

Stop Building AI Models. Start Building AI Products

Value comes from products. Not from models.

Fabrizio Lazzaretti

Senior Managing Consultant @ Wavestone | CNCF Ambassador | Speaker | Author

AI Summit Europe, Riga

07. May 2026

WAVESTONE



Disclaimer: I Love AI



At work

Currently leading multiple AI initiatives: LLMs, multi-agent systems and AI driven development.

A collection of technology logos arranged in a grid. The logos include: Kubernetes (blue ship's wheel), Prometheus (orange flame), Cilium (hexagonal pattern), Talos Linux (red and orange vertical bars), Ceph (red concentric circles), YOLO (black text), NVIDIA (green eye), and AMD Radeon Graphics (red and black).

kubernetes Prometheus
cilium Talos Linux ceph
YOLO NVIDIA RADEON GRAPHICS AMD

At home

Self-hosted Kubernetes cluster with GPU support running YOLO image classifier.



But

Business goal \neq technical detail. In production, optimize for business success.

I love AI. I see huge potential.
The talk that follows is about what often goes wrong on the way there and how to fix it.

95%

of GenAI projects fail to deliver measurable business returns.

It's not a new pattern

Different decade. Different model. Same shape of failure.

2018

GARTNER

85% of AI projects deliver erroneous outcomes.

Bias in data, algorithms, or the teams managing them.

Gartner forecast

Nov 2021

ZILLOW

Home-buying algorithm couldn't predict prices.

\$300M+ write-down. Zillow Offers shut down.

GeekWire

2008-15

GOOGLE FLU TRENDS

Worked 3 years, drifted 4 more before retirement.

"The parable of Google Flu: traps in big data analysis."

Lazer et al., Science 2014

Same shape every time: the model worked. The product didn't.

“

We're going faster. But not always in the right direction.

The faster we go, the further off course.

Some requests for AI I get are not founded in business needs



1

No optimization target

Ideas of using AI for process steps that can't be measured/optimized.

2

No training signal

Ideas of using a model without enough (or any) real training data.

3

No bottleneck

The step they want to optimize with AI is not the real problem, it will not bring value.



No business value → no ROI on the product
→ no happy customer.

“

*The model isn't the problem.
The undefined goal is.*

Sound familiar?

The reflex predates AI. We've seen this exact pitch before.

2018

“We need blockchain in our app.”

Why? *It's the future.*

Solving what? *Unclear.*

Status today? *Failed*

2024+

“We need a chatbot.”

Why? *It's the future.*

Solving what? *Unclear.*

Status today? *But it has to use AI.*

It's the same pattern, new tool

“We need an MCP server.”

Why? Everyone's building one. For what? TBD. Mgmt wants one.

SOAP / XML

1998 - SOAP, WSDL

But why?

REST / JSON

2005+ - REST, OpenAPI

But why?

Events

2016+ - Kafka

But why?

MCP

2025+ - Model Context Protocol

But why?

“We need an MCP server.”

Why? Everyone's building one. For what? TBD. Mgmt wants one.



98.7%

token reduction by replacing direct MCP tool calls with code execution

Anthropic Engineering, Nov 2025 · anthropic.com/engineering/code-execution-with-mcp

Four eras. Same reflex. — “A fool with a tool is still a fool.”

“

Let's measure the outcome

So, we just quickly define 1-2 KPIs?

A story of wrong KPIs.

Even when the metric is defined, the metric might not be the value.

THE CASE

Online portal (B2C) – We need to grow!

KPI management has set: Amount of registered users.

Result we hit: the KPI — every quarter.

Result we wanted: more customers that will use the portal and use the services.

We attracted the wrong people. The KPI was rising. The business wasn't.

Metric



Business
value




Goodhart applies.

“When a measure becomes a target, it ceases to be a good measure”


Business case has no AI in it

AI Chatbot




I want customers to get cheap, quick, highly customized support

MCP for Service X



I want X to have up-to-date data for use case Y

Business case has no AI in it



EXCEPTION


Real transformative use cases

- Pair-programmers like Cursor and Copilot.
 - Search that answers like Perplexity.
- => The product the customer buys is genuinely different.

Most products with AI are not like this.

From outside, AI changes nothing.

It's a black box — and that's the point. Same request in, same decision out.



From the customer's view

The customer should not be able to tell where the AI is.
The **request goes in; the decision comes out.**

Marketing thinking vs product thinking

“AI camera!” — the customer wanted a better photo.
Marketing changed; the product didn't.

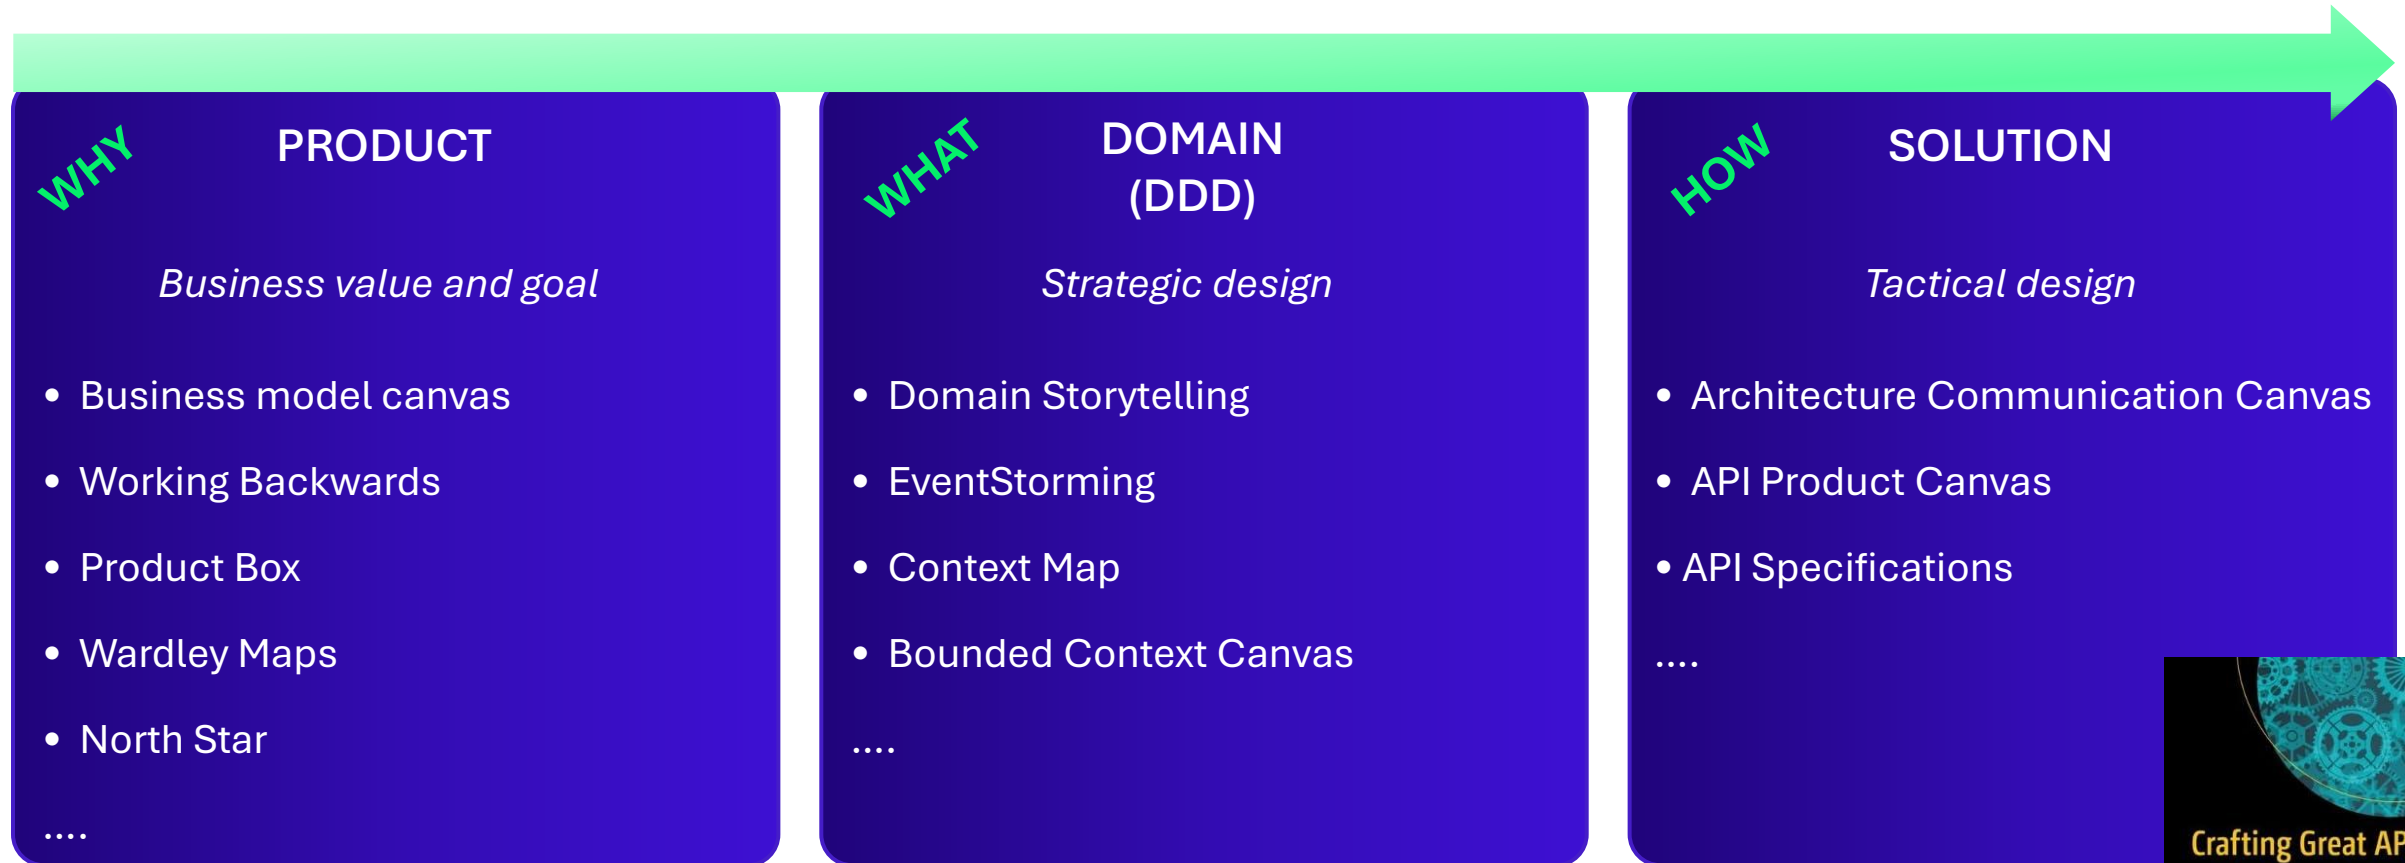
“

You can't optimize what you can't define.

