



UNLOCKING AI AGENT SUCCESS

with LangGraph vs. Camunda 8

viadee 

Who we are

Mario Micudaj



BPM & AI Consultant

- Data Science, BPM and Process Mining
- Camunda in the insurance industry
- Camunda as AI orchestrator in industrial industry

Andre Strothmann



BPM Consultant

- Process Engineer, Camunda Champion
- *Camunda* and *RPA* in the insurance industry

Who we work for

viadee Unternehmensberatung AG



founded 1994 in Muenster, Germany



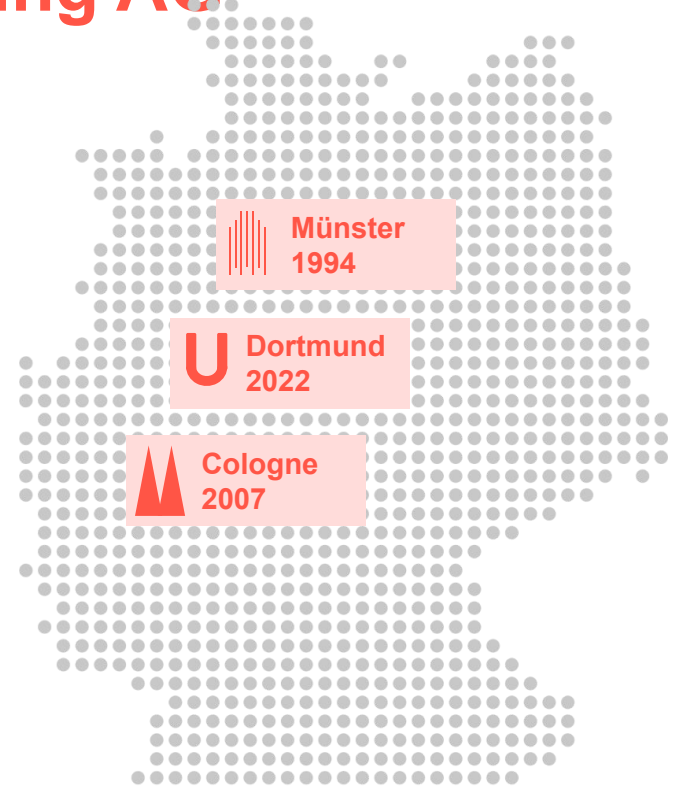
> 250 consultants



> 40 Mio. EUR revenue p.a.



> 3.500 consulting projects



Areas of competence



CLOUD ARCHITECTURES
& PLATFORMS



IT SECURITY



DATA & AI



BUSINESS ANALYSIS



SOFTWARE DEVELOPMENT
& ARCHITECTURE



ORGANISATIONAL
DEVELOPMENT



BPM & PROCESS
MANAGEMENT



QUALITY
ENGINEERING



PROJECT
MANAGEMENT

Agenda



01 Customer Case and Architecture

02 Implementation with LangGraph

03 Implementation with Camunda 8

04 Comparison

05 Q & A



Customer Case and Architecture

Service Desk Agent

Customer Case from retail industry



20,000

employees require support from the service desk around the clock.

> 2,000

tickets must be processed manually by hotline agents each month.

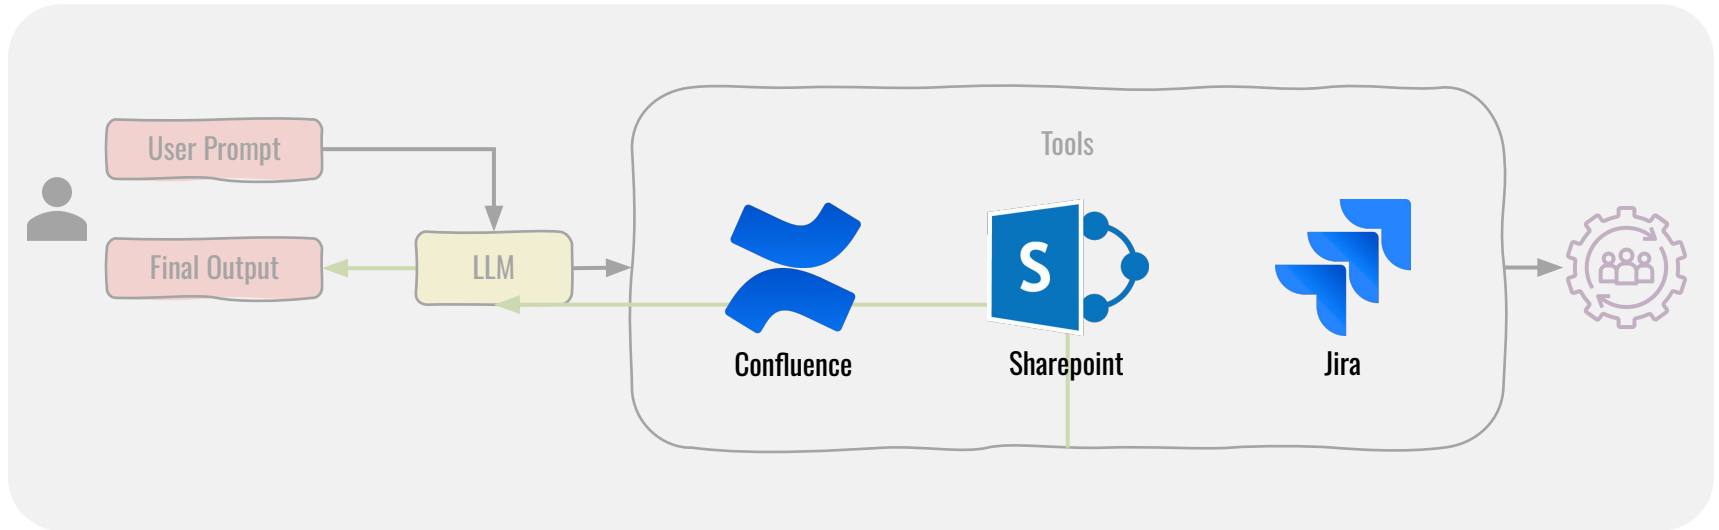
> *“The keyboard of my laptop is not responding to my typing.”*

> *“How do I recover deleted emails in Outlook?”*

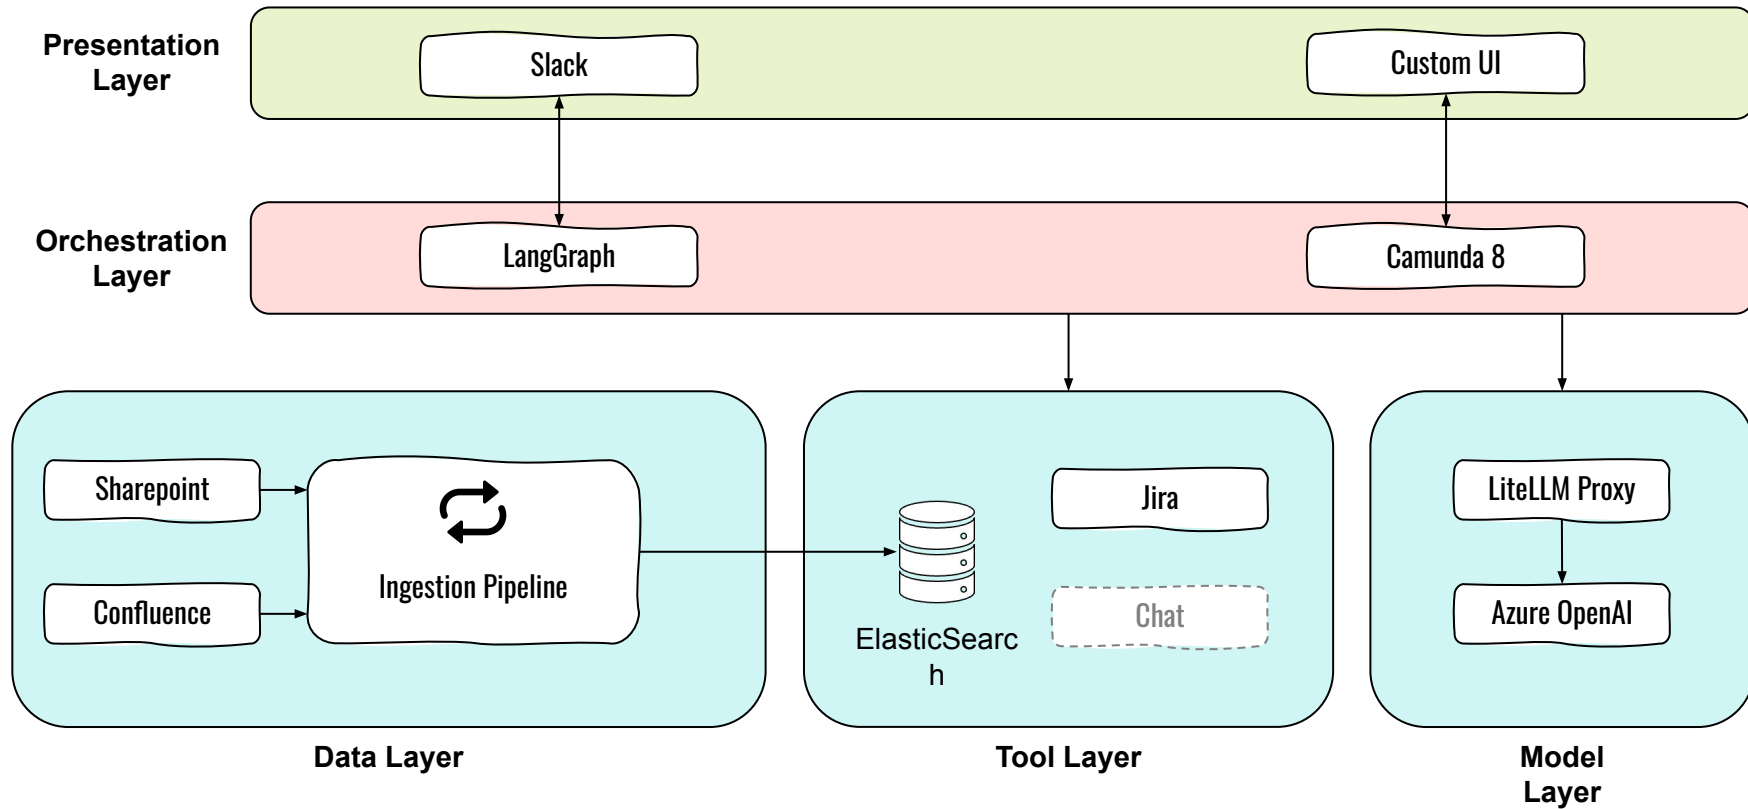
Customer Case from retail industry



Develop an intelligent chatbot to reduce support tickets and **relieve the service desk**, enabling **employees to resolve issues on their own**.



