



Azure AI Foundry and the dev toolchain to infuse AI in all your apps

Callum Anderson, Global Director, DevOps & SRE, Dentsu
Marlene Mhangami, Senior Developer Advocate, Microsoft
Trupti Parkar, Product Manager II, AI Platform, Microsoft



About me

Sarin Suriyakoon

Deputy Head of Tech

PALO^{IT}



PALO IT

A **global tech innovation consultancy** dedicated to achieving your business goals through next-level product-centric software delivery.

We provide expertise across the full product journey from ideation to launch and scale, using Sustainable IT practices as a commitment to crafting tech as a force for good.

50
Nationalities

18
Offices

10
Countries



25%
Organic Growth

\$70M
Turnover

700+
Technology experts from across the globe

100%
Independently-owned

0
Debt


Information Classification: CONFIDENTIAL (sensitive business information, the level of protection is dictated



GitHub Verified Partner



DOWNLOAD
PRESENTATION



What slows down Generative AI adoption?

Getting Started

The state of the art is evolving so quickly, it makes it difficult to decide what to use. Along with that, guidance and documentation is hard to find.

Development

Applications often require multiple cutting-edge products and frameworks which requires specialized expertise and new tools to stitch these components together.

Context

Large Language Model doesn't know about your data

Evaluation

It is hard to figure out which model to use and how to optimize for their use case.

Operationalization

Concerns around privacy, security, and grounding. Developers lack the experience and tools to evaluate, improve and validate the solutions for their Proof of Concepts, and to scale and operate in production.

What is Operationalization?

1



For Business

It's about streamlining essential capabilities within AI practice, including the integration of tools and processes necessary for the **successful delivery of AI-driven business outcomes**.

It's also about successfully **embracing change to propel organizational success**, capitalizing on the advantages offered by outcome LLM capabilities.

