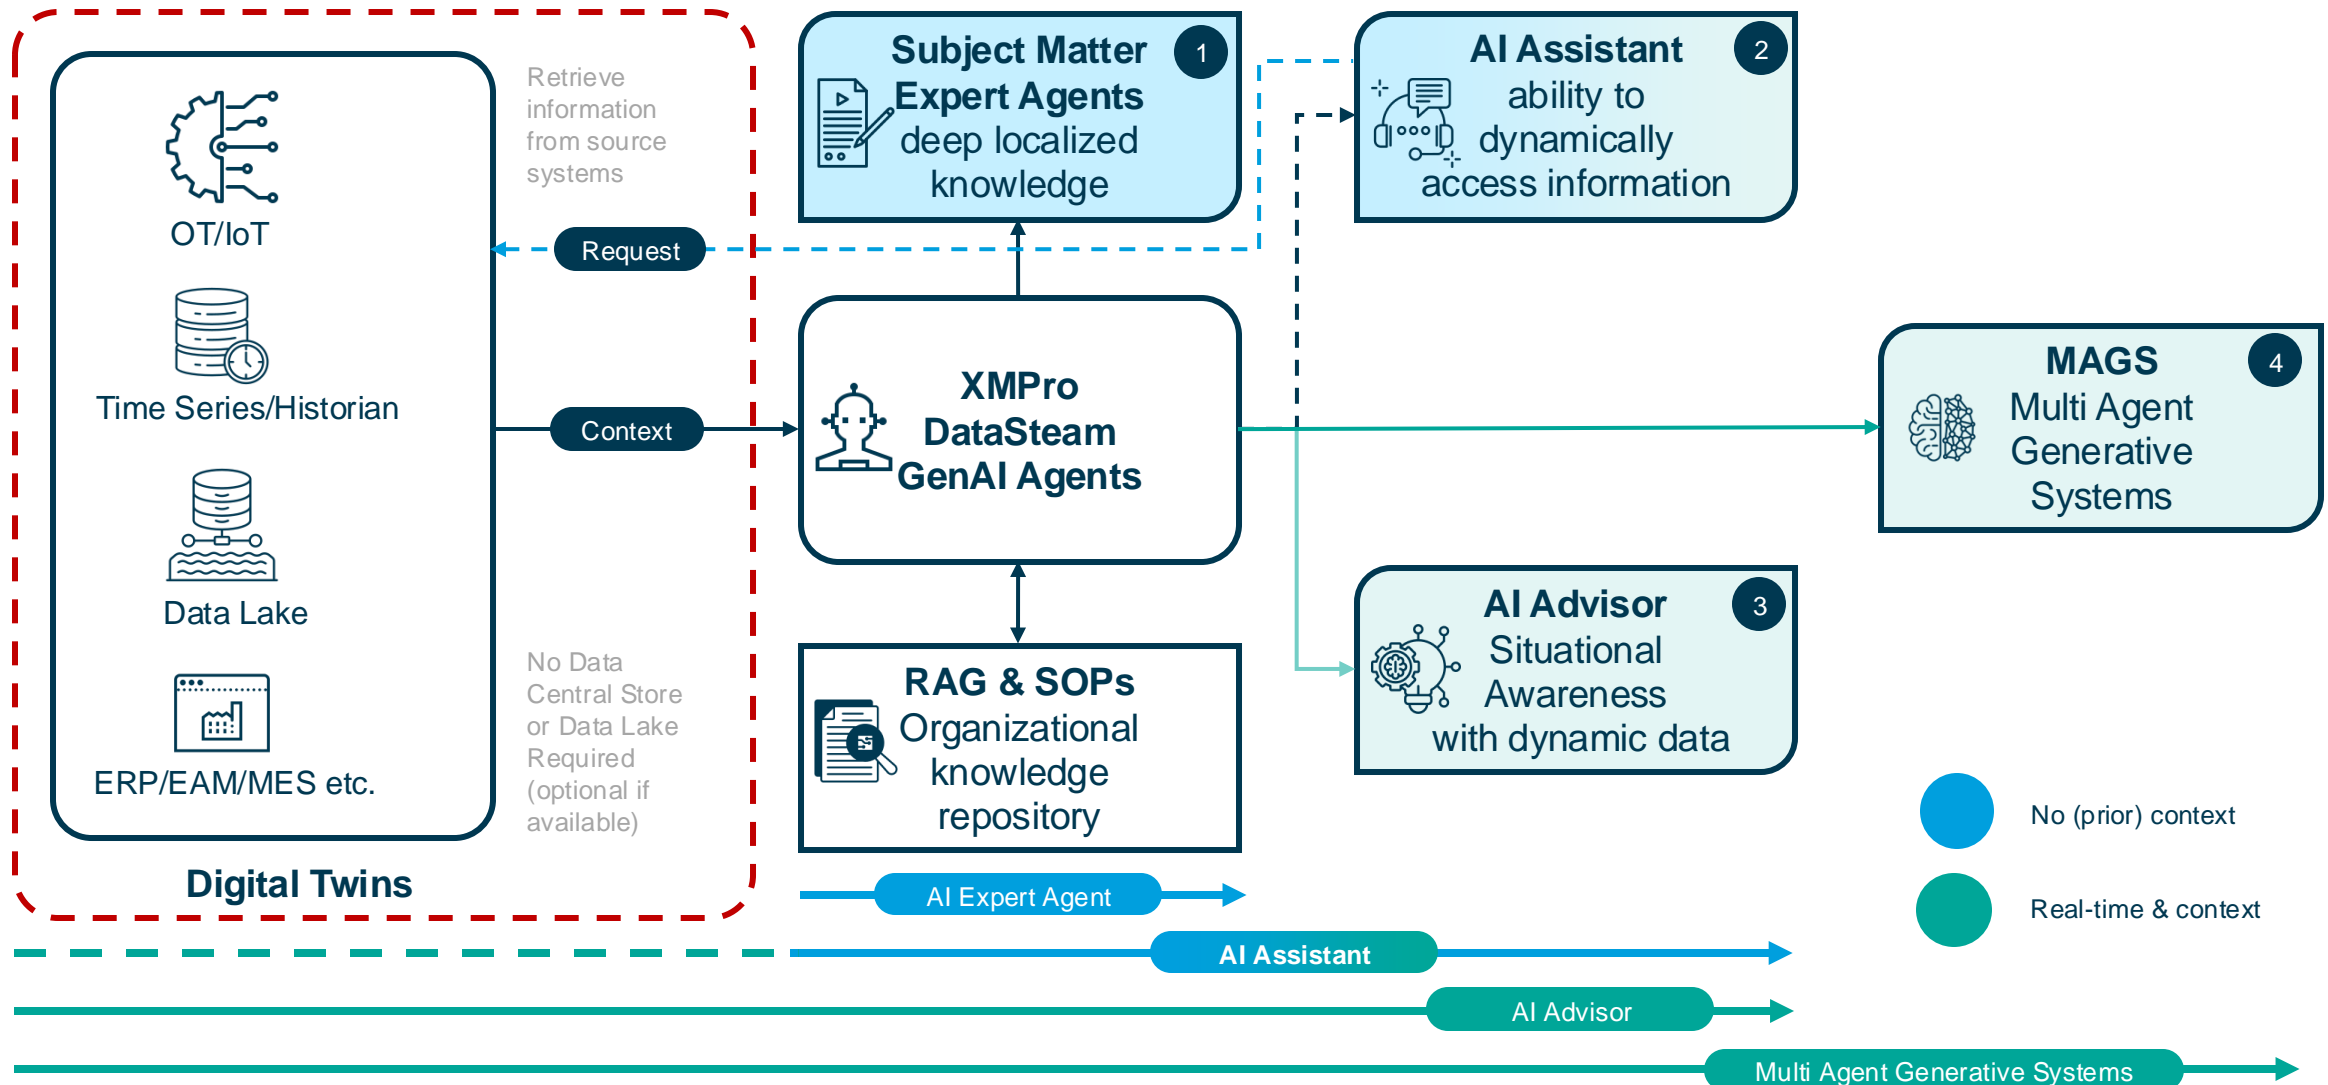
# Solution Stack for Intelligent Operations



# Digital Twin Capabilities Periodic Table v1.1


|  |  |   |   |  |  |   |   |
|--|--|---|---|--|--|---|---|
| <b>DS.AI</b><br>Data Acquisition & Ingestion | <b>DS.SG</b><br>Synthetic Data Generation          | <b>IR.ET</b><br>Enterprise System Integration | <b>IC.SR</b><br>Search                    | <b>IC.PR</b><br>Prediction                   |  | <b>UX.BV</b><br>Basic Visualization               | <b>UX.DB</b><br>Dashboards                  |
| <b>DS.ST</b><br>Data Streaming               | <b>DS.ON</b><br>Ontology Management                | <b>IR.EG</b><br>Eng. System Integration       | <b>IC.CC</b><br>Command & Control         | <b>IC.AI</b><br>Artificial Intelligence      |  | <b>UX.AV</b><br>Advanced Visualization            | <b>UX.CI</b><br>Continuous Intelligence     |
| <b>DS.TR</b><br>Data Transformation          | <b>DS.RP</b><br>Digital Twin (DT) Model Repository | <b>IR.IO</b><br>OT/IoT System Integration     | <b>IC.OS</b><br>Orchestration             | <b>IC.PS</b><br>Prescriptive Recommendations |  | <b>UX.RM</b><br>Real-time Monitoring              | <b>UX.BI</b><br>Business Intelligence       |
| <b>DS.CX</b><br>Data Contextualization       | <b>DS.IR</b><br>DT Instance Repository             | <b>IR.DT</b><br>Digital Twin Integration      | <b>IC.AL</b><br>Alerts & Notifications    | <b>IC.FL</b><br>Federated Learning           | <b>IC.BR</b><br>Business Rules                       | <b>UX.ER</b><br>Entity Relationship Visualization | <b>UX.BP</b><br>BPM & Workflow              |
| <b>DS.BP</b><br>Batch Processing             | <b>DS.DS</b><br>Domain Specific Data Management    | <b>IR.CL</b><br>Collab Platform Integration   | <b>IC.RP</b><br>Reporting                 | <b>IC.SM</b><br>Simulation                   | <b>IC.DL</b><br>Distributed Ledger & Smart Contracts | <b>UX.XR</b><br>Extended Reality (AV/VR/MR)       | <b>UX.GE</b><br>Gaming Engine Visualization |
| <b>DS.RT</b><br>Real-time Processing         | <b>DS.SA</b><br>Data Storage & Archive Services    | <b>IR.AS</b><br>API Services                  | <b>IC.AA</b><br>Data Analysis & Analytics | <b>IC.MA</b><br>Mathematical Analytics       | <b>IC.CS</b><br>Composition                          | <b>UX.GM</b><br>Gamification                      | <b>UX.3R</b><br>3D Rendering                |
| <b>DS.AS</b><br>Asynchronous Integration     | <b>DS.SR</b><br>Simulation Model Repository        | <b>MG.DM</b><br>Device Management             | <b>MG.EL</b><br>Event Logging             | <b>TW.EC</b><br>Data Encryption              | <b>TW.SC</b><br>Security                             | <b>TW.SF</b><br>Safety                            | <b>TW.RP</b><br>Responsibility              |
| <b>DS.AG</b><br>Data Aggregation             | <b>DS.AR</b><br>AI Model Repository                | <b>MG.SM</b><br>System Monitoring             | <b>MG.DG</b><br>Data Governance           | <b>TW.DS</b><br>Device Security              | <b>TW.PR</b><br>Privacy                              | <b>TW.RL</b><br>Reliability                       | <b>TW.RS</b><br>Resilience                  |

# The Roles that XMPro GenAI Agents Supports




# Continuum of AI-based Decision Augmentation


|                      |                              |                       |                                |                           |
|----------------------|------------------------------|-----------------------|--------------------------------|---------------------------|
| Scope                | Query Local Expert Knowledge | Query Enterprise Data | Specific entity state guidance | Manage tasks autonomously |
| Mode/Trigger         | User                         | User                  | User with Automation           | Automated Event           |
| User Interface       | Conversational               | Conversational        | Conversational                 | Composite                 |
| Live Data            | No                           | On Request            | Yes                            | Yes                       |
| Contextual Awareness | No                           | On Request            | Yes                            | Yes                       |



