# SAPINSIDER SUMMITSERIES Technology & Innovation Architecture Connection of the Conne

October 24-25, 2024 | Dallas

# Al Transformation of your Business

Frank Stienhans
SVP Solution Architecture

Babu Ayyalu
SAP on Cloud Solutions Architect

lemongrass

SAP on Cloud: Hyperscale. Hyperspeed. Hypercare.

# We are Lemongrass

Digitally enabled SAP on any Cloud















SAP on Cloud Roadmap



Cloud Patterns,
Architecture and Skills









Automated Service
Catalog with
Innovation-Centric Culture



SAP Data Analytics Accelerator



Lemongrass Cloud Platform (LCP)



© Lemongrass. All Rights Reserved.

# Representative Clients

Our client roster is made up of market leaders across a variety of industries including Consumer Products, Energy, Finance, Manufacturing, Agriculture, and more.

### Representative clients include:































































# What We'll Cover

- Al Transformation **Technology**
- Al Transformation **Process**
- Al Transformation People
- Al Transformation Solution

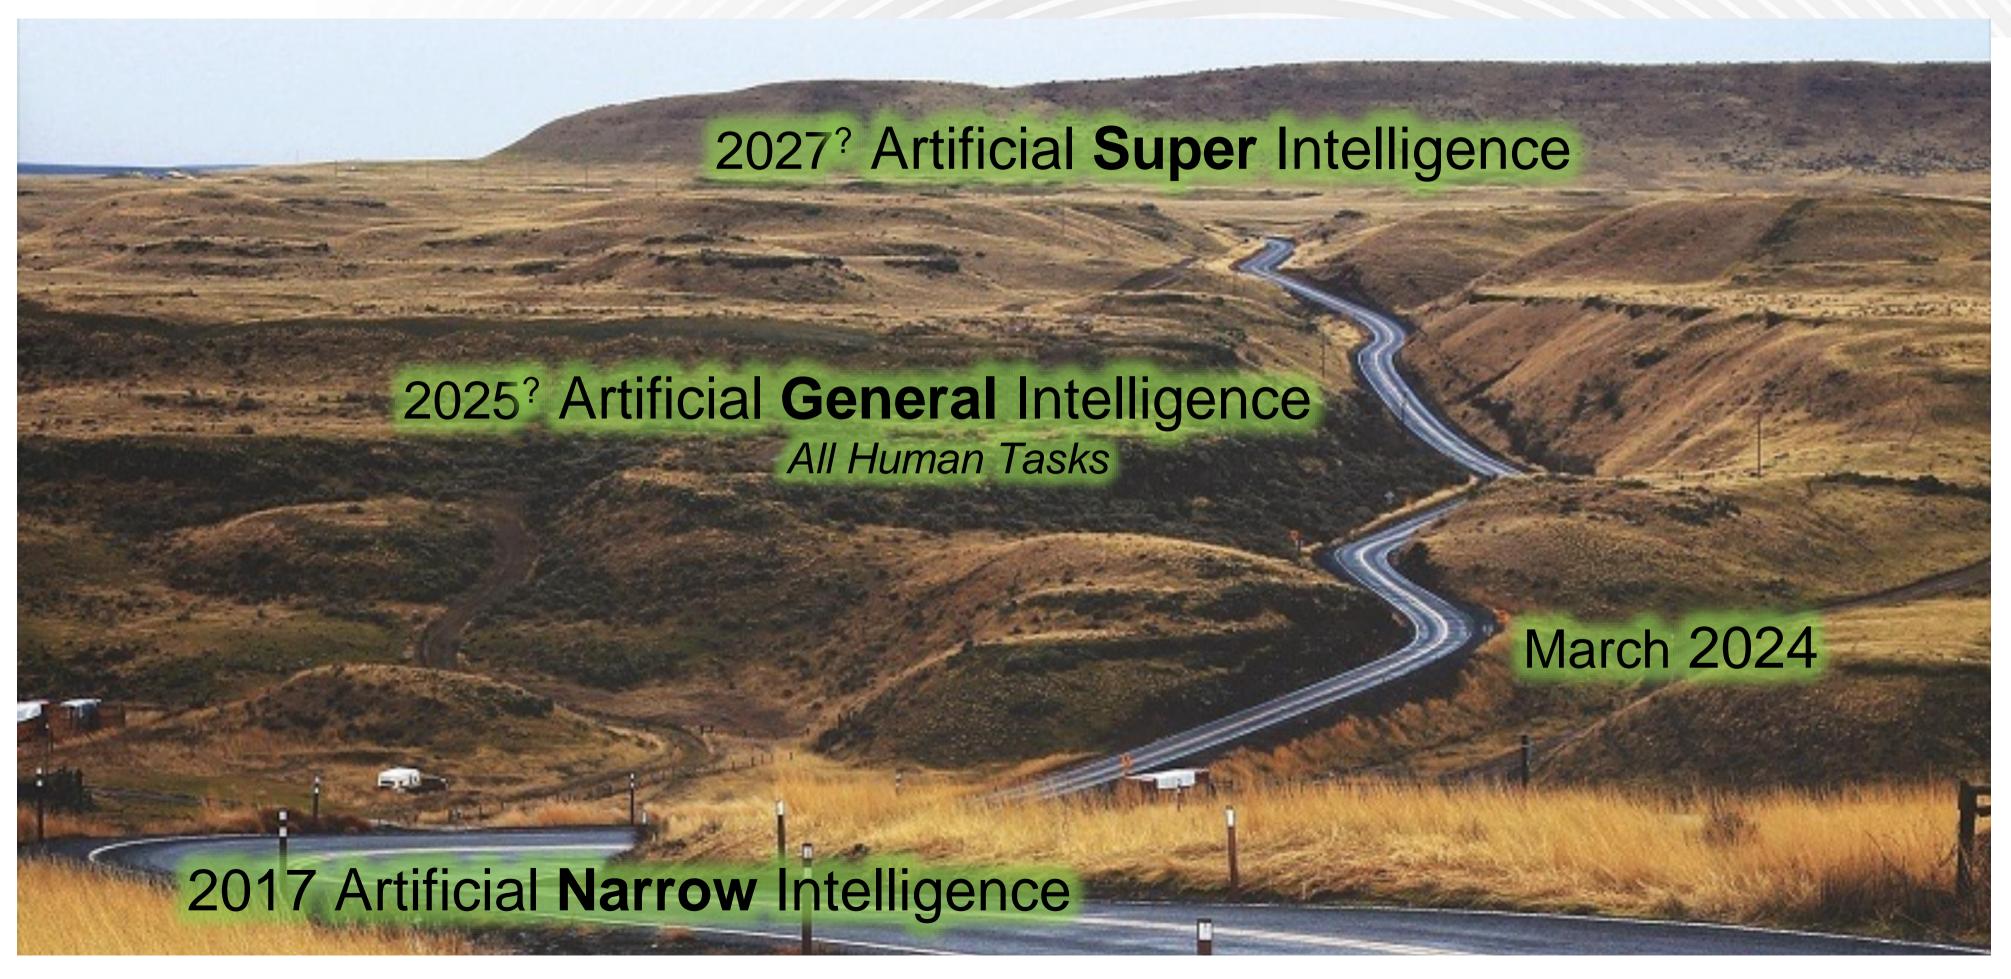
# Al Transformation - Technology

# Last 12 months as of May 2024

Cost	Input	Intelligence	Performance	Provid	ders	Hardware
6x reduction	4x increase	50% increase	2x increase	5x increase		2.5x throughput
\$ 30 <b>→</b> \$ 5	32K → 128K			OpenAI →	OpenAl Google Anthropic Meta	Nvidia Hopper → Blackwell
GPT 4 (March 2023) → GPT 4o (May 2024)					X.ai	

# Next 12 months as of May 2024

Cost	Input	Intelligence	Performance	Providers	Hardware
Increasing Competition	16x increase 128K → 2M	GPT 5	Faster Inference Hardware	OpenAl Google Anthropic Meta X.ai  Microsoft Amazon	Nvidia AMD Intel Google Amazon Microsoft



Definitions: http://www.eitc.org/research-opportunities/new-media-and-new-digital-economy/ai-machine-learning-deep-learning-and-neural-networks/ai-research-and-applications/stages-of-ai

# **Artificial General Intelligence**

Governance

Self-Govern

Memory **Self-Managing** Knowledge **Self-Learning Data Sources Self-Expanding** Tools **Self-Extending** Pipelines / Workflows **Self-Defining Foundation Models** 

Transactional, Short-Term

Long-Term, Archives, Right to forget?

Private & Public, Structured & Unstructured, Realtime & Historic

Web, Services, ... Capabilities

e.g. DevOps, SAP

GPT, Claude, Gemini, ...