Architecture

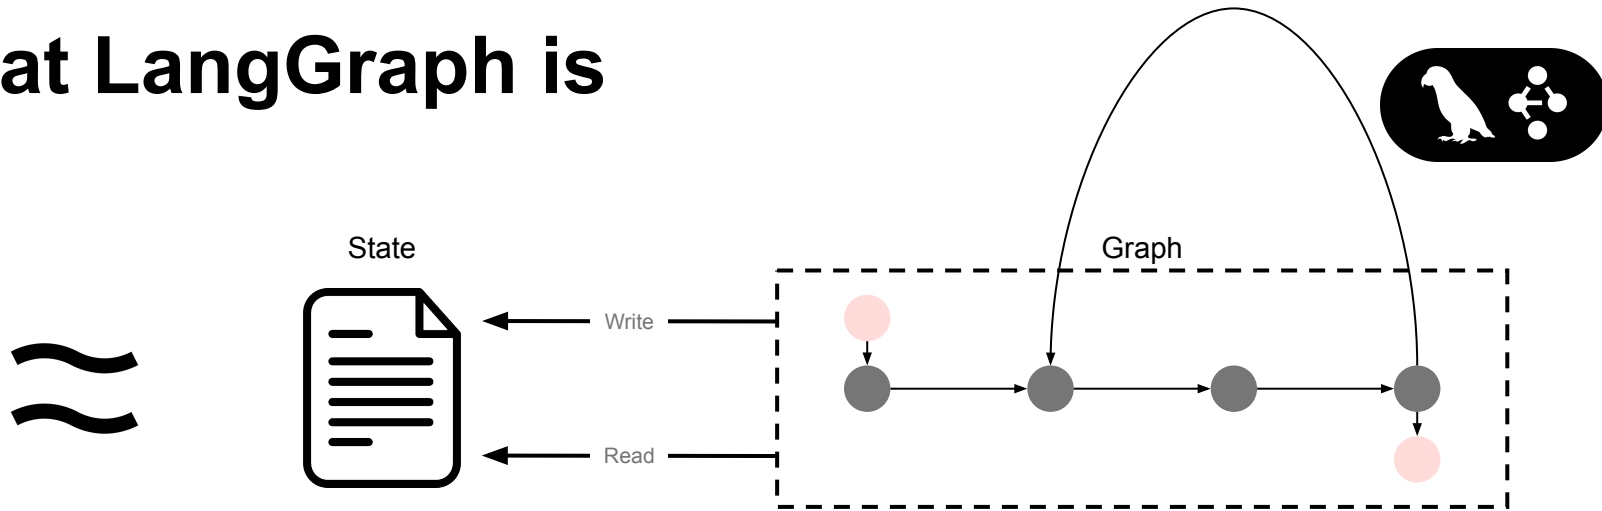




Implementation with LangGraph

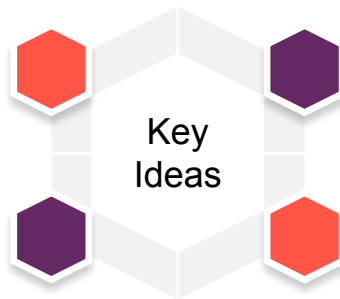
Service Desk Agent

What LangGraph is



Extension to LangChain with
graph-based orchestration on top

Framework for building AI agents
using graph-based workflows

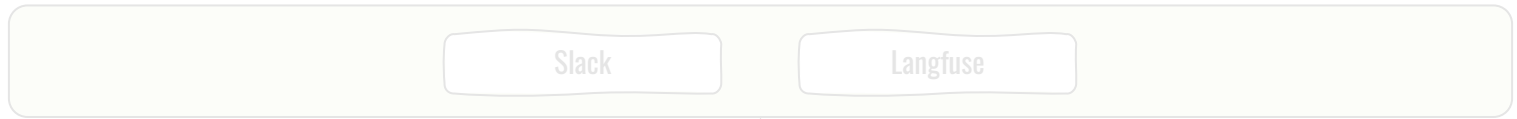


Graph constructs enable **if/else**,
parallel branches, and **loops**

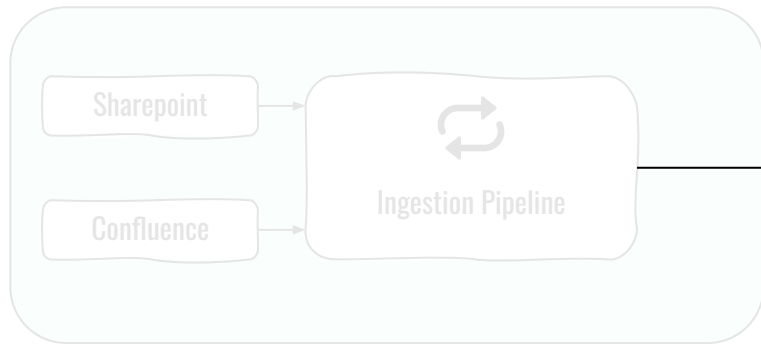
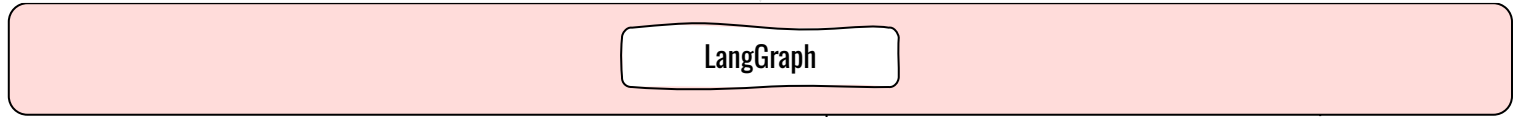
Open-source, community-driven project
that welcomes external contributions

Implementation of service desk agent

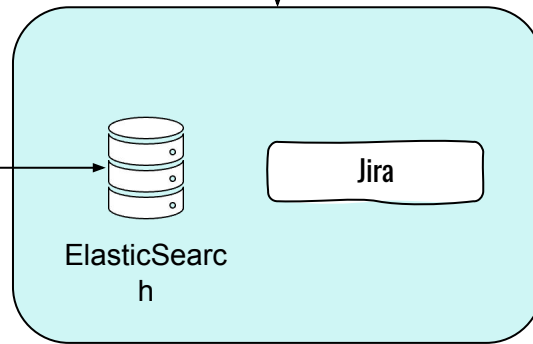
agent
Presentation
Layer



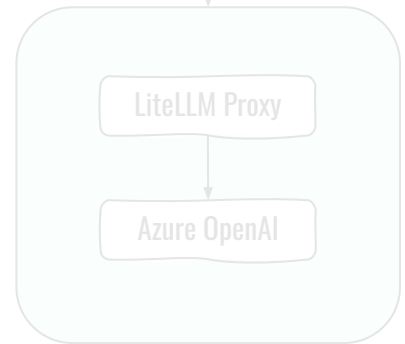
Orchestration
Layer



Data Layer



Tool Layer



Model
Layer



Implementation of service desk agent

... in three steps:

- 1 Definition of available tools

- 2 Definition of nodes and edges

- 3 Definition of graph and agent

Implementation of service desk agent

... in three steps:

1

Definition of available tools

```
Decorator ----- @tool
def ask_user_for_ticket_creation(request_text):
    """
    This tool can be used to create a Jira ticket for the user at the service desk so that someone can handle the request afterwards.
    """
    summary = get_llm().invoke(f"Generate a short, but descriptive title (do NOT exceed 250 characters) for a jira issue based on the following request: {request_text}").strip('')
    description = get_llm().invoke(f"Generate a compact description for a jira issue with all details available from the following request: {request_text}")
    ticket = create_jira_ticket(summary, description)
    return f"Successfully created a jira ticket with the following details: {ticket}."

@tool
def query_knowledge_base(request_text):
    """
    Use this tool to search the knowledge database for the AI Agents Inc context.
    """
    query = get_llm().invoke(f"Generate a search query to search for relevant information in the knowledge database based on the following request: {request_text}")
    search_result = answer_query_against_rag(query)
    return f"Searching the knoweldge base has resulted in: {search_result}."
```

Implementation of service desk agent

... in three steps:

2

Definition of nodes and edges

```
def tool_node(state: dict):
    """Performs the tool call."""
    result = []
    for tool_call in state["messages"][-1].tool_calls:
        tool = tools_by_name[tool_call["name"]]
        observation = tool.invoke(tool_call["args"])
        result.append(ToolMessage(content=observation, tool_call_id=tool_call["id"]))
    return {"messages": result}
```

```
def should_use_tool(state: MessagesState) -> str:
    """
    If the LLM has made a tool call, we are at the tools node.
    If the llm has not made a tool call, we want to escape the loop.
    """
    messages = state["messages"]
    last_message = messages[-1]
    return TOOLS_NODE if last_message.tool_calls else END
```

Implementation of service desk agent

... in three steps:

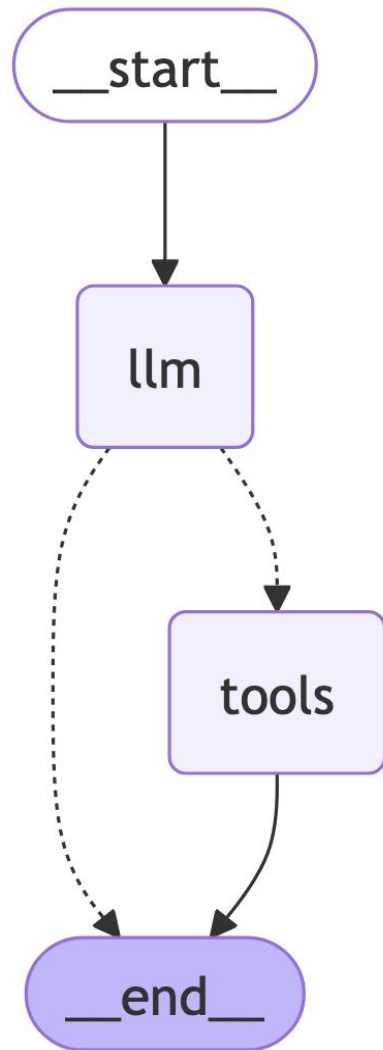
3

Definition of **graph** and **agent**

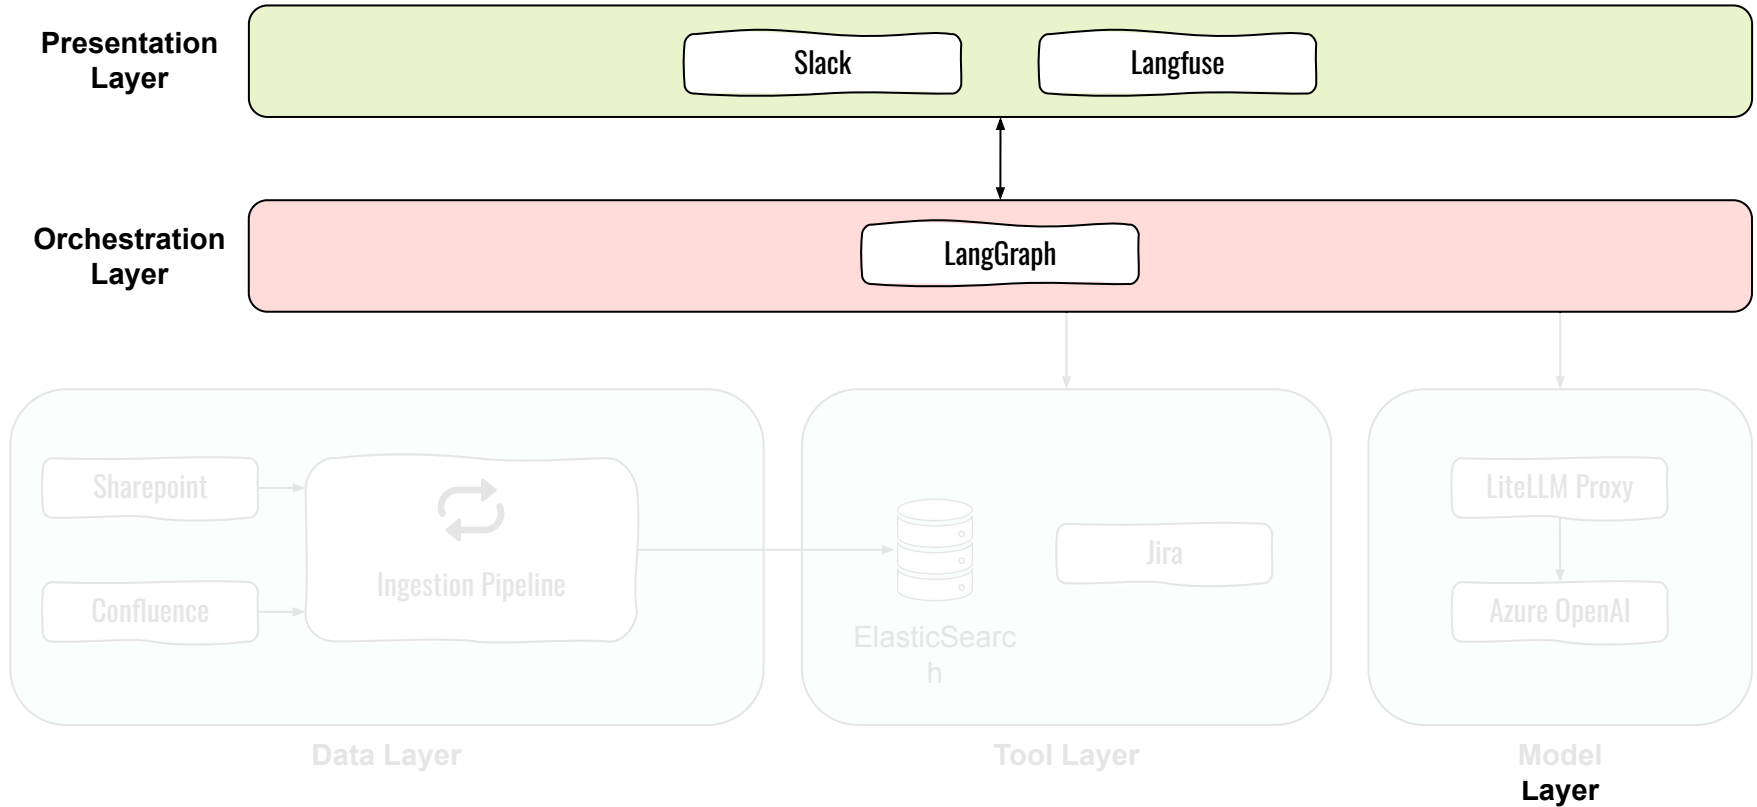
```
agent_builder = StateGraph(MessagesState)
agent_builder.add_node(LLM_NODE, llm_node)
agent_builder.add_node(TOOLS_NODE, tool_node)

# Connect the nodes
agent_builder.add_edge(START, LLM_NODE)
agent_builder.add_conditional_edges(
    LLM_NODE,
    should_use_tool,
    {TOOLS_NODE: TOOLS_NODE, END: END}
)
agent_builder.add_edge(TOOLS_NODE, LLM_NODE)

# Compile the agent
agent = agent_builder.compile().with_config({"callbacks": [langfuse_handler]})
```



Demo and monitoring



AA

AI ...

Home

Threads

Huddles

Drafts & sent

Directories

Starred

Channels

Direct ...

Apps

Slackbot

LangGraph AI ...

Service Desk A...

Add apps

Service Desk Agent

Chat History About

Service Desk Agent APP

How to use Service Desk Agent?

B I S | | | | |

Reply...

+ Aa © © | | |

AI Agents Inc. / service-desk-langgraph

Tracing

Traces Observations

Search...

IDs / Names

Past 30 min

Env 1

Filters

Table View 0

Columns 14/26

Timestamp	Name	Input	Output
No results.			

Rows per page 50 Page 1 of 0

Monitoring with Langfuse

The screenshot displays the Langfuse monitoring interface for a specific trace. The top bar shows the trace ID: 1693768f50f29d905bad14f43ec396b6. The left sidebar contains a search bar and a tree view of the trace components. The main area is divided into two panels: a 'LangGraph' panel on the left and a 'LangGraph ID' panel on the right.

LangGraph Panel (Left): Shows a hierarchical tree of the trace components. The root is 'LangGraph' (11.21s), which branches into 'llm' (1.73s) and 'tools' (2.95s). The 'llm' component further branches into 'ChatOpenAI' (1.73s) and 'should_use_tool' (0.00s). The 'tools' component branches into 'query_knowledge_base' (2.95s) and 'OpenAI' (0.56s). The 'OpenAI' component branches into 'search' (0.32s).

LangGraph ID Panel (Right): Displays the execution details for the 'LangGraph' component. It shows the execution time (11.21s), total cost (\$0.001095), and the number of items (5,914 → 491). Below this, there are sections for 'Input' and 'Output'.

Input Section: Shows a table with 'Path' and 'Value' columns. The 'messages' path is expanded, showing 7 items.

Path	Value
messages	[..., ..., ..., ...6 more]
> 0	7 items
> 1	7 items
> 2	7 items
> 3	7 items
> 4	7 items
> 5	7 items
> 6	7 items
> 7	7 items
> 8	7 items

Output Section: Shows a table with 'Path' and 'Value' columns. The 'messages' path is expanded, showing 10 items.

Path	Value
messages	[..., ..., ..., ...11 more]
> 0	7 items
> 1	7 items
> 2	7 items
> 3	7 items
> 4	7 items
> 5	7 items
> 6	7 items
> 7	7 items
> 8	7 items
> 9	10 items
> 10	0 items

At the bottom of the interface, there is a diagram showing the flow of the execution. It starts with a green box labeled '__start__', which points to a blue box labeled 'llm'. The 'llm' box points to a blue box labeled 'tools'. The 'tools' box points back to the 'llm' box, forming a loop. Finally, the 'llm' box points to a red box labeled '__end__'.

