Holistic AI

Designing Ethical AI at Scale

Design Systems & Product **Leadership**

by Sim Deol

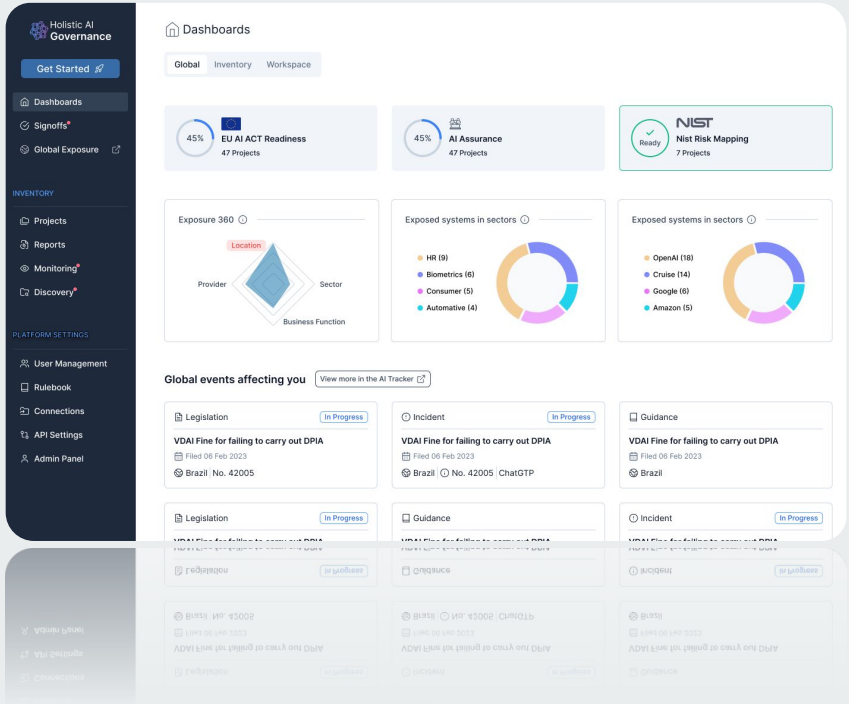

Outline

- Context and goals
- The problem
- My role
- Approach
- Scaling design
- Products and deliverables
- Rollout, impact and reflection
- Further exploration

Context

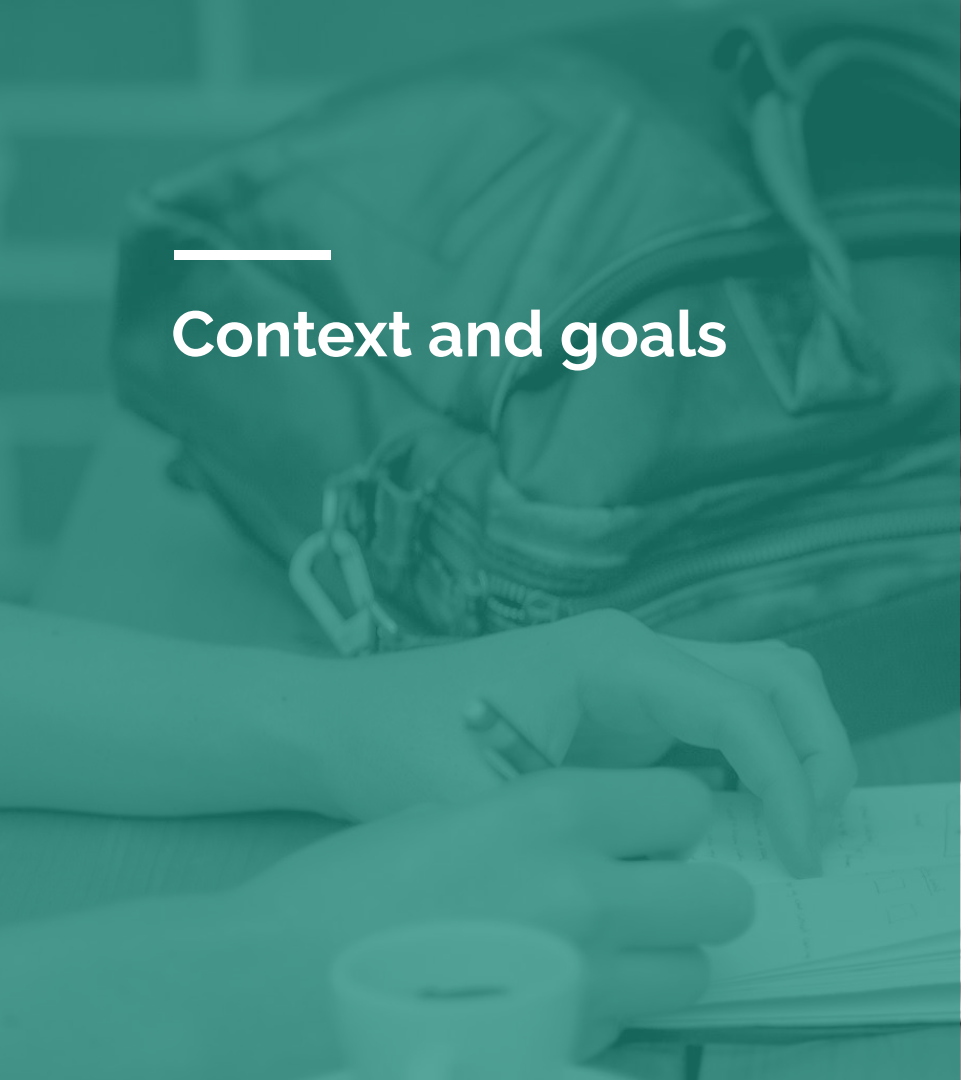

Context and goals

Company - **Holistic AI** - enterprise platform for AI governance, risk and compliance

Initial State - One legacy Angular platform, custom components, poor scalability, no unified UX, inconsistent UI

Goal - Modernise the core platform and enable rapid expansion into new products

Mission - Redesign infrastructure, scale UX, accelerate delivery

The Governance Platform - Centralised Ethical AI Management

The original core platform was robust but rigid.

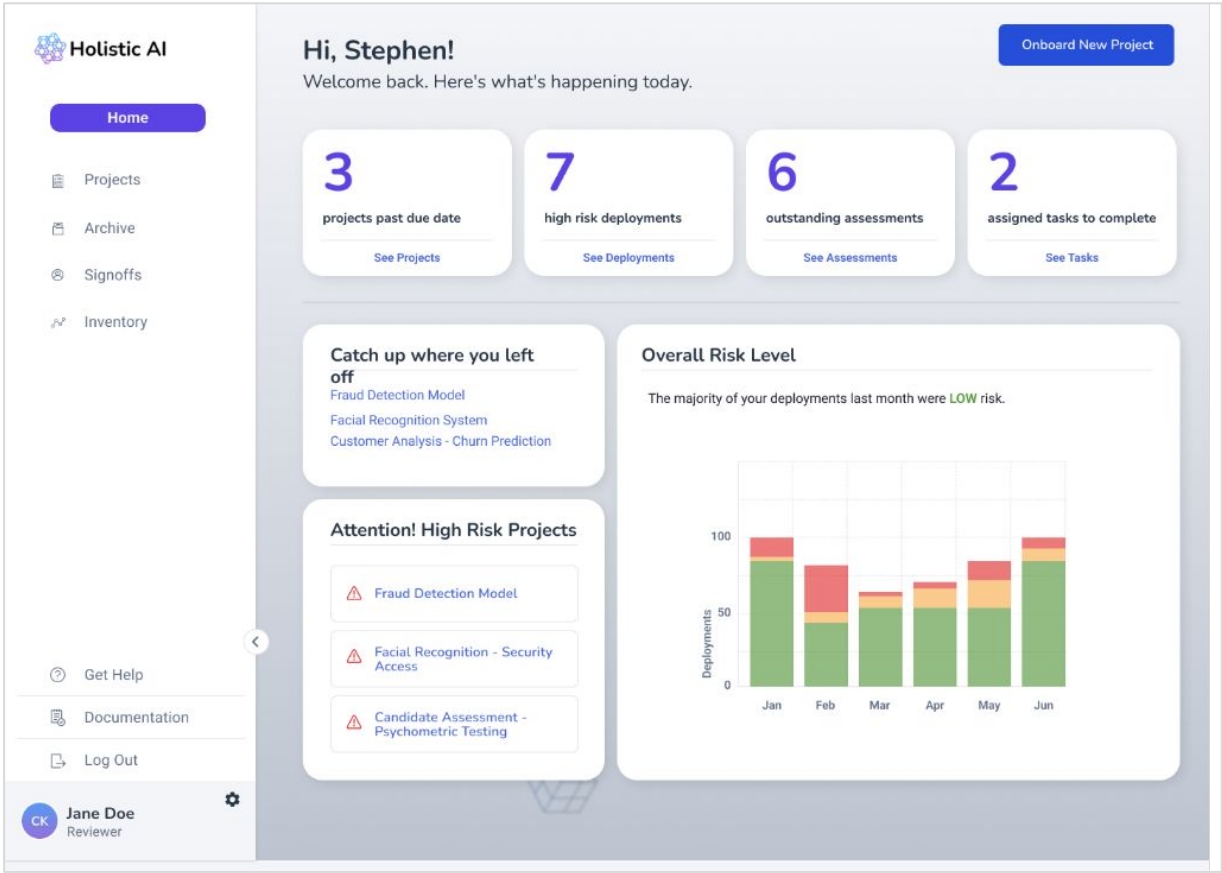

The Challenge

- Outdated tech (Angular + bespoke components) hindered dev velocity
- Difficult for engineers to onboard and iterate quickly
- No reusable design patterns, inconsistent UI, high design/development overhead
- Needed to scale into two new products while redesigning the original

My role

Sim Deol

Sr Product Designer

- Sole Product Designer across all product streams
- Led UI audits, design system strategy, and component library development
- Partnered closely with engineers and product managers to align on priorities, technical frameworks and implementation approach
- Assisted in roadmapping, sprint leadership
- Conducted research and collated insights from users
- Led design reviews & stand-ups,

Approach

Strategic & Technical Approach

- Audited UI and front-end workflows to identify friction points, as well as a comprehensive review of existing UI and UX patterns across the product.
- Worked with engineering to implement a React + Tailwind CSS stack, leveraging shadcn/ui prebuilt components to accelerate delivery despite limited front-end capacity.
- Designed and built a tokenised component library in Figma to mirror implementation structure
- Introduced UX conventions, naming logic, and documentation to ensure ease of adoption

Cross-functional collaboration

- Aligned with PMs on scope and client priorities (e.g., Unilever retention, Mapfre onboarding).
- Partnered with engineers, ensuring design intent met technical feasibility.
- Collaborated with legal, policy, and data science teams to embed compliance and trust into UX.