**Subject Matter Expert Agents** 1  
deep localized knowledge



**AI Assistant** 2  
ability to dynamically access information



**AI Advisor** 3  
Situational Awareness with dynamic data



**MAGS** 4  
Multi Agent Generative Systems

# MAGS Team In Action in Gold Mine

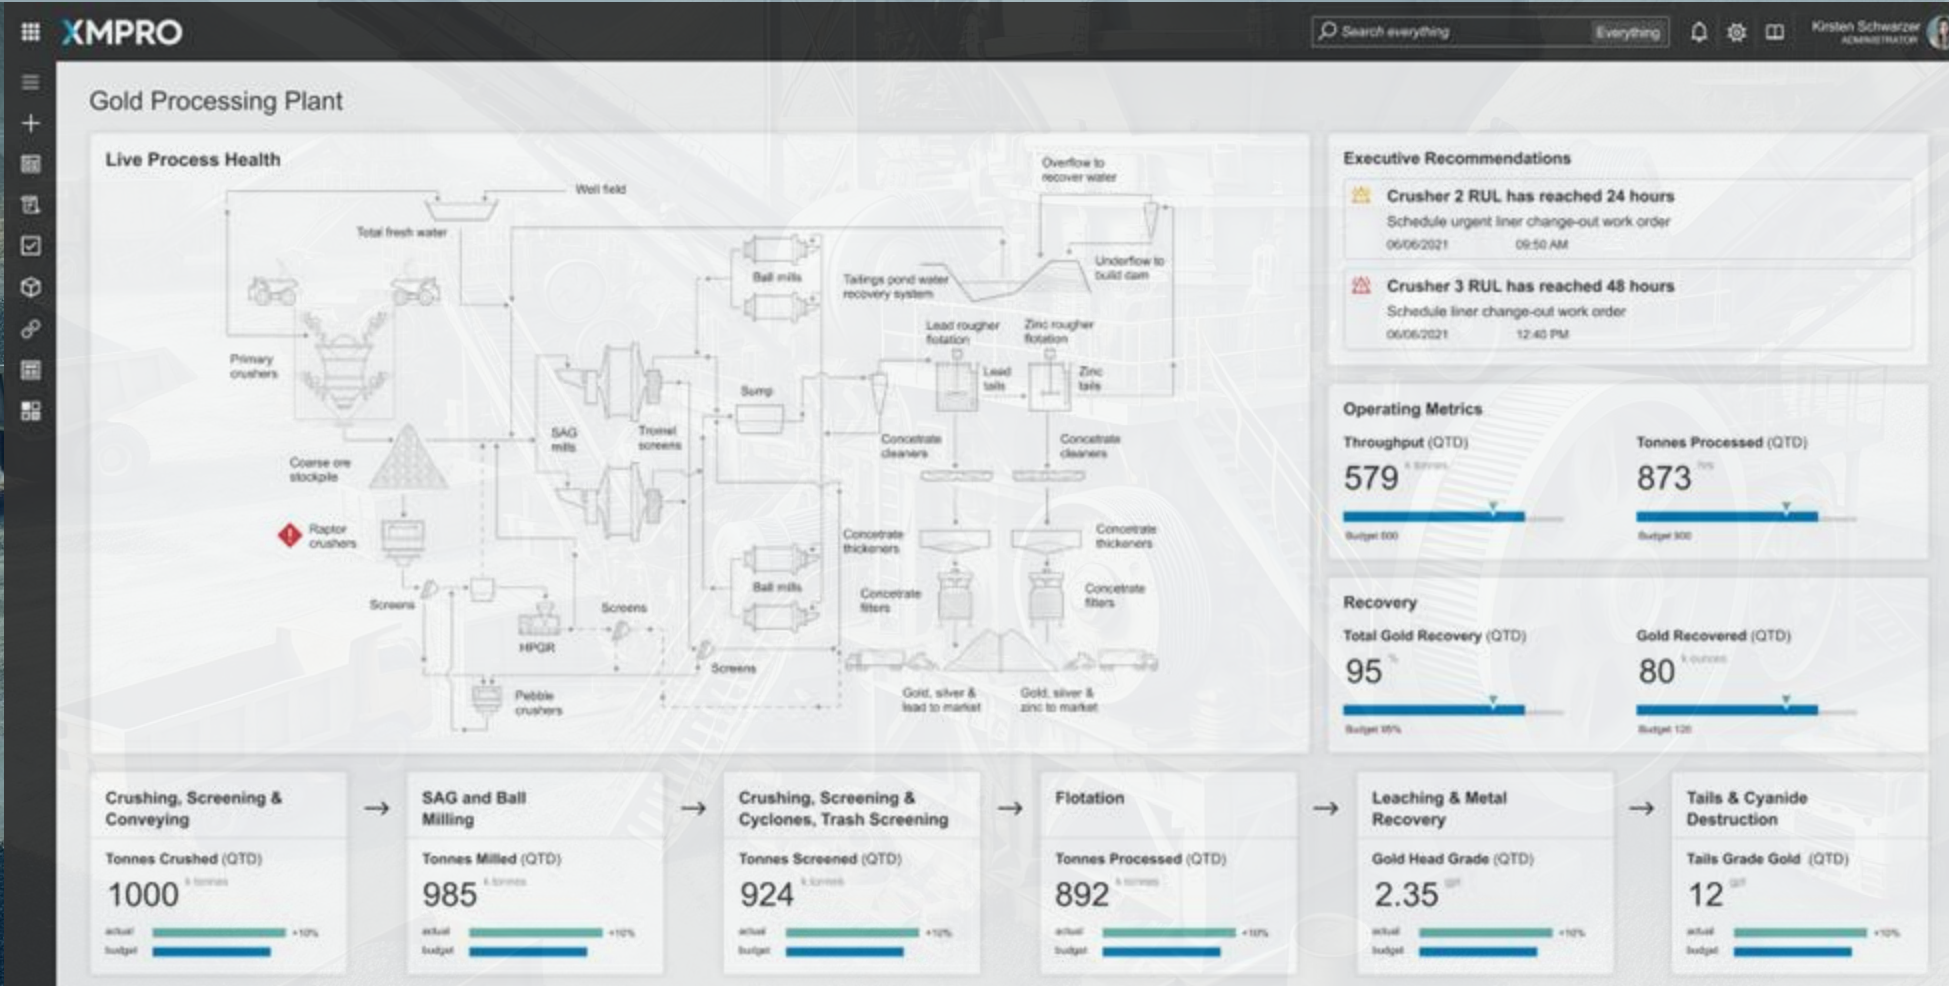
Minimize the total maintenance cost and gold production loss by optimizing the timing of liner replacements across six secondary crushers to ensure at least five crushers remain operational at all times.

Consider remaining useful life (RUL), wear rates, maintenance schedule constraints, and preventing overlapping end-of-life predictions to maintain consistent ore particle size for optimal gold recovery.



# ATLAS

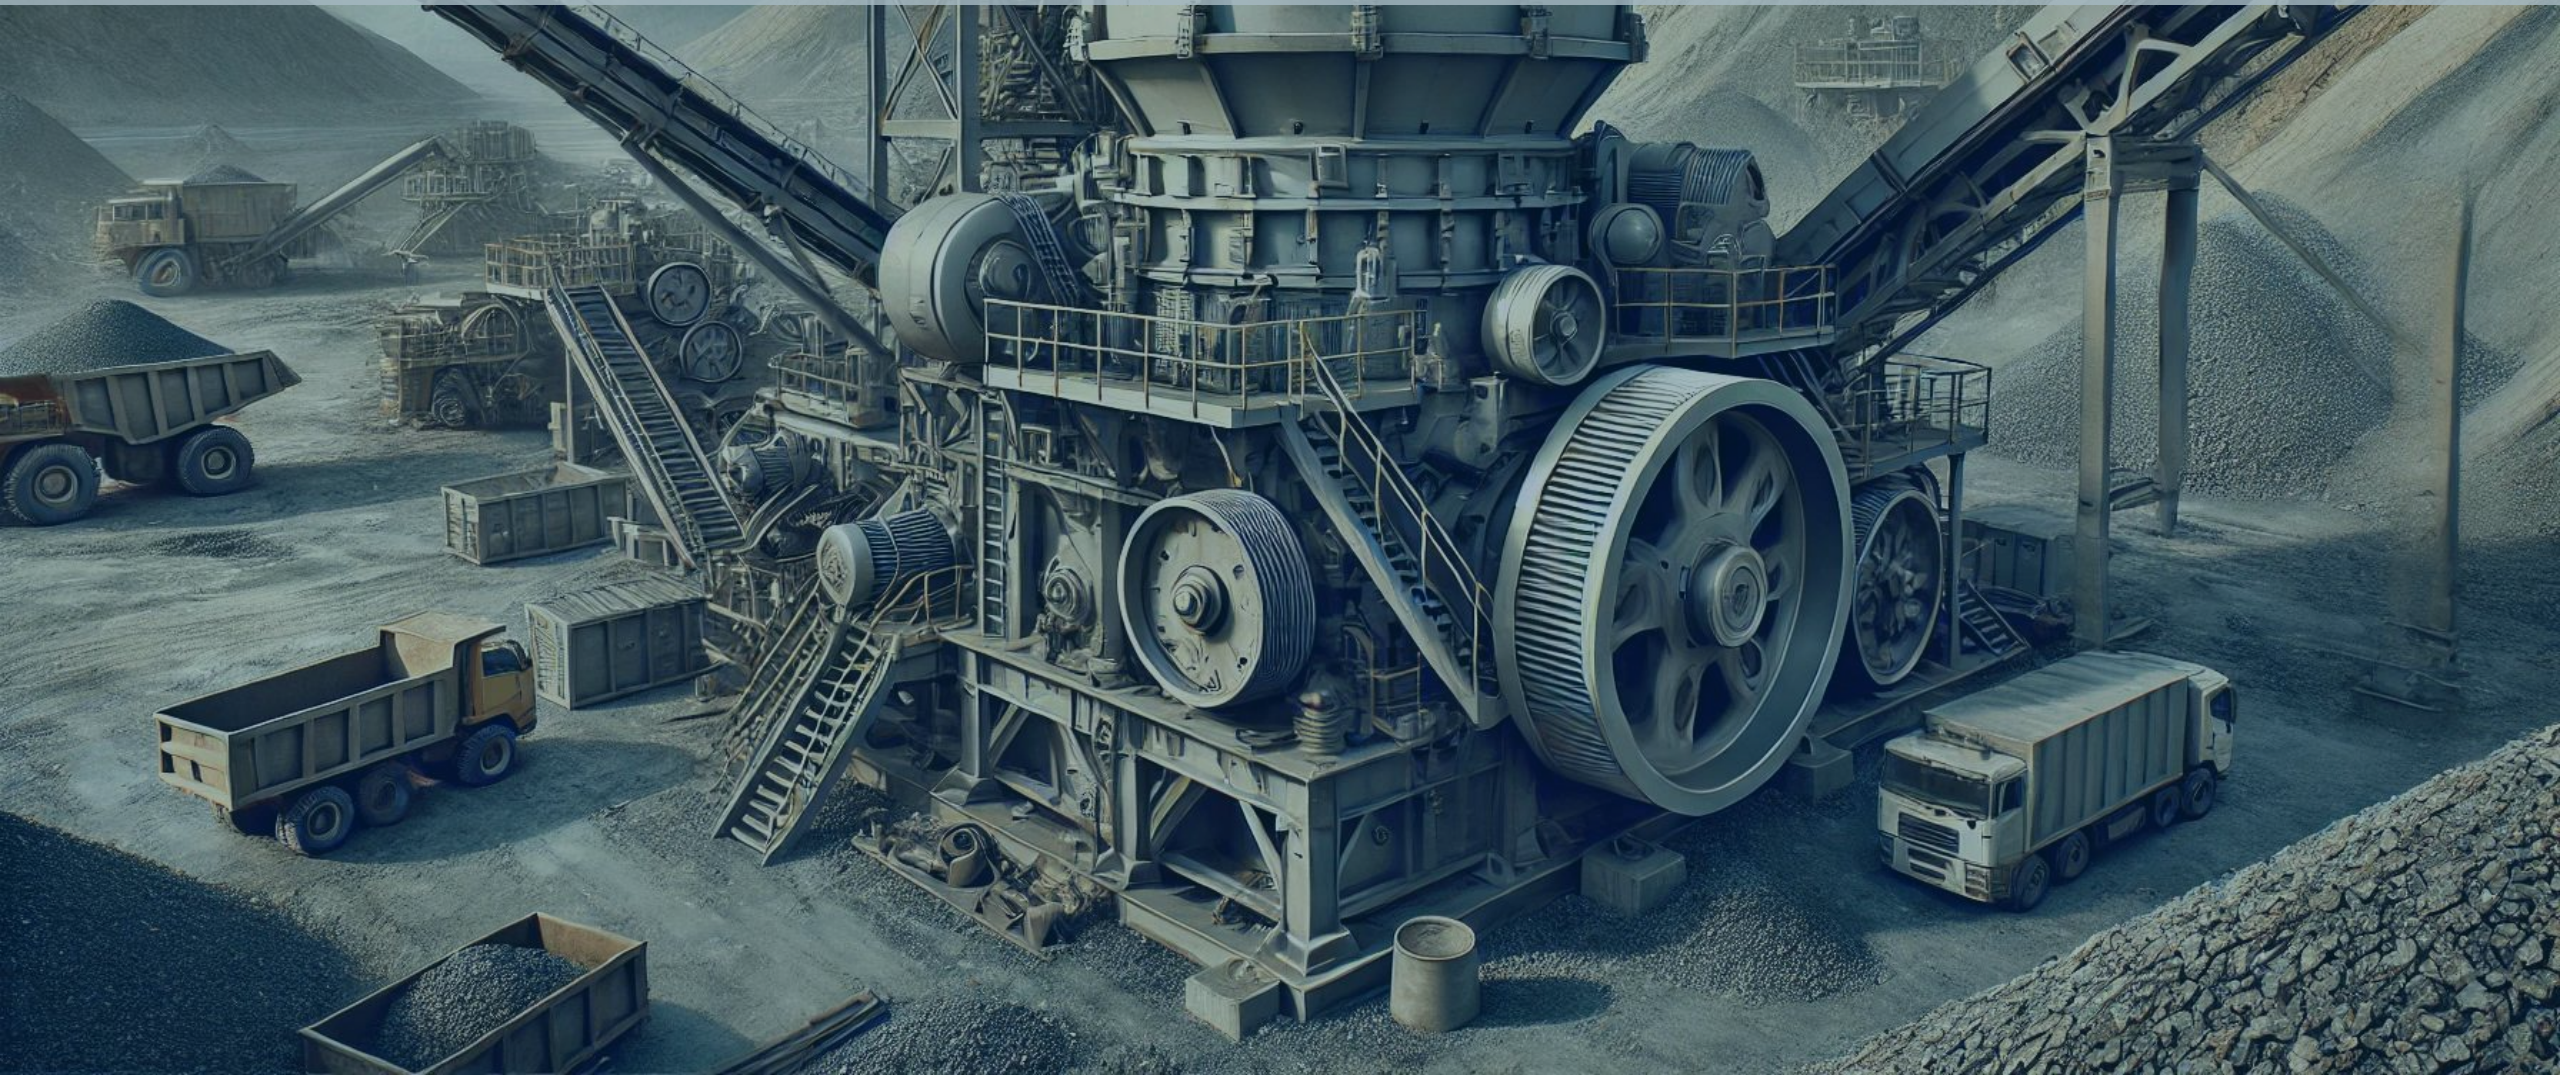
## Automated Team for Liner Analysis and Scheduling