Anyone who claims otherwise is just overfitting to a proxy metric.

How do we get from a business goal to a working AI product?

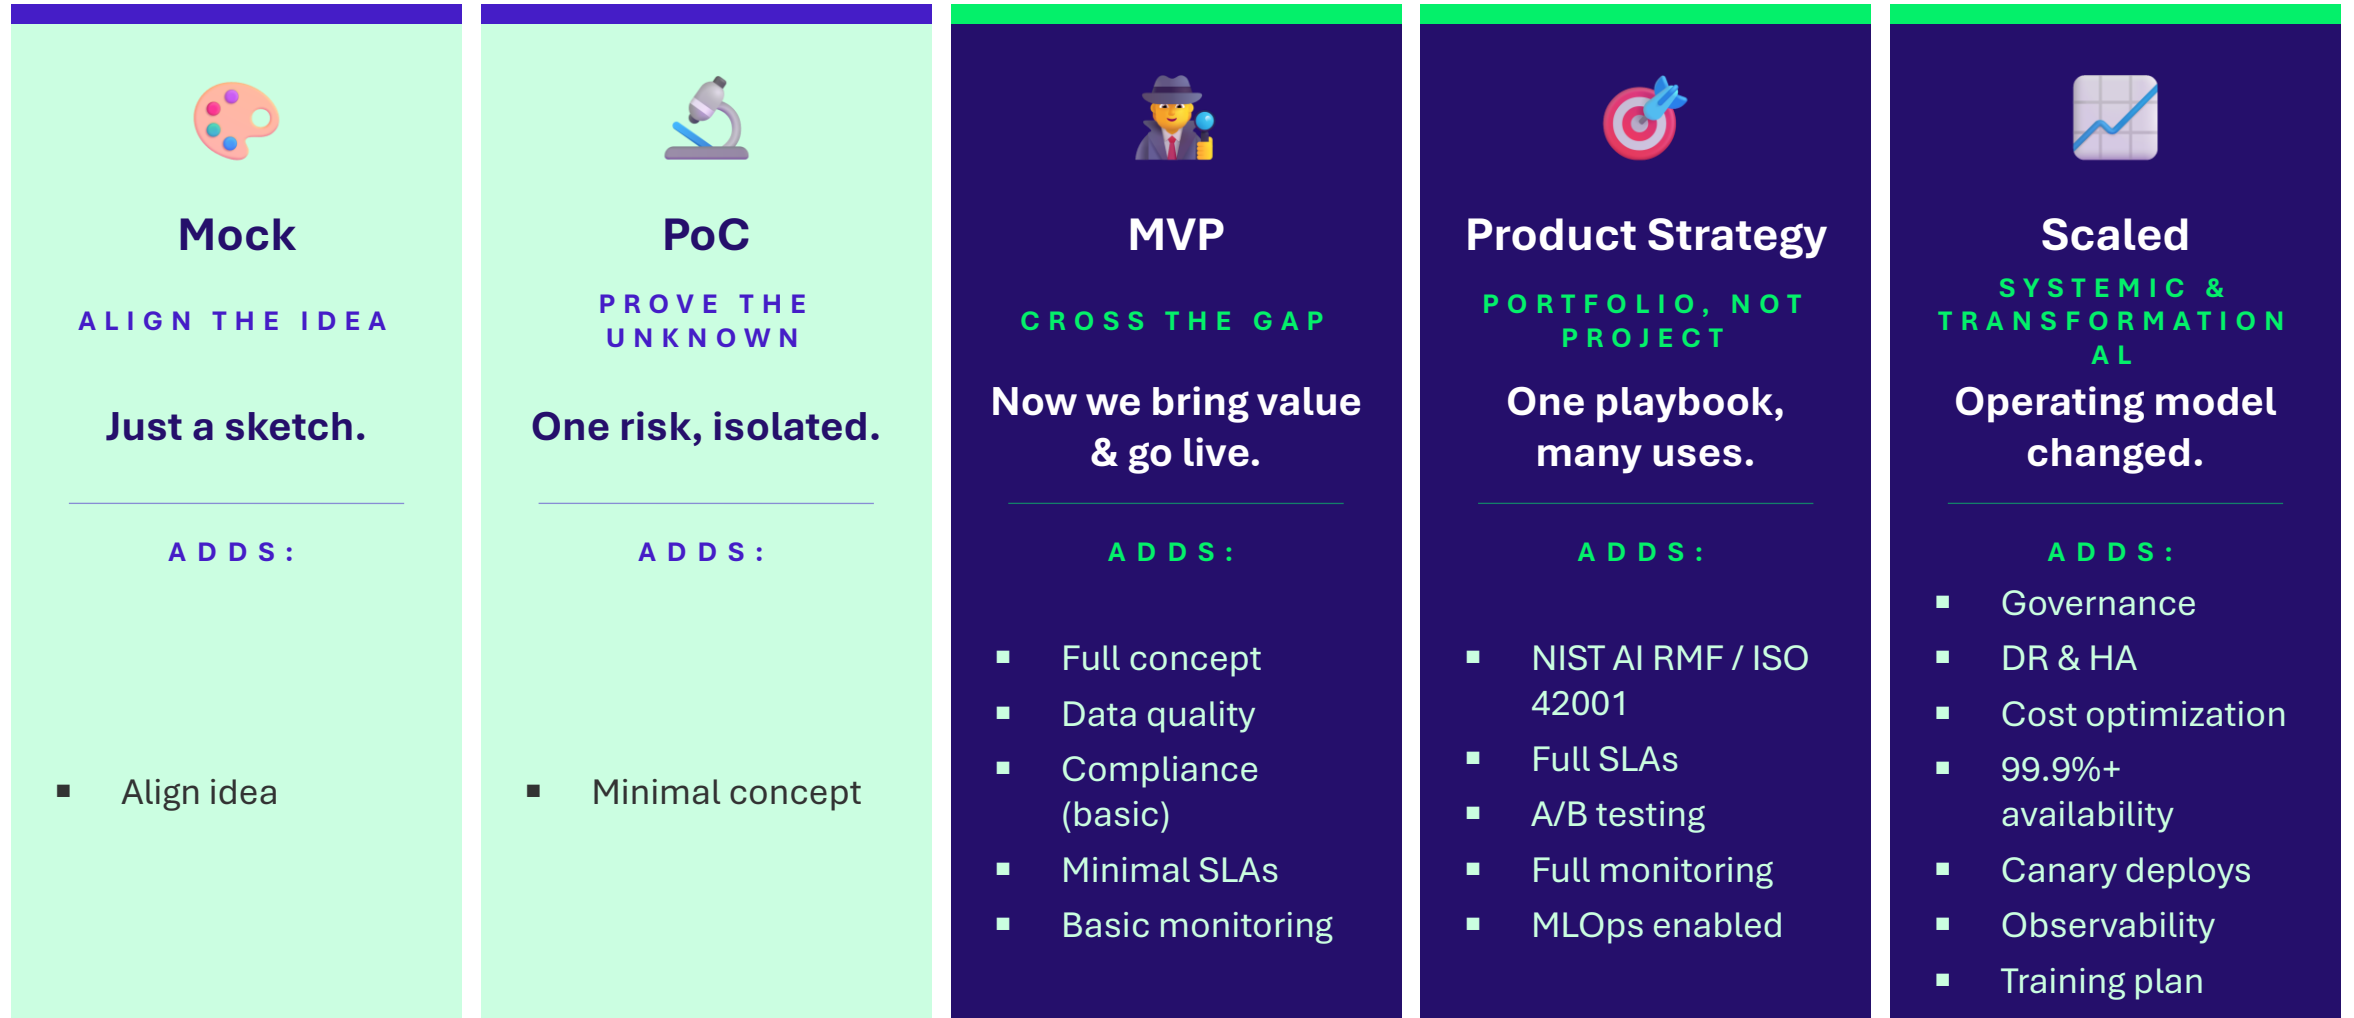
Three layers: From value to build



**Crafting Great APIs
with Domain-Driven
Design**


Craftsmanship of Asynchronous
and Synchronous APIs
—
Annegret Junker
Fabrizio Lazzaretti