Start-up/Scale-up vs Large structures

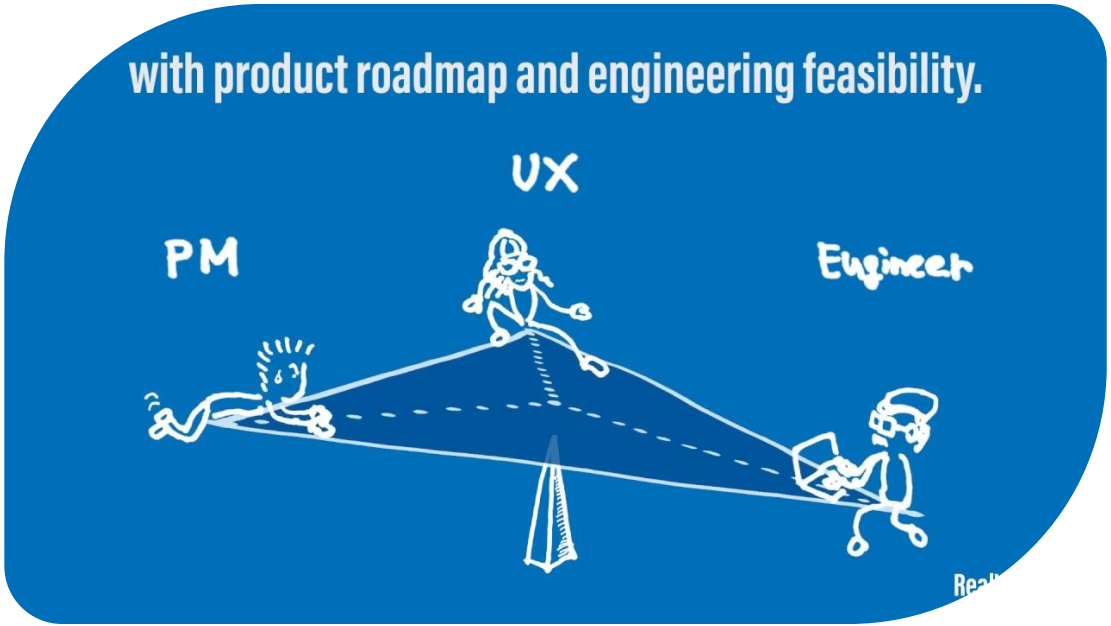

Start-up

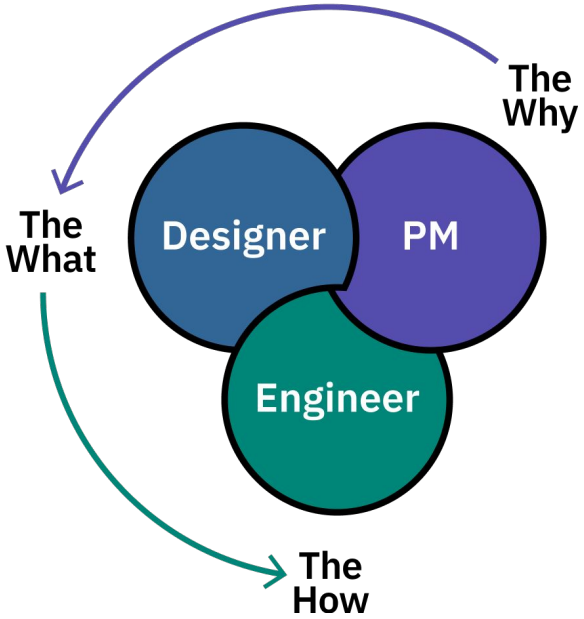

Larger structures

Work-in-progress

ADMIN

Who they are

Senior-level employees or leadership team

ROLES

Chief Data Officers (CDOs)

AI Governance Leads

Compliance Officers

other executives.

Overview

- These users oversee the AI systems and are responsible for ensuring compliance, reducing risks, and improving the overall AI quality.
- They may or may not have technical and regulatory knowledge
- Familiarity with organizational goals, risk management frameworks, and business strategies.
- Strong decision-making skills with the ability to balance technical complexity and business needs.

Goals

- Ensure the organization's AI systems align with regulatory and ethical standards.
- Oversee AI governance frameworks and ensure compliance across all projects.
- Mitigate risks associated with AI (e.g., bias, data security, model robustness).
- Continuously improve the quality, transparency, and accountability of AI systems across the enterprise.
- Efficiently manage resources and team capacity for AI-related projects.

Needs

- High-level dashboards with key metrics and insights on AI system performance, risks, and compliance.
- Regular reports that show trends, potential risks, and strategic recommendations.
- Clear visibility into the status of ongoing AI projects, audit logs, and compliance checks.
- Tools to define and adjust governance policies and risk frameworks at an organization-wide level.
- Ability to delegate tasks and monitor team progress efficiently.

Pain Points

- Difficulty overseeing multiple AI projects and teams at scale.
- Challenges in keeping up with evolving regulations and ensuring compliance across diverse AI models.
- Limited visibility into the root cause of AI failures or bias across projects.
- Lack of alignment between technical data and business impact when making decisions.
- Ensuring effective collaboration between cross-functional teams (AI engineers, legal, compliance, etc.).

Tasks

- Reviewing AI governance reports and compliance dashboards.
- Adjusting risk and governance frameworks as needed.
- Making decisions based on AI performance insights and compliance outcomes.
- Monitoring the progress of AI projects and ensuring resources are allocated properly.
- Approving key governance policies or escalating issues to senior leadership.
- Overseeing audits, model evaluations, and critical incident responses.

MEMBER (USER)

Who they are

Mid-level to junior-level employees

ROLES

Data Scientists

AI Engineers

Compliance Specialists

AI Project Managers

Overview

- These users are responsible for implementing and maintaining AI models within the organization and ensuring they adhere to governance and compliance rules.
- More focused on day-to-day operational tasks and project execution.
- Familiar with specific AI tools and frameworks, but less concerned with overall organizational strategy.
- Some understanding of compliance and regulatory standards as they relate to their projects, but may not be experts in governance.

Goals

- Successfully complete AI projects while ensuring they meet governance and compliance standards.
- Improve the performance, fairness, and reliability of AI models they work on.
- Identify and address risks, including bias and model drift, in their daily work.
- Collaborate with other team members to execute tasks efficiently.

Needs

- Clear and actionable tasks that directly relate to their AI projects.
- Detailed insights on their AI models' performance, fairness, and compliance.
- Notifications and alerts when their models exhibit potential risks or non-compliance issues.
- Guidance on how to resolve governance issues or optimize their AI systems.
- Access to project-specific documentation, guidelines, and best practices.

Pain Points

- Frustration with unclear or changing compliance requirements.
- Difficulty managing multiple AI models and projects while ensuring adherence to governance policies.
- Overwhelmed by manual compliance checks or model evaluations.

Tasks

- Running bias detection and mitigation tests on their models.
- Ensuring compliance with governance policies.
- Responding to alerts about non-compliance, bias, or performance degradation in models.
- Collaborating with other members on AI project execution and reporting.
- Attending regular check-ins to provide updates on project status and risks.