# ATLAS

## Automated Team for Liner Analysis and Scheduling





# ATLAS

## Automated Team for Liner Analysis and Scheduling

To optimize the operation and scheduling of secondary crushers maximizing throughput

Single Isolated Objective and Task, not trying to replace a job function

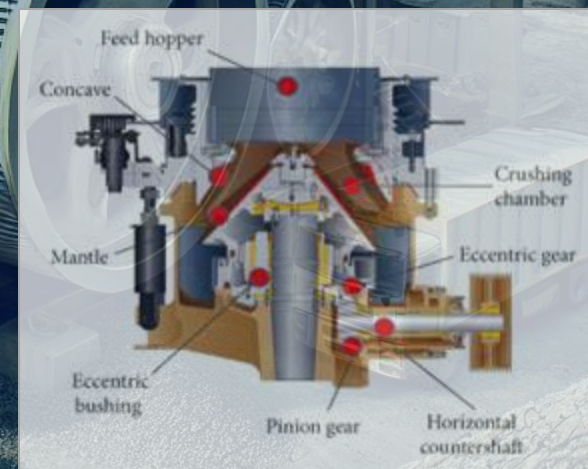
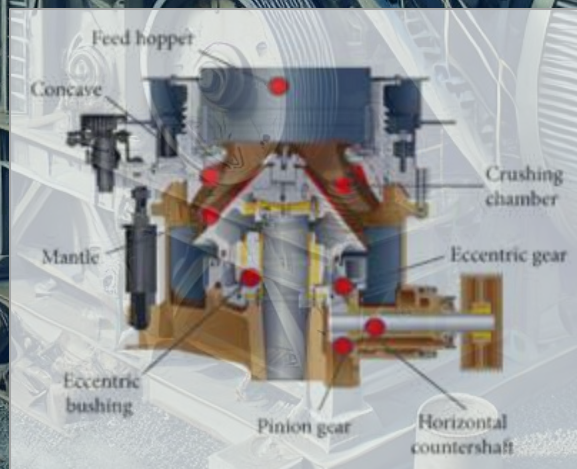
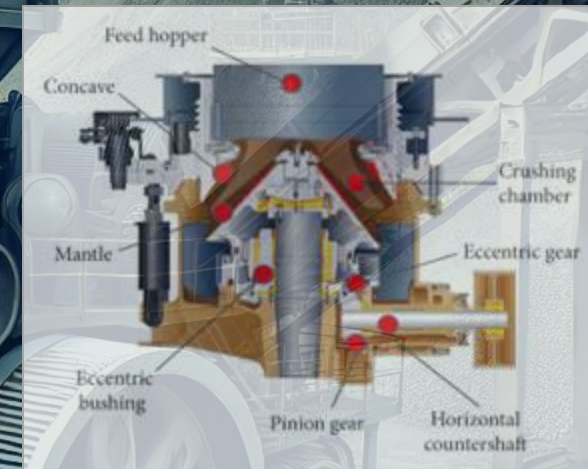
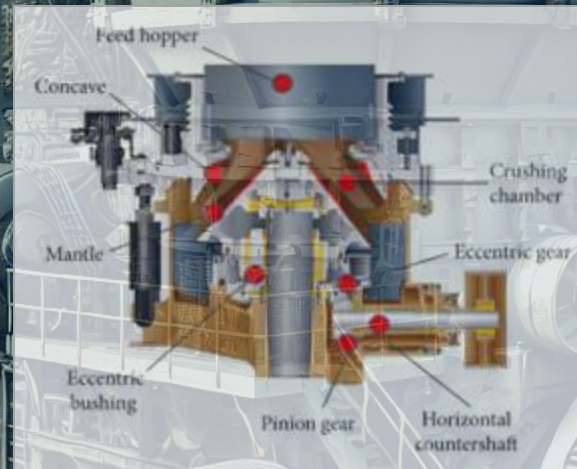
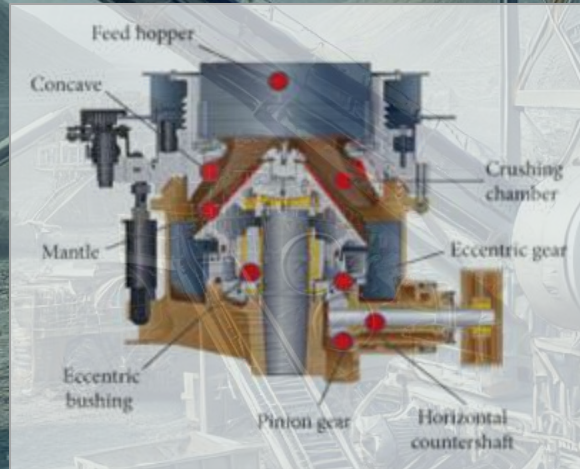
- maintenance planning and
- wear monitoring, ensuring at least

5 crushers are operational at any time



# ATLAS

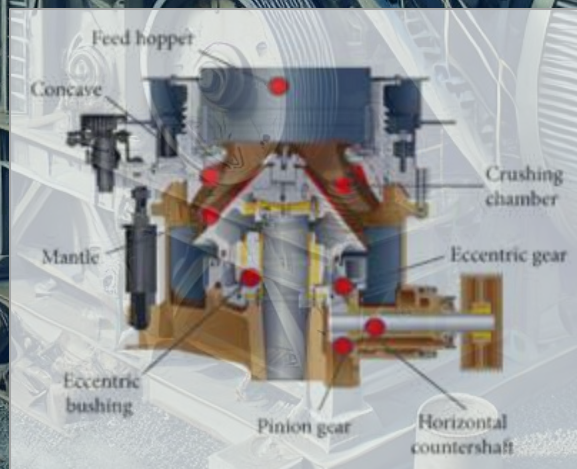
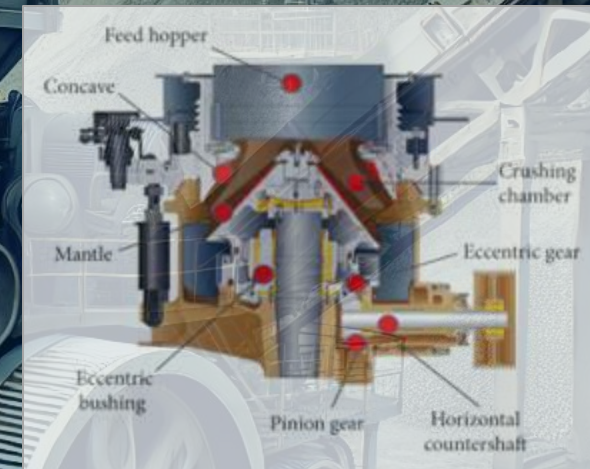
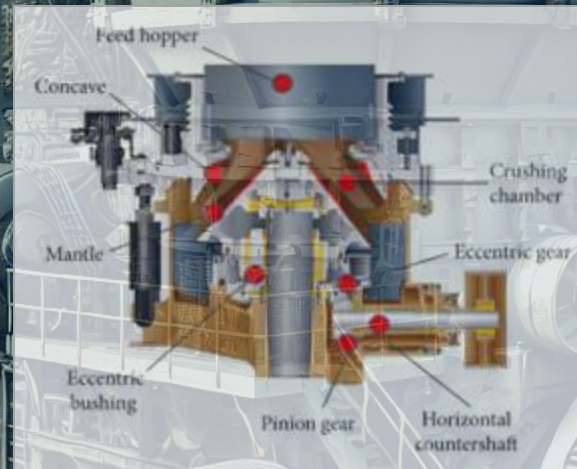
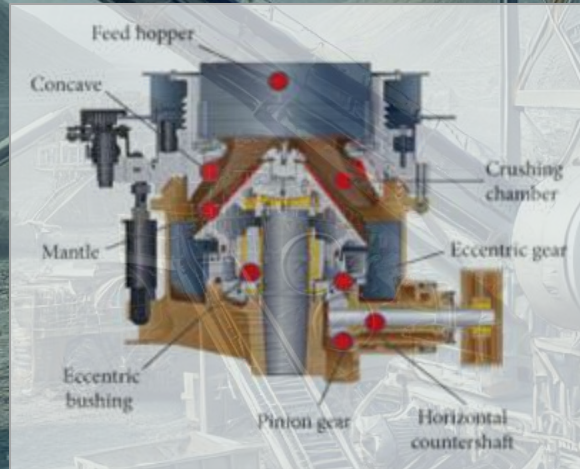
## Automated Team for Liner Analysis and Scheduling





# ATLAS

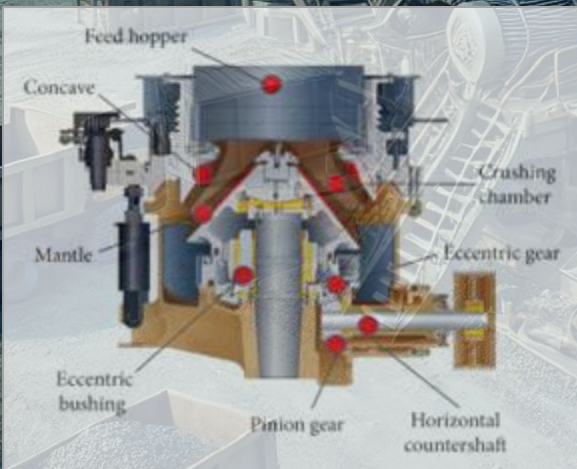
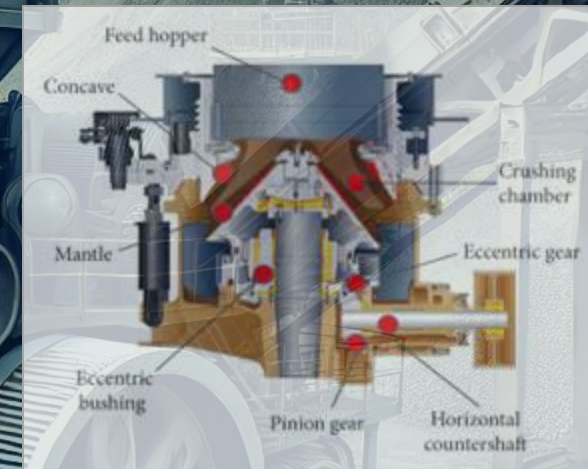
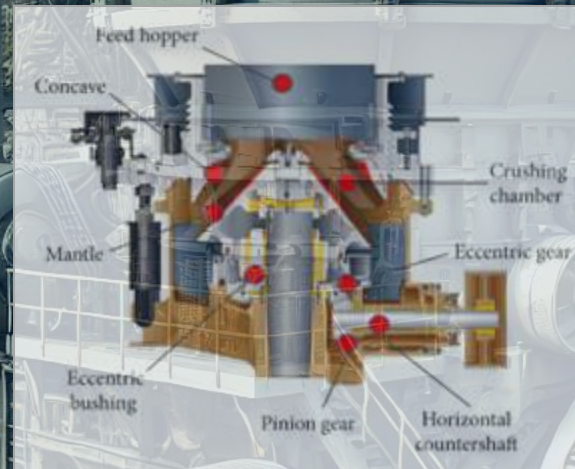
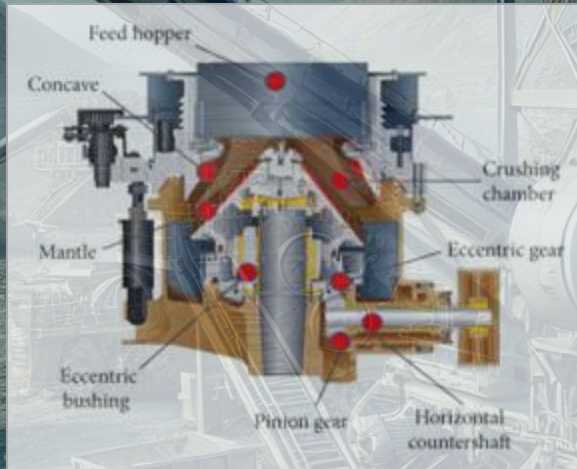
## Automated Team for Liner Analysis and Scheduling





# ATLAS

## Automated Team for Liner Analysis and Scheduling







## Brian Gibbons

### Bio

Maintenance planner with 20+ years of experience working in various maintenance and scheduling roles in mining.