3

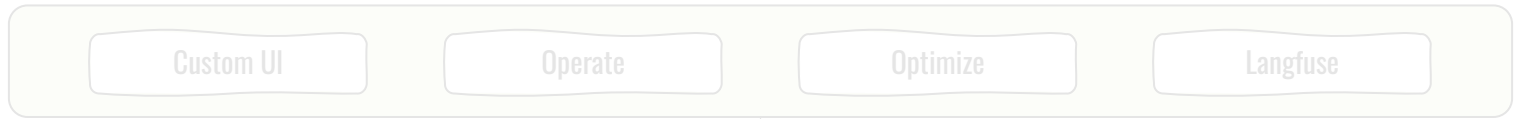


Implementation with Camunda 8

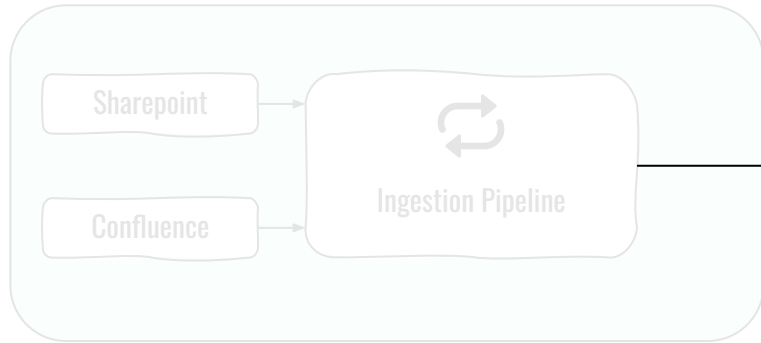
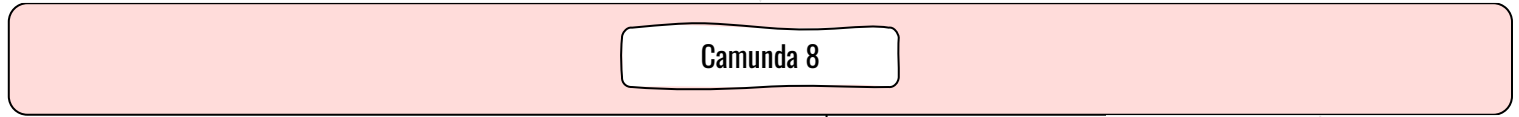
Service Desk Agent