### Demo Assistant Scenario II

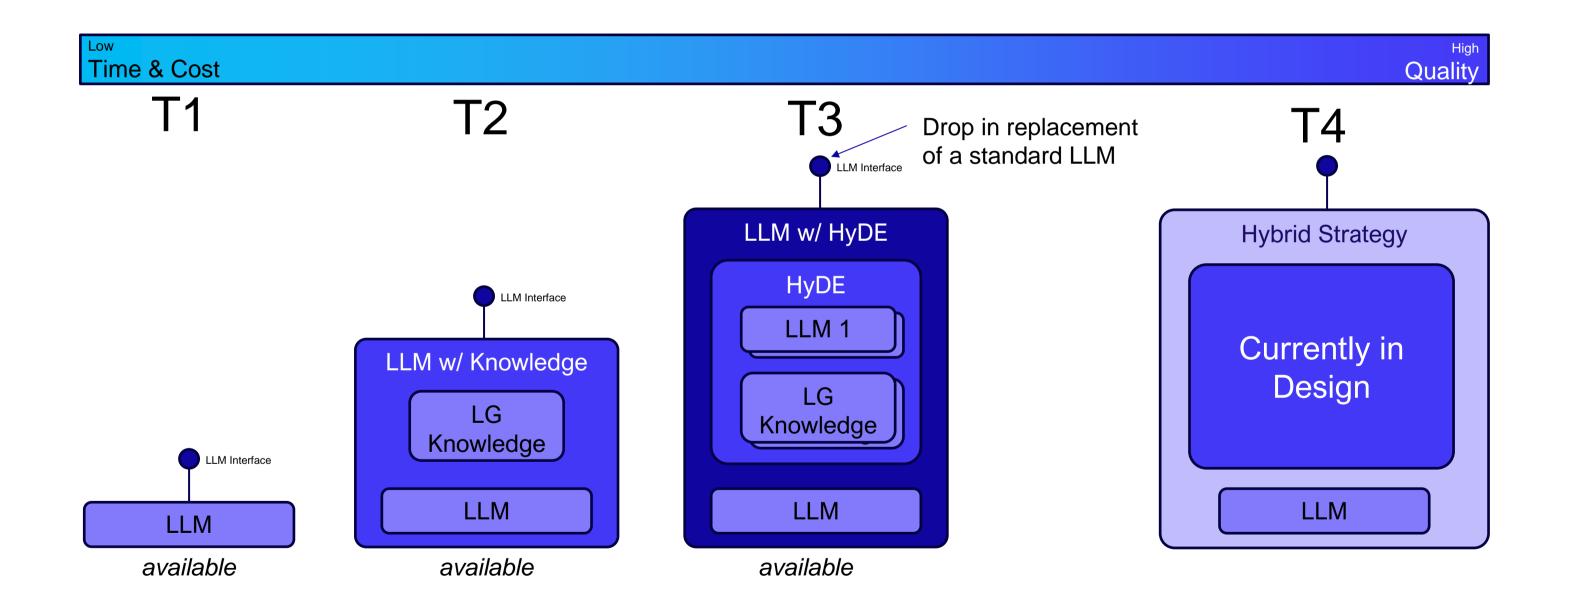
# Lemongrass Internal Al Knowledge August 2024

Knowledge Base					
Bundle	Files	Pages	Tokens		
	3	5,258	2,139,382		
	3	6,359	2,265,264		
	35	239	326,988		
	2,522	31,819	19,865,533		
	17,012	179,311	106,246,583		
	42	571	153,418		
	5	3,491	1,735,521		
Diverse	15	2,361	1,194,215		
	2,504	29,919	13,210,437		
Knowledge	6	1,718	691,708		
	21,748	21,748	561,015		
	151	402	1,407,964		
	3,950	12,761	45,622,012		
	62,567	64,269	37,482,710		
	9,227	9,335	11,379,146		
	8	9	15,130		
	4,851	4,974	5,473,104		
	6,980	7,114	8,303,227		
J	1,927	2,418	8,584,750		
TOTAL	133,556	384,076	266,658,107		

Infrastructure Footprint: 3.5 GB RAM

Knowledge Lookup: 0.004 to 0.3 seconds

### **Composite Large Language Models**

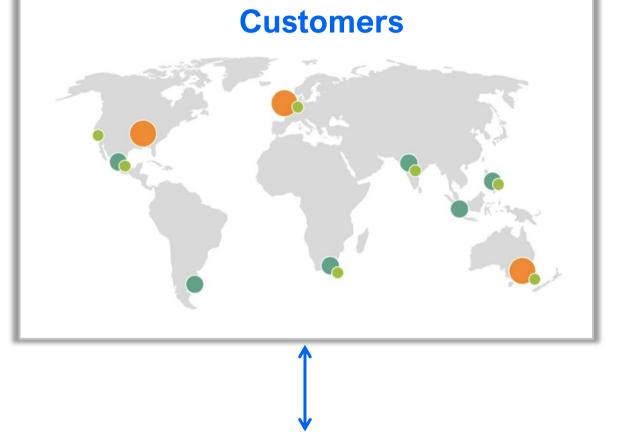


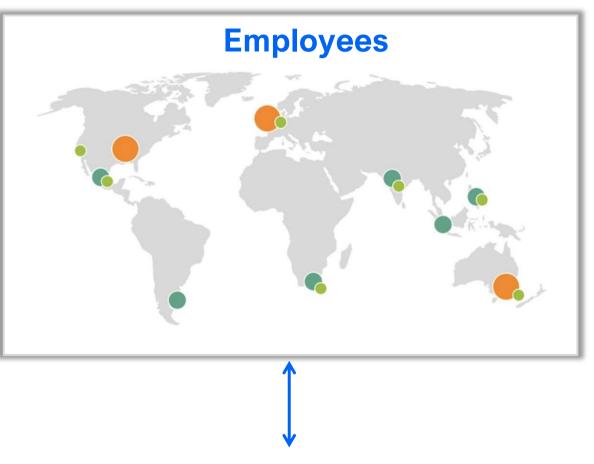
HyDE: Precise Zero-Shot Dense Retrieval without Relevance Labels <a href="https://arxiv.org/abs/2212.10496">https://arxiv.org/abs/2212.10496</a>