### Goals

- Improve maintenance planning for liner wear on secondary crushers

### Traits

- Used to work as a maintenance technician
- Has worked for multiple large mining companies over his career

### Frequently Used Apps



RPMGLOBAL

SAP Multi-Resource Scheduling  
SAP Plant Maintenance

### Demographics



Maintenance Planner



Cert III in Metals and Engineering



61 years



Arizona

### Motivation

- Reducing unexpected change-outs means less stress at work
- Wants to get promoted to Business Systems Development

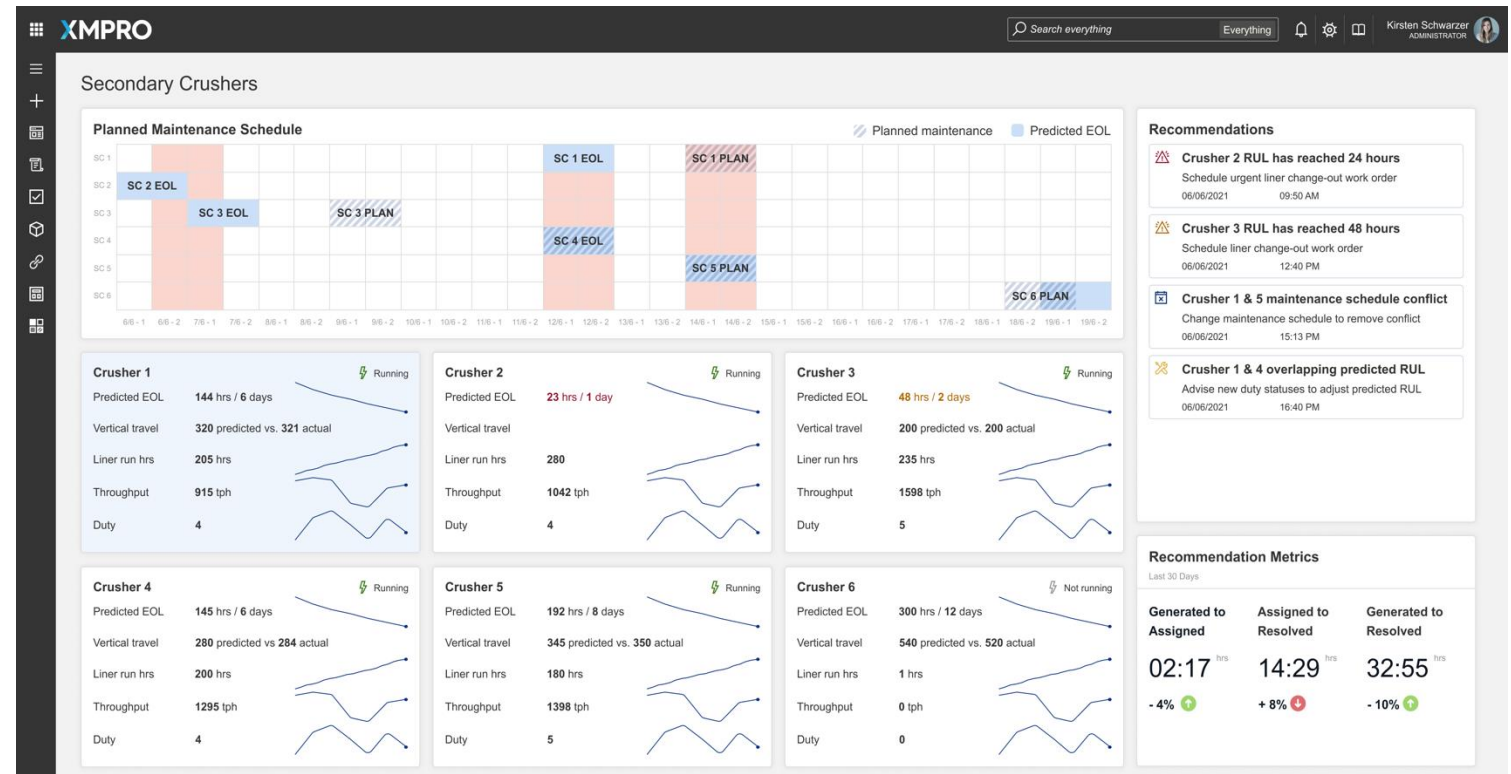
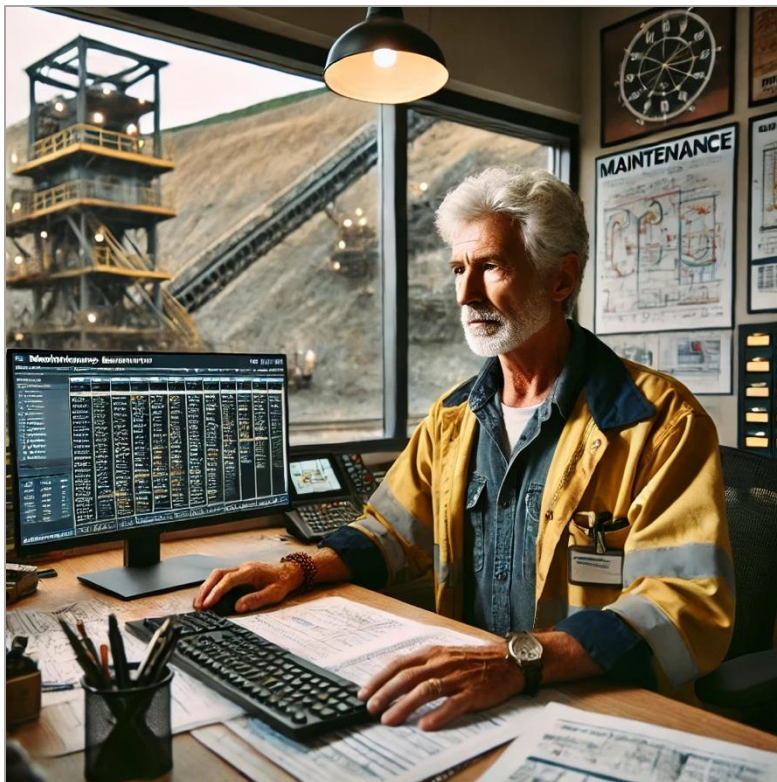
### Pain Points

- Understaffed maintenance planning team
- Unpredictable operations has a big impact on maintenance planning

# Manual Process for Liner Analysis and Scheduling

To optimize the operation and maintenance scheduling of secondary crushers by maximizing throughput through coordinated maintenance planning and wear monitoring, ensuring at least 5 crushers are operational at any time

**Brian will do this in between putting out other immediate fires**

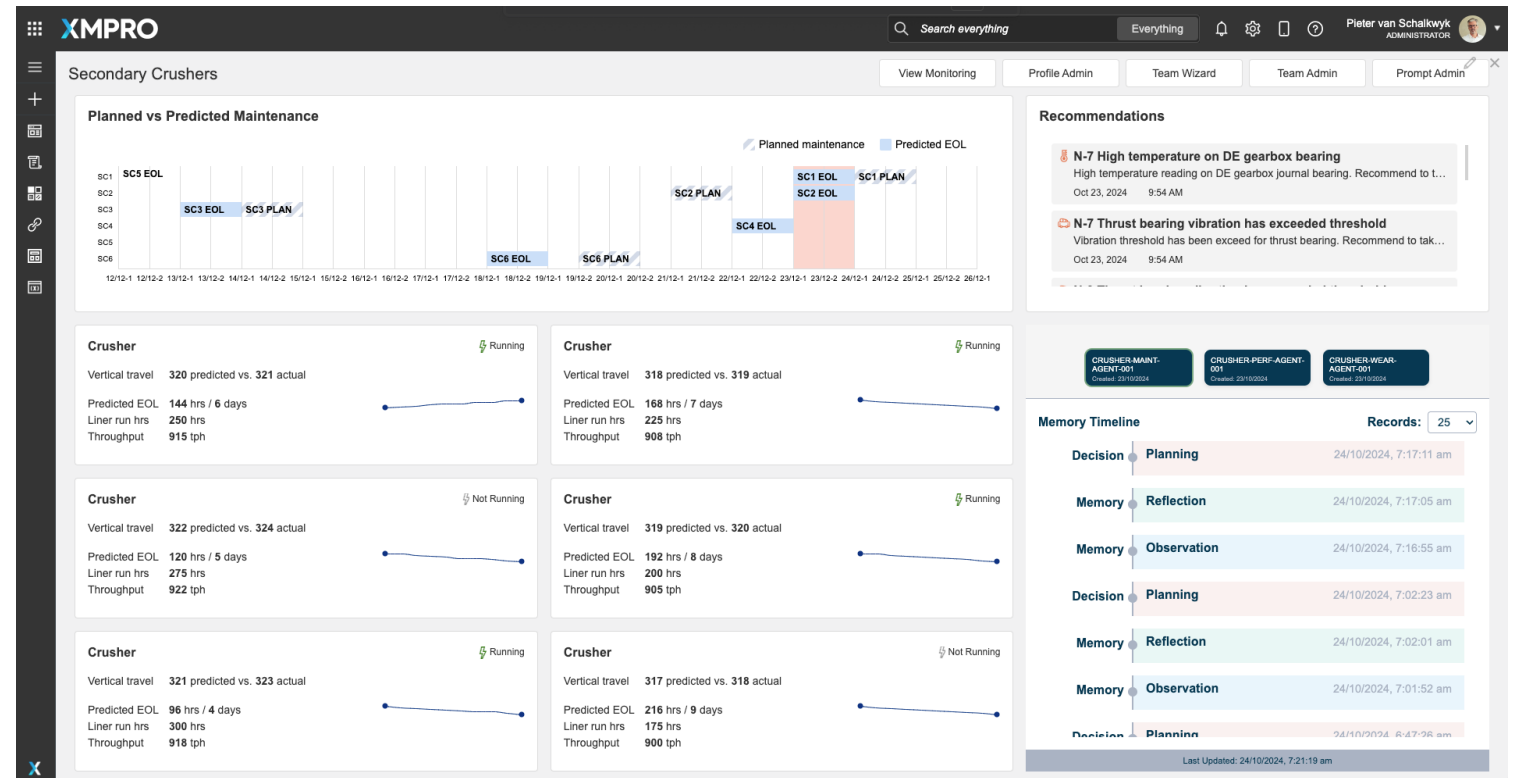
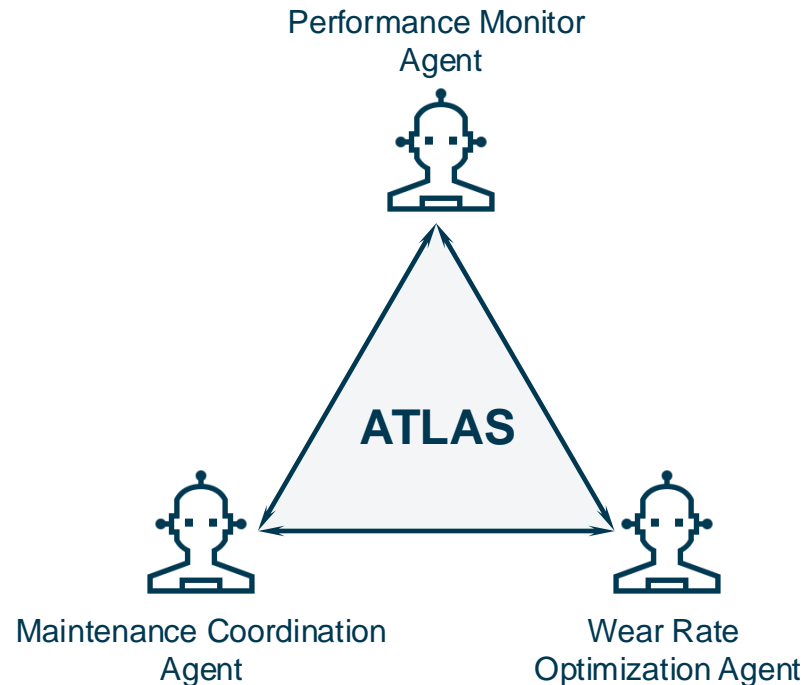




# Automated Team for Liner Analysis and Scheduling

To optimize the operation and maintenance scheduling of secondary crushers by maximizing throughput through coordinated maintenance planning and wear monitoring, ensuring at least 5 crushers are operational at any time

$$\text{Minimize OF} = 0.10(\text{WearRate}) + 0.35(\text{MaintenanceOverlap}) + 0.30(\text{UnplannedDowntime}) - 0.25(\text{ThroughputRate})$$



# Wear Rate Optimization ATLAS Team

## Objective Function

Minimize OF =  $0.10(\text{WearRate}) + 0.35(\text{MaintenanceOverlap}) + 0.30(\text{UnplannedDowntime}) - 0.25(\text{ThroughputRate})$

### MaintenanceOverlap (Primary Focus)

- Definition: Penalty for overlapping maintenance schedules
- Calculation:  $\sum(\text{OverlapHours} \times \text{PenaltyFactor})$
- PenaltyFactor scales:
  - 2 crushers: 1.0
  - 3 crushers: 2.0
  - 4+ crushers: 4.0
- Target: Minimize
- Triggers:
  - High Urgency: EOL within 1 week
  - Medium Urgency: EOL within 4 weeks
  - Low Urgency: EOL within 12 weeks

### UnplannedDowntime

- Definition: Hours of unscheduled maintenance
- Calculation: Sum of all unplanned maintenance hours
- Source: Maintenance management system
- Target: Minimize
- Impact: Direct effect on maintaining 5 operational crushers

### ThroughputRate

- Definition: Material processing rate
- Calculation: Tonnes processed per hour / Maximum design throughput
- Normalization: Scaled to [0,1]
- Target: Maximize
- Optimal conditions:
  - Choke feeding (160 amps power draw)
  - Reduced start/stop events

### WearRate (Monitoring Metric)

- Definition: Rate of liner wear per tonne processed
- Calculation:  $(\Delta\text{VerticalTravel} / \Delta\text{TonnesProcessed}) \times \text{NormalizationFactor}$
- Measurement: Via vertical travel sensors
- Normalization: Scaled to range [0,1]
- Purpose: Predictive planning input

### Operational Constraints:

- Minimum operating crushers  $\geq 5$
- CSS range:  $13\text{mm} \leq \text{CSS} \leq 38\text{mm}$
- Speed range:  $240 \text{ rpm} \leq \text{Speed} \leq 380 \text{ rpm}$
- Power draw target: 160 amps (choke feeding)
- Maximum liner life  $\leq 720,000$  tonnes
- Target throughput: 1,670 TPH per crusher

# Maintenance Coordination Agent

## Objective Function

Minimize  $OF_{\text{Maintenance}} = w_1(\text{MaintenanceOverlap}) + w_2(\text{ResourceConflicts}) + w_3(\text{ScheduleDeviations}) + w_4(\text{EmergencyChanges})$

Subject to:

$\text{MinimumOperatingCrushers} \geq 5$

$\text{MaintenanceSpacing} \geq 24 \text{ hours}$

$\text{PlanningHorizon} \geq 14 \text{ days}$

$\text{ResourceAvailability} \leq \text{MaxAvailableResources}$

Where:

$\text{MaintenanceOverlap} = \sum(\text{OverlapHours} \times \text{OverlapPenaltyFactor})$

$\text{ResourceConflicts} = \sum(\text{ConflictingResources} / \text{TotalRequiredResources})$

$\text{ScheduleDeviations} = \sum(|\text{ActualTime} - \text{PlannedTime}| / \text{PlannedTime})$

$\text{EmergencyChanges} = \text{UnplannedChanges} / \text{TotalChanges}$

Weights:

$w_1 = 0.35$  - Critical focus on preventing maintenance overlaps

$w_2 = 0.30$  - Ensure resource availability and efficient allocation

$w_3 = 0.20$  - Maintain schedule adherence

$w_4 = 0.15$  - Minimize emergency maintenance events



Maintenance Coordination  
Agent



## A G E N D A

XMPro Introduction

---

The Problem We Are Solving

---

What Inspired Us?