Research Snippet

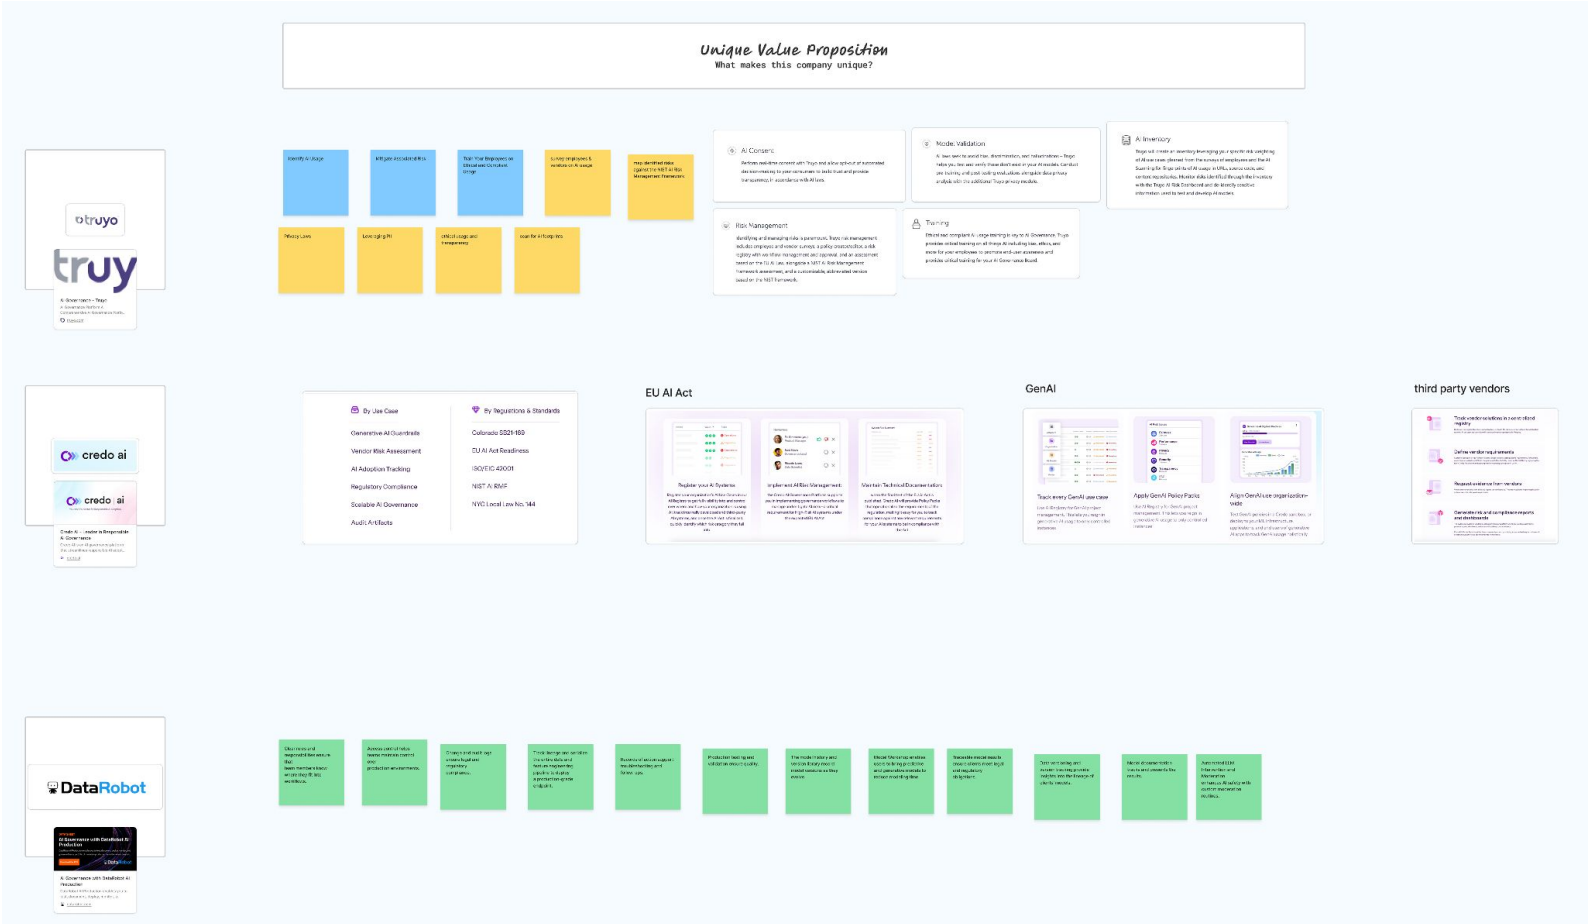

Workshops snippet

findings from the meeting

painpoint: end user comes to their flow to submit all answers, and they need to do it again on HAI for assurance and EU AI act

needs: ask all information upfront, additional questions are asked if when it is necessary

painpoint: they have to come to HAI to archive so many projects cause there are no activies.

whenever there is an AI inventory, they always come to answer the questions in the Microsoft form

questions are not for assurance but for internal goverance

the intention behind this is they don't people come to HAI before ideation is done, other stages are PoC, pilot and Live.

they need to make sure they have enterprise architecture to align with their strategy

PRA ID?

also ask for legel team has been involved or not

insights: with questions that not get answered, people may come back to answer them, or drop.

painpoint: they have to come to HAI to archive so many projects cause there are no activies.

assurance stage is manually updated for reporting

comments area to update the status

Gen AI greenlist to see if a tool is approved by U, it is a first step, still need AI assurance

ideation

avoid answer same questions repeatly

clearly state that the project needs to be Poc or after to process

build a good way for them to communicate with enterprise architecture and legal people

nice to have a comment/ note feature to add extra info/tags

only ask necessary questions, limit the number of questions

separate projects which are at ideation stage and the stages after

automate some process, eg. greenlist/redlist apps, if redlist, provide explanations and usage instructions

questions in the form are editable, and should be easy to edit. We should save history version for them.