2



For IT

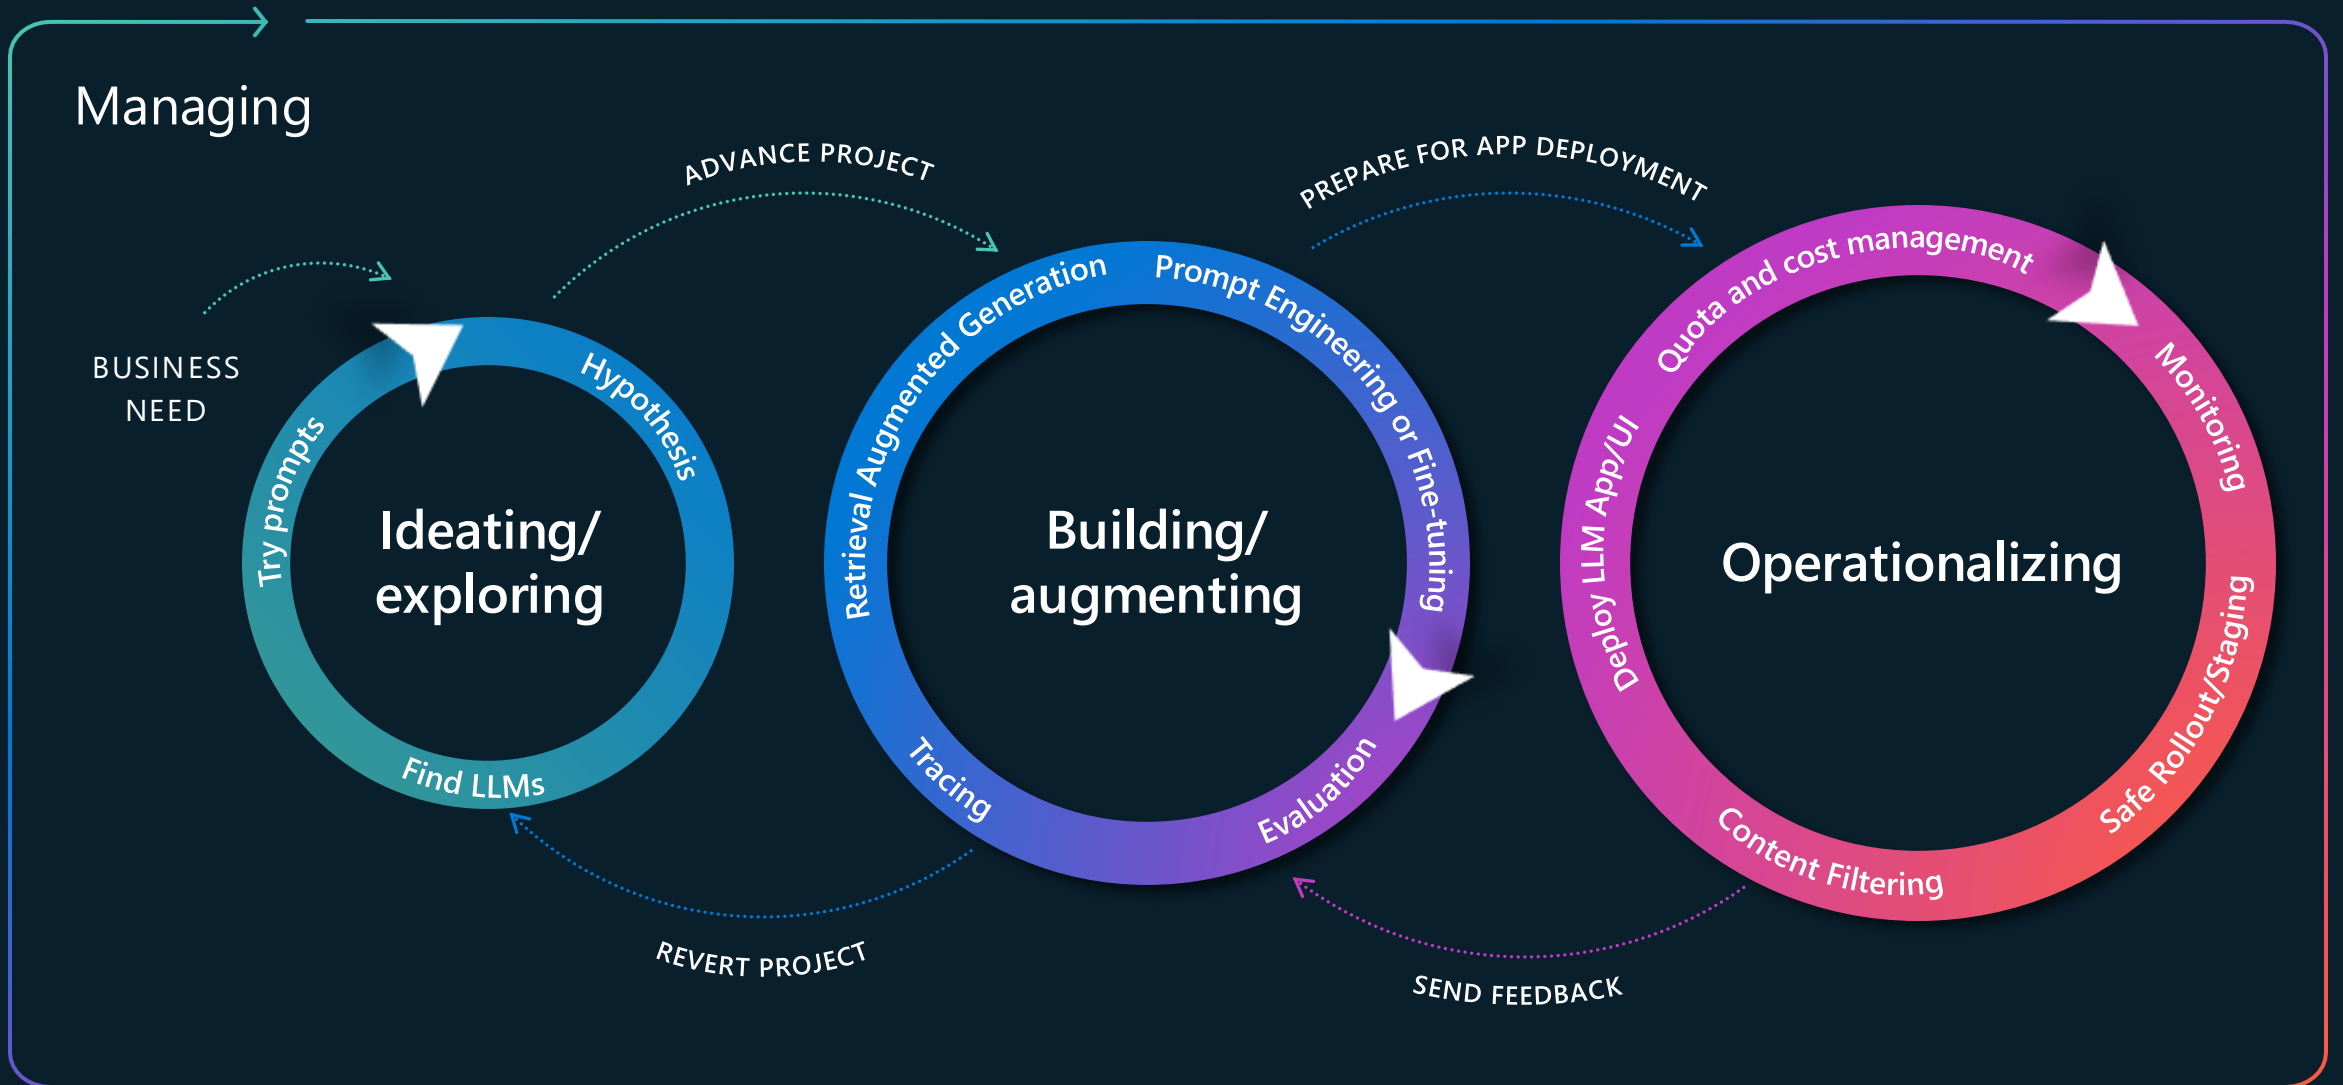
It's about **navigating technical complexities** while integrating the LLMs into real-world solutions.