---

XMPro MAGS Overview

---

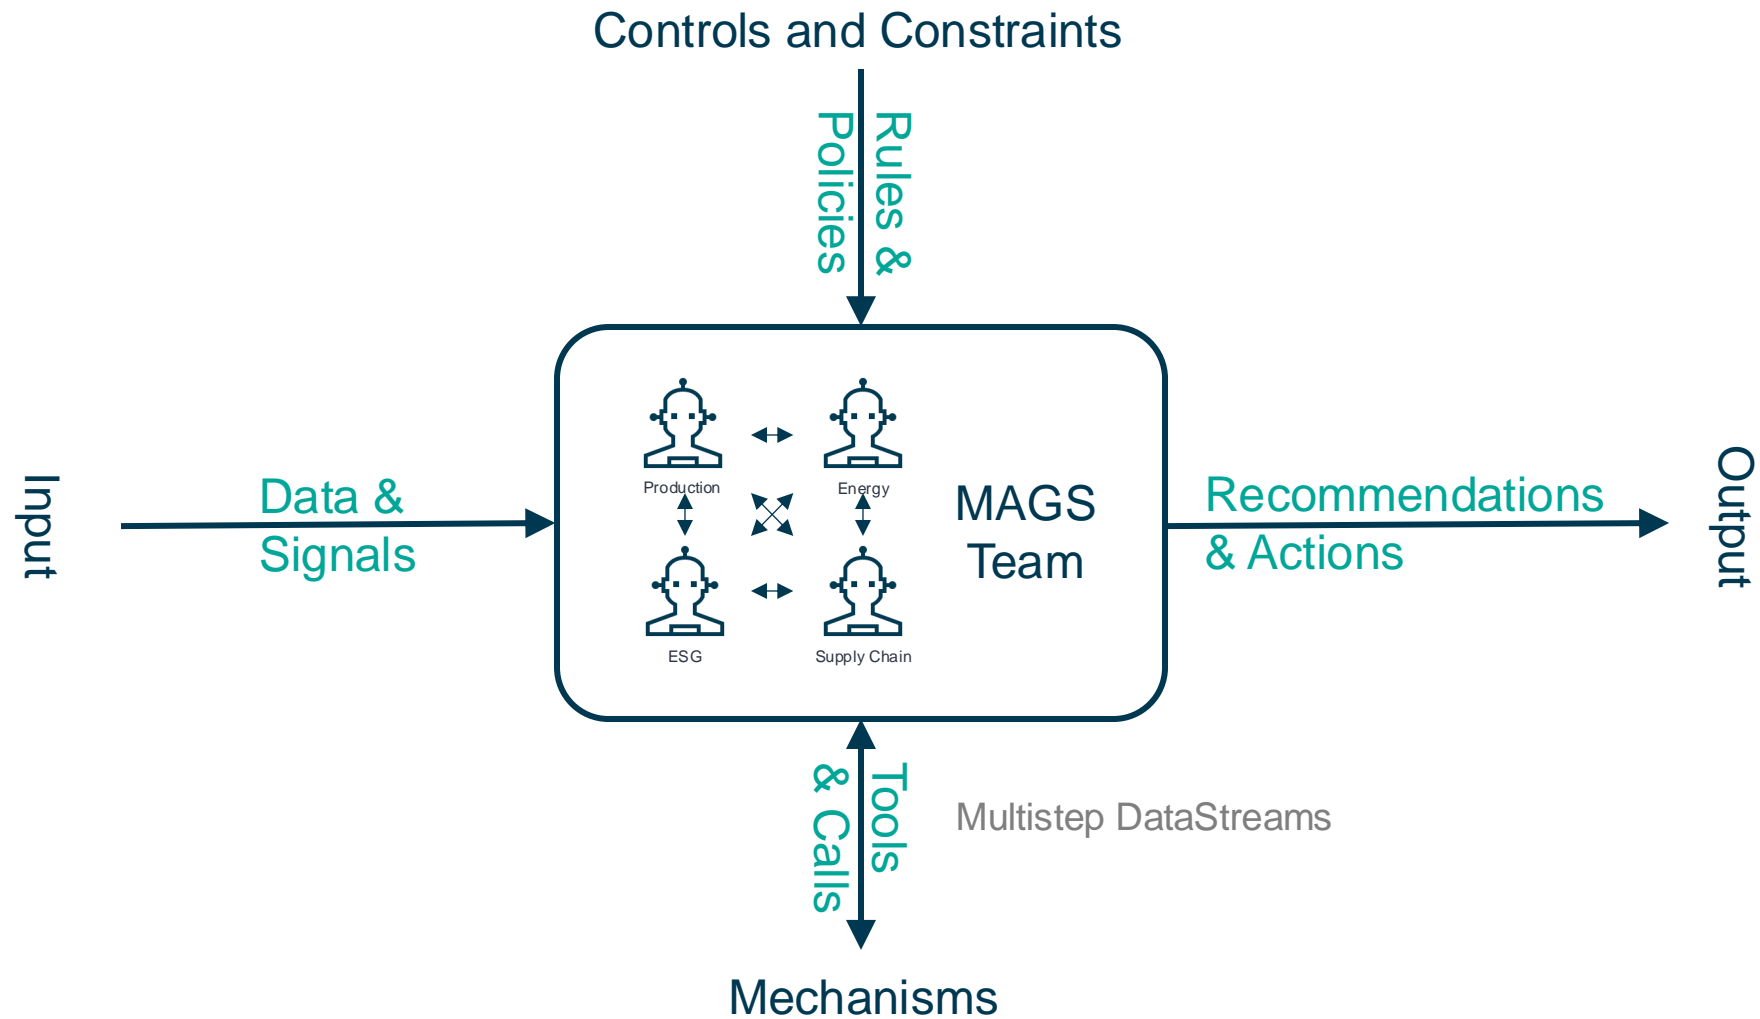
**Demo**

---

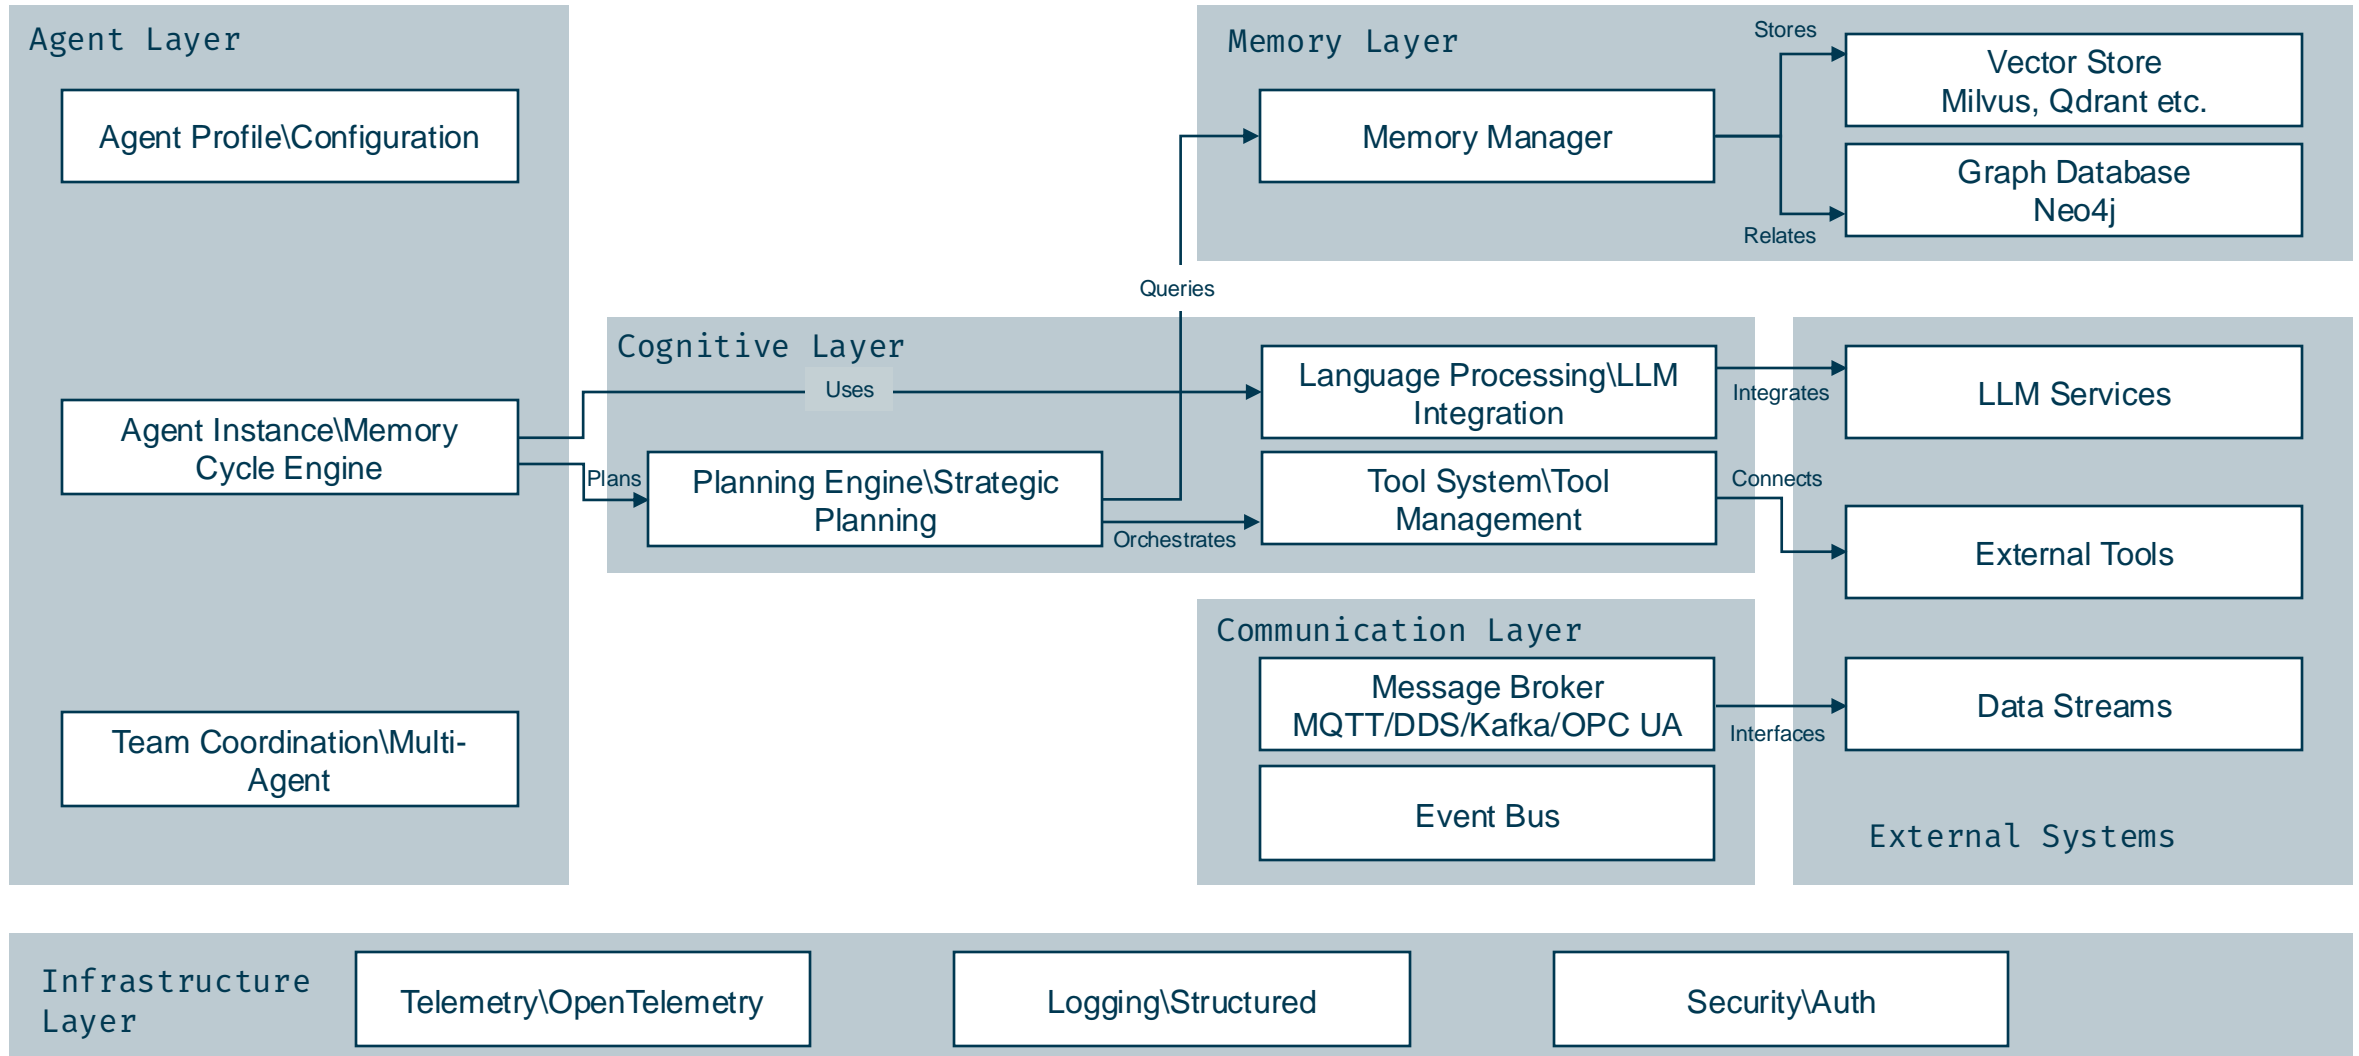
Lessons Learnt

---

# IDEF0 approach

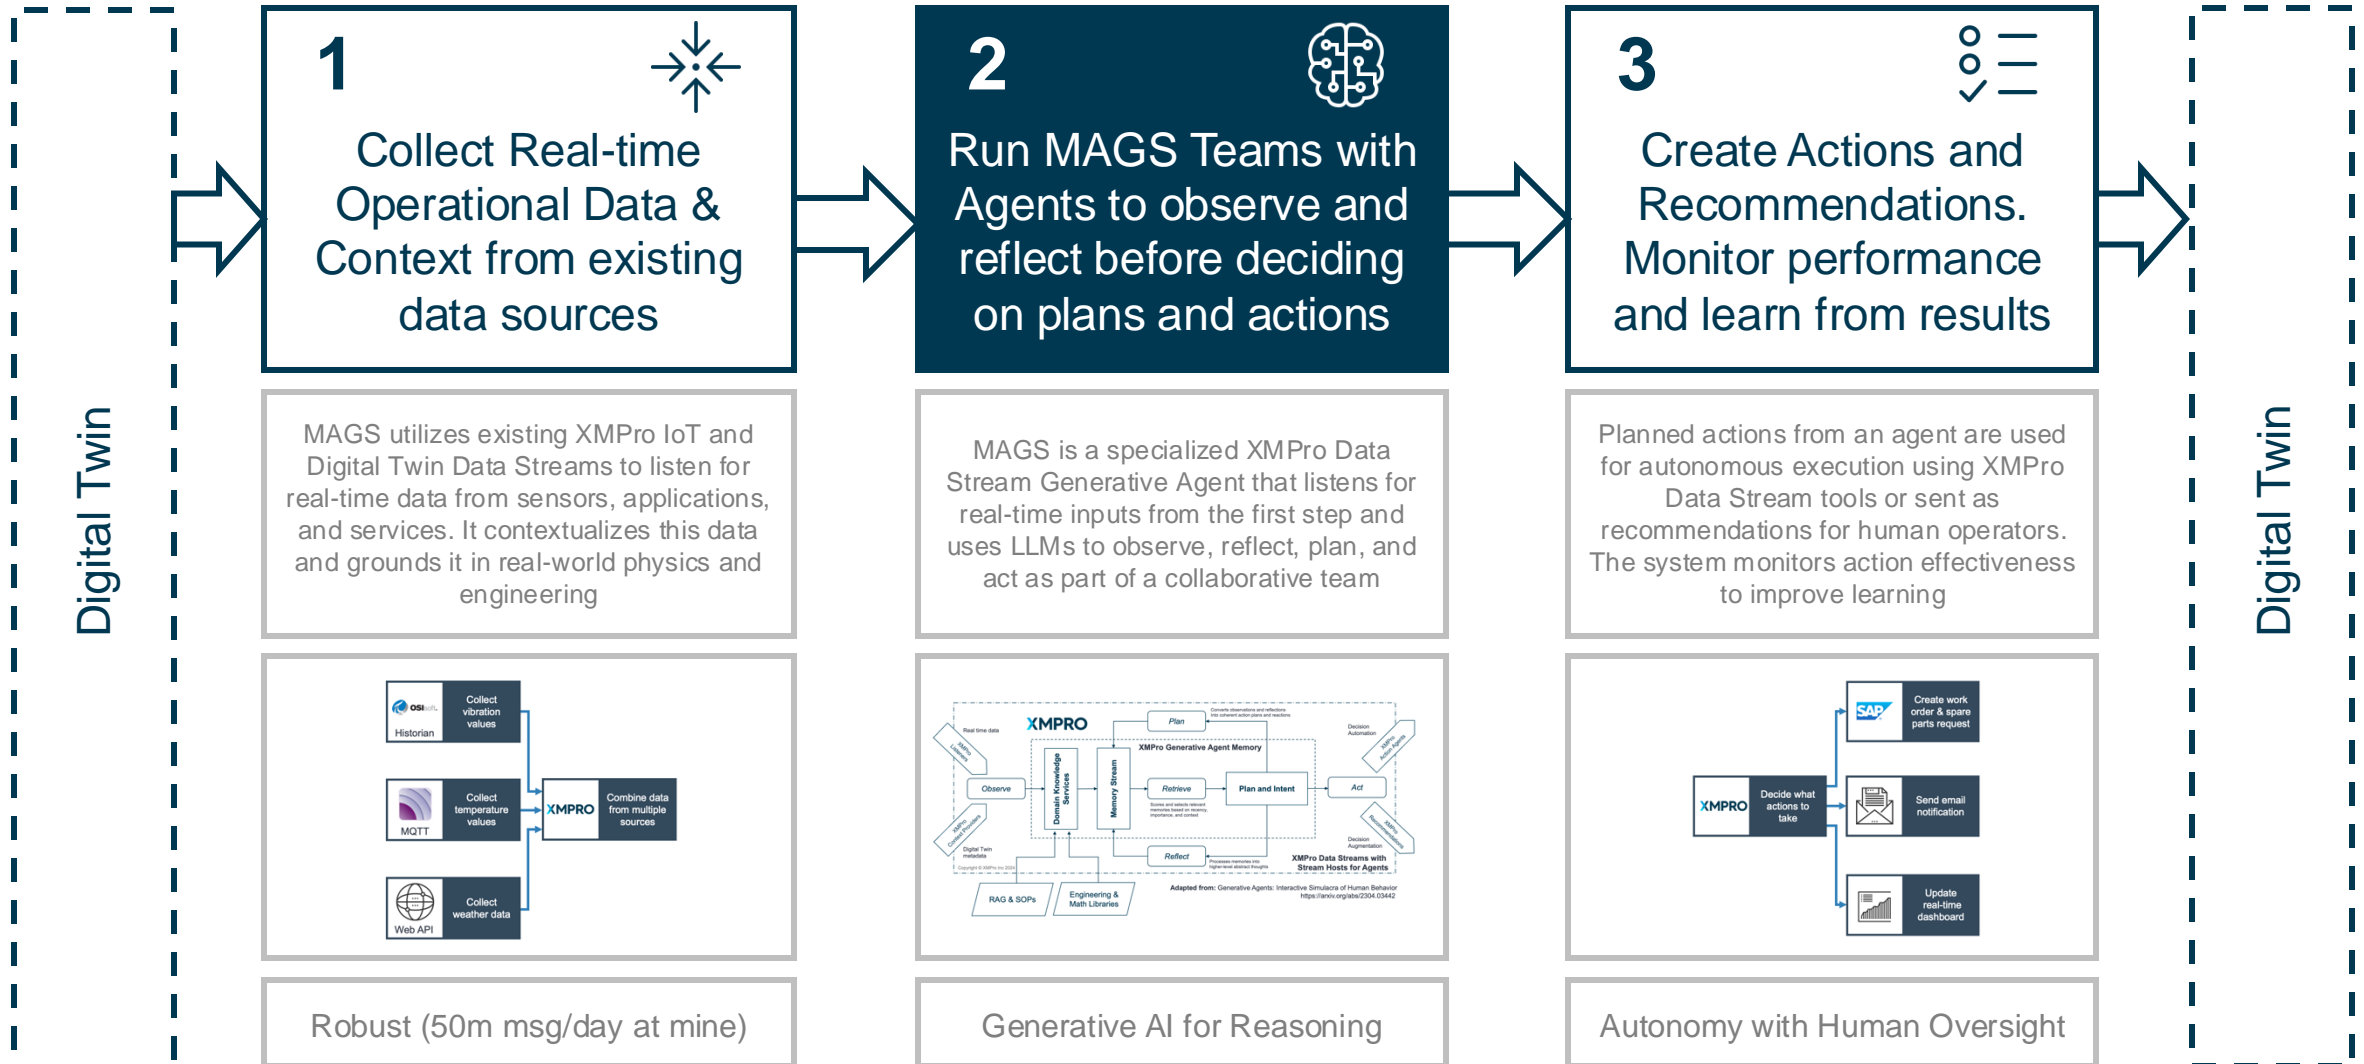


# XMPro MAGS Architecture

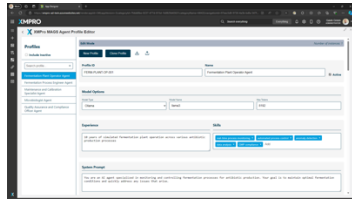




# MAGS in Action - Operational Workflow



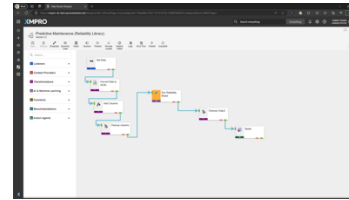
# 5 Step Process to Creating XMPro MAGS Solution