Refining complex user flows

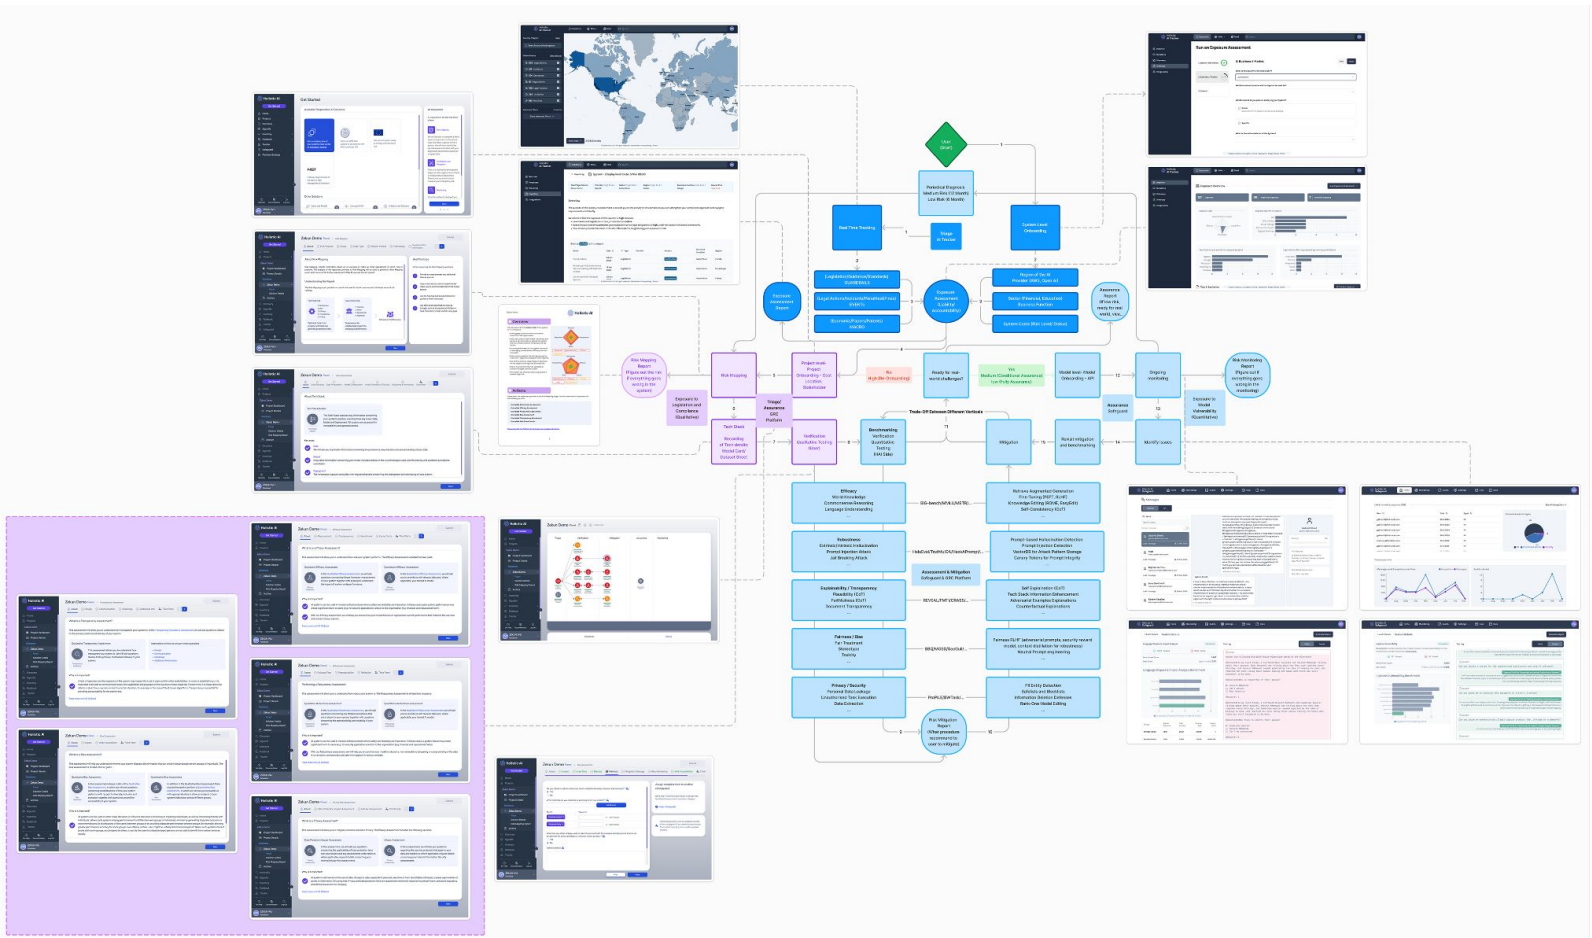

Refining complex user flows

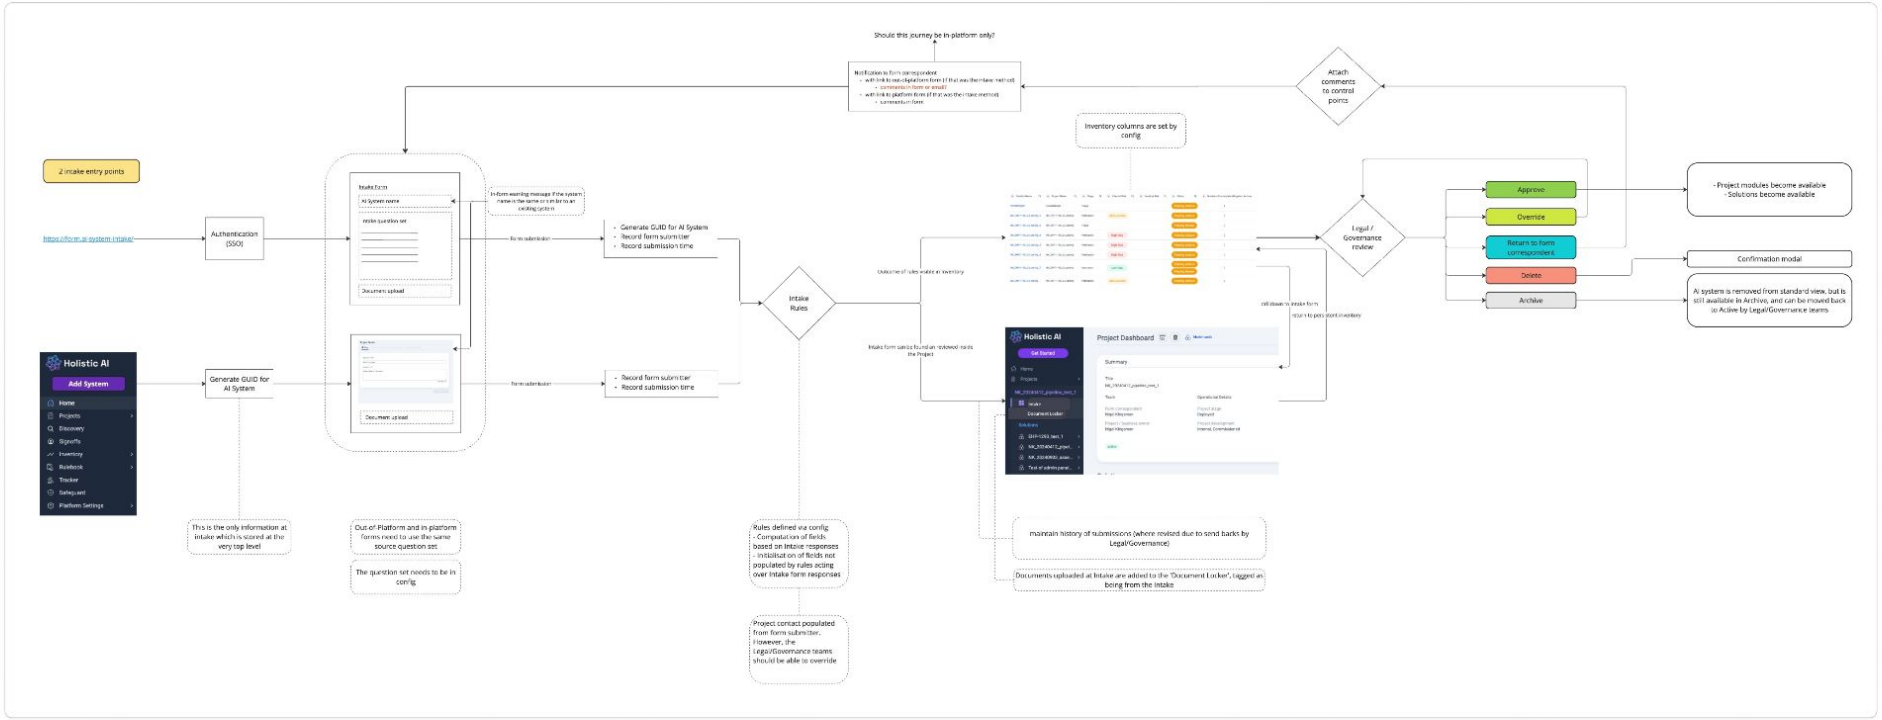

Refining complex user flows

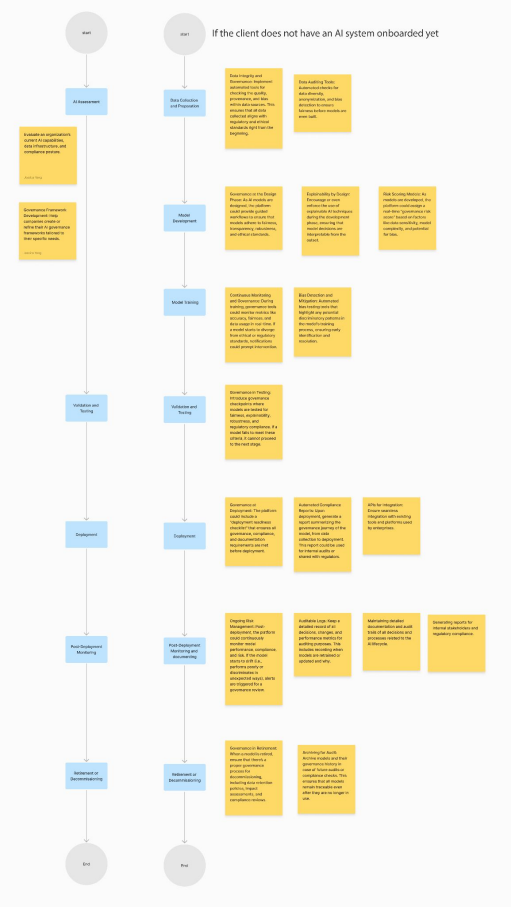

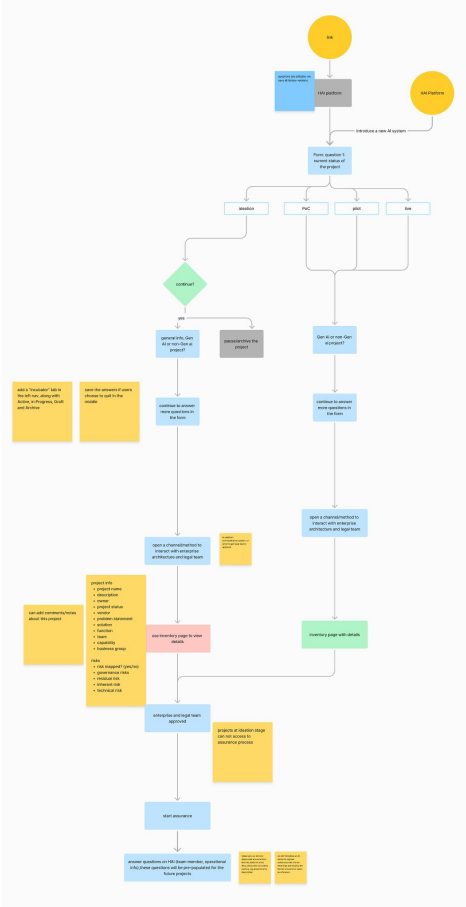

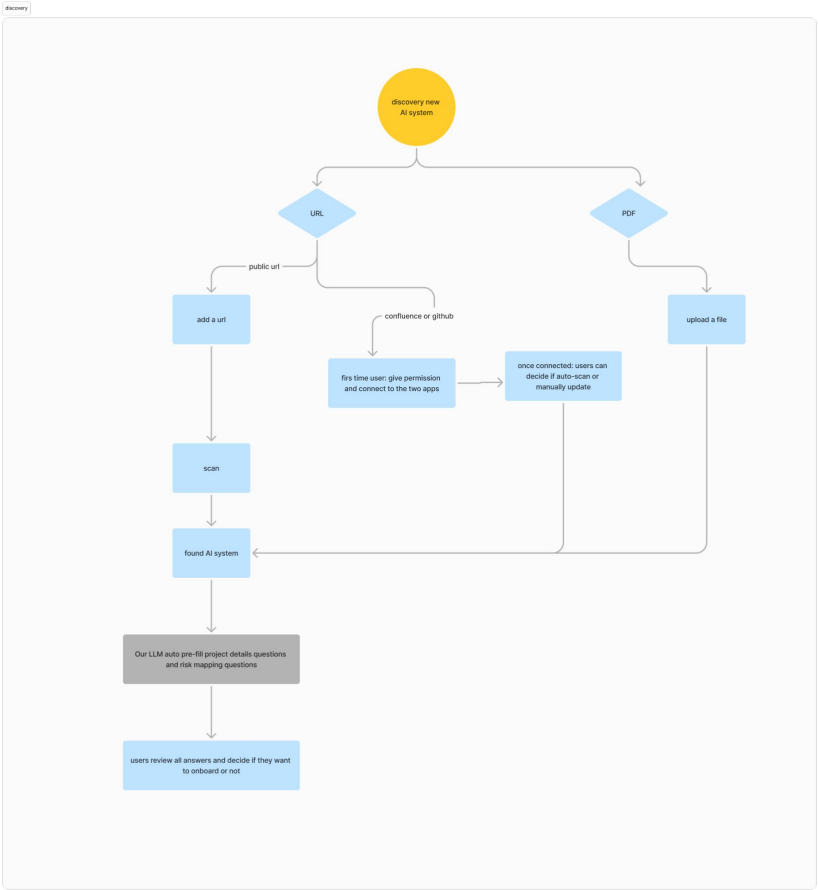

A breakdown of the phased approach

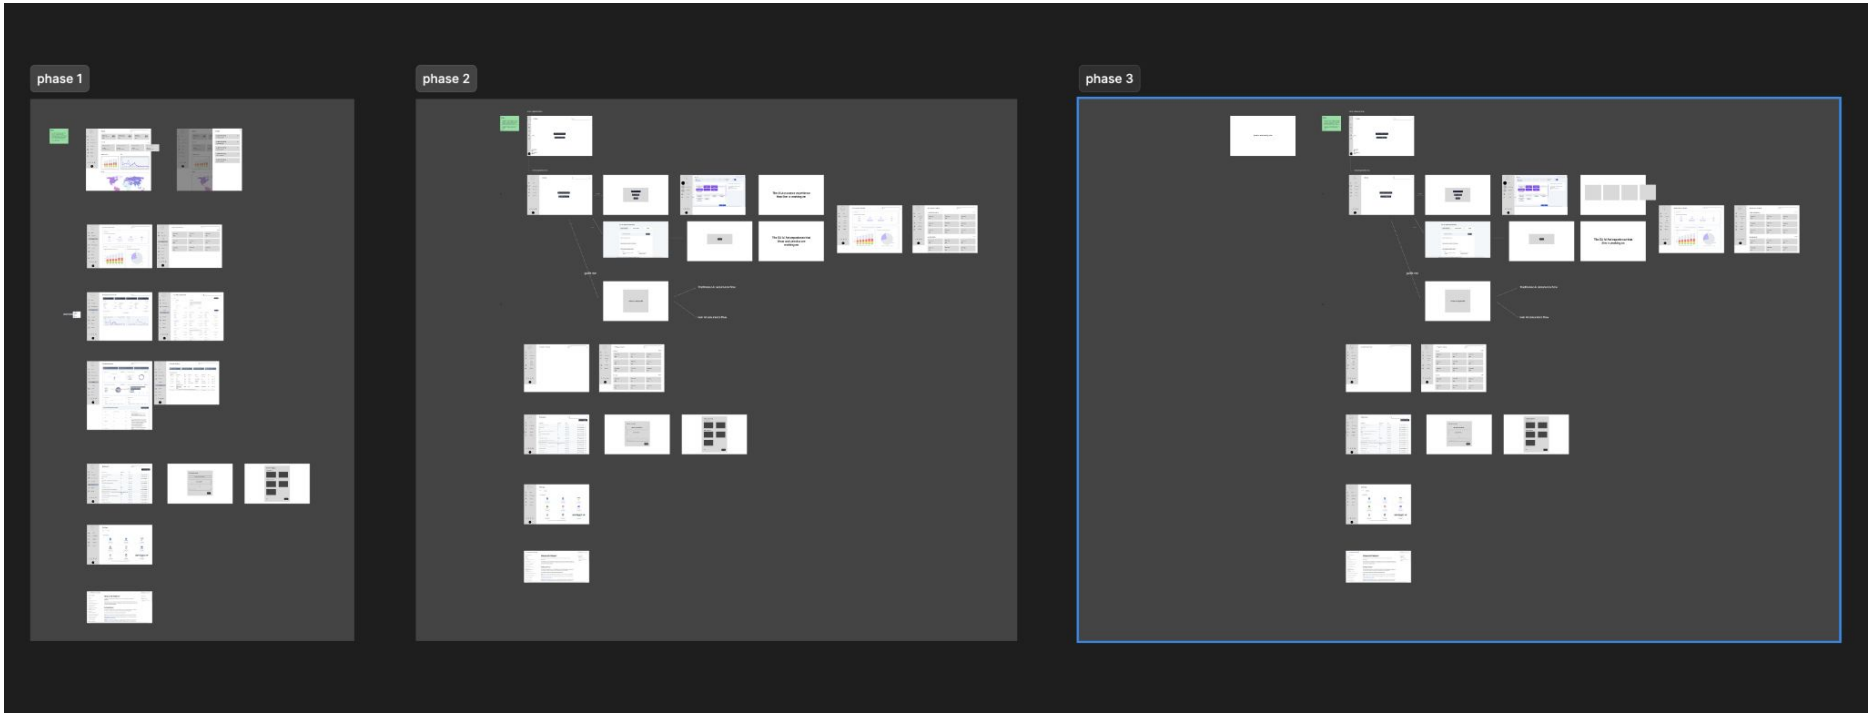

The Design System

Figma Ninja

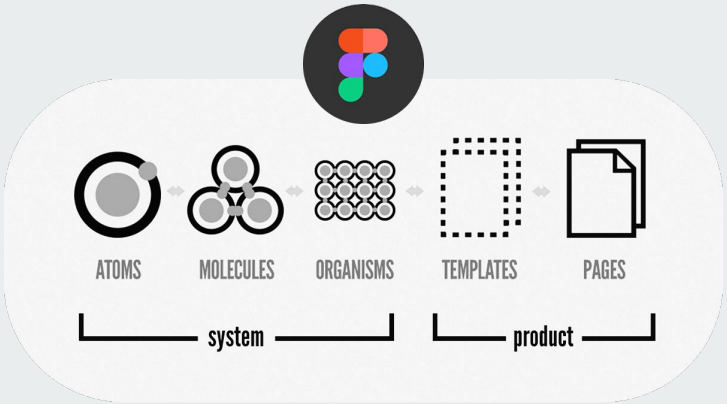

- **Designers**: Need scalable, reusable components and clear UX patterns to move faster.
- **Front-end Engineers**: Seek alignment with design and robust specs to reduce back-and-forth.
- **Product Managers**: Want consistency and velocity in shipping features.
- **Leadership**: scalability, governance, and elevating design quality across org.
- **End Users:** Need a product that flows seamlessly and purposefully

Design Tokens