Implementation of service desk agent

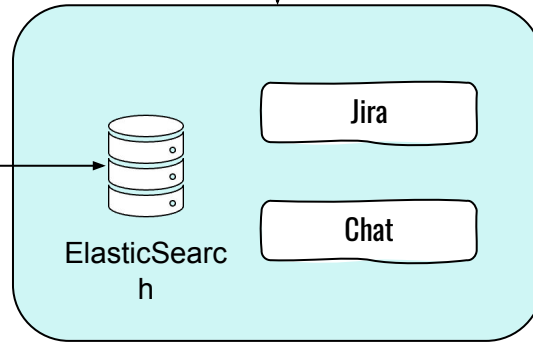
agent
Presentation
Layer



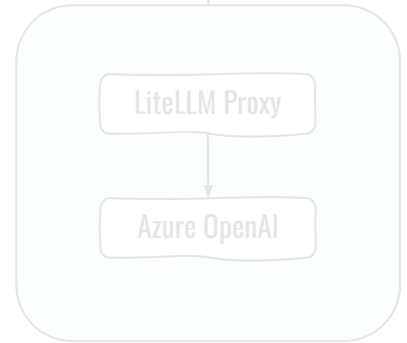
Orchestration
Layer



Data Layer



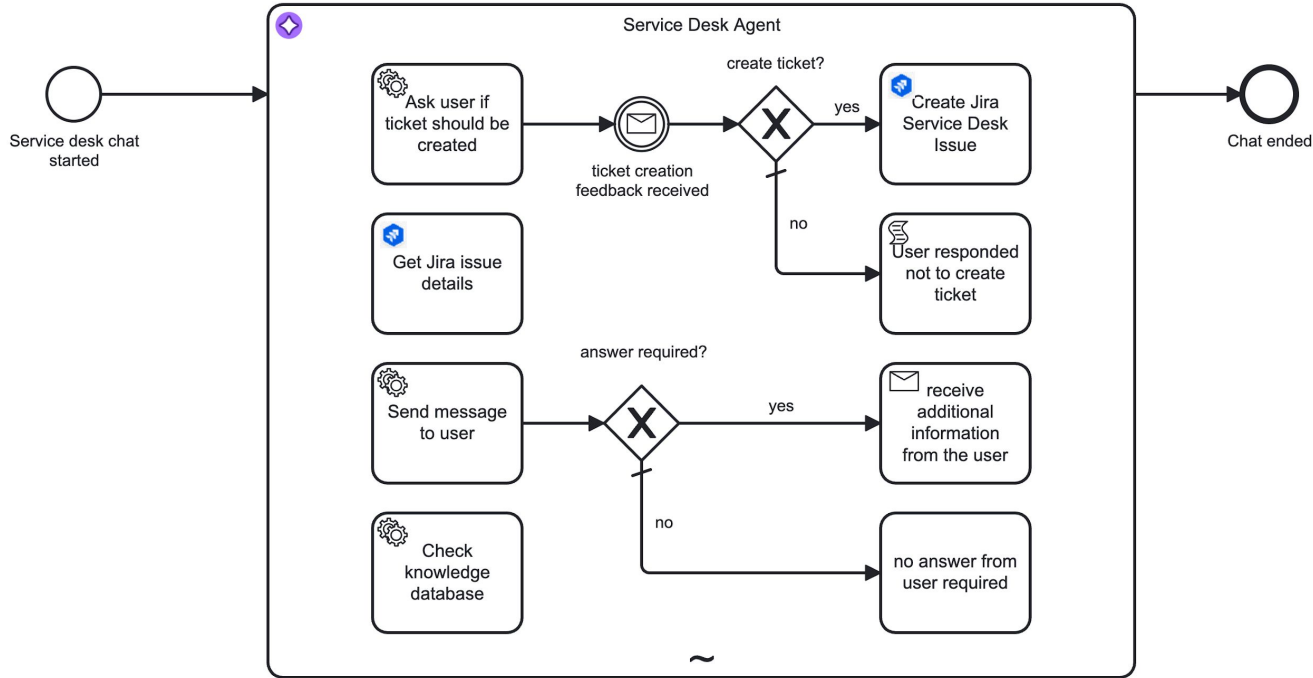
Tool Layer



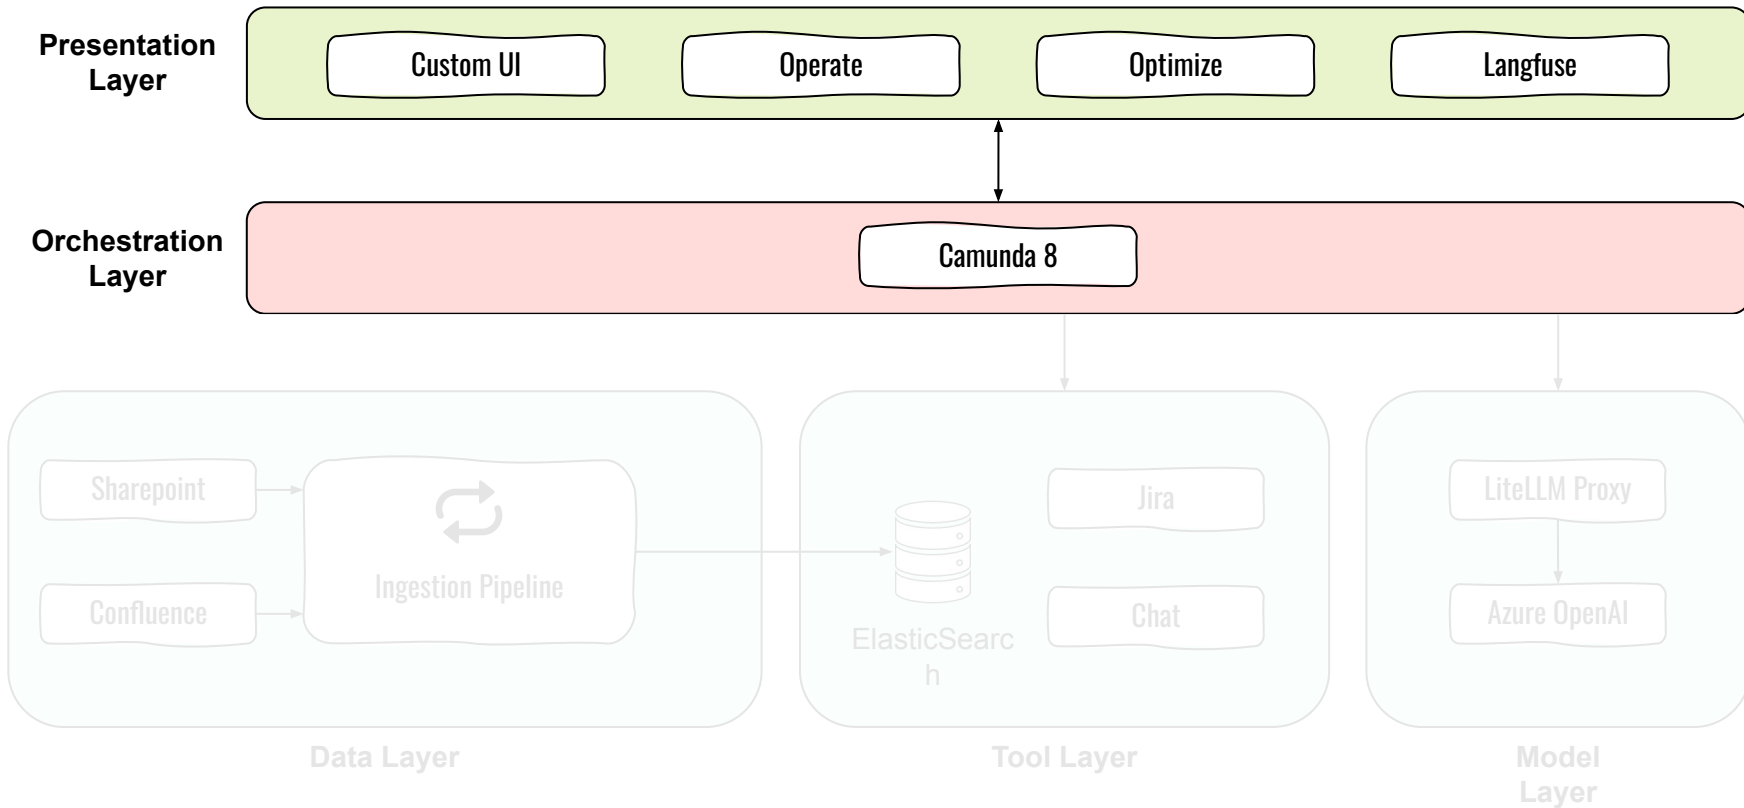
Model
Layer



AI Agent in Camunda



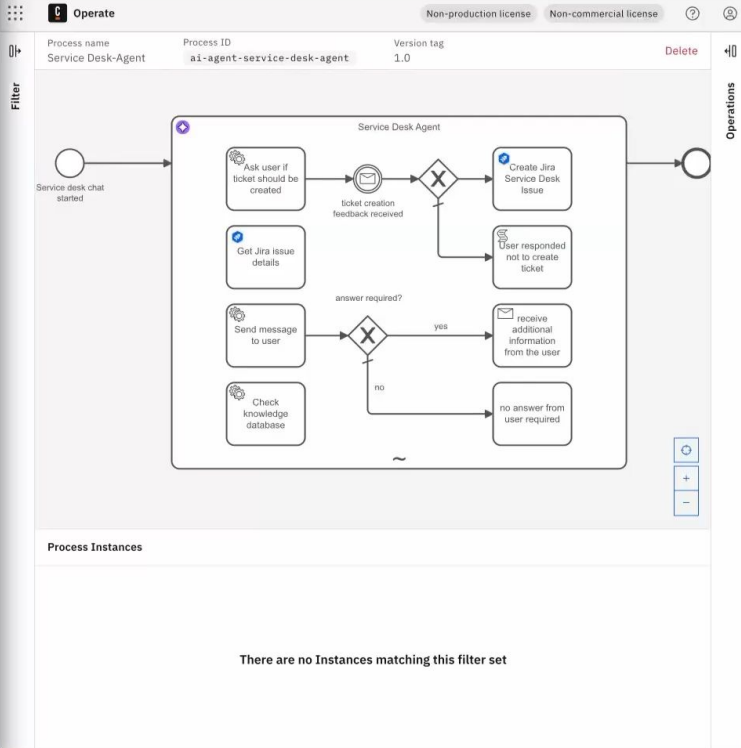
Demo and monitoring



SD Service Desk Agent

Type a message...

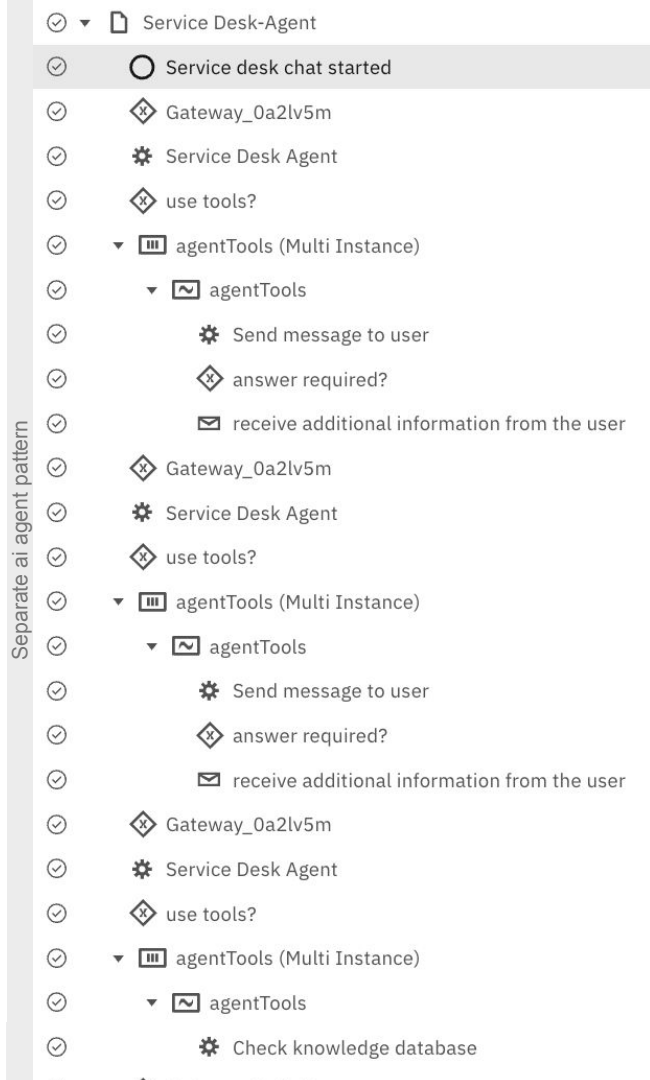
Send



Operational transparency

Operate

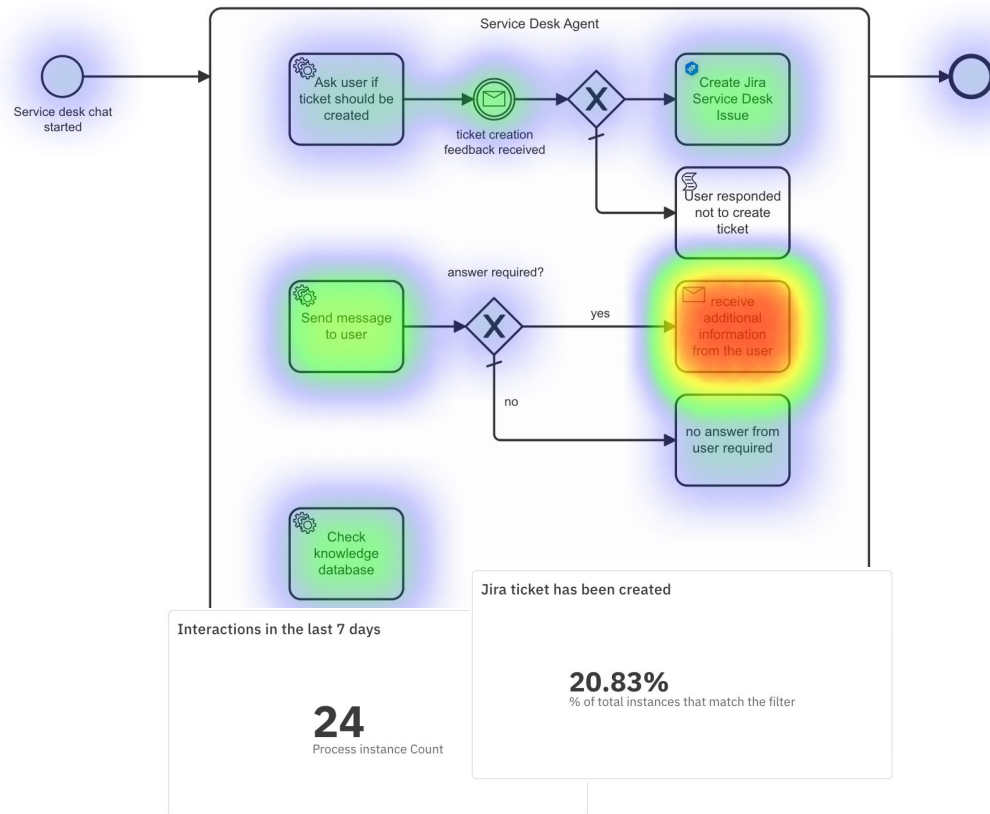
- Complete flow of each agent interaction visible
- Tool sub-step level (deterministic steps)