### Demo Assistant Scenario I

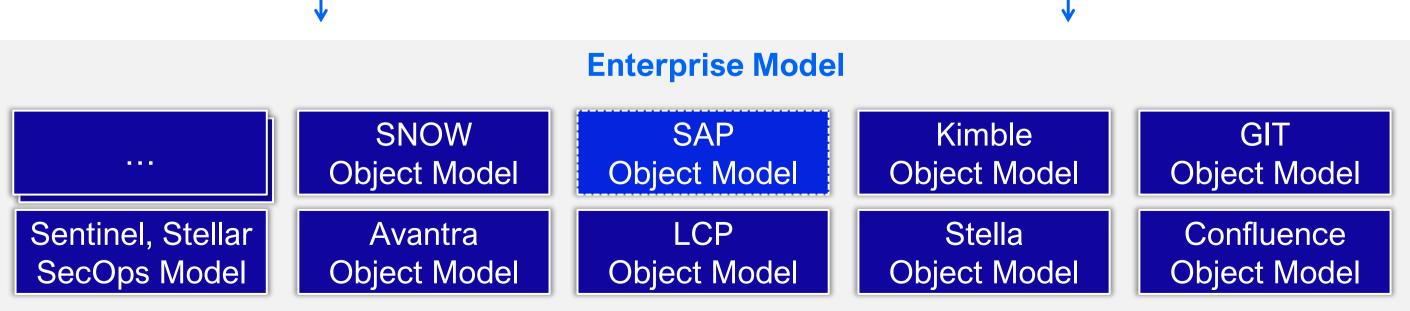
# **Challenge 1: The Great Divide**

**The Questions** 





**The Answers** 



### Demo Assistant Scenario II

# **Generating Solution Architectures**

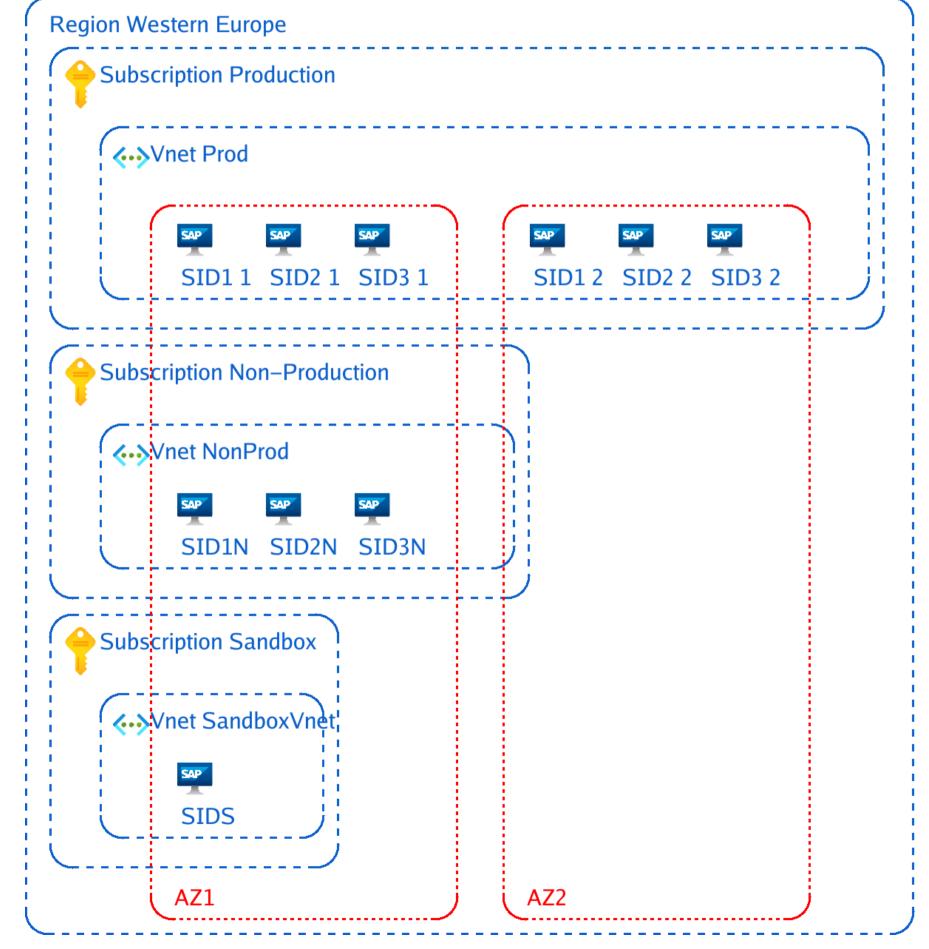
6:03 AM

 $\odot$ 

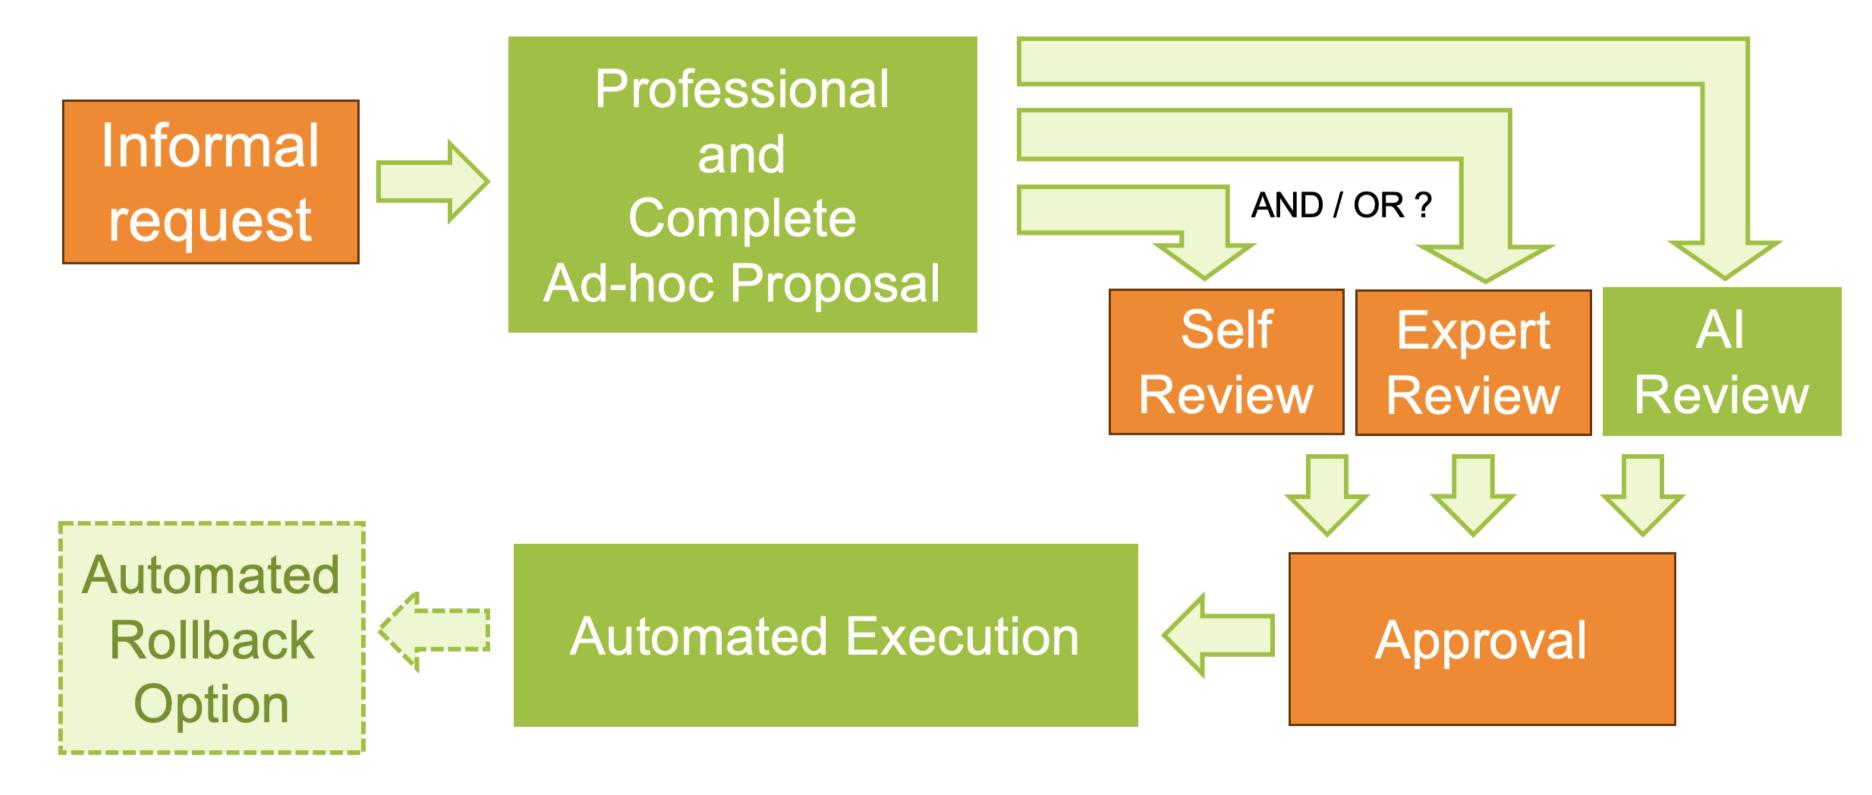
architect an Azure Solution Architecture with

- Production 3 HA SIDs
- Non-Production: Matching SIDs without HA
- Sandbox: a single SID