The 5 phases: What actually changes



Features can change depending on use case

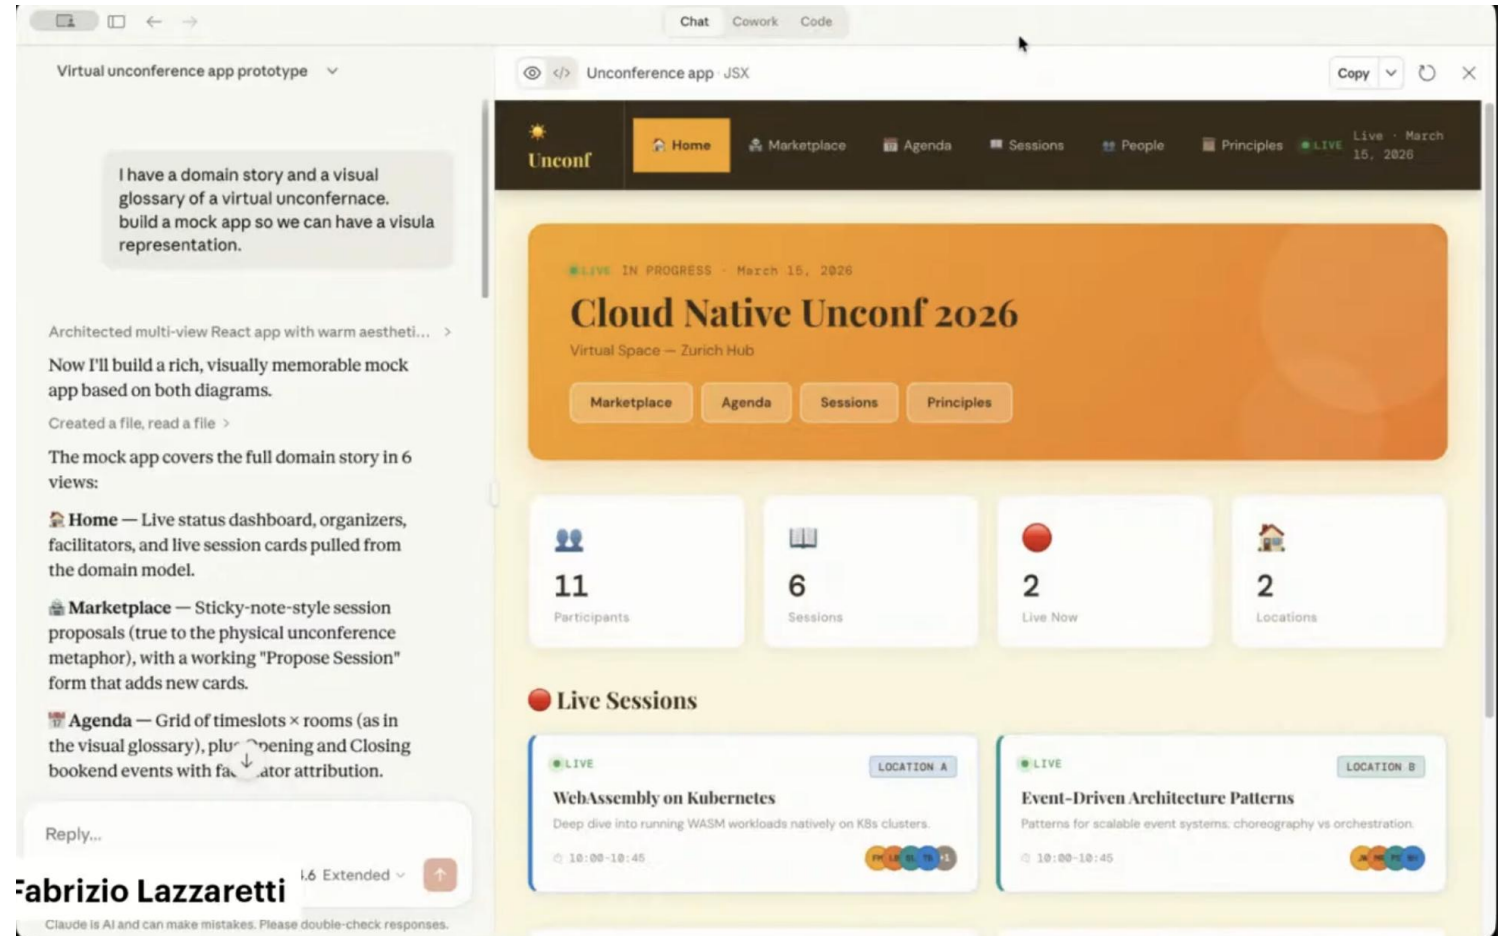
Phase 1: 🎨 Mock



Quickly make the idea visible

Discuss:


- UX: Does this feel like the product they actually want?
- Which data is needed?
- What success would look like — before any code is written (except AI generated frontend code).



Further reading: <https://lazzaretti.me/post/talks/2026/loosely-coupled-domain-storytelling/>

Phase 2: PoC

Prove the unknown. One risk at a time. Define success before you start.



Q 1

What do we want to prove?

One thing. Technical, data-side, or product-side.

Q 2

When have we proven it?

What's in scope, what's out of scope.

Q 3

How will we measure success?

A number, not a feeling.


WHEN A PoC EARNS ITS PLACE

Run a PoC only when the unknown is technical or data-side. Multiple PoCs may be needed — each proves a different thing.

If everything is clear, skip the PoC. It's a demo, not a proof, if you can't define what success looks like.

Phase 3: MVP

The PoC worked. Now the real questions start.



BUDGET

- Initial cost
- Operational cost
- Iterative update cost
- What are the costs if it fails?

SUCCESS

- When is it successful?
- How is quality defined?
- Which KPIs do we have, and how measured?

ARCHITECTURE

- Loosely coupled?
- Clean APIs between components?
- Sensible service cuts?
- Security?

RESILIENCE

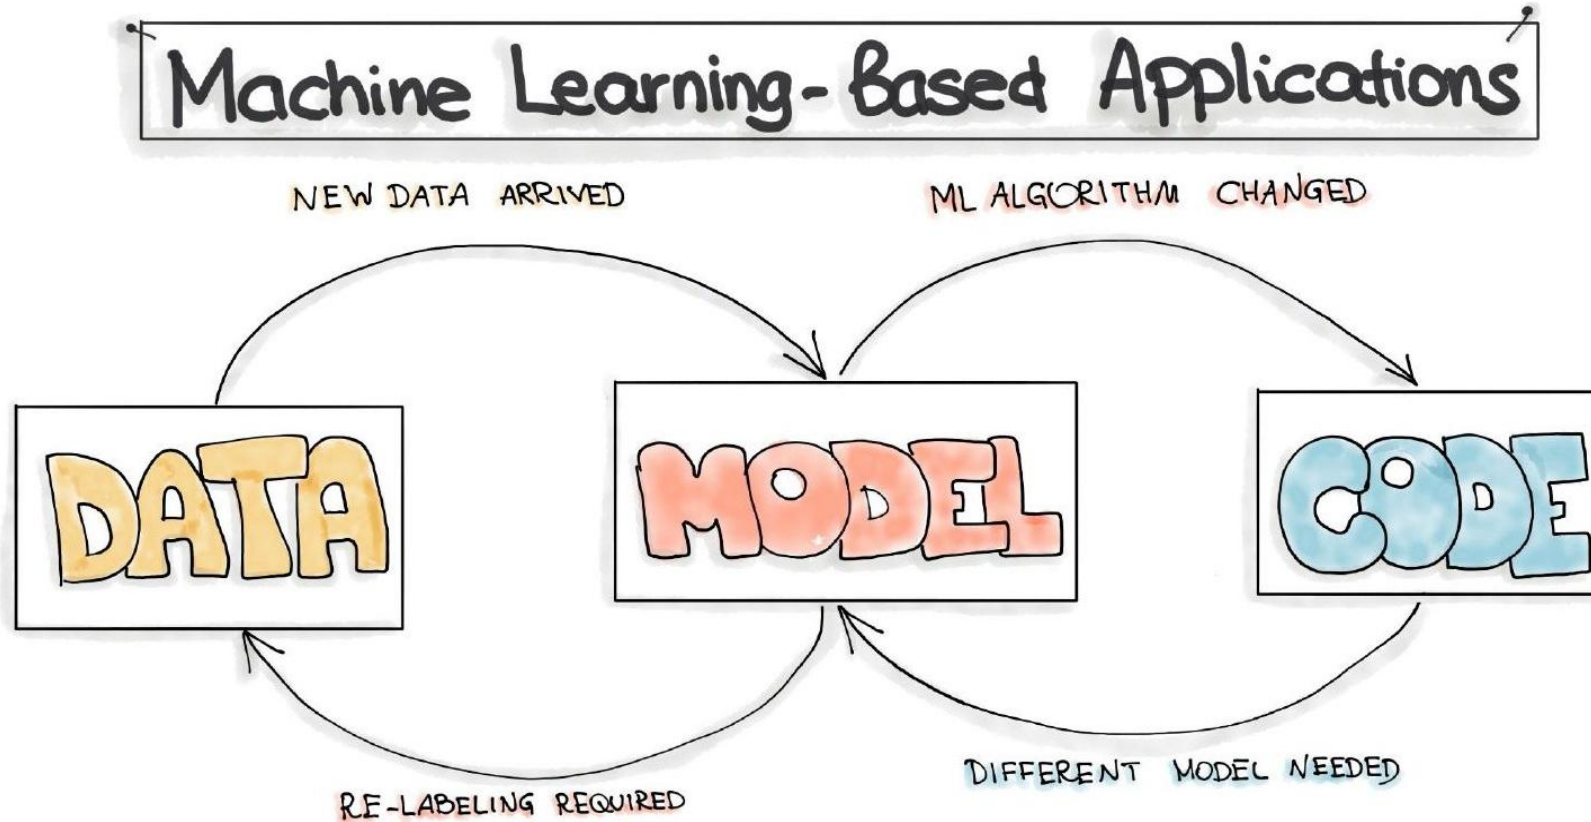

- What if it doesn't work?
- Recovery time after failure
- Stochastic failures — plan for them

SECURITY & COMPLIANCE

Security risk, business risk, risk acceptance. · What if data leaks? · And... who pages on incidents at 3am?

Phase 4: 🎯 Product Strategy


The MVP is live. Now let's build a long-lasting strategy.