Operational transparency

Optimize

- Process metrics
- Heatmaps of tool usage

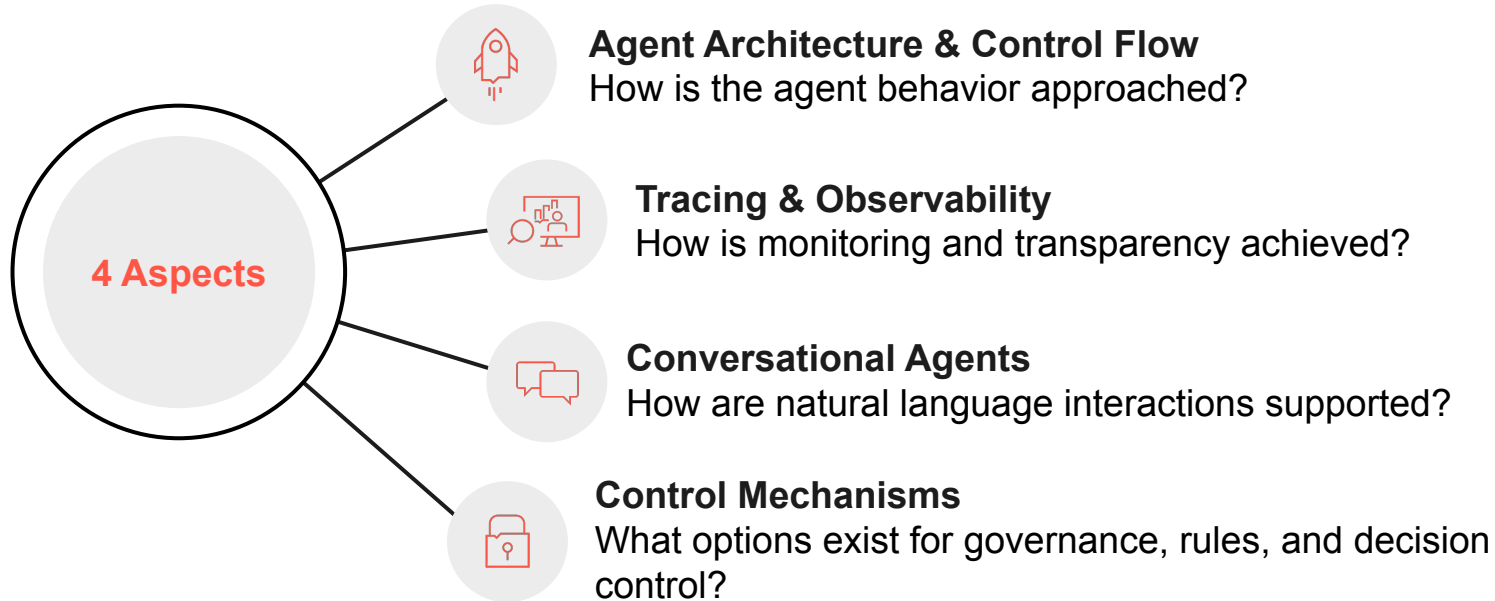


4 Comparison



LangGraph vs. Camunda 8

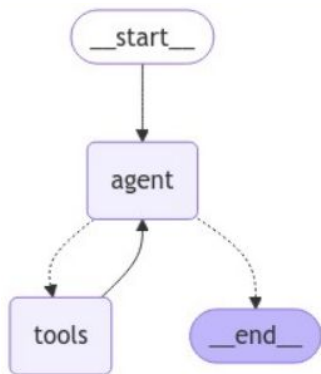
Key Aspects for Comparison



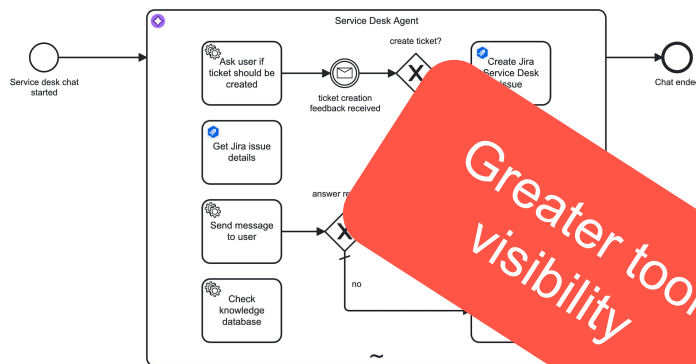
Agent Architecture & Control Flow



Code-first modular agent architecture that allows **dynamic control flow**



Process-first architecture with **static control flow** embedded in ad-hoc subprocess



Greater tool visibility

Example: Tool definition



```
@tool
def ask_user_for_ticket_creation(request_text):
    """
    This tool can be used to create a Jira ticket for the user at the service desk,
    """
    summary = get_llm().invoke(f"Generate a short, but descriptive title (do NOT exc
    description = get_llm().invoke(f"Generate a compact description for a jira issue
    ticket = create_jira_ticket(summary, description)
    return f"Successfully created a jira ticket with the following details: {ticket}"
```



SERVICE TASK
Ask user if ticket should be created

General

Name
Ask user if ticket should be created

ID
ask_user_for_ticket_creation

Documentation

Element documentation
This tool can be used to create a Jira ticket for the user at the service desk, so that someone can handle the request afterwards.

Template [+ Select](#)

Task definition

Job type *fx*
ask_user_for_ticket_creation

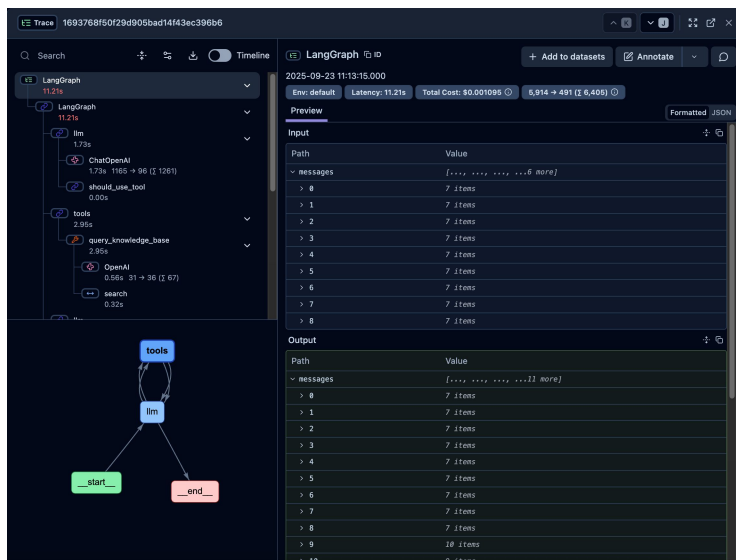
Retries *fx*

Inputs

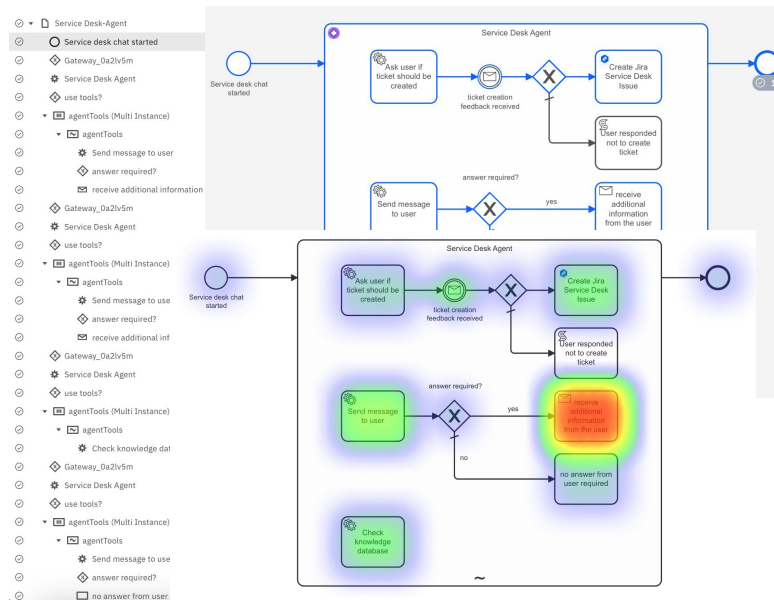
summary
Local variable name
summary
Variable assignment value *fx*
= fromAi(toolCall.summary, "A consive title for the request of the user", "string")

Variable: toolCallResult

Tracing & Observability



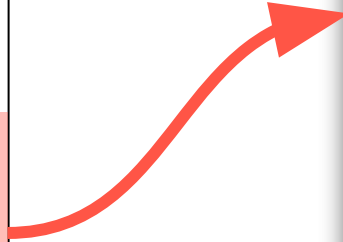
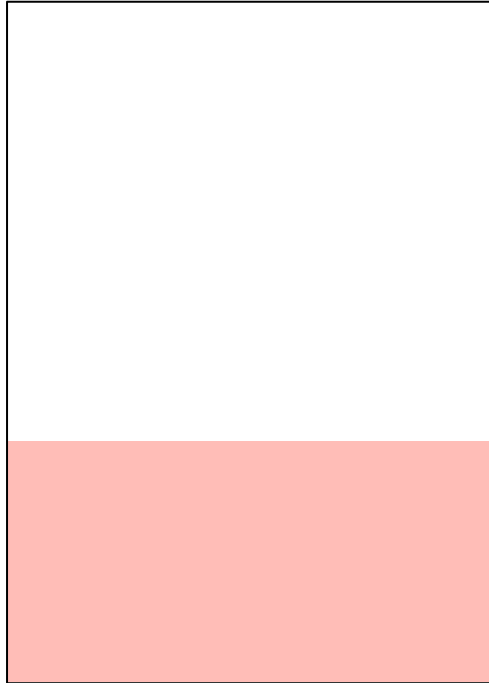
Graph-Level transparency



Process-level transparency



LLM observability in Camunda 8



Trace f85f8a0f-952b-4f37-a7c2-b3bfa4365a3f

Search [] Timeline []

Hotline Agent Camunda 8 3m 10s Conciseness: 0.70

- litellm:camunda8 1.30s 1196 → 33 (I 1229)
- litellm:camunda8 0.73s 1251 → 55 (I 1306)
- litellm:camunda8 0.81s 1330 → 25 (I 1355)
- litellm:camunda8 0.61s 7 → 0 (I 7)
- litellm:camunda8 2.61s 2100 → 122 (I 2222)
- litellm:camunda8 2.09s 2256 → 80 (I 2336)
- litellm:camunda8 1.62s 2301 → 144 (I 2445)
- litellm:camunda8 0.87s 2383 → 85 (I 2468)
- litellm:camunda8 1.05s 1712 → 76 (I 1788)
- litellm:camunda8 0.74s 1649 → 54 (I 1703)

litellm:camunda8 ID [] + Add to datasets [] Annotate [] Playground []

2025-09-05 08:34:40.348

Latency: 0.81s Time to first token: 0.81s Env: default \$0.000236

1330 prompt → 25 completion (I 1355) gpt-4o-mini temperature: 0.5 stream: false max_retries: 0

extra_body: [] system_fingerprint: fp_efad92c60b

Preview Formatted JSON

- ATP Web Portal: <https://atp-web.ewer.local/>

The answer must always be generated in the english language.