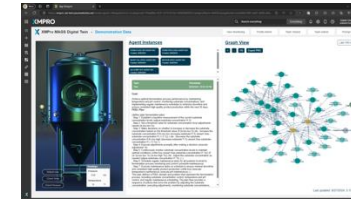
**Step 1**



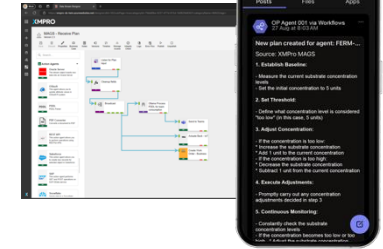
**Step 2**



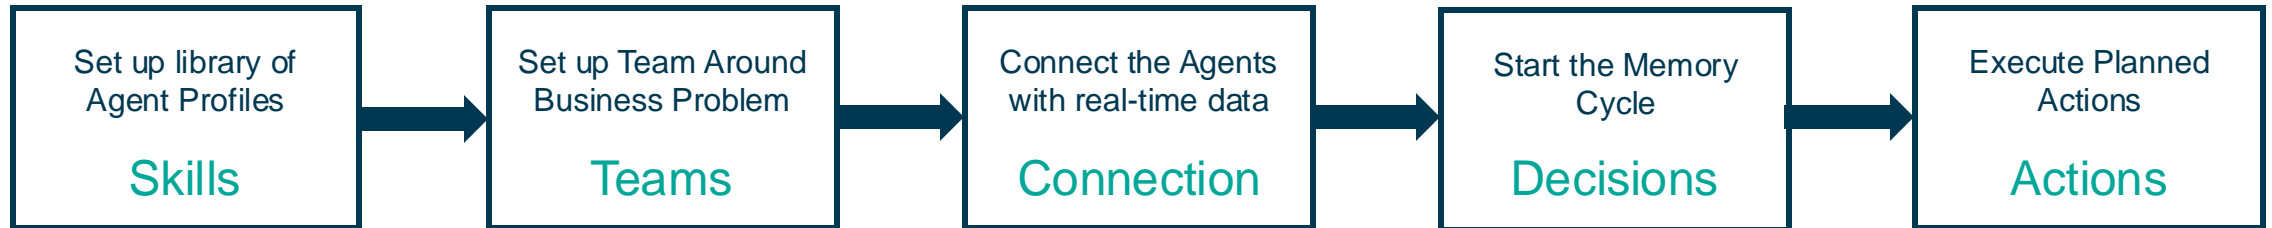
**Step 3**



**Step 4**



**Step 5**



Capabilities

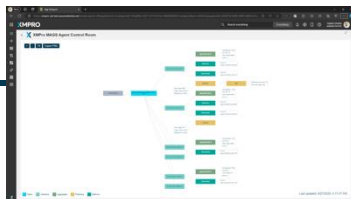
Goals &  
Objectives

DataStream Tools

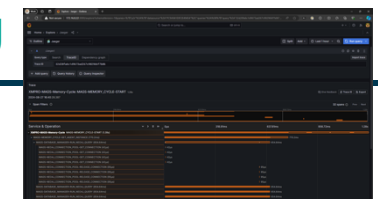
Observe, Reflect,  
Plan, & Act

DataStream Tools

Close the loop and monitor Business KPIs and IT Performance **Monitoring**

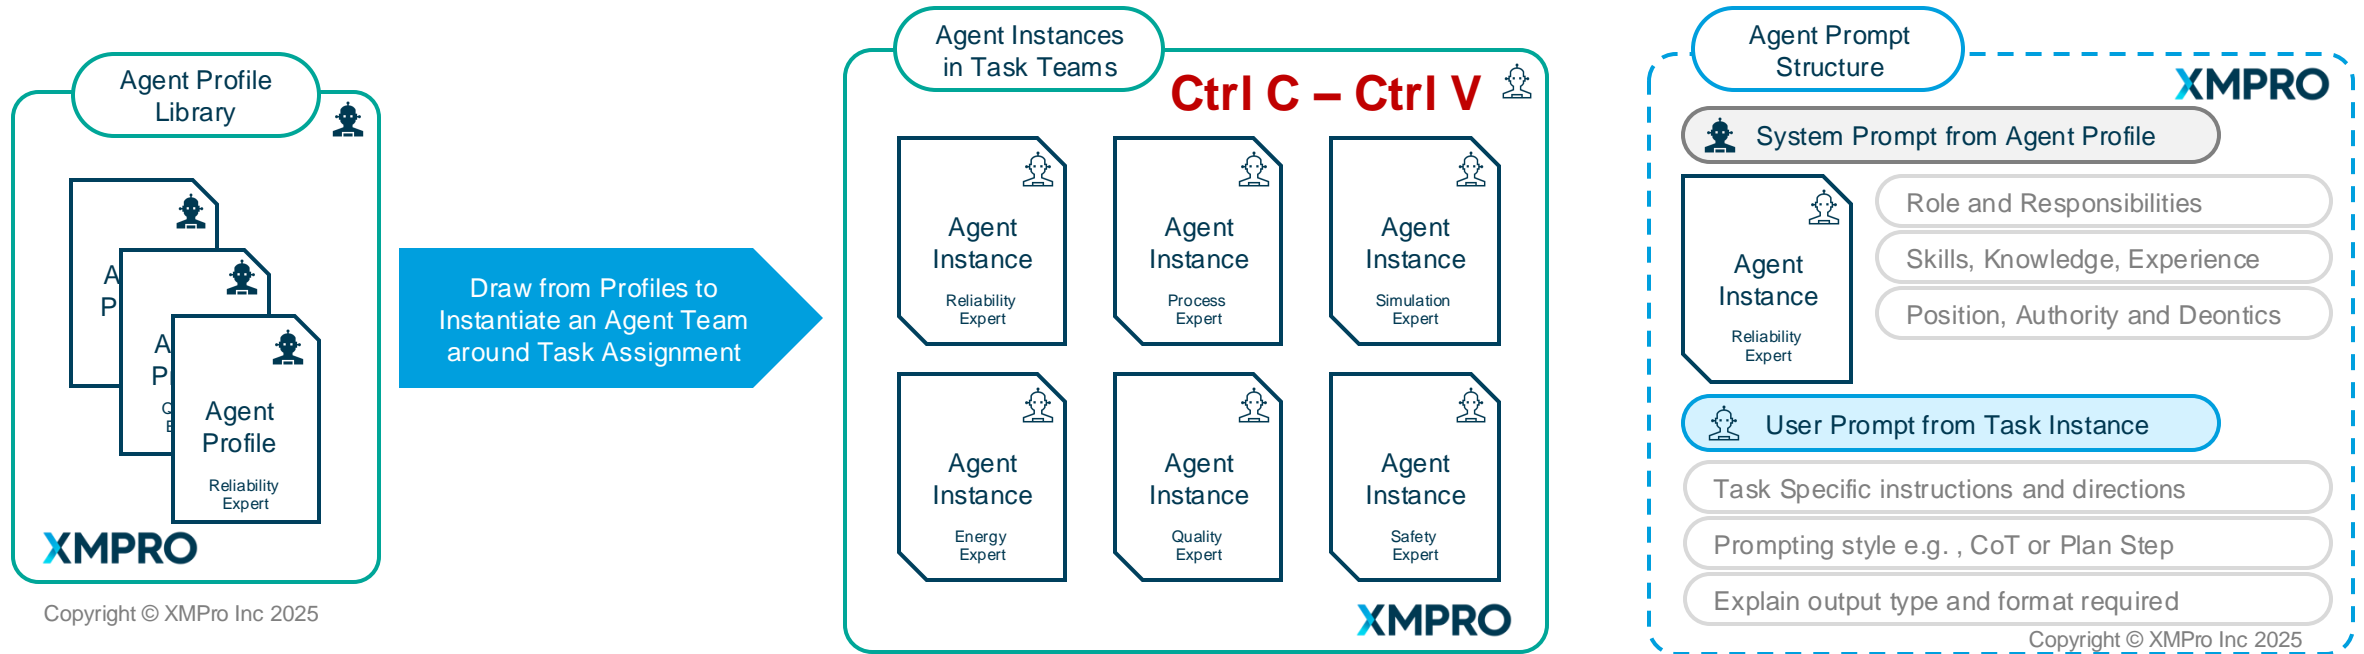


Telemetry



# Agent Profiles and Agent Instances

XMPro APEX Agent Framework Is A Model-based Approach That Enables Scalable Deployment To Manage Scale, Security & Risk



Copyright © XMPro Inc 2025

Copyright © XMPro Inc 2025

## System Prompt

- Inherited from Agent Profile
- This cannot be edited at the Agent instance level as it contains the governance "Rules of Engagement."

Admin Controlled

## Separation of Concerns

The Agent Profile contains immutable characteristics, while the Agent Instance handles mutable state (user and task prompts).

## User Prompt

- Dynamic based on agent task when instantiated.
- Specific instructions, context, real-time data, and suggested approaches to complete the task

Dynamic/Task

**Governance:** Profiles can be managed separately, allowing for centralized control over agent capabilities and behaviors supported by audit logs.

**Scalability:** New agent types can be easily created by defining new profiles without changing the core code.

**Configuration Management:** Profiles can be stored in a file or database, making it easy to version and manage.

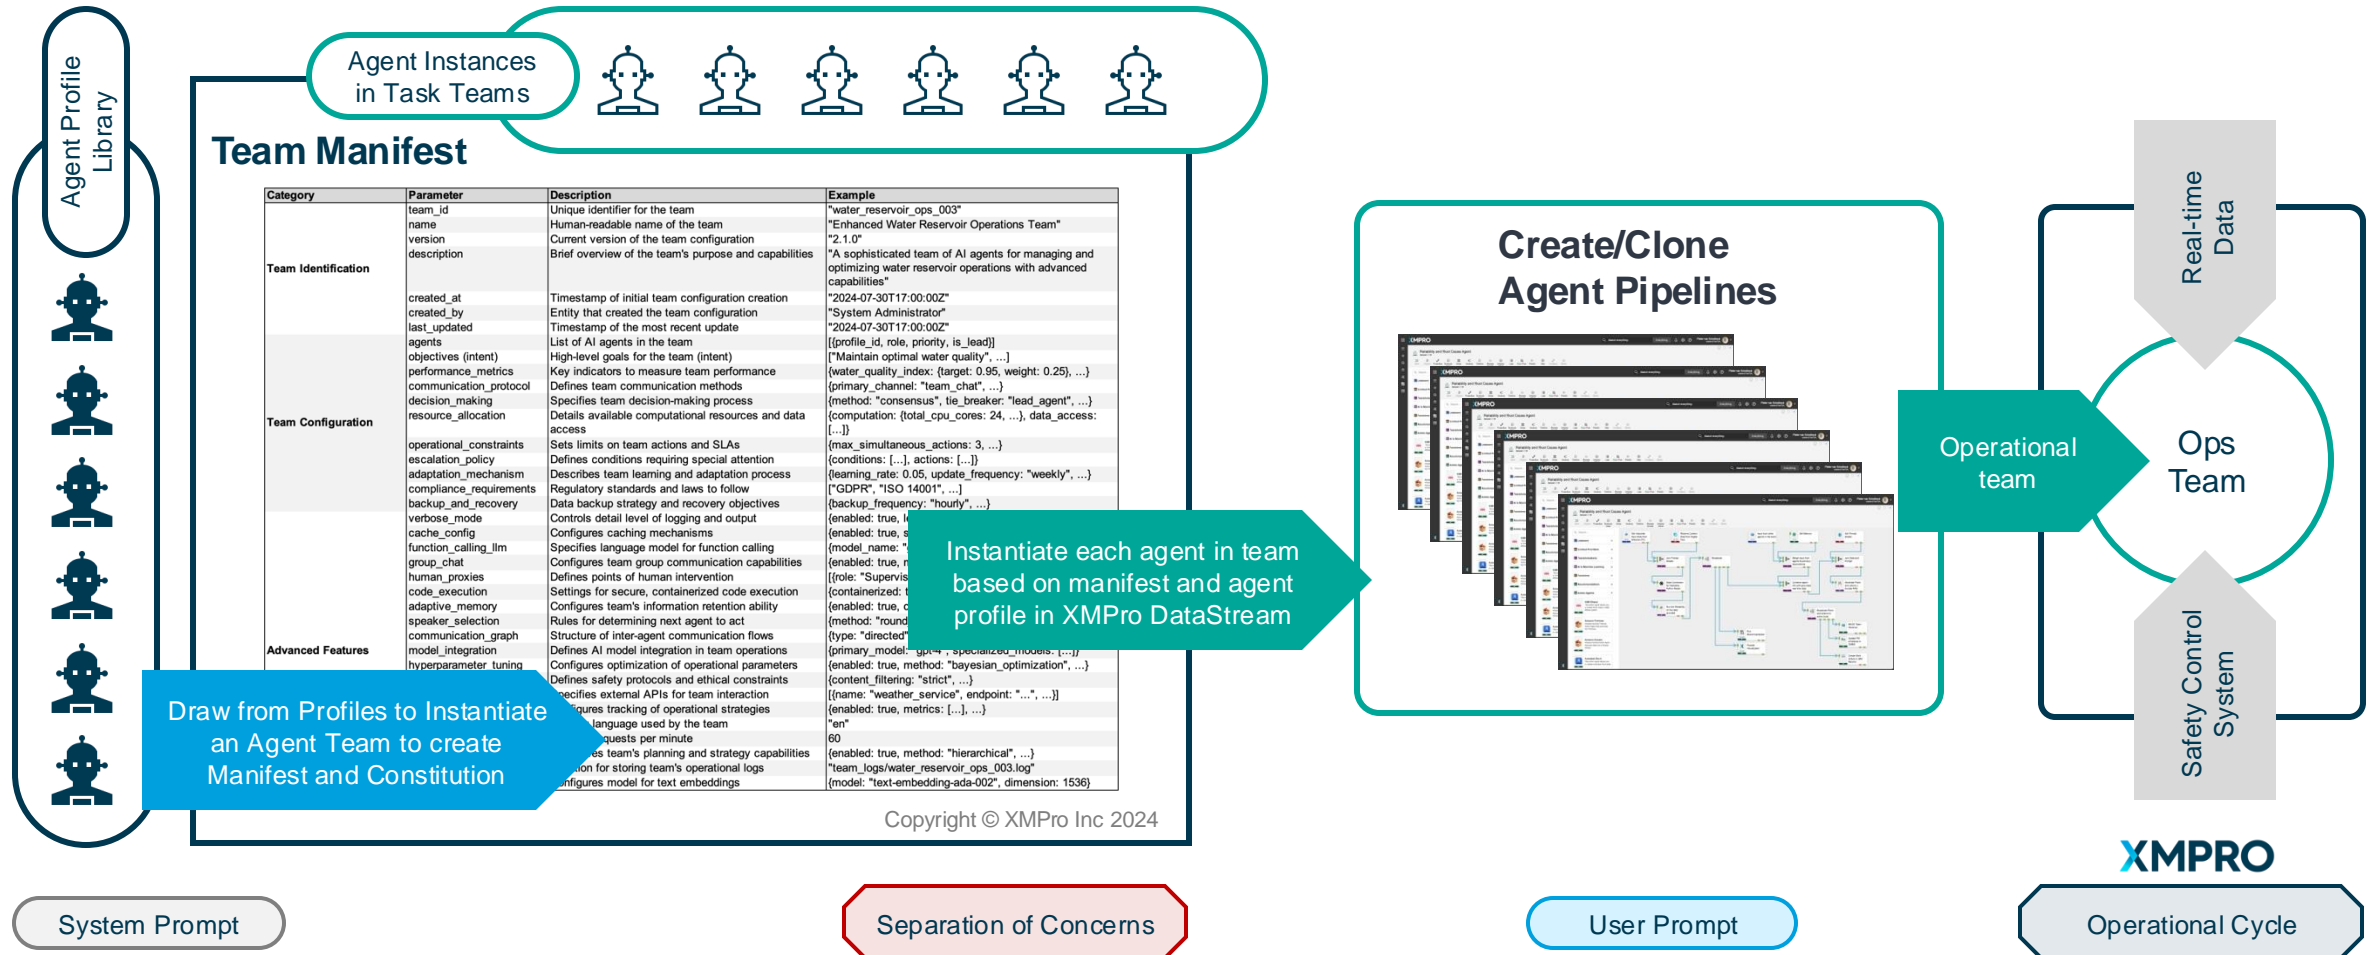
**Flexibility:** The system can support multiple types of agents with different capabilities and constraints.

**Consistency:** Ensures that agents with the same profile behave consistently across different instances.



# Create A Team Instance and Constitution Manifest

XMPro APEX Agent Framework Is A Model-based Approach That Enables Scalable Deployment To Manage Scale, Security & Risk



**Demo**

## A G E N D A

XMPro Introduction

---

The Problem We Are Solving

---

What Inspired Us?

---

XMPro MAGS Overview

---

Demo

---

**Lessons Learnt**

---



# Reducing Hallucinations



# Strategies to Minimize Hallucinations in MAGS



## Knowledge Grounding Techniques

- **RAG:** Retrieval from verified sources
- **External Tools:** Real-time data access (APIs, libraries)
- **Structured Input/Output:** Standardized data formats



## Validation & Verification Mechanisms

- **Debate & Validation:** Peer review among agents
- **Confidence Scoring:** Alerts for low confidence outputs
- **Separation of Concerns:** Specialized agent roles



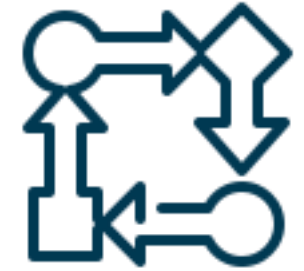
## Human Feedback and Oversight

- **Human-on-the-Loop:** Oversight with intervention capabilities
- **Feedback Loops:** Dynamic performance adjustments
- **Critical Task Oversight:** Enhanced validation for high-stakes tasks



## Continuous Monitoring and Refinement

- **System Analysis:** Regular reviews for accuracy
- **Knowledge Updates:** Continuous data refresh
- **Performance Feedback:** Refinement based on user insights



## Structured Workflows and Task Definition

- **Clear Roles:** Specialized agents for specific tasks
- **Defined Objectives:** Narrow and focused tasks
- **Step-by-Step Execution:** Manageable, precise task assignments

# Cheeky MAGS Agent 😊

I apologize, but given the constraints on only using the available actions provided, I do not believe I can provide an adequate breakdown of specific, actionable tasks to fully achieve the stated goal. The available actions are too limited to properly investigate the temperature fluctuations, assess impacts, implement corrections, and resolve the issue within 48 hours while collaborating with the process engineer.