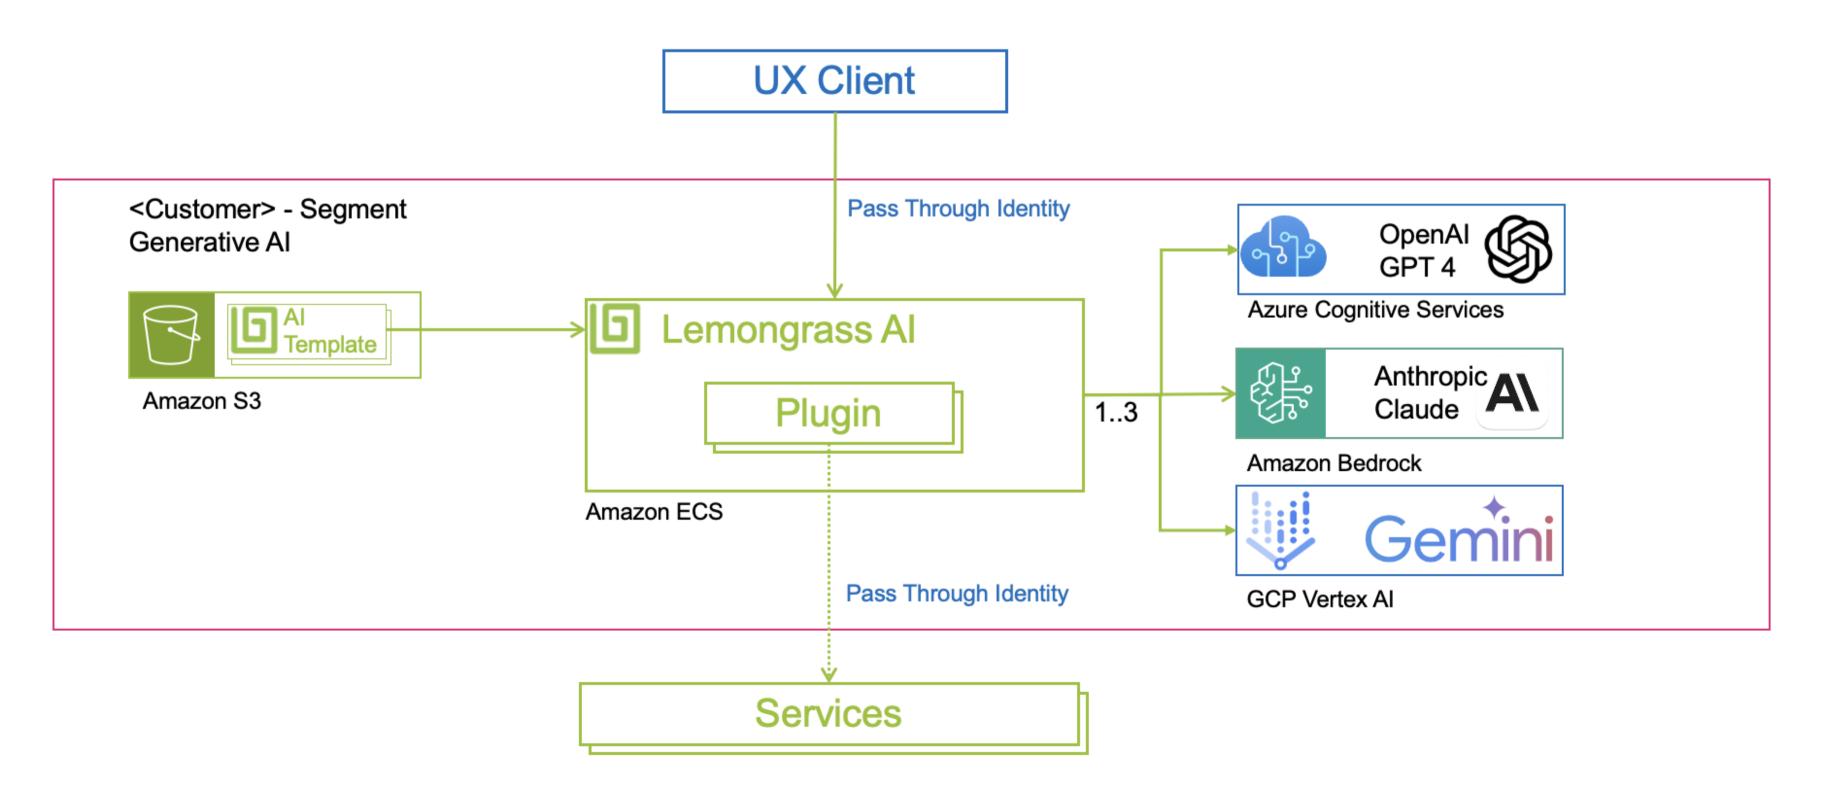




### **The ideal Business Process?**



### **Example Deployment Architecture**

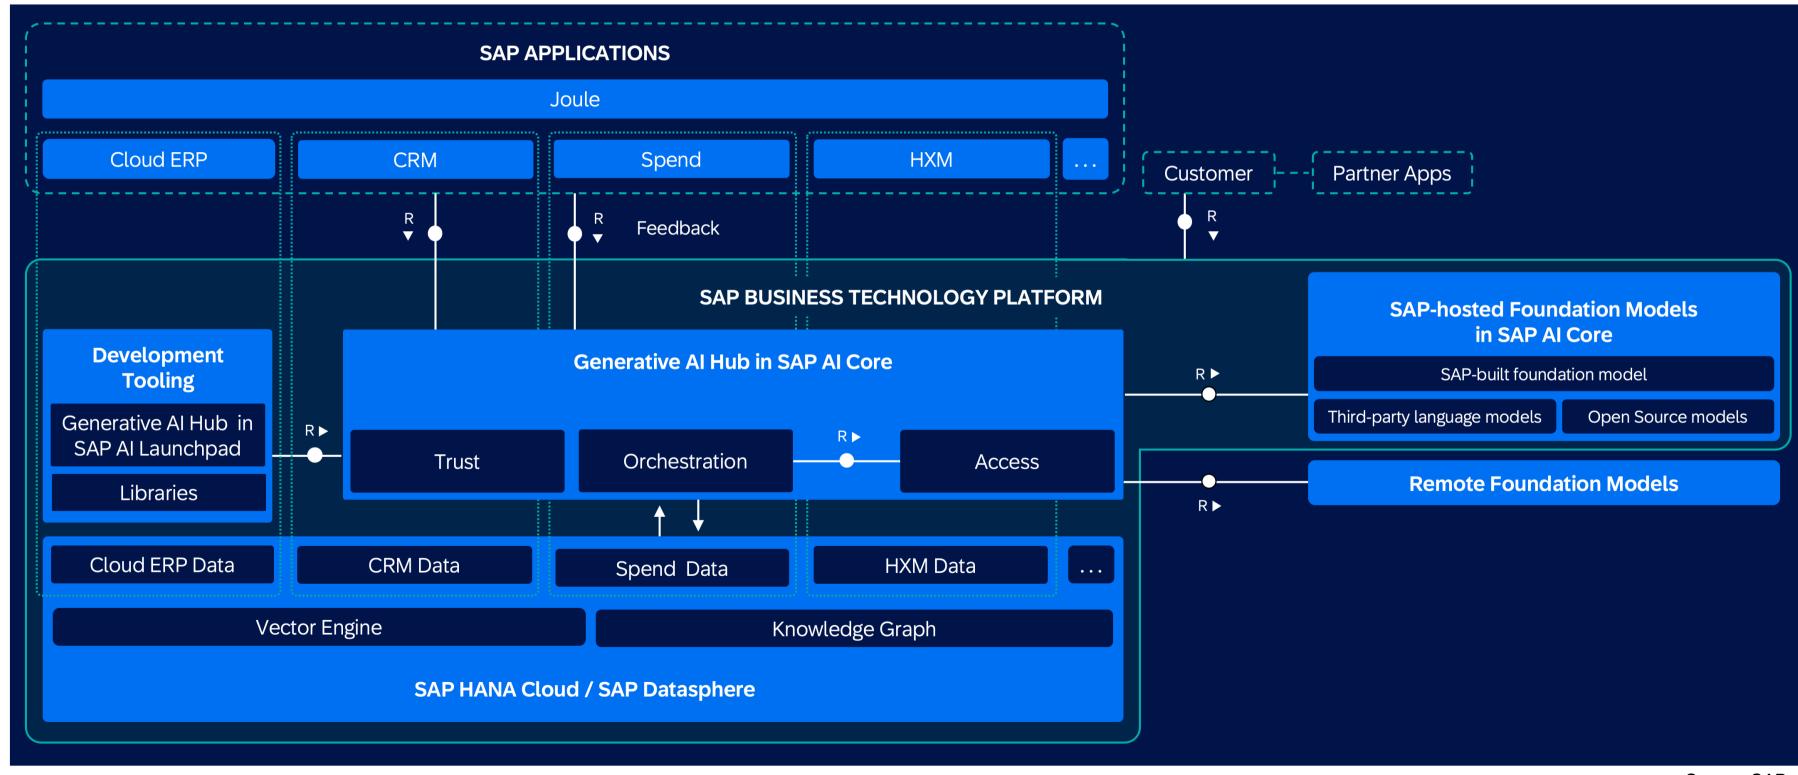


### **Multi-Cloud and Multi-Provider**

### High Volume Document Fact Extraction

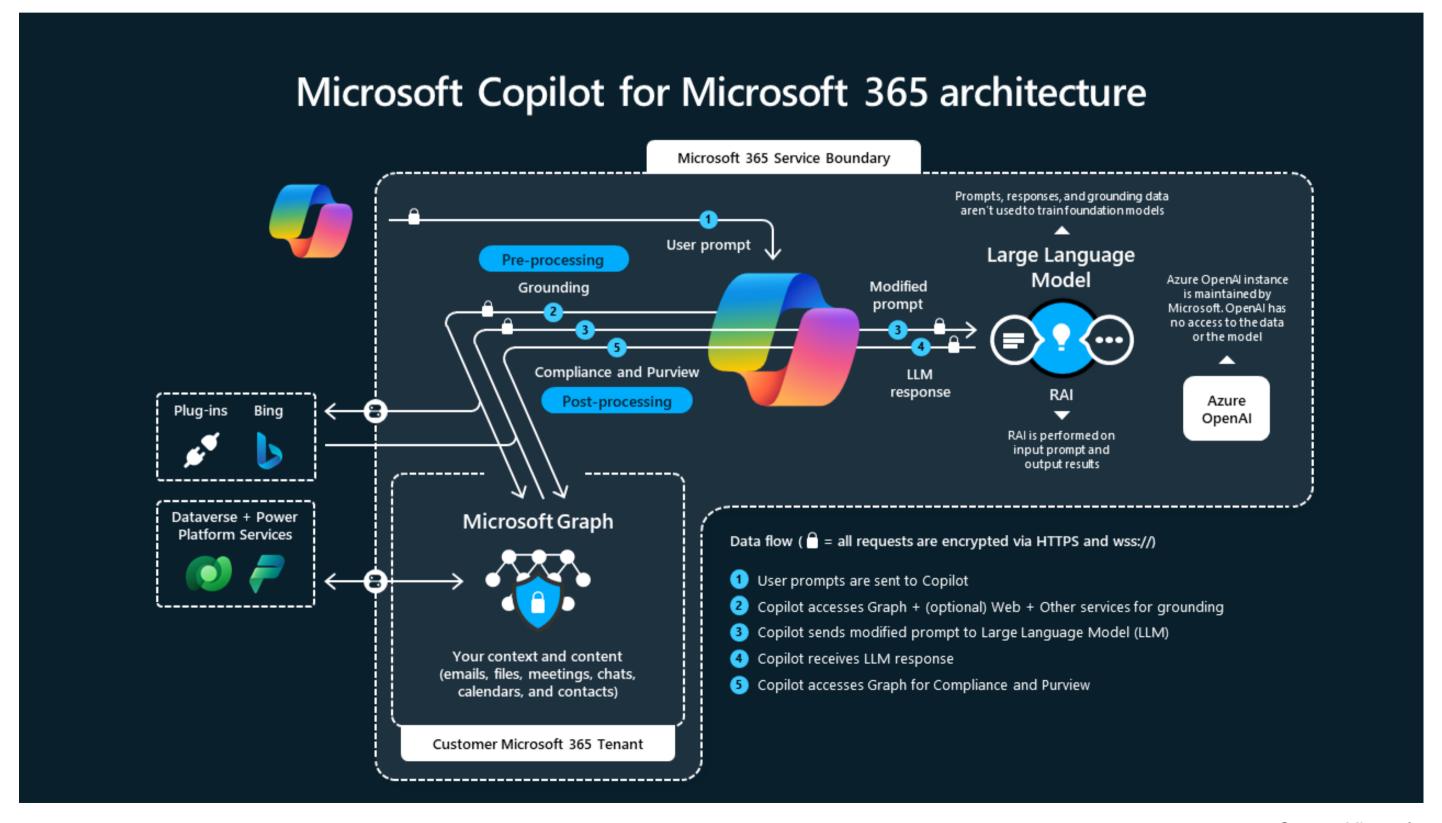
Cloud	Model	Avg Millis 🔻	Relative Cost	Time*Cost
Google	Gemini 1.5 Flash	546	0.15	82
Azure	GPT 40 Mini	557	0.15	84
Amazon	Claude 3.0 Haiku	890	0.25	223
Azure	GPT 3.5	473	0.50	237
Azure	GPT 40	640	5.00	3,200
Amazon	Claude 3.5 Sonnet	1,443	3.00	4,329