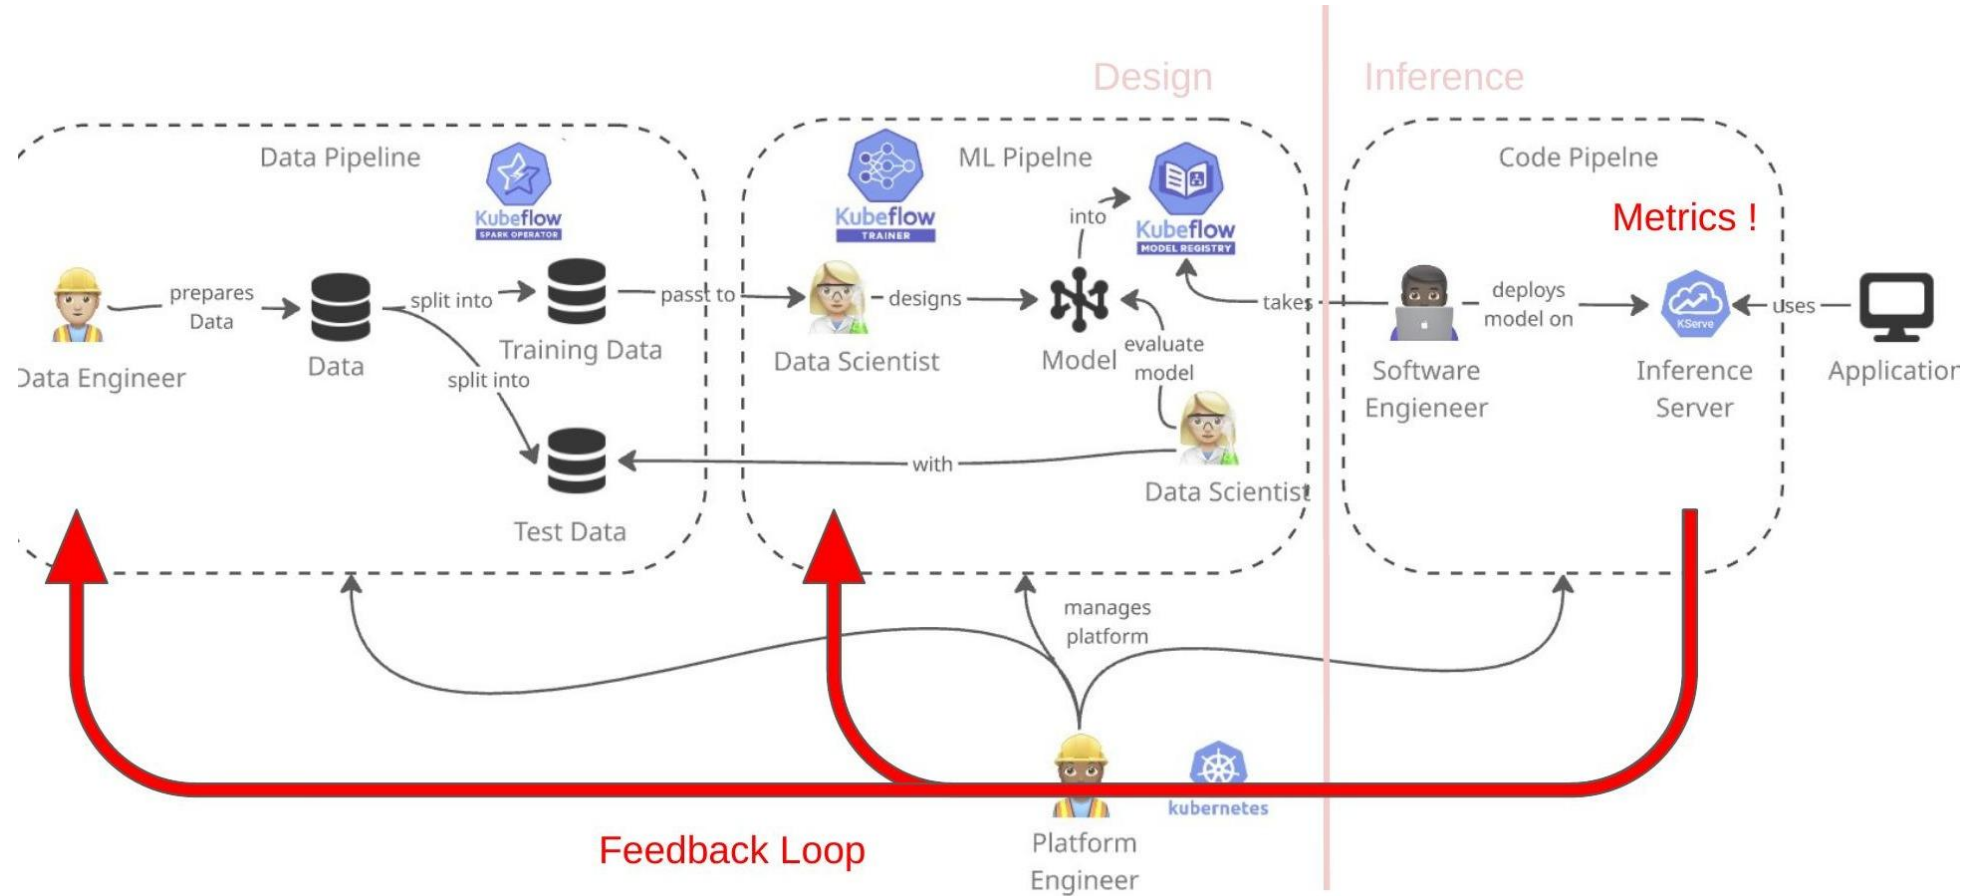
PLUS

Compliance, Full SLAs, A/B testing, Full monitoring,

Phase 5: Scaled



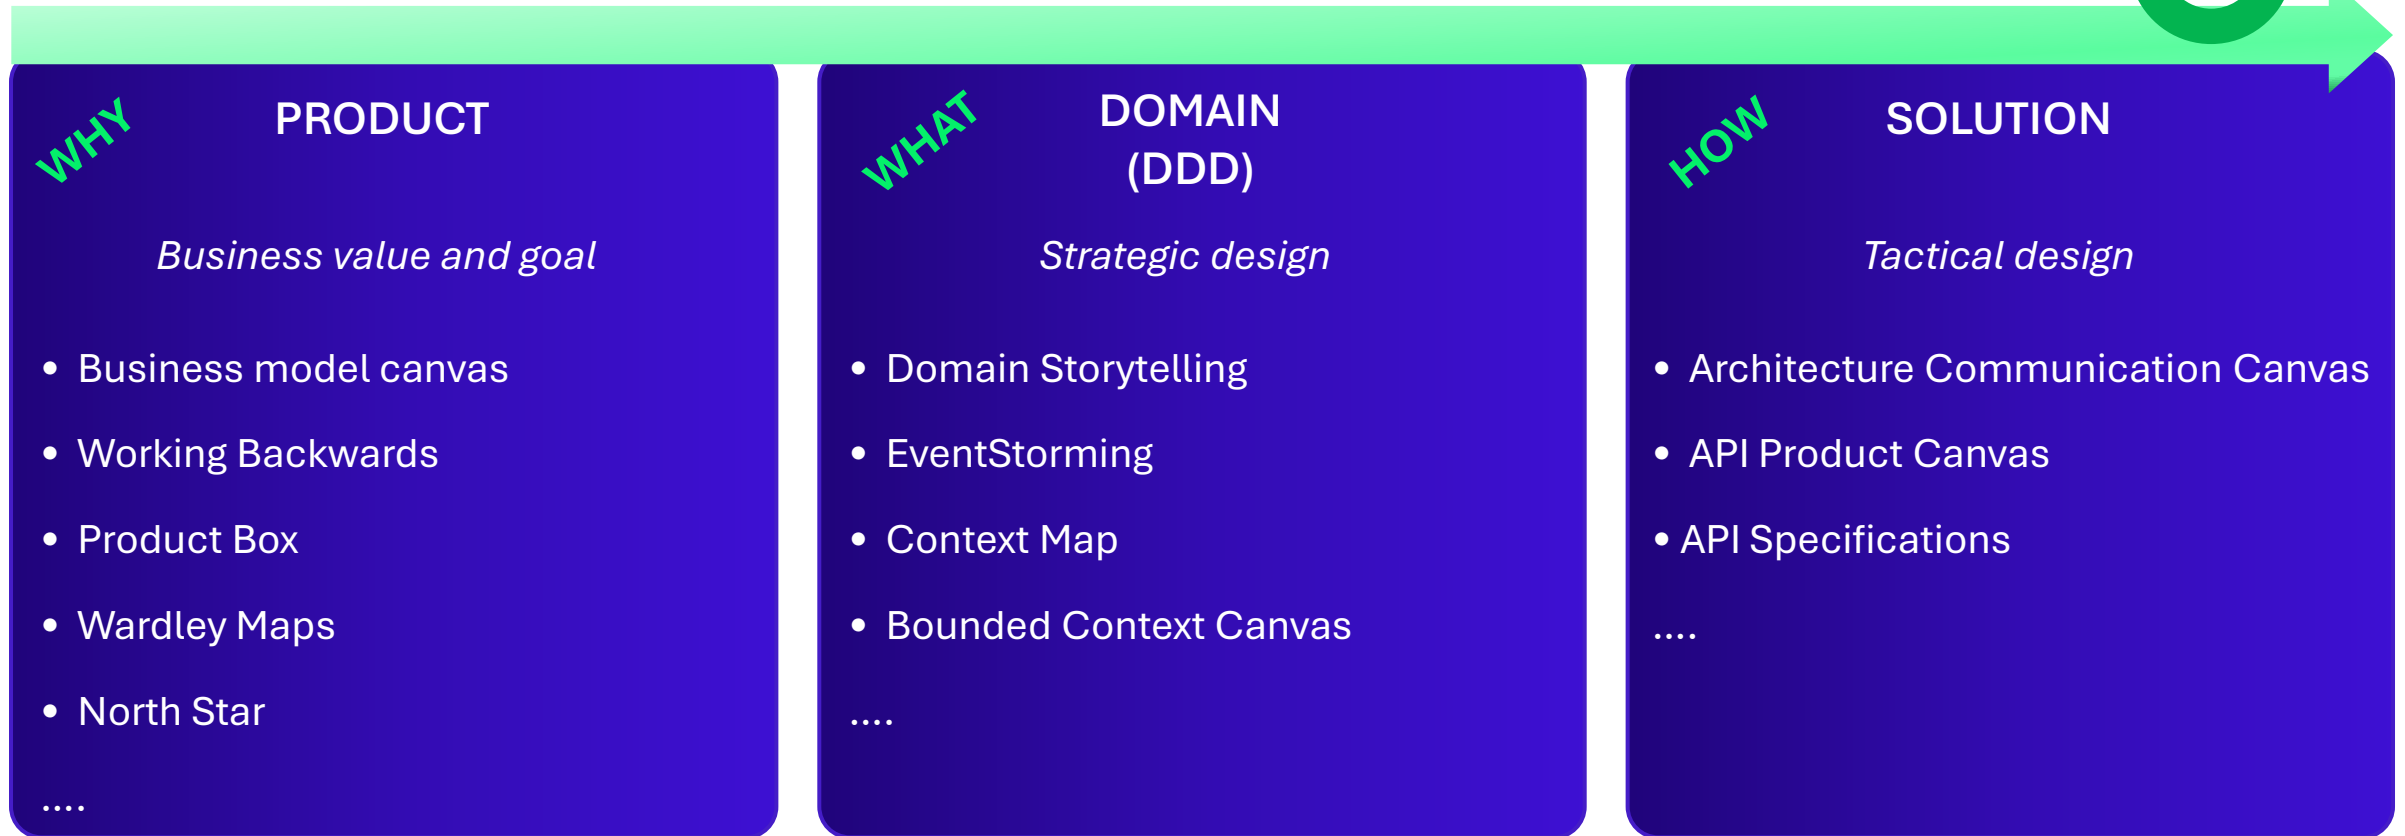
Production metrics feed back into the data pipeline. Without the loop, MLOps is just deployment.