Show 4 more ...

Tool

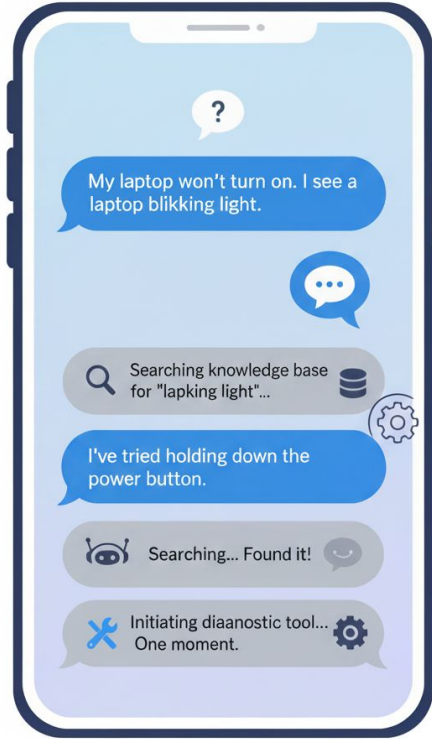
the feedback from the user was: The esc key is not owkring

Path	Value
json	1 items
tool_call_id	"call_j31A76vJtW3E1Z3VTEpaZxF"

Assistant

Path	Value
tool_calls	[{"function": ...}]
	3 items
function	2 items
arguments	1 items
queryKnowledgeBase	"keyboard issues, ESC key not working"
name	"query_knowledge_database"
id	"call_GgELfcvo7LAf00PXs6wxCz3o"
type	"function"
function_call	null
annotations	empty list

Conversational Agents



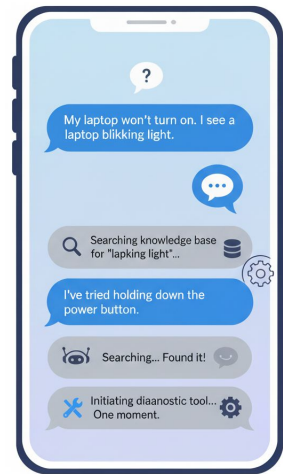
Conversational Agents



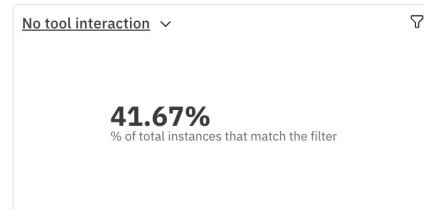
- Natively conversational
- **Natural language** is the **control interface**



- BPMN not natively dialog-oriented
- Agent can become conversational by using **LLM-based tools**
- Less natural for **free-flowing, dynamic dialogue** compared to LangGraph



Prompting important,
but also not reliable



Control Mechanisms



- Governed by **schema validation**, **tool restrictions**, and **moderation layers**
- Allows **dynamic guardrails** during execution, maintaining control
- More flexible, but requires **custom policy design** and guidance



- Uses **BPMN** and **DMN** for a structured flow of control during execution
- Strong support for **process-level policies**, approvals, and roles
- Easy to integrate with **enterprise governance** and compliance systems



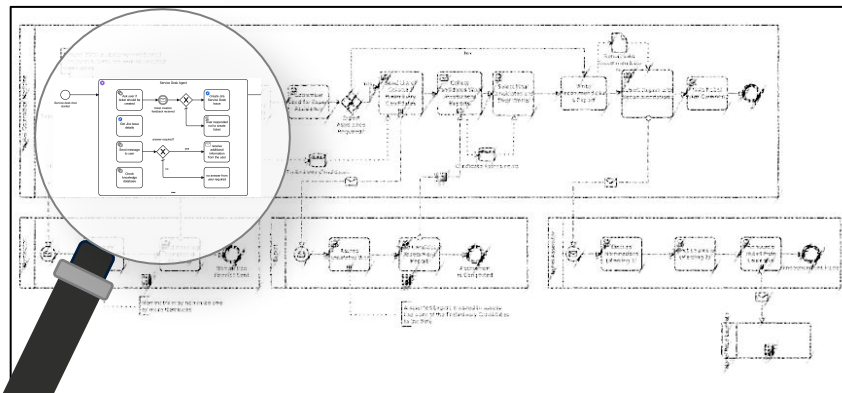
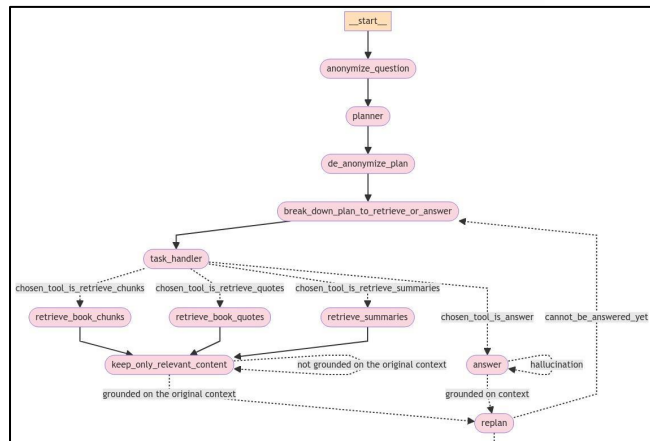
Both platforms support guardrails -
but in different ways



Comparing Apples & Oranges?



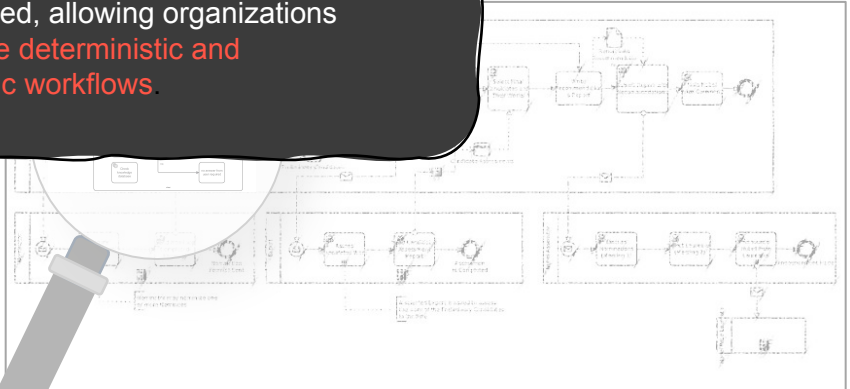
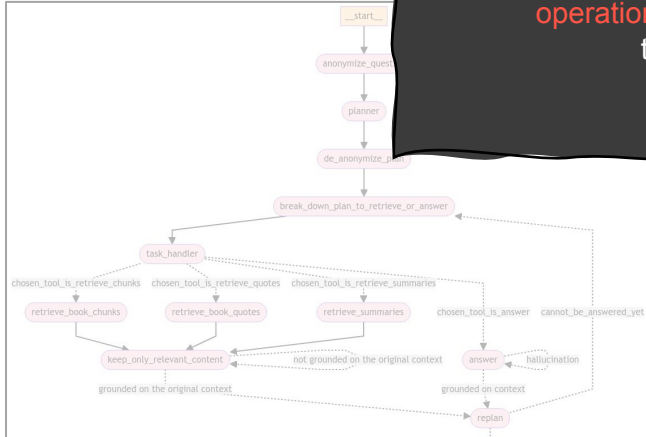
VS



Comparing Apples & Oranges?



LangGraph is specialized to build agents, whereas Camunda 8 is designed to automate end-to-end business processes and operationalize agents as needed, allowing organizations to seamlessly combine deterministic and non-deterministic workflows.





**IT COMES DOWN TO USE
CASES!**

viadee 

Thank you!

Questions?



Andre Strothmann
Andre.Strothmann@viadee.de

viadee Unternehmensberatung AG
Anton-Bruchausen-Str. 8
48147 Münster
Tel: +49 251 77777 377
www.viadee.de



Mario Micudaj
Mario.Micudaj@viadee.de

viadee Unternehmensberatung AG
Anton-Bruchausen-Str. 8
48147 Münster
Tel: +49 251 77777 138
www.viadee.de