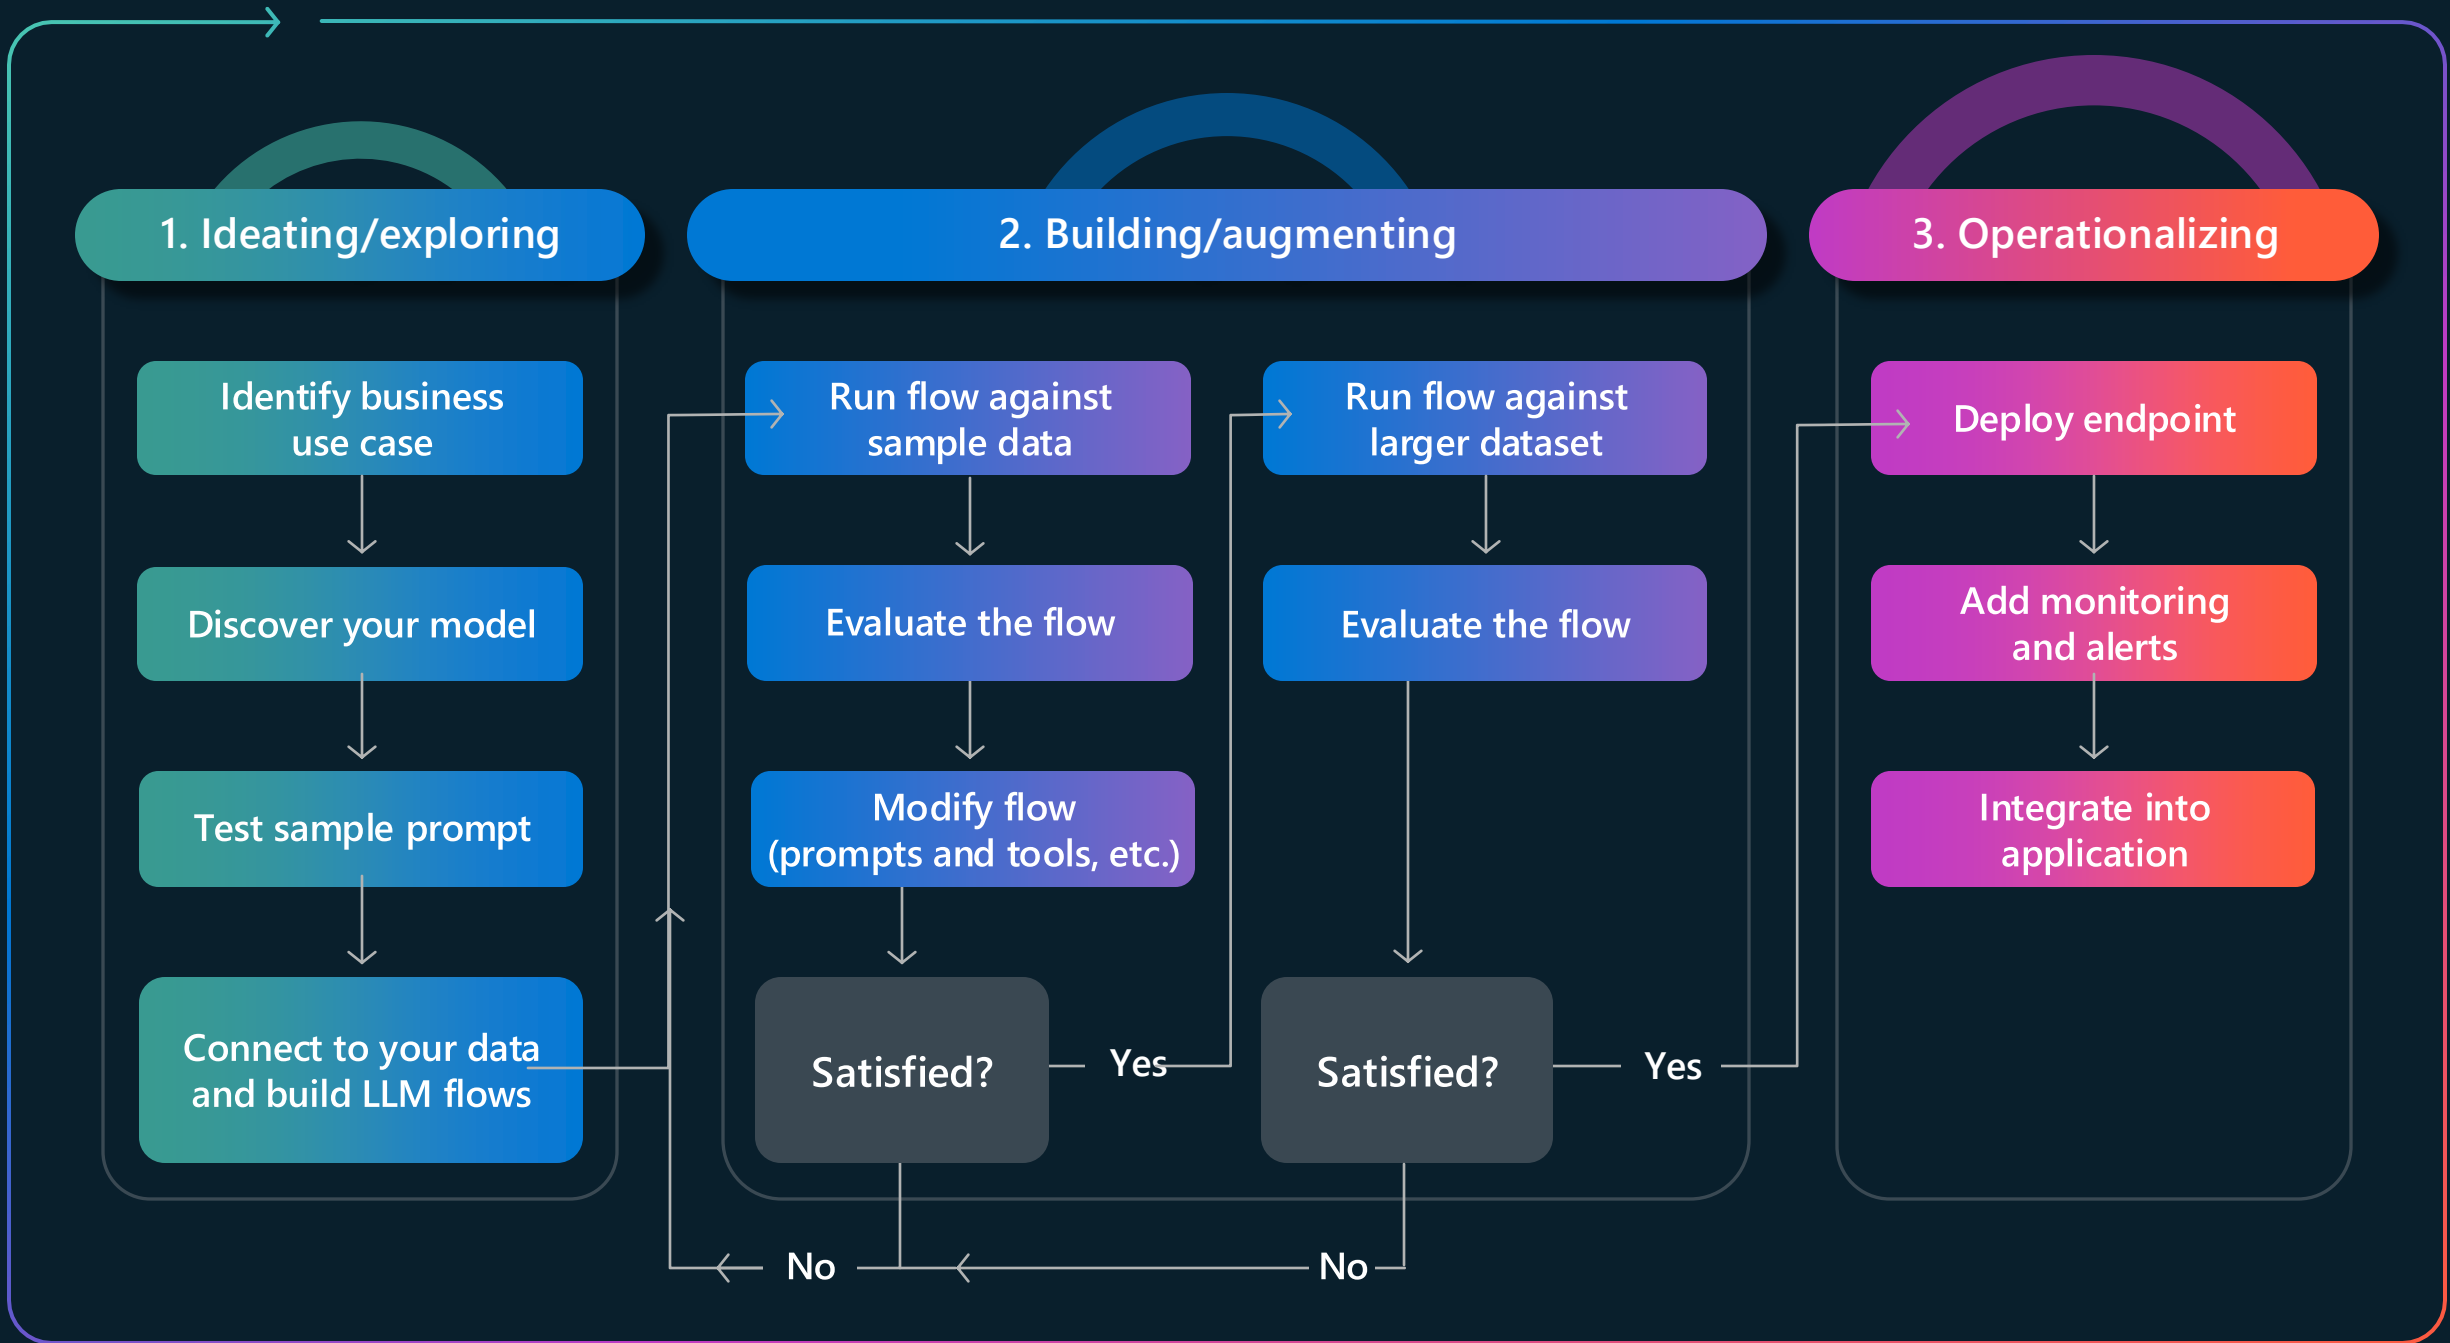
It's also about harmonizing non-deterministic "learners" associated with AI, with traditional **Solution Lifecycle and Service Management** principles.