PLUS

Governance, DR & HA, cost optimization, 99.9%+ availability, canary deploys, observability, training plan, ...

From value to build, phase by phase



Mock

MVP

PoC

Product Strategy

Scaled

Let's do an example

— The Online Library



**Crafting Great APIs
with Domain-Driven
Design**

Craftsmanship of Asynchronous
and Synchronous APIs

—
Annegret Junker
Fabrizio Lazzaretti

Apress®

The Product

SCOPE

Community driven online library getting money by donations

GOAL

We want to make the book licensing process simpler, better, and less time consuming

AI?

Maybe AI can help?



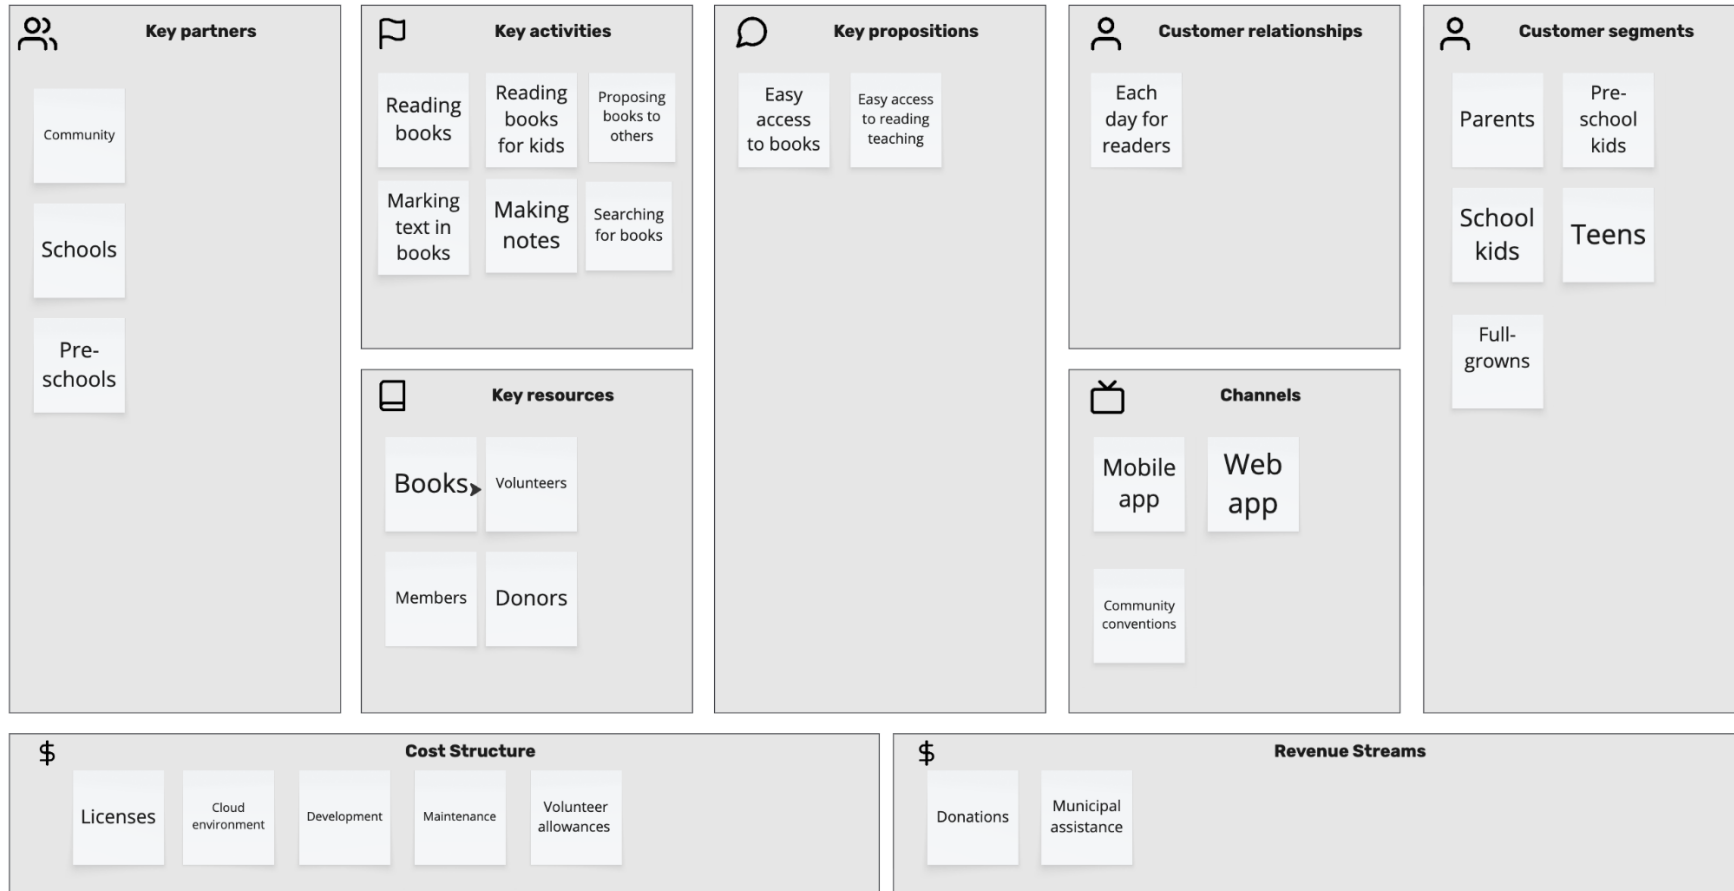
**Crafting Great APIs
with Domain-Driven
Design**

Craftsmanship of Asynchronous
and Synchronous APIs

—
Annegret Junker
Fabrizio Lazzaretti

The Product: Business Model Canvas

The Business Model Canvas



What is the product?
When is the product successful?