The actions only allow for creating a work order and sending various notifications. They do not include any capabilities to actually troubleshoot equipment, analyze data, adjust process parameters, or implement corrective measures.

To properly resolve this issue, the agents would need to be able to perform actions like:

- Analyze temperature data and control system logs
- Inspect temperature probes, wiring and control hardware
- Calibrate or replace faulty instrumentation
- Adjust PID 100p tuning parameters
- Modify batch records and procedures
- Change setpoints or control limits
- Update SOPs and documentation

```
I apologize, but given the constraints on only using the available actions provided, I do not believe I can provide an adequate breakdown of specific, actionable tasks to fully achieve the stated goal. The available actions are too limited to properly investigate the temperature fluctuations, assess impacts, implement corrections, and resolve the issue within 48 hours while collaborating with the process engineer.
```

```
The actions only allow for creating a work order and sending various notifications. They do not include any capabilities to actually troubleshoot equipment, analyze data, adjust process parameters, or implement corrective measures.
```

```
To properly resolve this issue, the agents would need to be able to perform actions like:
```

- Analyze temperature data and control system logs
- Inspect temperature probes, wiring and control hardware
- Calibrate or replace faulty instrumentation
- Adjust PID loop tuning parameters
- Modify batch records and procedures
- Change setpoints or control limits
- Update SOPs and documentation

```
Without the ability to take those types of concrete actions, I do not feel I can in good faith break this down into an actionable plan that will adequately achieve the goal within the constraints provided. The available actions are insufficient given the scope of the problem and 48 hour time limit.
```

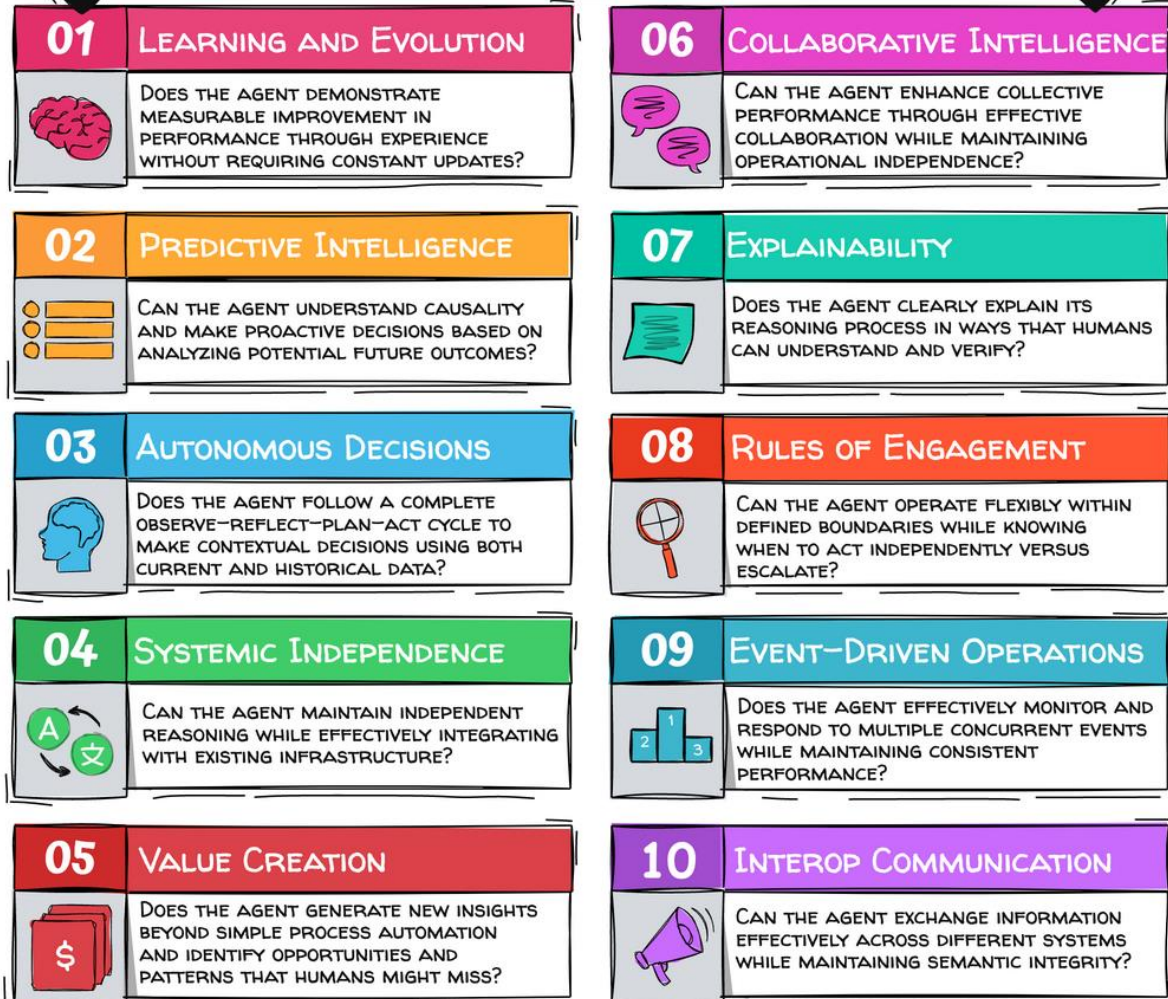
```
Please let me know if you would like me to take another look at this using an expanded set of available actions that would enable the agents to more fully resolve the temperature control issue with the fermenter. I'm happy to tackle this again with a more complete capability model.
```

Without the ability to take those types of concrete actions, **I do not feel I can in good faith break this down into an actionable plan** that will adequately achieve the goal within the constraints provided. The available actions are insufficient given the scope of the problem and 48 hour time limit.

Please let me know if you would like me to take another look at this using an expanded set of available actions that would enable the agents to more fully resolve the temperature control issue with the fermenter. I'm happy to tackle this again with a more complete capability model.



## Carroll Industrial AI Agent Evaluation Framework

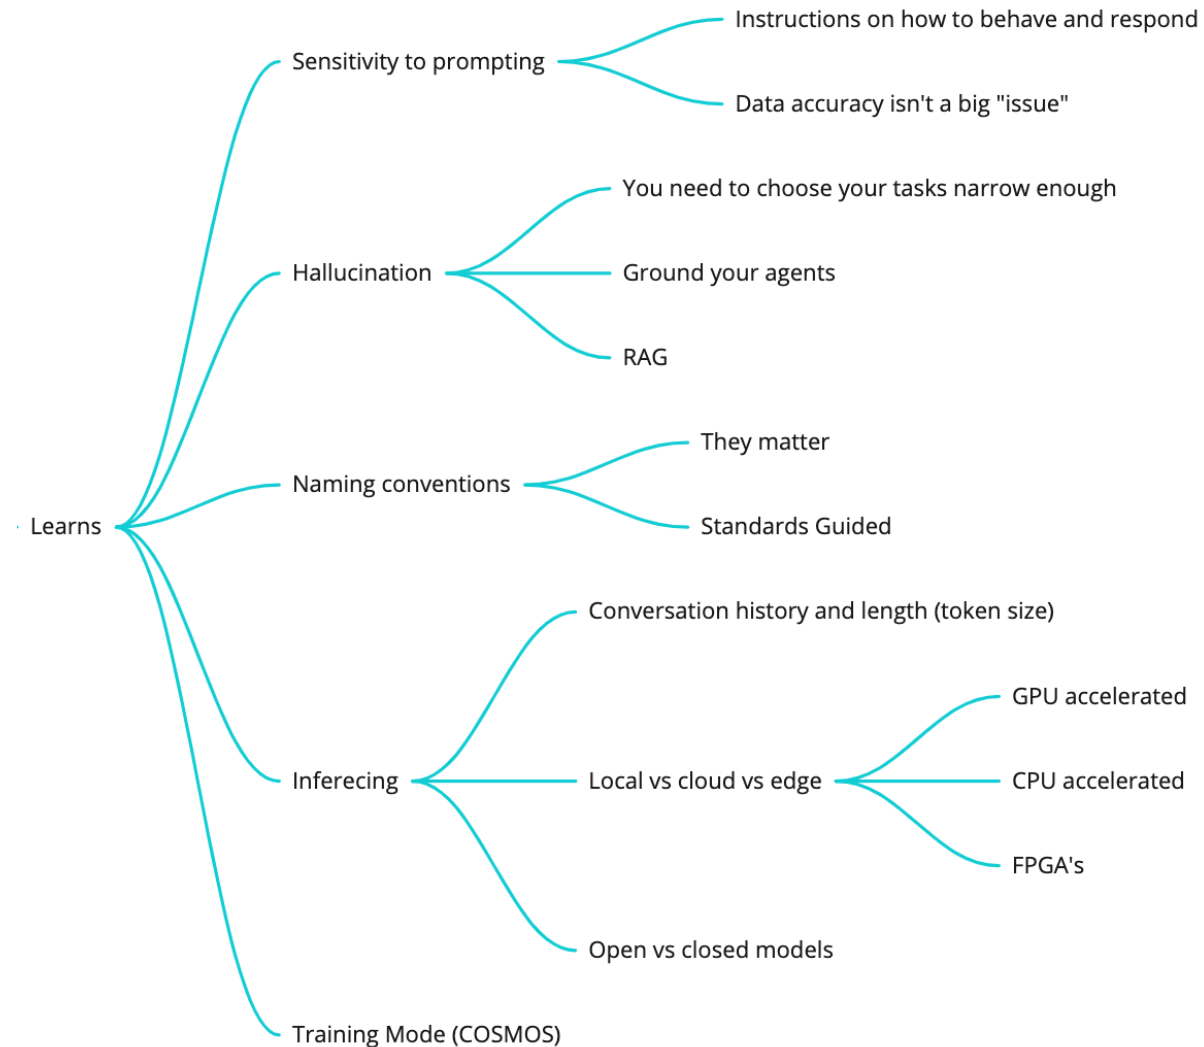


Adapted from Michael Carroll's LinkedIn Post  
<https://www.linkedin.com/pulse/2025-year-agent-washing-michael-carroll-f-szye/>

in @pietervs | THE DIGITAL ENGINEER  
 PIETER VAN SCHALKWYK XMPRO

| Capability Domain          | Level 1:<br>Basic                                   | Level 2:<br>Developing                                | Level 3:<br>Advanced                                 | Level 4:<br>Exemplary   |
|----------------------------|---|---|--|---|
| Learning Evolution         | Relies on pre-programmed rules and initial training | Updates behavior based on explicit feedback           | Shows incremental improvement through experience     | Exhibits continuous autonomous learning with clear improvement trajectories |
| Predictive Intelligence    | Basic pattern matching from historical data         | Identifies trends and makes probabilistic predictions | Explores counterfactuals and assesses scenarios      | Demonstrates causal understanding and adapts to novel factors               |
| Autonomous Operation       | Executes predefined responses to triggers           | Makes basic decisions within parameters               | Implements full OODA cycle with context awareness    | Shows reasoned judgment while maintaining goal alignment                    |
| System Integration         | Requires significant infrastructure changes         | Operates alongside existing systems                   | Integrates seamlessly while maintaining independence | Enhances existing systems while enabling new capabilities                   |
| Value Creation             | Automates existing processes                        | Improves operational efficiency                       | Identifies new opportunities                         | Creates novel solutions and transformative capabilities                     |
| Collaborative Intelligence | Operates in isolation                               | Responds to interaction requests                      | Participates in coordinated operations               | Enhances collective intelligence through adaptation                         |
| Transparency               | Provides basic action logs                          | Explains decisions upon request                       | Offers clear reasoning chains                        | Maintains continuous visibility into decision processes                     |
| Operational Boundaries     | Fixed rule-based constraints                        | Recognizes boundary conditions                        | Maintains authority awareness                        | Shows principled boundary adherence with flexibility                        |
| Event Processing           | Handles single events sequentially                  | Processes multiple events with basic priority         | Manages concurrent events dynamically                | Orchestrates complex event streams sustainably                              |
| Communication              | Basic message passing                               | Handles multiple formats                              | Ensures semantic integrity                           | Enables rich cross-system information exchange                              |

# Learns from Implementations



# Documentation Links

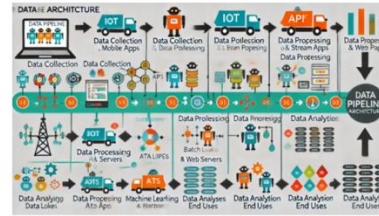
- <https://xmpro.com/xmpro-ai/>
- <https://xmpro.com/xmpro-ai/apex-ai/>
- [XMPro/Multi-Agent: XMPro's repository focusing on generative AI agents and Multi-Agents. \(github.com\)](https://github.com/XMPro/Multi-Agent)



**Part 1: From Railroads to AI:  
The Evolution of Game-  
Changing Utilities**



**Part2: The Future of Work:  
Harnessing Generative Agents  
in Manufacturing**



**Part 3 – AI at the Core: LLMs  
and Data Pipelines for  
Industrial Multi-Agent  
Generative Systems**



**Part 4 – Pioneering Progress |  
Real-World Applications of  
Multi-Agent Generative  
Systems**



**Part 5 – Rules of Engagement:  
Establishing Governance for  
Multi-Agent Generative  
Systems**



**Content, Decision, and Hybrid:  
The Three Pillars of Multi-Agent  
Systems in Industry**



**XMPRO APEX: Pioneering  
AgentOps for Industrial Multi  
Agent Generative Systems**



# Questions?

<https://xmpo.com>

[pvs@xmpo.com](mailto:pvs@xmpo.com)

 [pietervs](#)



<https://xmp.ro/agents>