The paradigm shift from MLOps to GenAIOps

	Traditional MLOps	GenAIOps
Target audiences	ML Engineers Data Scientists	AI Engineers App developers
Assets to share	Model, data, environments, features	LLM, agents, plugins, prompts, chains, APIs
Metrics/evaluations	Accuracy, Sensitivity	Quality: accuracy, similarity Harm: bias, toxicity Correct: roundedness Cost: token per request Latency: response time, RPS
ML models	Build from scratch	Pre-built, fine-tuned served as API (MaaS)

GenAI Lifecycle in the real world



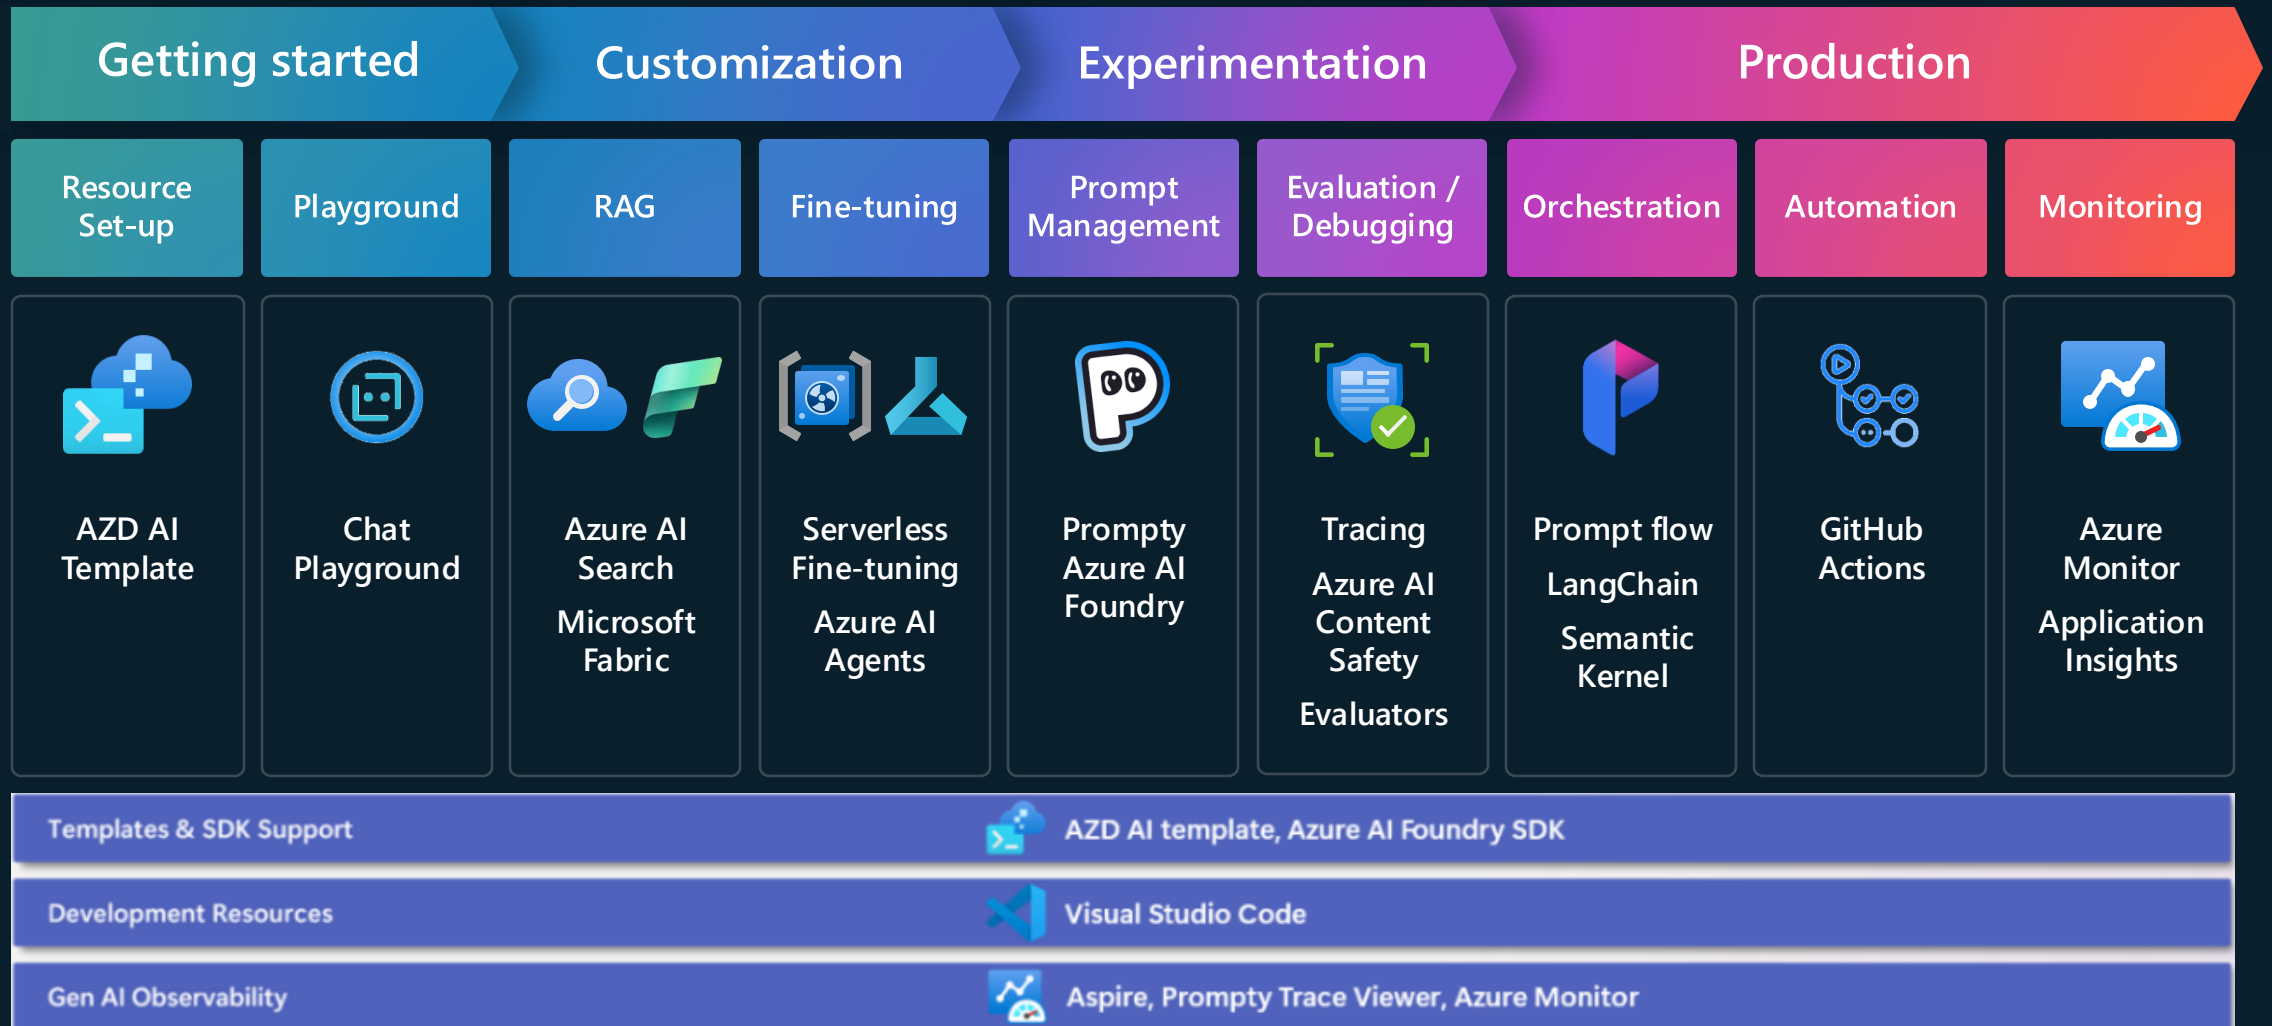


Updated

Where do I start?

Robust GenAIOps Toolchains in Azure AI

Bringing your AI apps to production



Building the application



Local dev environment to the cloud in a single step

Azure Developer CLI

```
azd init  
azd up  
azd monitor  
azd pipeline config
```

azd commands



Your app



azd up



The cloud

Q&A