Crafting Great APIs
with Domain-Driven
Design

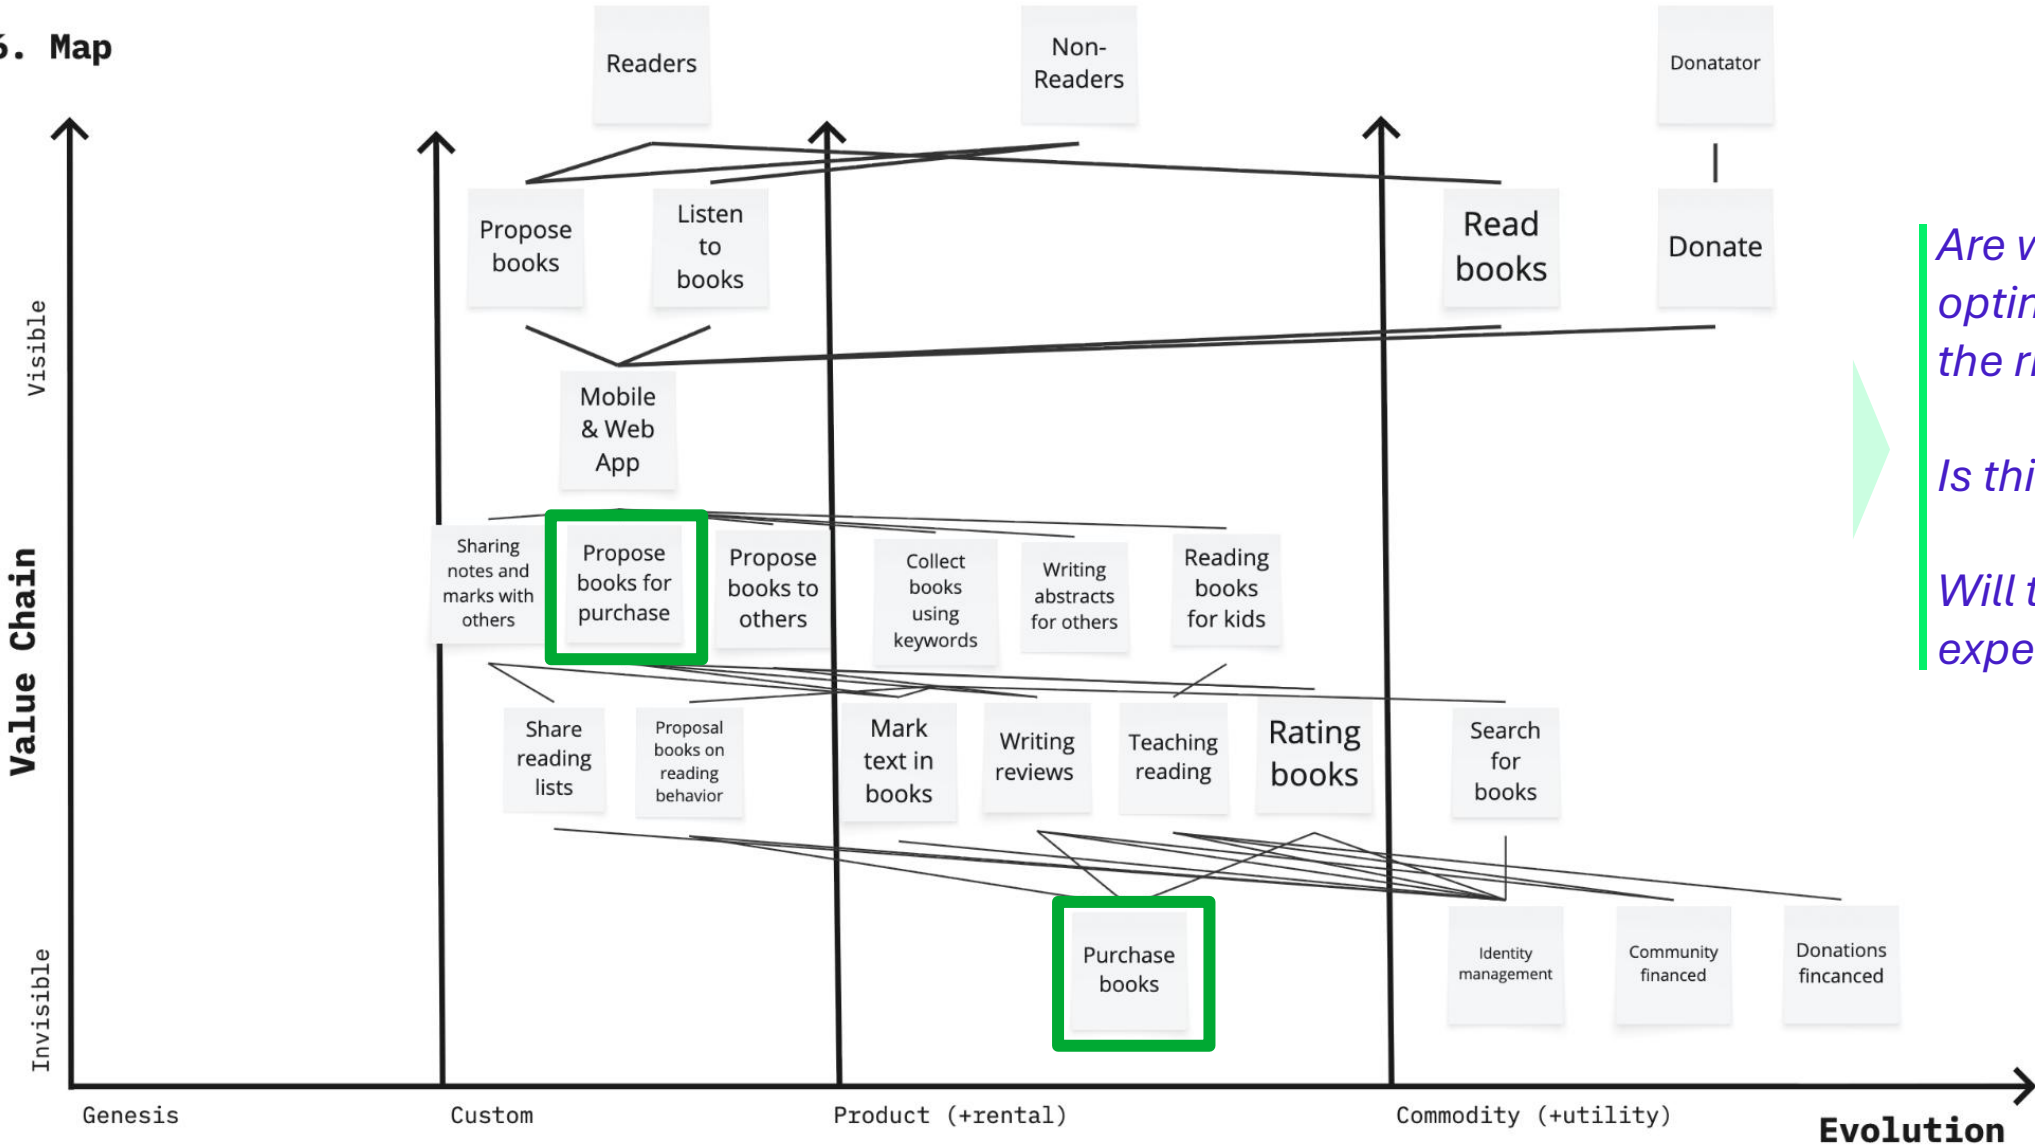
Craftsmanship of Asynchronous
and Synchronous APIs

—
Annegret Junker
Fabrizio Lazzaretti

Source: [Strategyzer AG](#) | License: [CC By-SA 3.0](#)

The Product: Wardley Mapping


6. Map



Are we optimizing/investing in the right place?

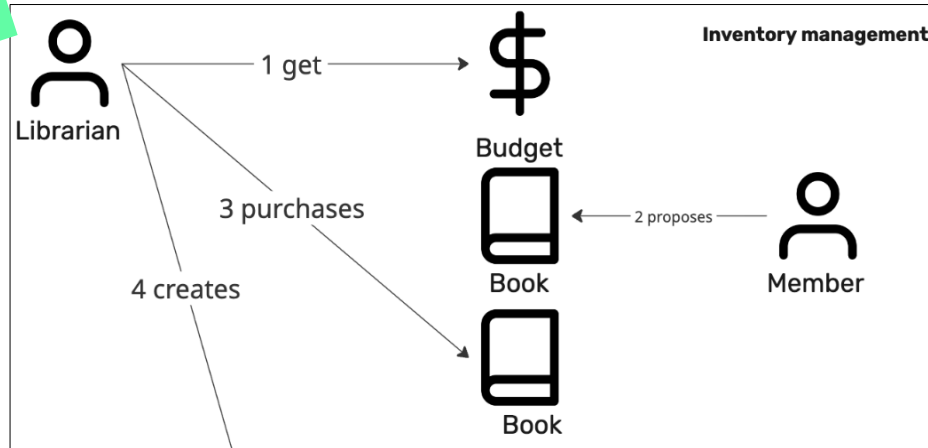

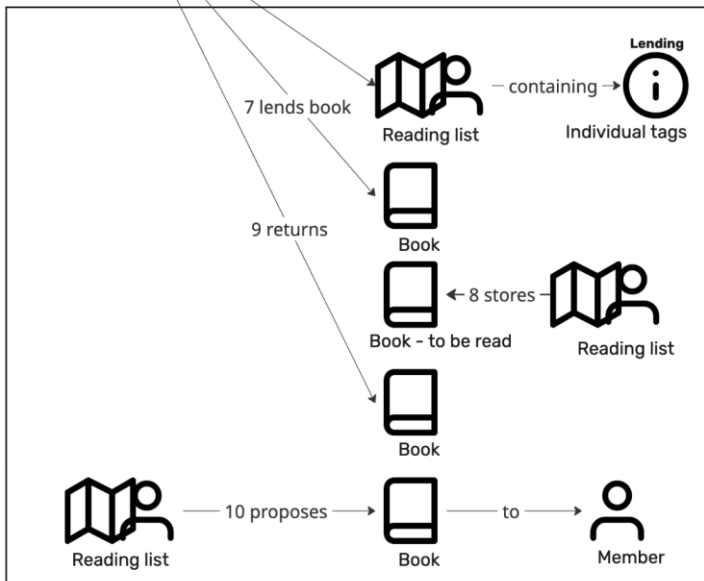

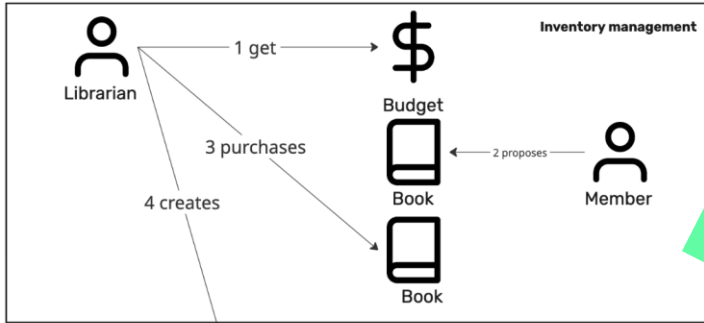
Is this a priority?

Will this improve the experience/product?



The Product: Domain Story

Searching and Lending



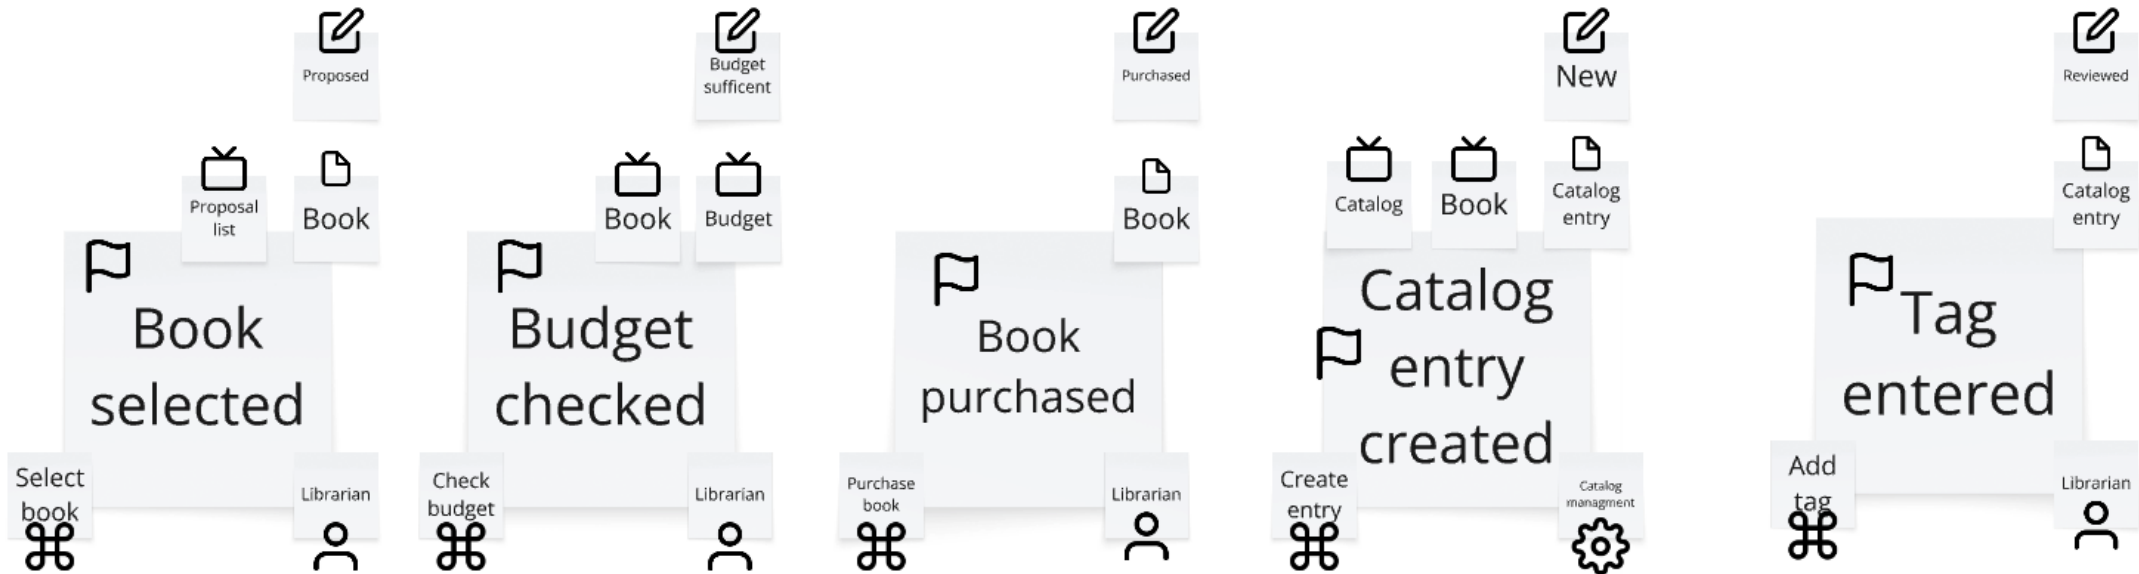
How is the workflow?