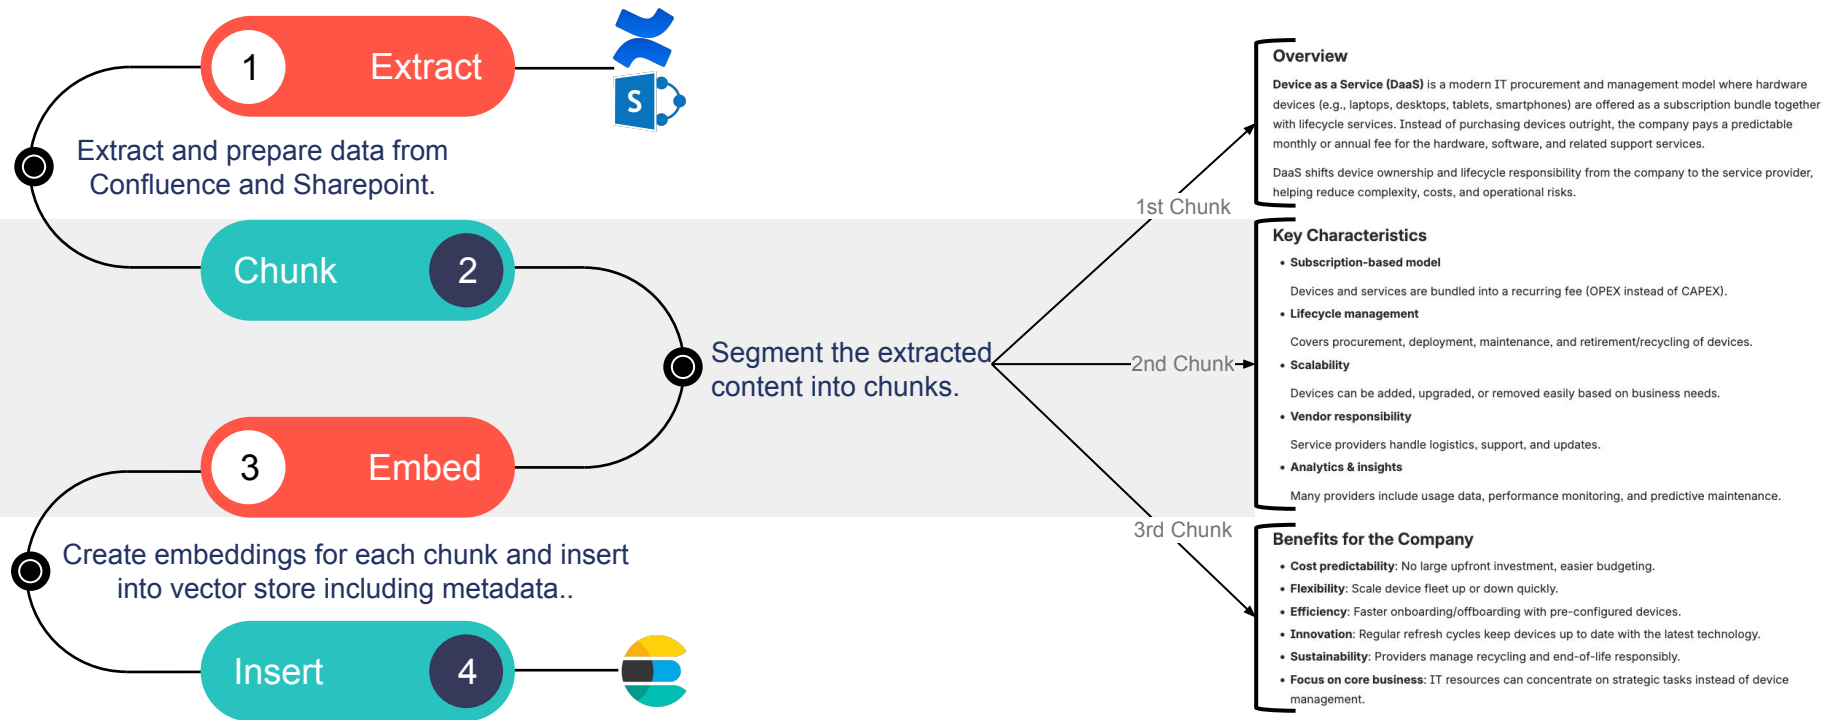


© viadee



viadee Unternehmensberatung AG
viadee.de

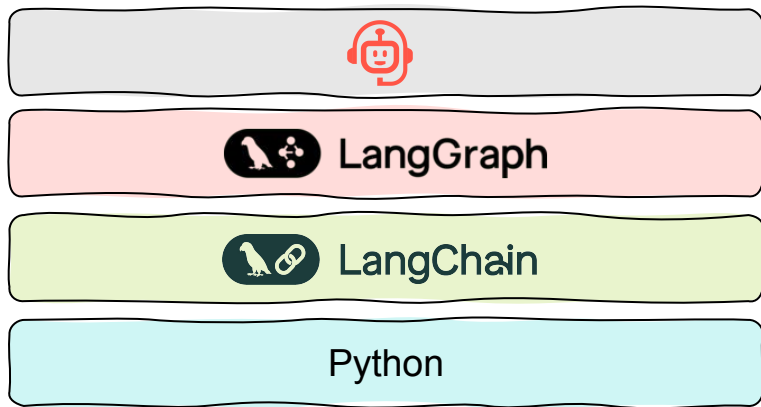
Preparing the knowledge base



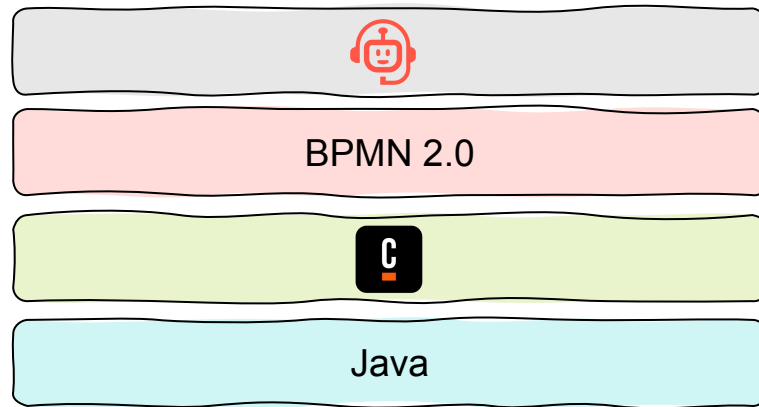
Technology Stack



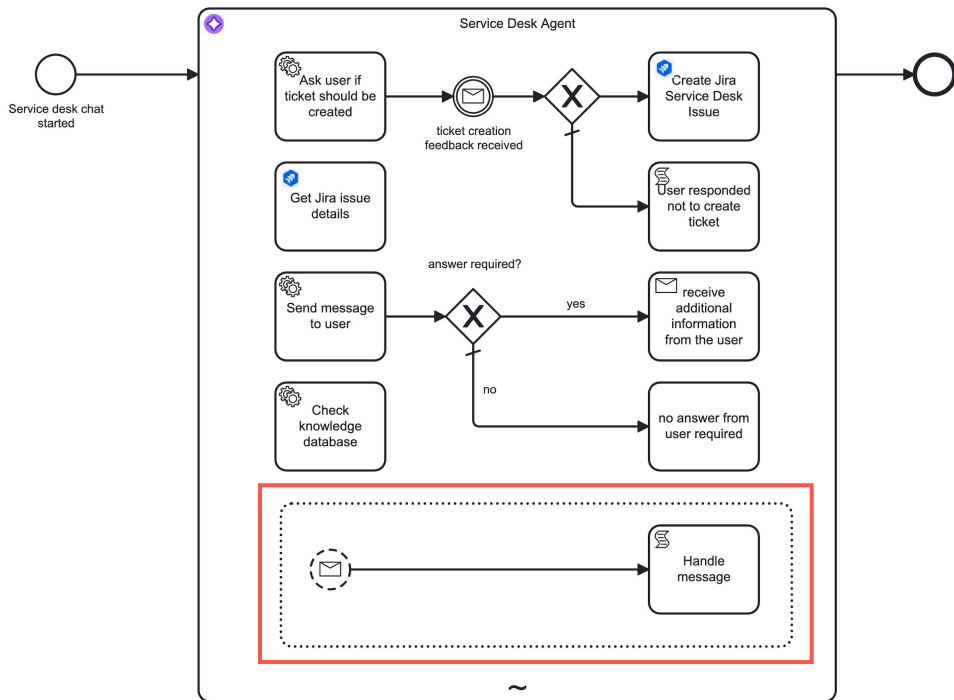
- low-level orchestration framework for building ai agents
- Code-oriented




- AI Agent connector based on LangChain4j
- Model-oriented



Event subprocess in AI Agent



 **AI AGENT**
Service Desk Agent

Limits

Event handling

Event handling behavior

☒ Wait for tool call results

☒ Interrupt tool calls

Support for „interrupting“
message input

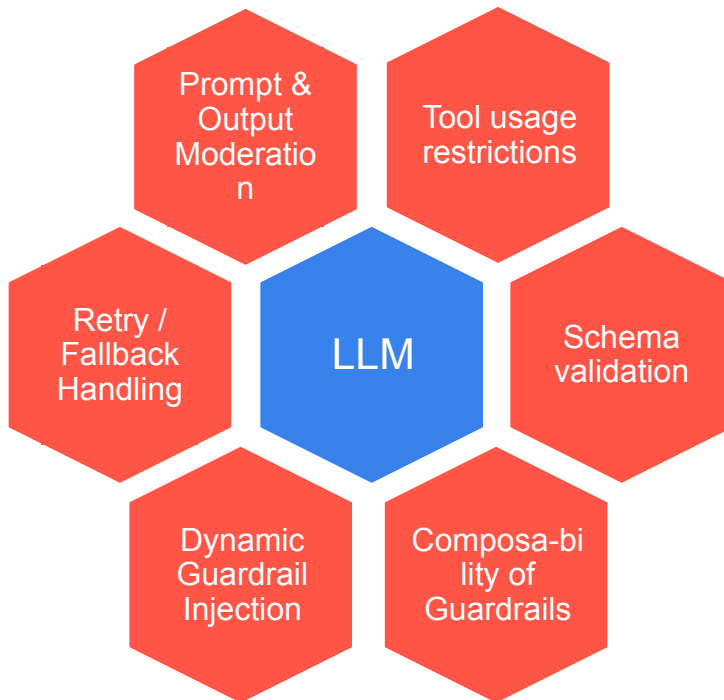
Guardrails in LLM Model provider



LLM guardrails

Guardrails
on tool
output

Guardrails
on graph
edges



LLM guardrails of
model provider

Limited by
modelling
possibilities