- 100+ components created and rolled into scalable tokenised system
- Accelerated dev workflows and improved design consistency
- Enabled us to launch new products with minimal UI rework
- Adopted across all three platforms

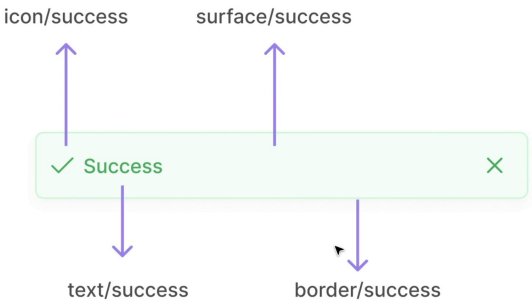

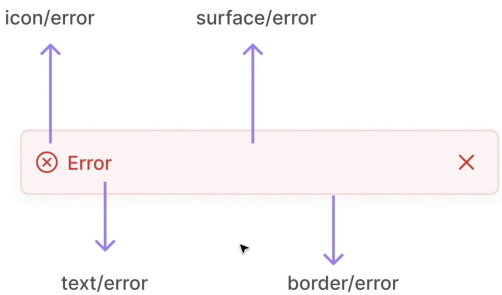

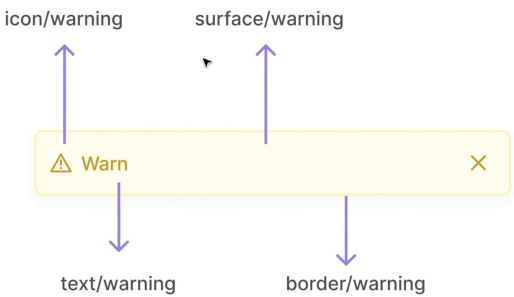

Example: A breakdown of the ‘success’ colour tokens

Brand

Name	Value
green	
100	BAFFB9
200	81F57F
300	7ED37C
400	53C351
450	53C351
500	11820F
600	067504
700	055603

Alias

Name	Value
success	
100	green/100
200	green/200
300	green/300
400	green/400
500	green/500
600	green/600
700	green/700

Mapped

Name	Value
text	
success	success/500
icon	
success	success/500
border	
success	success/600
surface	
success	success/100

Example: multi-brand and theming

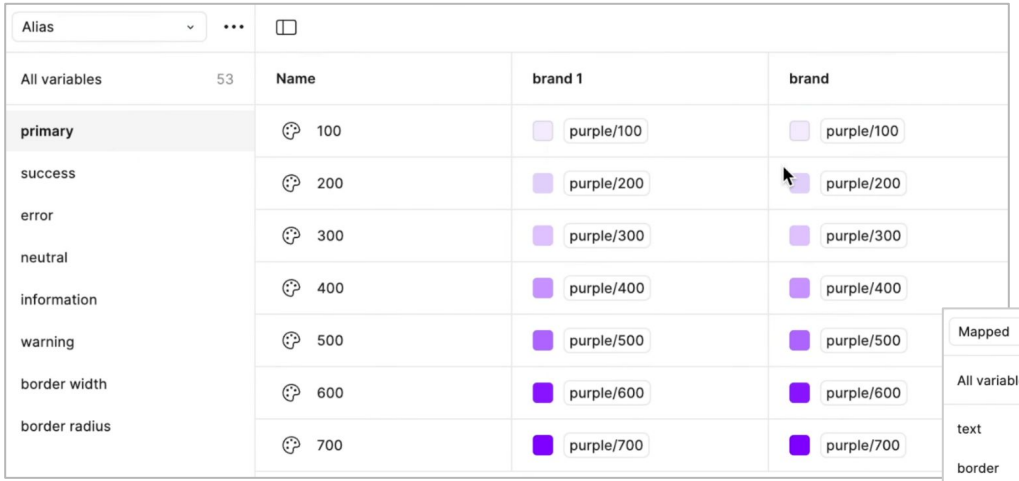

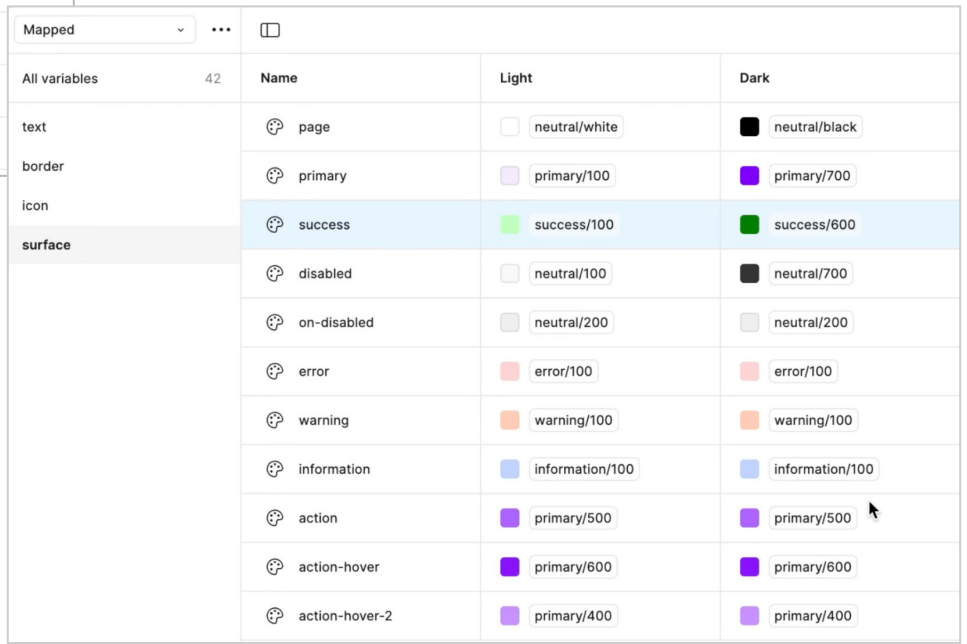

Why does this matter?