Do we all have the same understanding?

**Crafting Great APIs
with Domain-Driven
Design**

Craftsmanship of Asynchronous
and Synchronous APIs
—
Annegret Junker
Fabrizio Lazzaretti

The Product: Event Storming



What happens exactly?

What can go wrong?

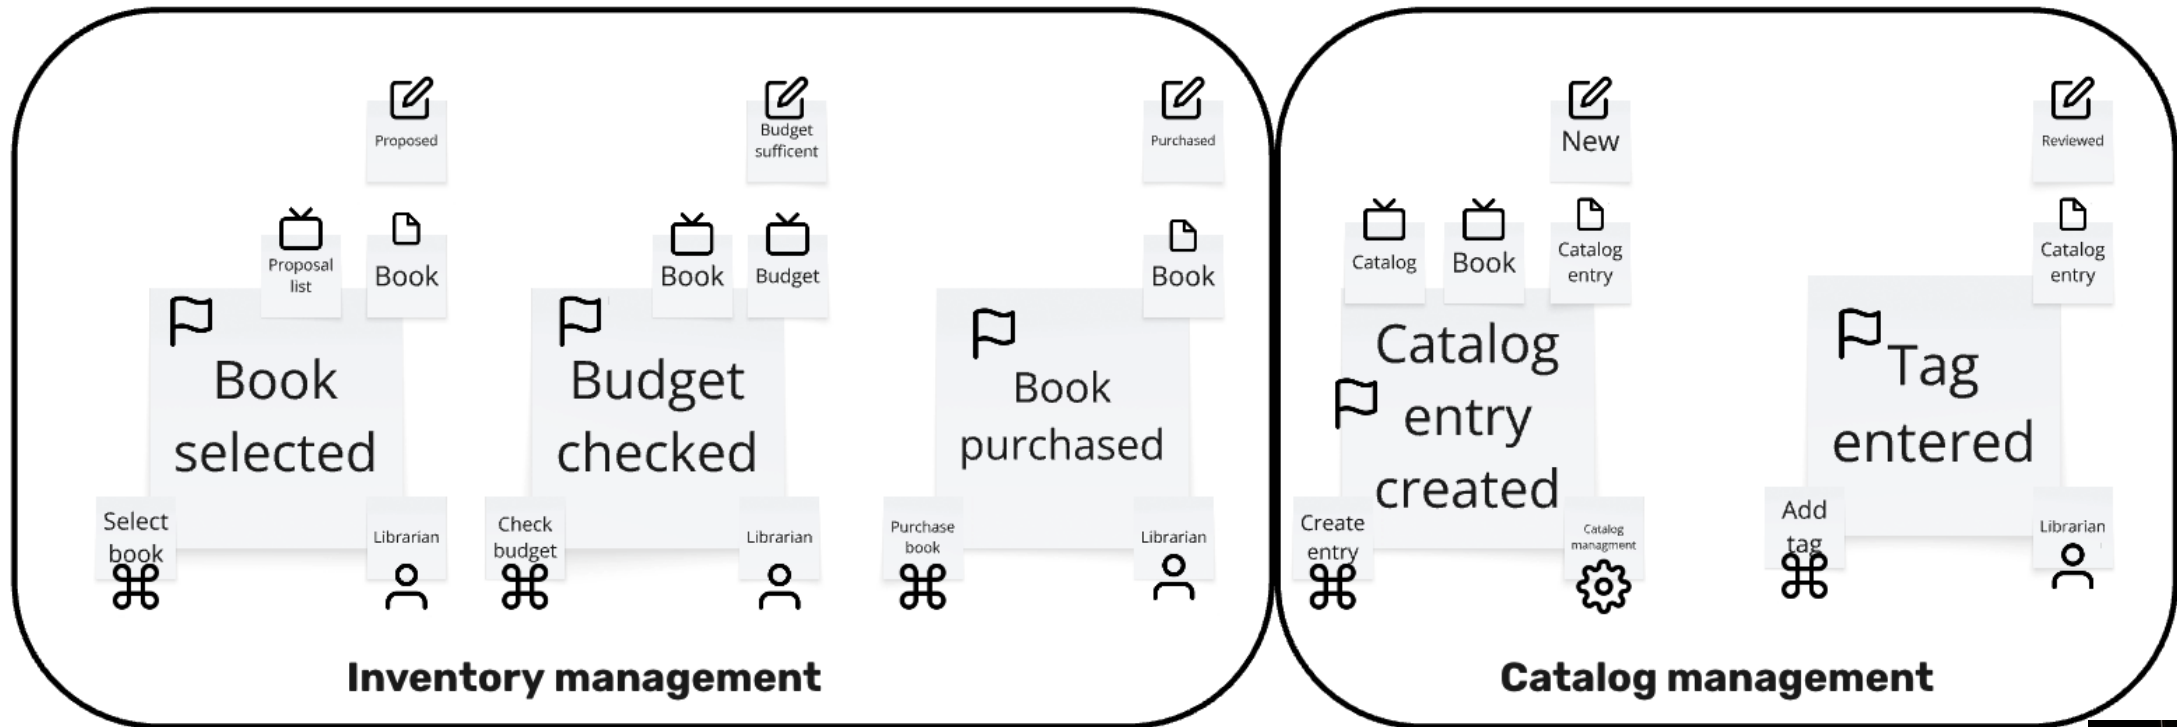
Does every step bring value?

**Crafting Great APIs
with Domain-Driven
Design**

Craftsmanship of Asynchronous
and Synchronous APIs

Annegret Junker
Fabrizio Lazzaretti

The Product: Bounded Contexts



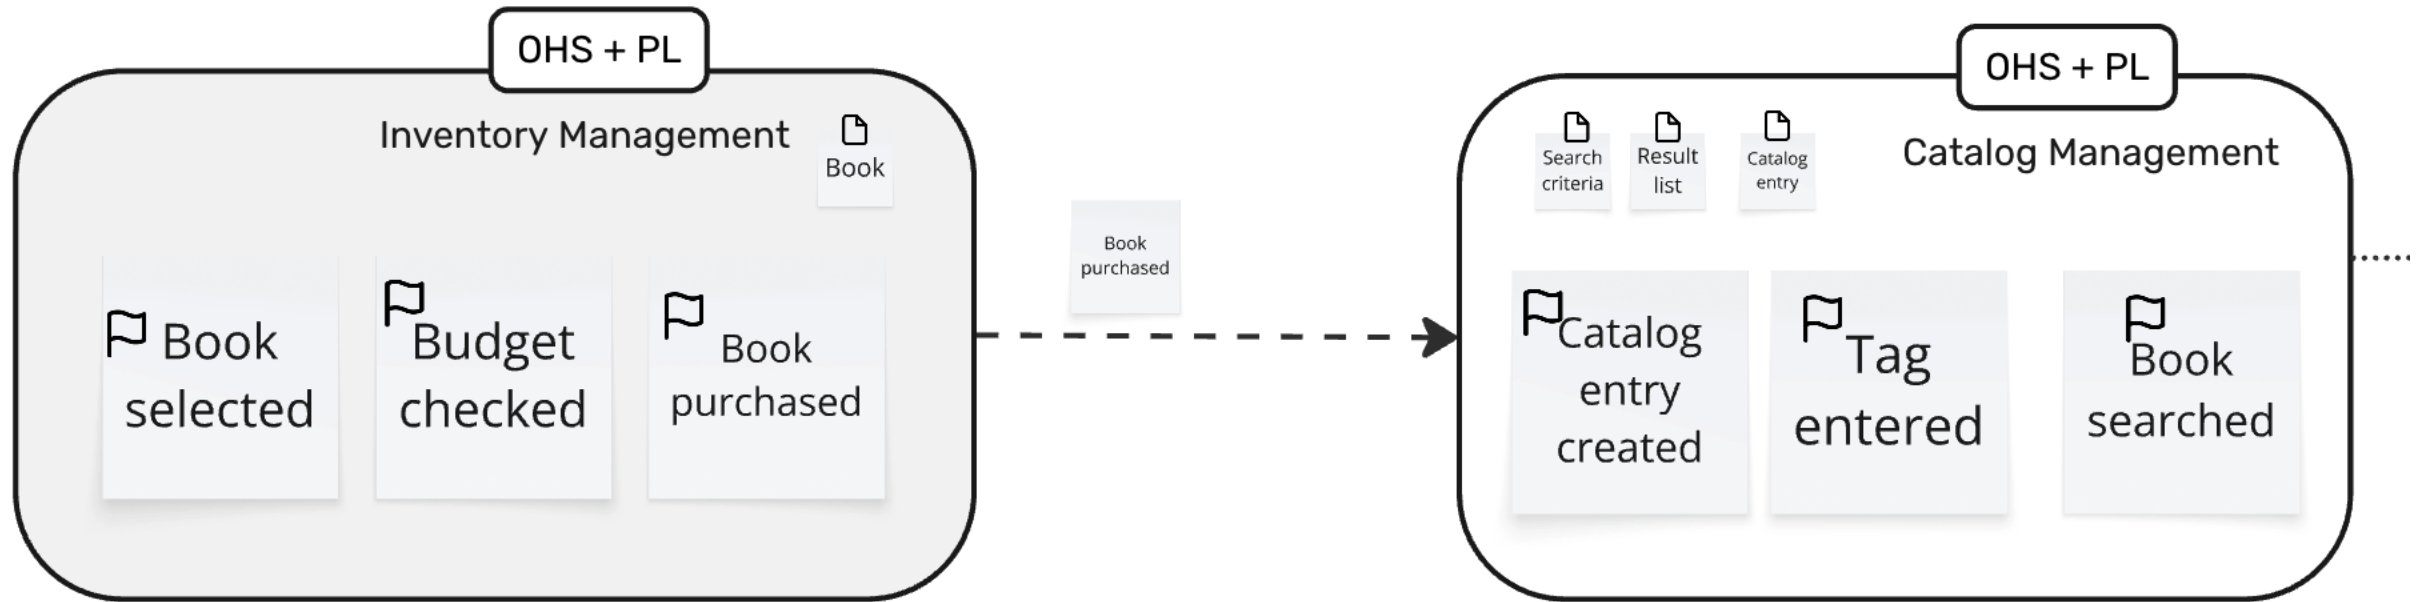
*How can we split the app into separate contexts?
Same language, same quality attributes, single team responsible*