### **SAP Joule**



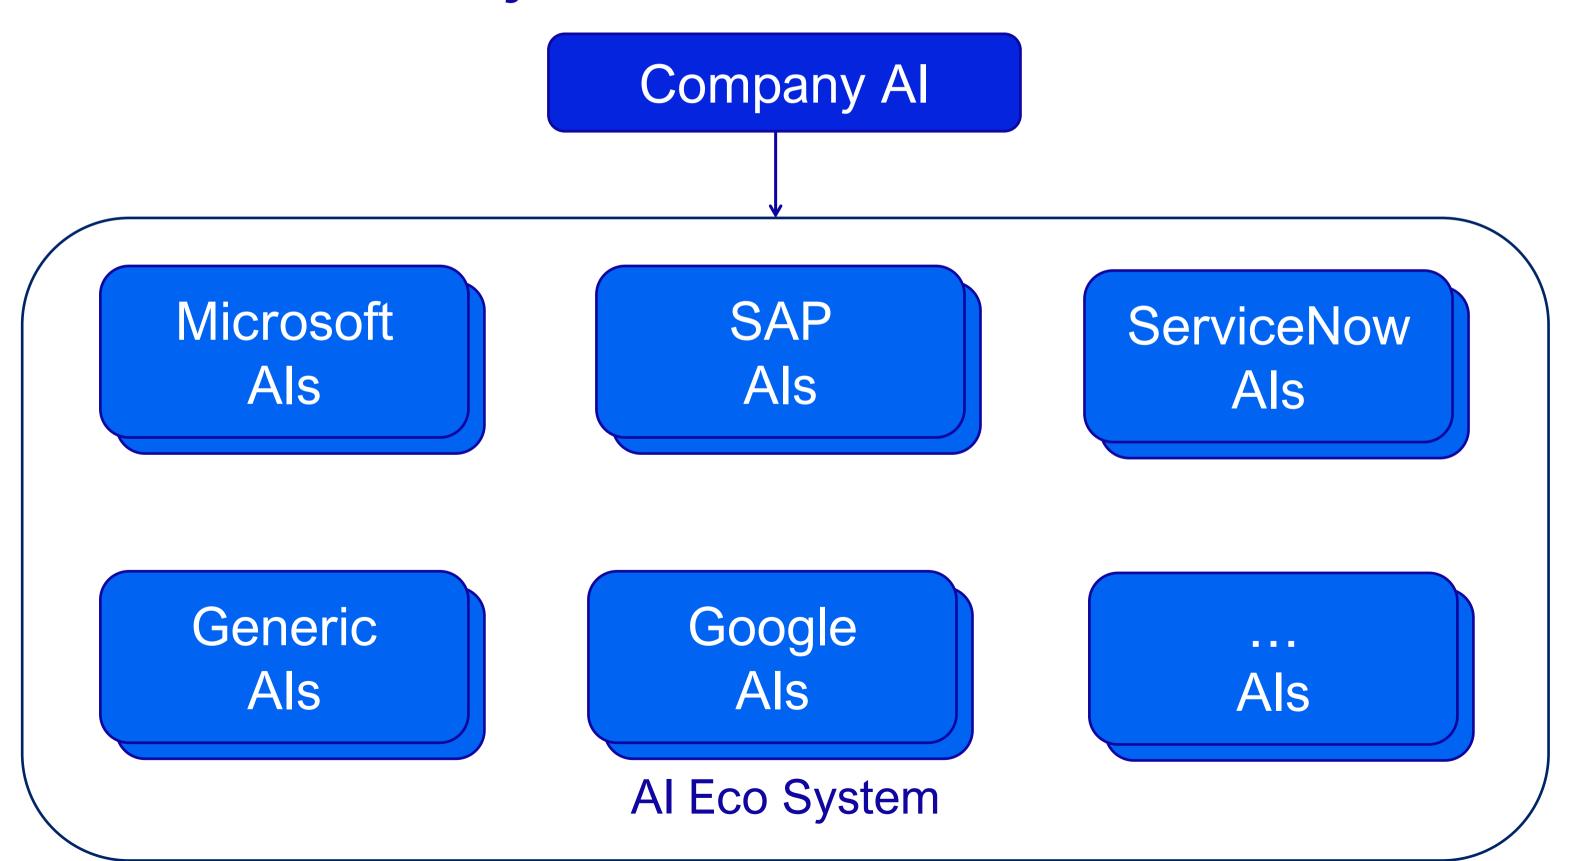
Source: SAP

# **MS** Copilot

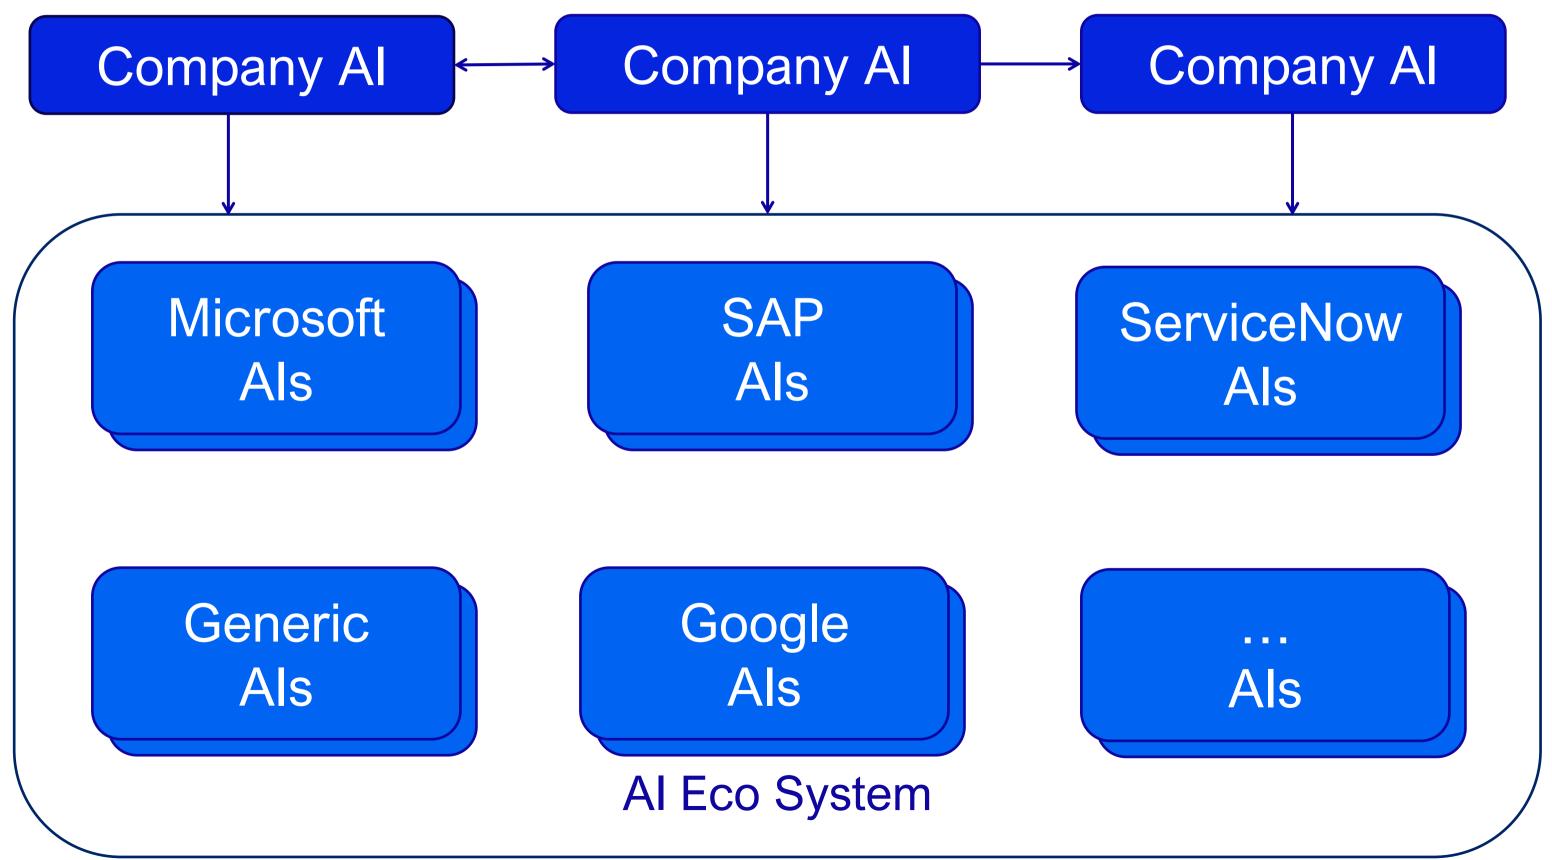


Source: Microsoft

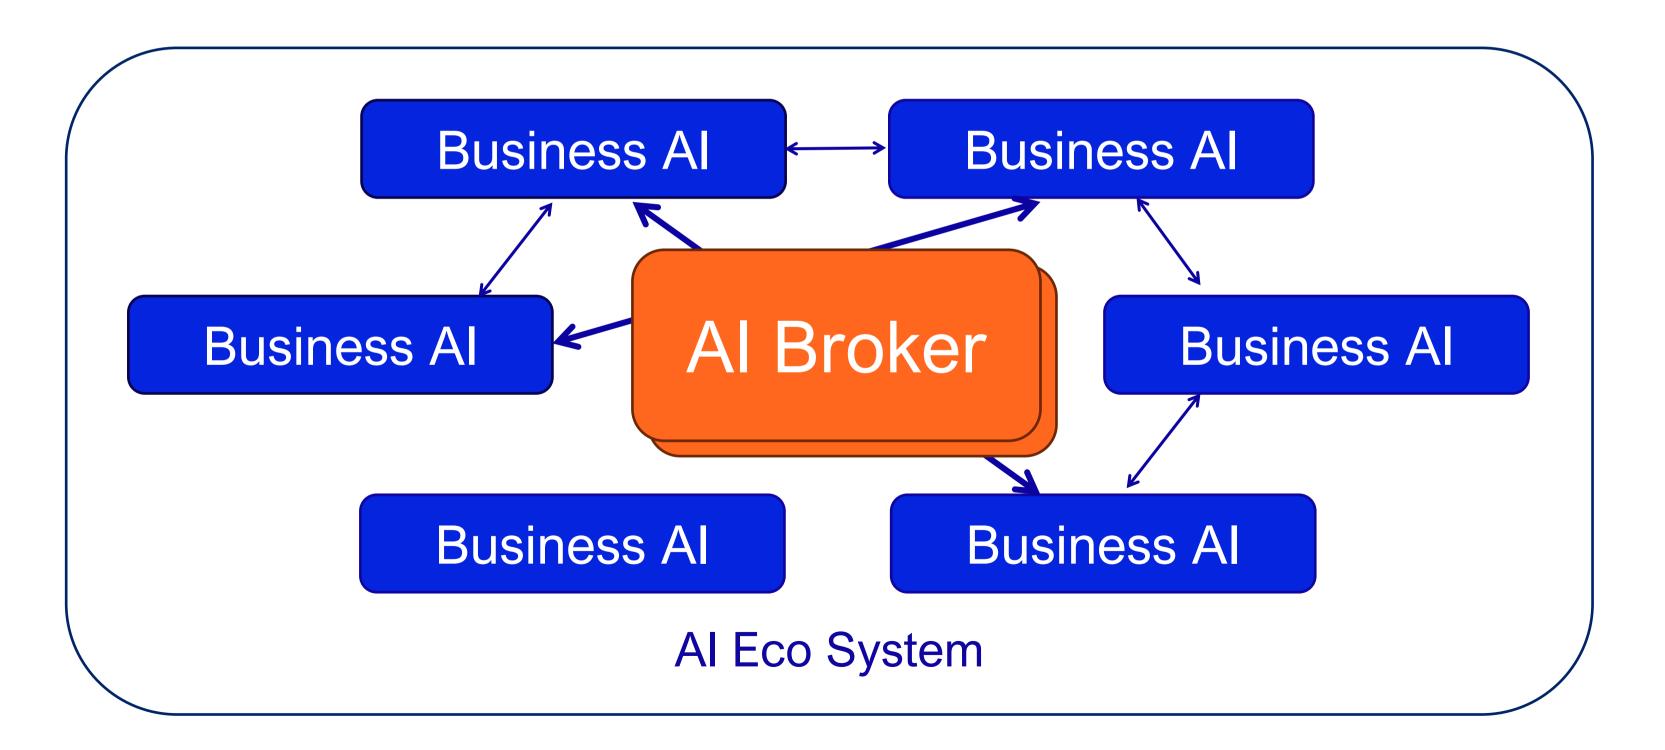
### Where are we today?



### Where are we tomorrow?



### Where might we be heading?

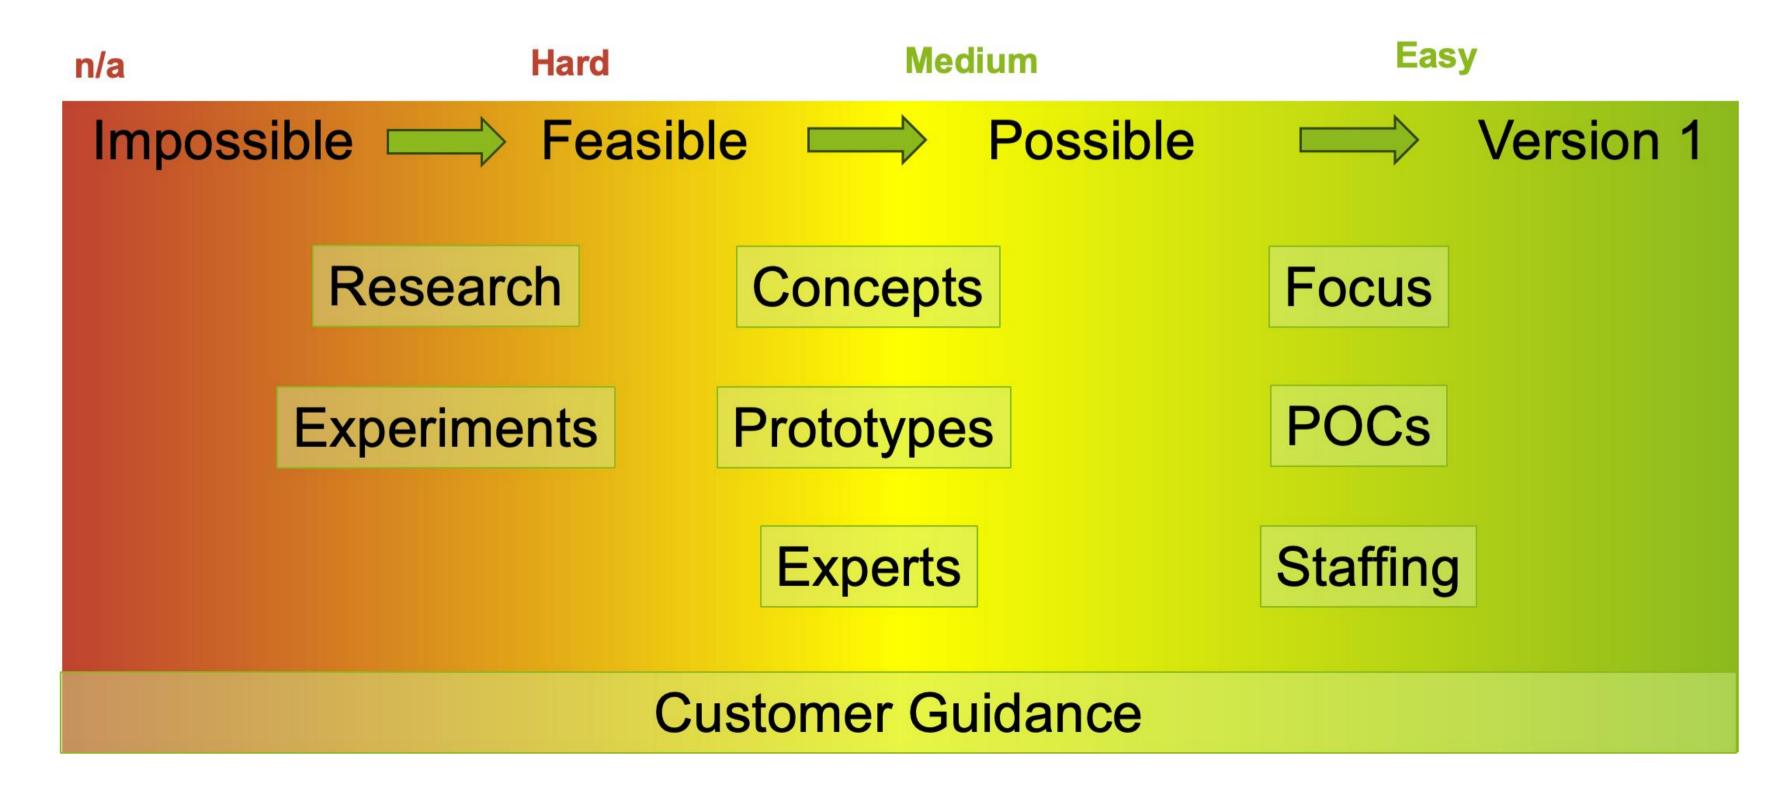