- Ensures wider team of designers in future have a consistent pattern to follow
- Much less likelihood of an inconsistent UI across teams
- Removes guesswork on components and styles
- Far easier to make system wide changes

The design system at a glance

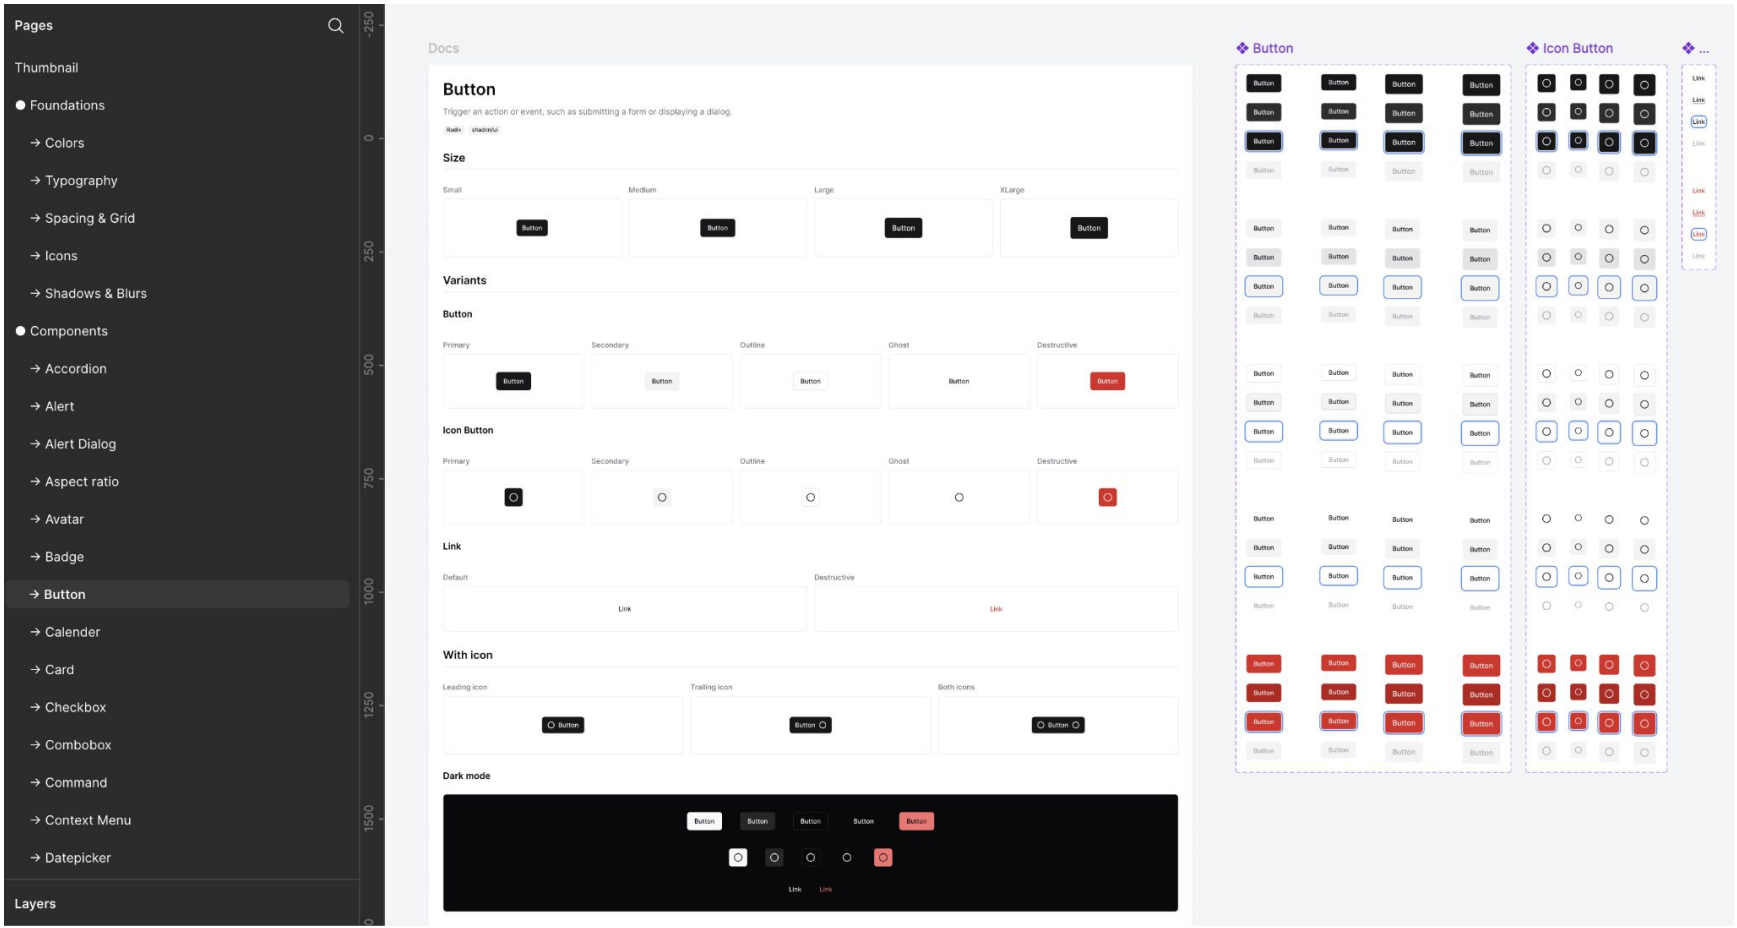

Auto-Layout

- Responsive by Design
- Scalable Component Libraries
- Mimics the flexbox/grid logic used by front-end engineers

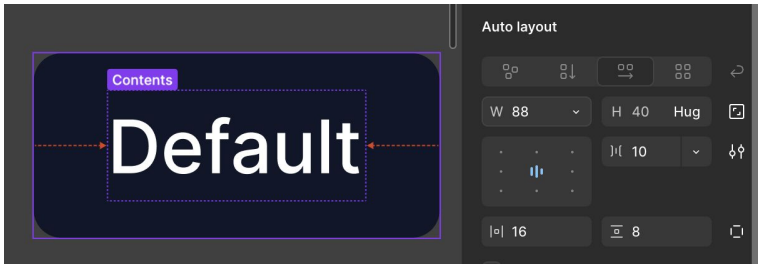

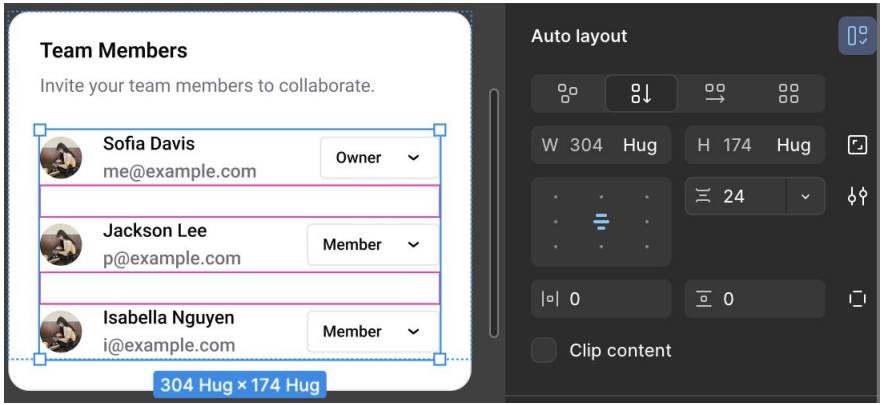

Atomic design system

- 100+ components created and rolled into scalable tokenised system
- Accelerated dev workflows and improved design consistency
- Enabled us to launch new products with minimal UI rework
- Adopted across all three platforms

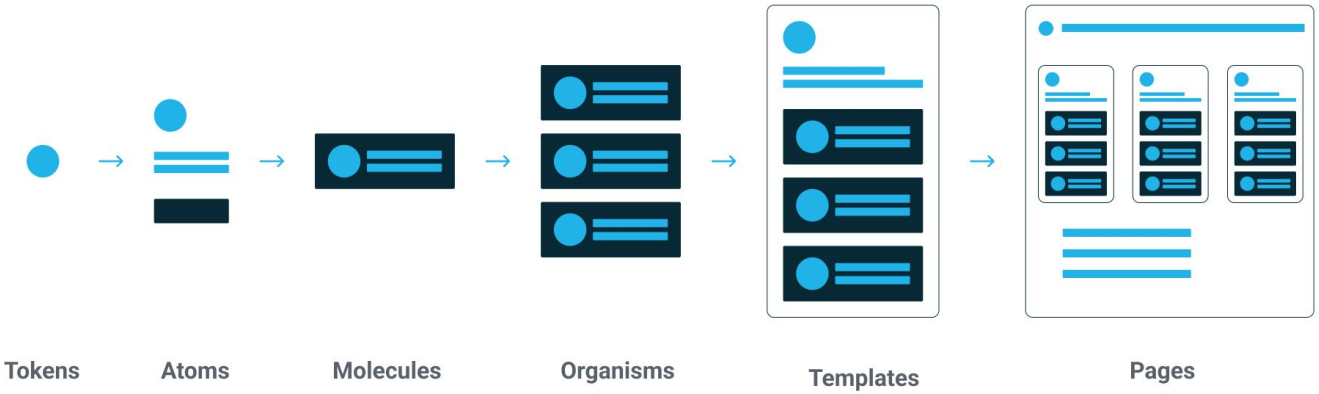

When to go Atomic

- Large organisations with complex design systems
- Teams that need a shared language for scalability
- Creating reusable components for different contexts

- Smaller teams or projects with limited scope
- Teams that need a shared language for multi-product projects
- When a simple, practical component library is sufficient

Build from scratch or build off an existing DS?

From scratch:

- Unique product domain or UI patterns
- When there's inconsistency or fragmentation across the product
- Targeting multiple squads or regions, needing consistency at scale

Existing DS:

- Startup or MVP and speed matters more than scale or 'beauty'
- Design needs are fairly standard
- Working inside a platform that already uses one

In the case of Holistic AI:

- Design system was built using shadcn as a base (standard functionality, less focus on UI flair more focus on functionality. We can add custom components on top: imaging viewer, workflow builder, etc.)
- Atomic design (scalable, thinking of product suite in future) if resources allow
- Tokenised library

The new platform(s)

The Core platform redesign

Centralised Ethical AI Management - The original core platform was robust but rigid. I evolved it for enterprise scalability. Just some of the improvements:

- Proposed and designed a new intake flow that automated onboarding via file uploads and integrations with Confluence, GitHub

- Designed the workflow and UX around key frameworks (e.g., EU AI Act), giving organizations a transparent view of their compliance and risk posture.

Result: significant decrease in manual effort and user frustration (Unilever reported a 60% increase in monthly AI system onboarding), helping retain Unilever as a long-term client and winning two new clients.