Crafting Great APIs with Domain-Driven Design

Craftsmanship of Asynchronous and Synchronous APIs

Annegret Junker
Fabrizio Lazzaretti

The Product: Context Map



Can we build separate services out of them?

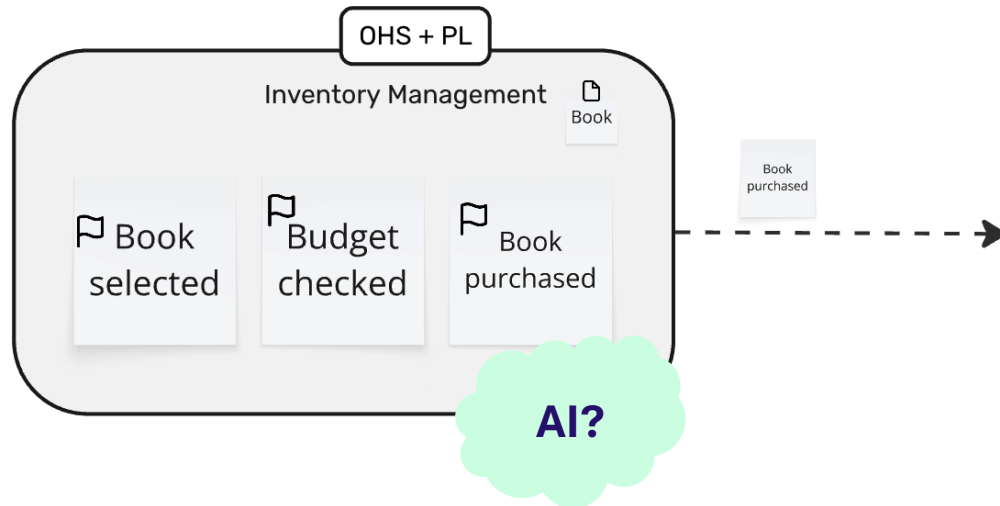
How will they communicate with each other while staying decoupled?

**Crafting Great APIs
with Domain-Driven
Design**

Craftsmanship of Asynchronous
and Synchronous APIs

—
Annegret Junker
Fabrizio Lazzaretti

The Product: Let's ask if AI helps



What do we want to improve? (Speed? Better results?)

Will AI help here? (Can it help with the given steps?)

→ Probably yes: Research about books, compare offers,...

→ AI in Coding is not AI in the product! → Often code is better than AI in execution

How do we measure success?

→ loans per book, dissatisfaction rate, time used to buy book licenses ...

Do we need a PoC?

→ What can judge the result?

Can we do an MVP right away?

→ Maybe? Show if book selection is better, with more loans?

“

“But agents do it for me now.”

“But agents do it for me now.”

MYTH

“Just put an agent in the center.”

REALITY

An agent monolith is still a monolith. Same coupling, same failure modes, harder to debug and often slow and expensive to run.

MYTH

“Agents handle the design for us.”

REALITY

Agents are themselves separated by roles; this is what makes them more precise. And should one have the responsibility to do everything?

MYTH

“We don't need APIs anymore — only MCP.”

REALITY

MCP is also an API. Same separation of concerns still applies. The protocol changed — but is this what was really needed? Look at the interaction first!



Deeper MCP and interaction architecture critique:

WeAreDevelopers 2026 talk — “Two Architects, One Mission: Rethinking How Systems Talk.”, by me and Patrick Müller

Take home.

AI is a tool you might use to deliver it.

Build products that happen to use AI — not AI that hopes to become a product.

01

Define the product before the model.

If you can't fill the product box or write the press release, you don't have a product. From the outside, AI almost always changes nothing.

02

Question the KPI before celebrating it.

A rising metric isn't proof of value.

03

Distinguish marketing thinking from product thinking.

Customers want outcomes, not labels.

Wavestone Services: From Foundations to Transformation: We turn AI into Business Value



BUILD THE **FOUNDATIONS**
FOR SCALABLE GROWTH

- Design, select, deploy **Data, AI** and **Agentic Platforms**
- Prepare **data quality, industrialization** and **scaling**
- Build the **AI ecosystem**, vendors strategy and **partnerships**



TARGET TANGIBLE **VALUE**
AND ORGANIZE

- Development of **strategies, maturity assessment, business cases** and **value measurement framework**
- Definition of **AI governance, operating model**, and **“trust” framework**



TRANSFORM BUSINESS
WITH **AI & AGENTIC**

- Optimize and redefine **business processes** with **AI** and **Agentic**
- Build **AI models and Agents, AI products, agentic factories**
- Manage **change, upskilling**, AI for all.

BUSINESS

We bring **deep sector knowledge** and **domain expertise**, enabling us to steer your AI transformation towards the new AI-first business models and processes that will generate the greatest value.

TECHNOLOGY

We are **builders**. We deliver AI agents quickly while also establishing strong and scalable foundations to help you transition into an AI-driven company.

PEOPLE

While leveraging AI to augment capabilities, **we keep humans at the center**. We foster close collaboration with client teams and prioritize the impact of these technologies on employees and society.

Thank you. Find me here.



Fabrizio Lazzaretti

**Senior Managing Consultant
@ Wavestone**

→ Now we have time for questions

Get in touch via
QR Code



lazzaretti.me/aisummit26



Our assets

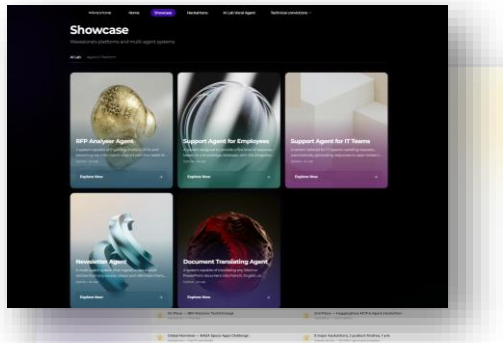
PARTNERSHIPS

Key international and European players:
Microsoft, Google, Mistral, AWS,
SAP, Salesforce, ServiceNow,
Snowflake, Cognigy, Databricks, Dataiku...

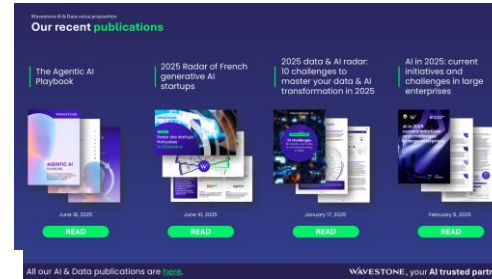
EXPERTISES

Business Analysts & AI Product Managers
Adoption & change managers
Organisation & TOM specialists
Data, AI & Agentic engineers
AI Act, Sovereignty, Cyber specialists
Architects

AI LAB




PUBLICATIONS & EVENTS



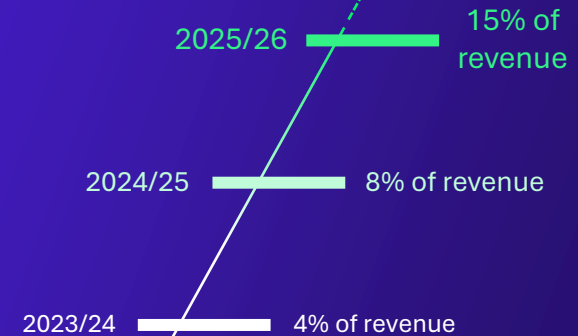
Les Echos

GROUNDING IN SUCCESSFUL EXPERIENCE WITH MAJOR INDUSTRY PLAYERS



International talent pool
~1000 consultants


4 main geographies with dedicated AI teams
France, US, UK, Germany
+ nearshore Poland, Romania, Morocco



A **SUCCESSFUL** POSITIONING IN AI,
CONTINUING **PROGRESS**

Alongside our credentials, we keep improving and sharing: our recent AI publications

Global AI survey
2025: The paradox of
AI adoption



October 15, 2025

READ

The Agentic AI
Playbook



June 18, 2025

READ

2025 Radar of French
Generative AI
startups



June 10, 2025

READ

2025 data & AI radar:
10 challenges to
master your data & AI
transformation in 2025



January 17, 2025

READ