# Al Transformation - Process

### **Idea - Checkpoints**

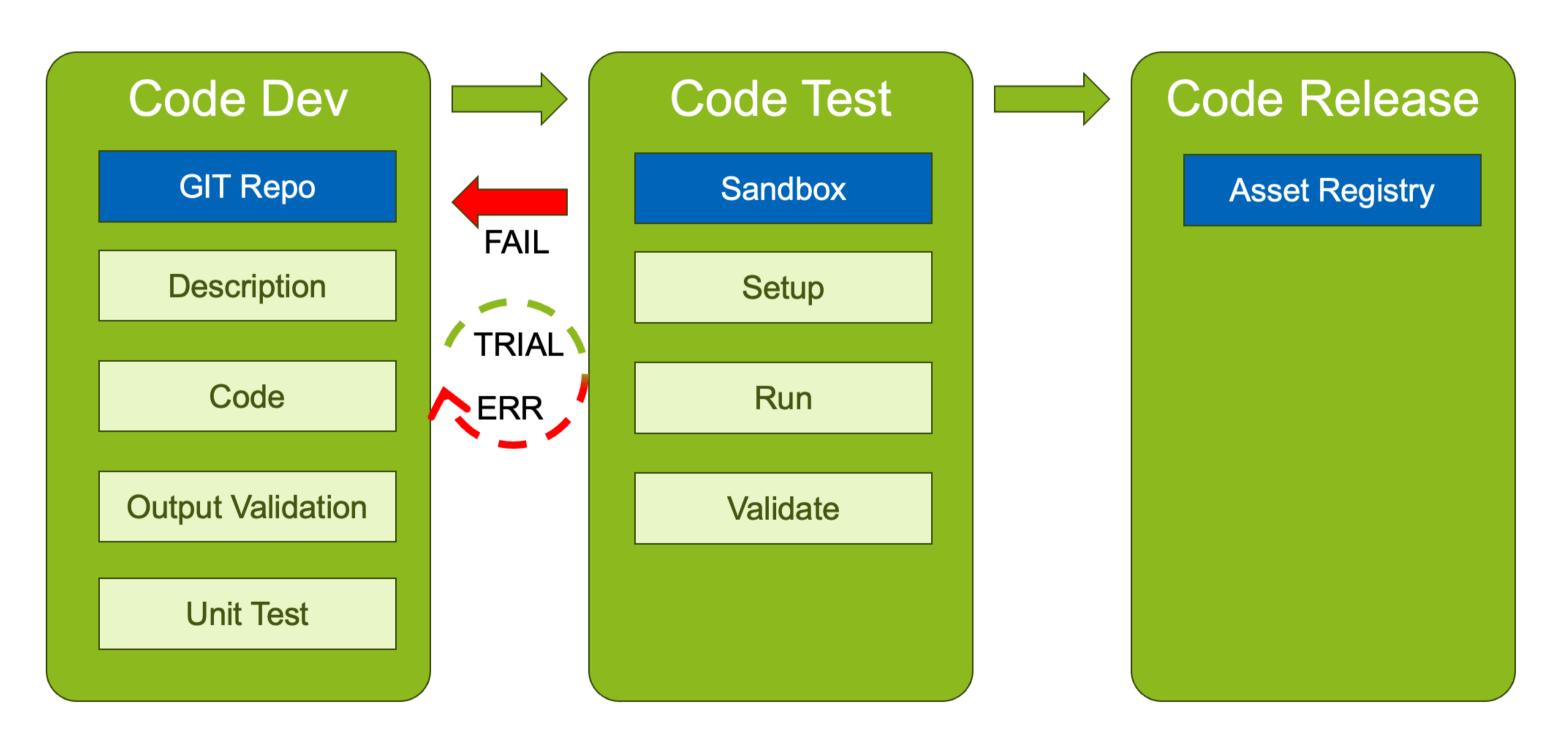
- 1. Significant Business Value
- 2. Is the AI smart enough?
- 3. Do we have the Domain Knowledge?
- 4. Do we have the Data?
- 5. Cost to the Business
- 6. Impact on the Organization