Before and after

Before

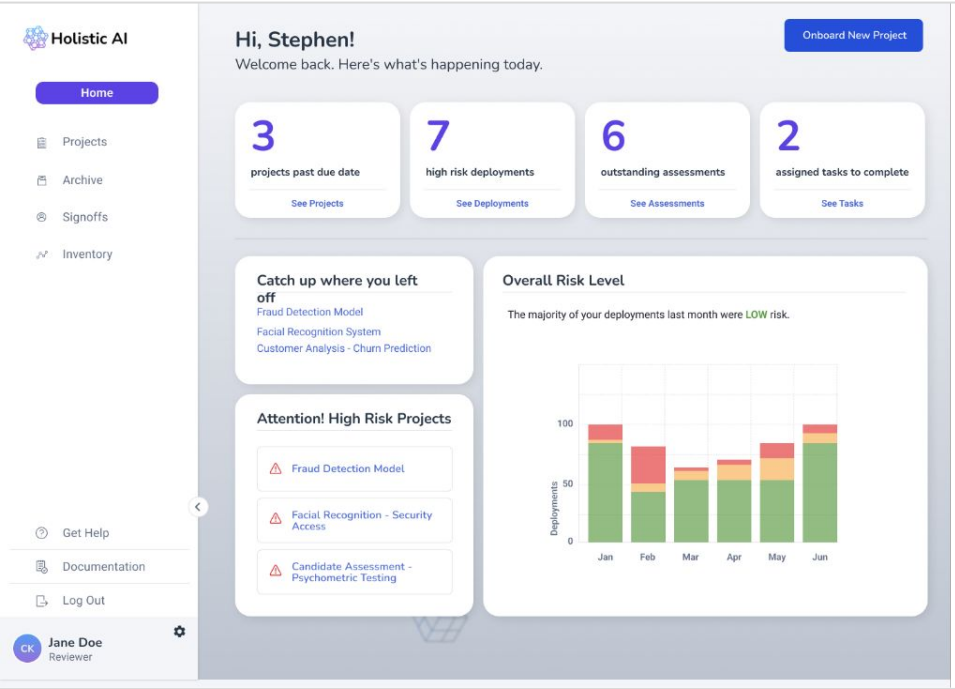

After

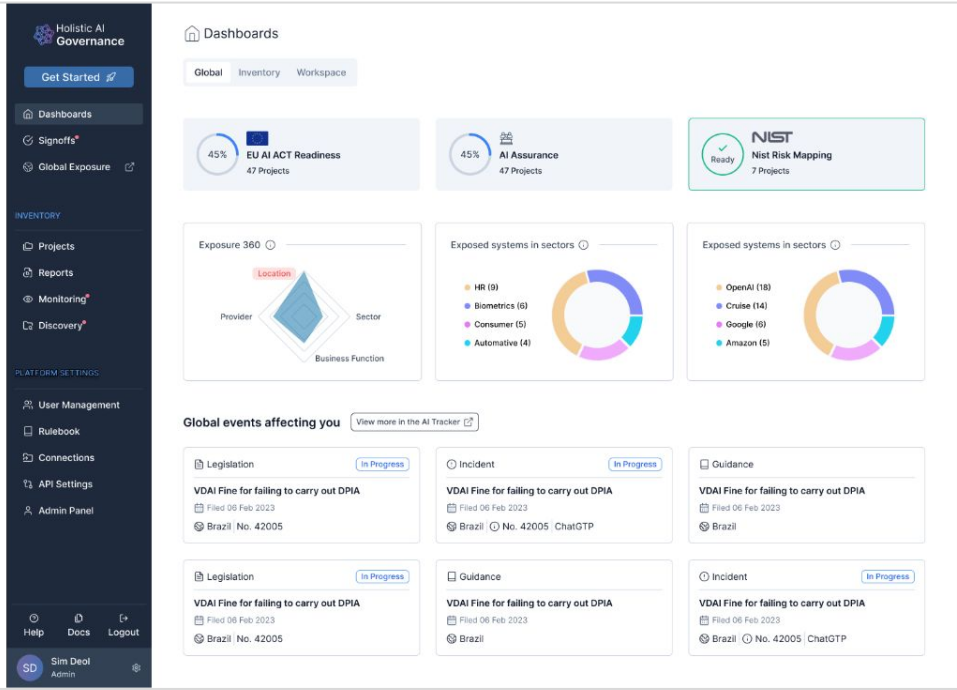

The ‘Get Started’ screen

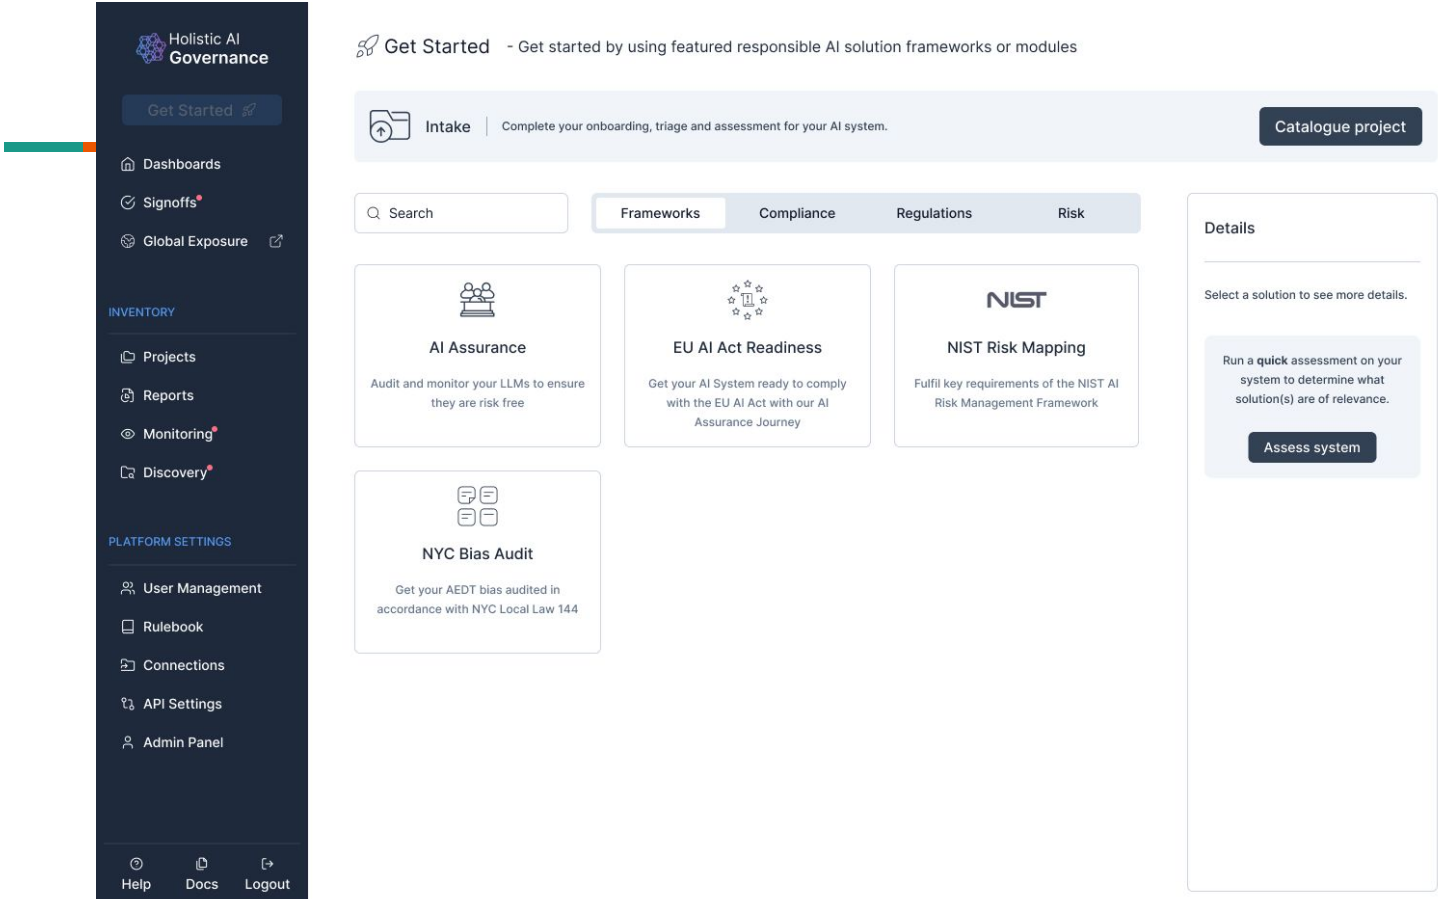

Simplifying system onboarding

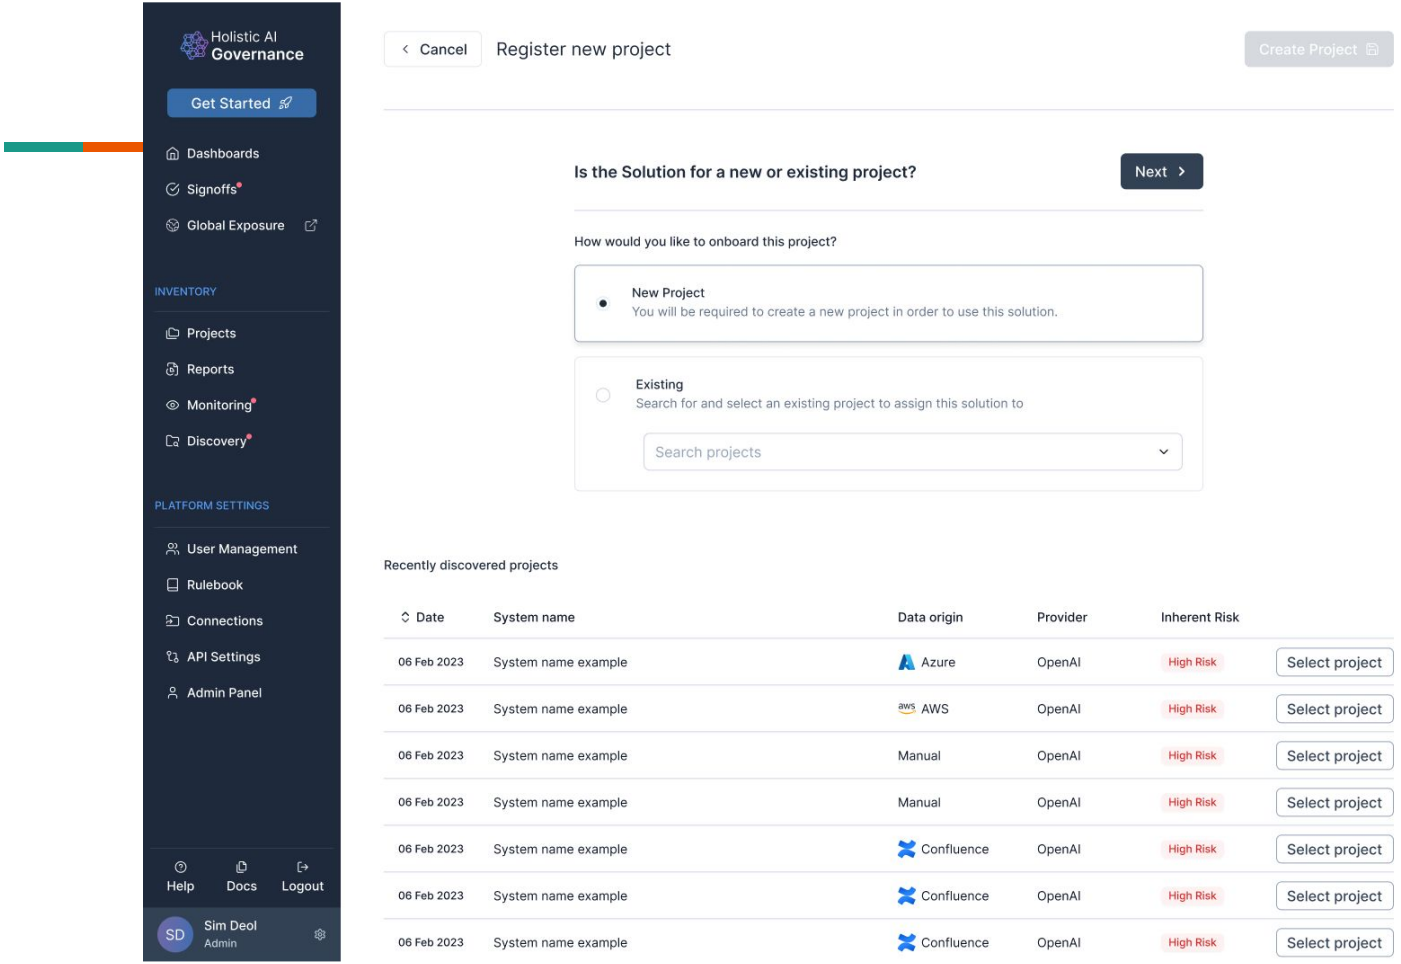

Collaborative view

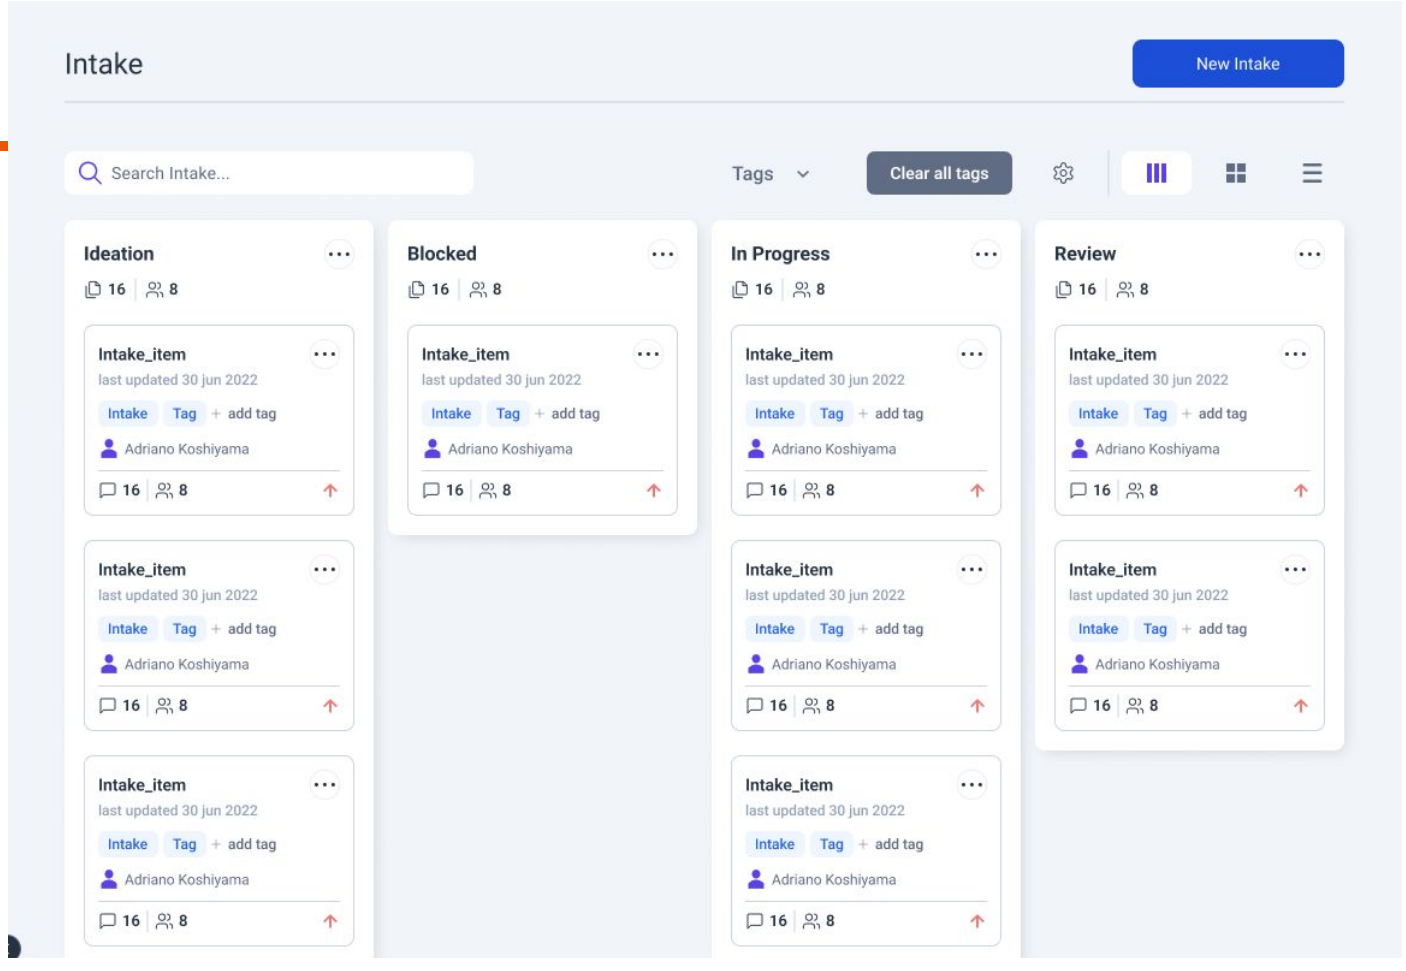

Collaborative view

Messages

Latest comments

Intake | Projects

Search projects or intake...

Intake_item
last updated 24 Sep 2024
Active | Tag | + add tag
Sim Deol — 16 — 3

Project_01
last updated 30 jun 2024
Active | Tag | + add tag
Adriano K — 8 — 3

Project_02
last updated 08 Apr 2023
Active | Tag | + add tag
Nigel Kingsman — 9 — 3

Project_03
last updated 24 Mar 2023
Active | Tag | + add tag
Sim Deol — 2 — 3

Project_04
last updated 30 jun 2022

Sim Deol REVIEWER — 24 Sep 2024 at 13:01
Commented on: Intake_item (Stage 2)
There seems to be a mistake on the project details. Please contact nigel.kingsman@holisticai.com.

Kasem Yassine ASSIGNEE — 24 Sep 2024 at 18:00
Commented on: Intake_item (Stage 5)
I still need to amend 3 fields on the form. Can I get back to this after I sort some other things out?

Sim Deol REVIEWER — 24 Sep 2024 at 18:32
Commented on: Intake_item (Stage 2)
Okay but I also added Nigel into the loop so he can take a look while you work on those