### **Any Innovation Idea**



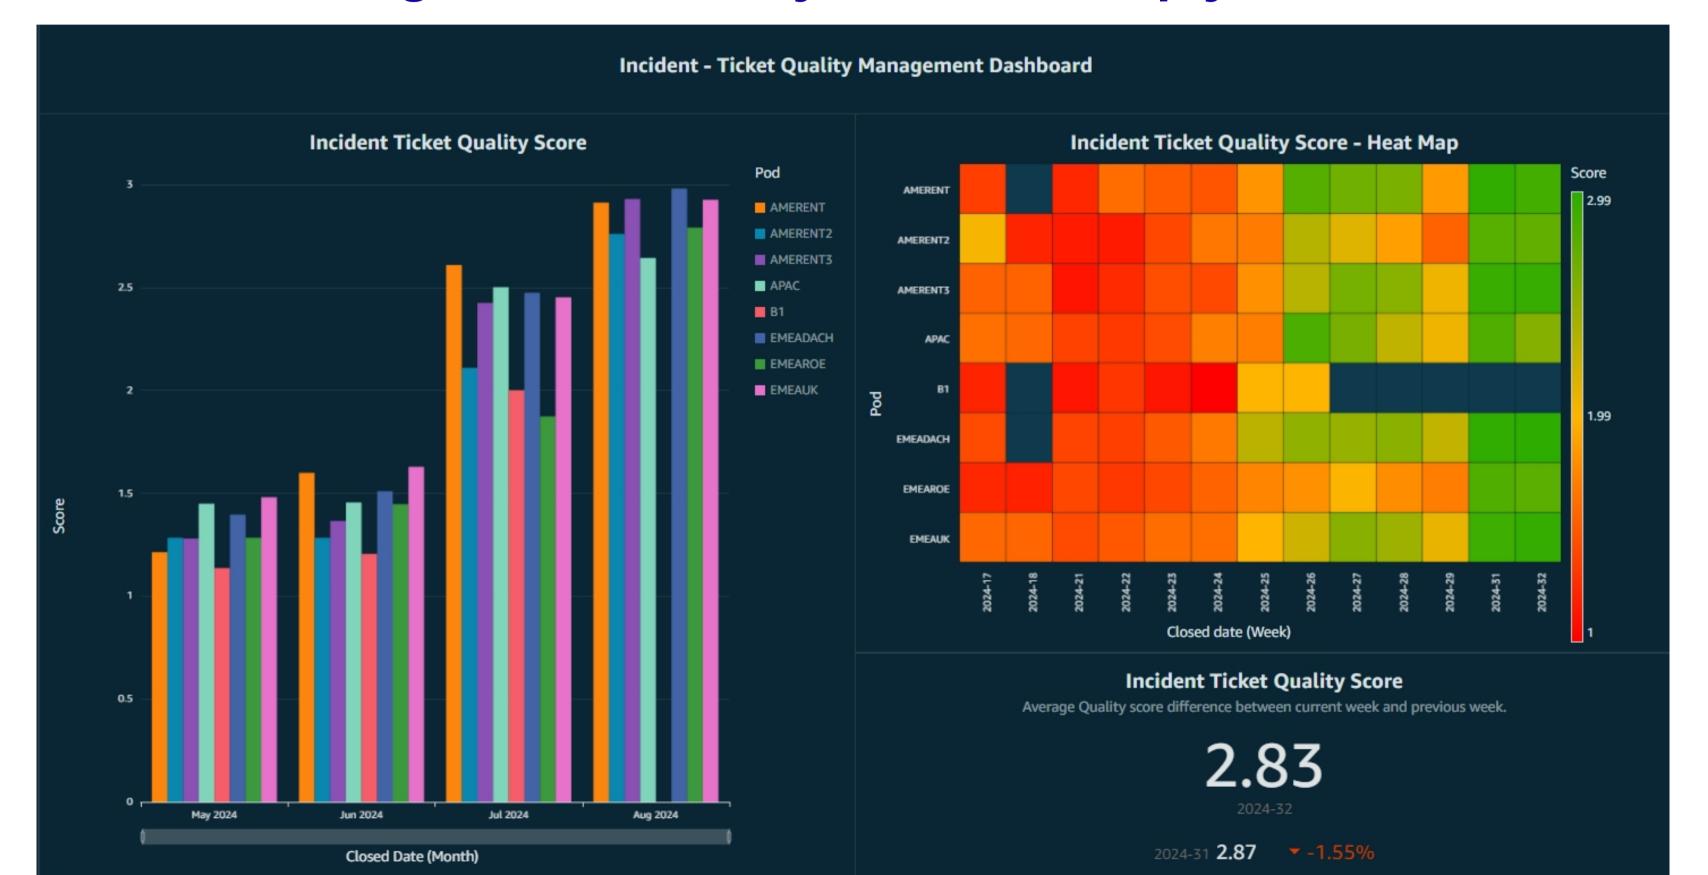
### Agent Scenario I

# Simple DevOps Process an Engineer follows

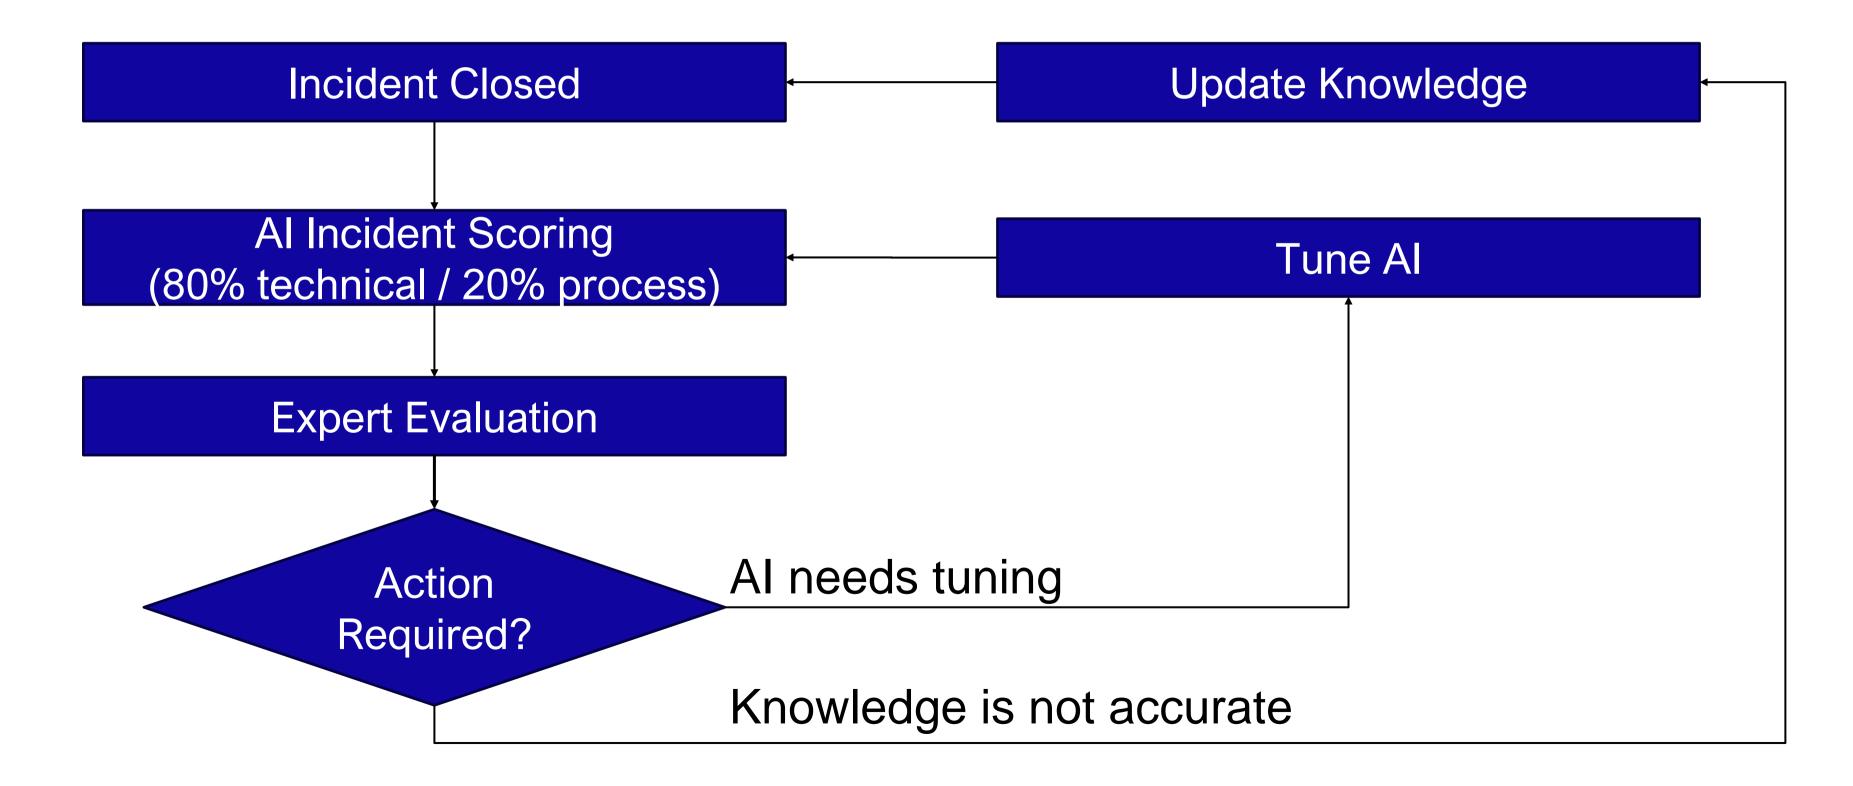


# Al Transformation - People

### **Incident Management – Quality of Human Reply over time**



### **Human Feedback - Al Training Loop**



# Incident Management – Correctness of Al Rating of Human Reply



# Al Transformation - Solution

# **Lemongrass Al Strategy**

SAP To Cloud Transformation for Customers

Advisory

Design

Implement

Operate

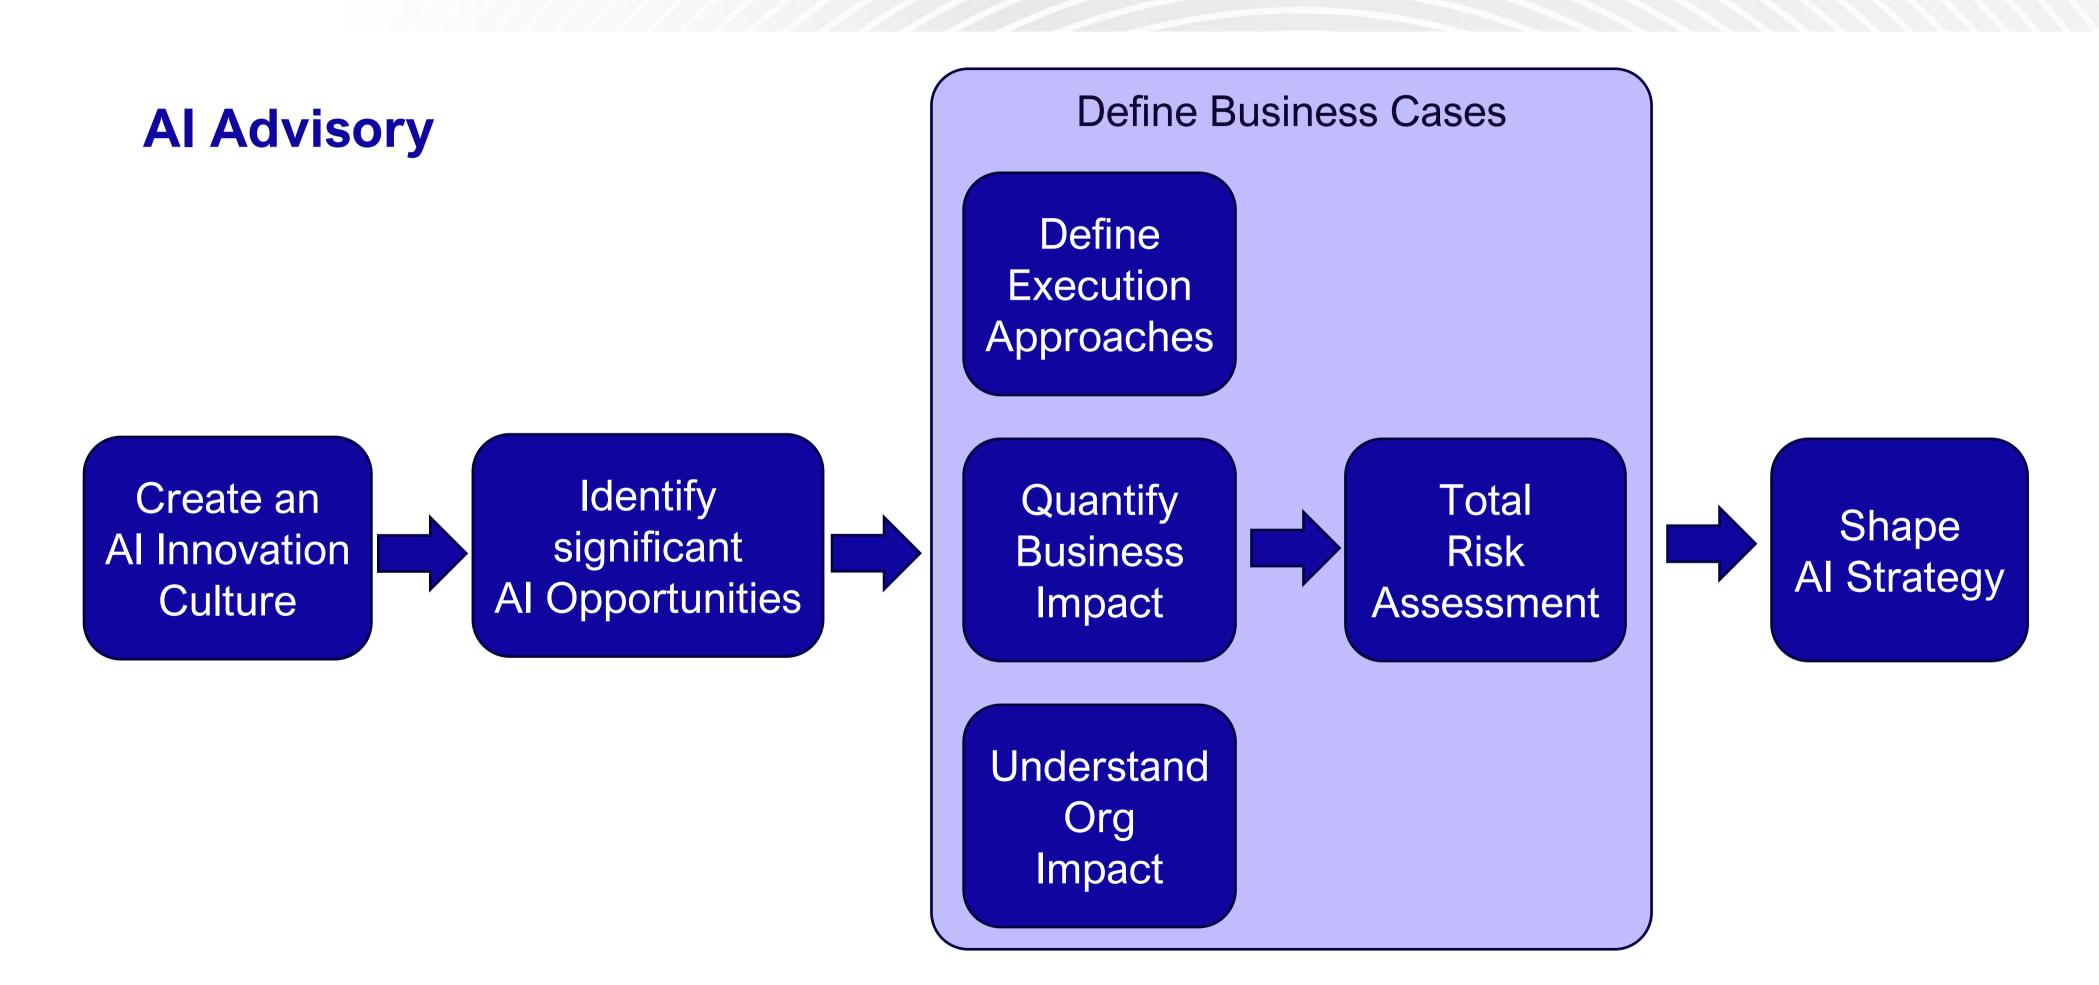
Al Transformation of Lemongrass

Advisory

Design

Implement

Operate





### Where to Find More Information

### ©Gartner 2024 Hype Cycle for Emerging Technologies

The 25 disruptive technologies to watch on the Gartner, Inc.

### SAP AppHaus • Business AI Explore Workshop – Identify AI use cases with business impact.

• Aims to help identify business challenges and opportunities that can be addressed using Artificial Intelligence (AI)

### https://www.microsoft.com/en-us/research/research-area/artificial-intelligence/

Microsoft Research publications pursuing computing advances on Al

### https://www.sap.com/products/artificial-intelligence/business-technology-platform.html

Using SAP BTP to infuse AI into SAP applications, extensions, and analytics

### https://aws.amazon.com/ai/generative-ai/?refid=8f1c8b7a-7d88-4fec-9624-839fb8e22c8f

Generative AI with AWS

### https://azure.microsoft.com/en-us/products/machine-learning/generative-ai

Generative AI with Azure

### **Generative AI | Google Cloud**

Generative AI with GCP

### Generative Al Accelerator | Data, Analytics & Al Services (lemongrasscloud.com)

Generative Al accelerators from Lemongrass

# **Key Points to Take Home**

- Now is the right time to start your Al journey.
- Identify tasks that are important, yet not enjoyed.
- Think about the organizational impact of Al.
- Use an architecture that keeps your business flexible to AI changes.
- Engage with Experts

# Thank you! Any questions?

### **Frank Stienhans**

www.linkedin.com/in/frankstienhans/

# Babu Ayyalu

www.linkedin.com/in/babu-ayyalu

Please remember to complete your session evaluation.

# SAPINSIDE

SAPINSIDER MEMBERSHIP MATTERS. YOU HAVE TO BE PART OF THE GLOBAL COMMUNITY.