Type comment — Submit

Team Details — Add user

Team Members

Sim Deol
REVIEWER
sim.deol@holisticai.com

Kasem Yassine
ASSIGNEE
kasem.yassine@holisticai.com

Intake Name
Intake_item

Last Updated
24 Sep 2024 at 18:32

Potential Risk
Low Risk

Scaling Design Across New Products

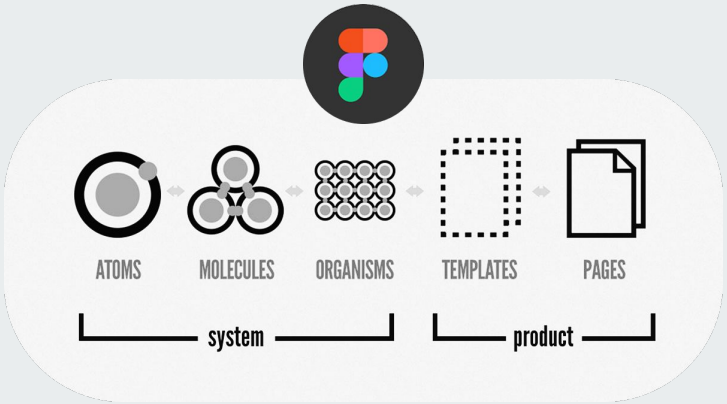

- With the new system in place, we expanded into two products in parallel:
    - **AI Tracker:** Policy monitoring for legal and compliance teams globally
    - **AI Safeguard**: LLM prompt auditing for risk, PII, hallucination, toxicity and much more
- System allowed immediate reuse and adaptation of components
- Designed with consistency, speed, and accessibility in mind

8 - The AI Tracker

Making Global AI Regulation Navigable - A real-time intelligence tool for monitoring AI legislation, policy, lawsuits, and more.

- Translated dense, multi-region AI law data into digestible UX
- Created dynamic filtering, search, and visualisation tools to help users digest large, complex data
- Introduced dynamic filters, visual maps, and intelligent grouping
- Designed for mixed audiences: legal users and operational decision-makers

Result: successful launch of a brand new product with over 1k sign-ups in one month.

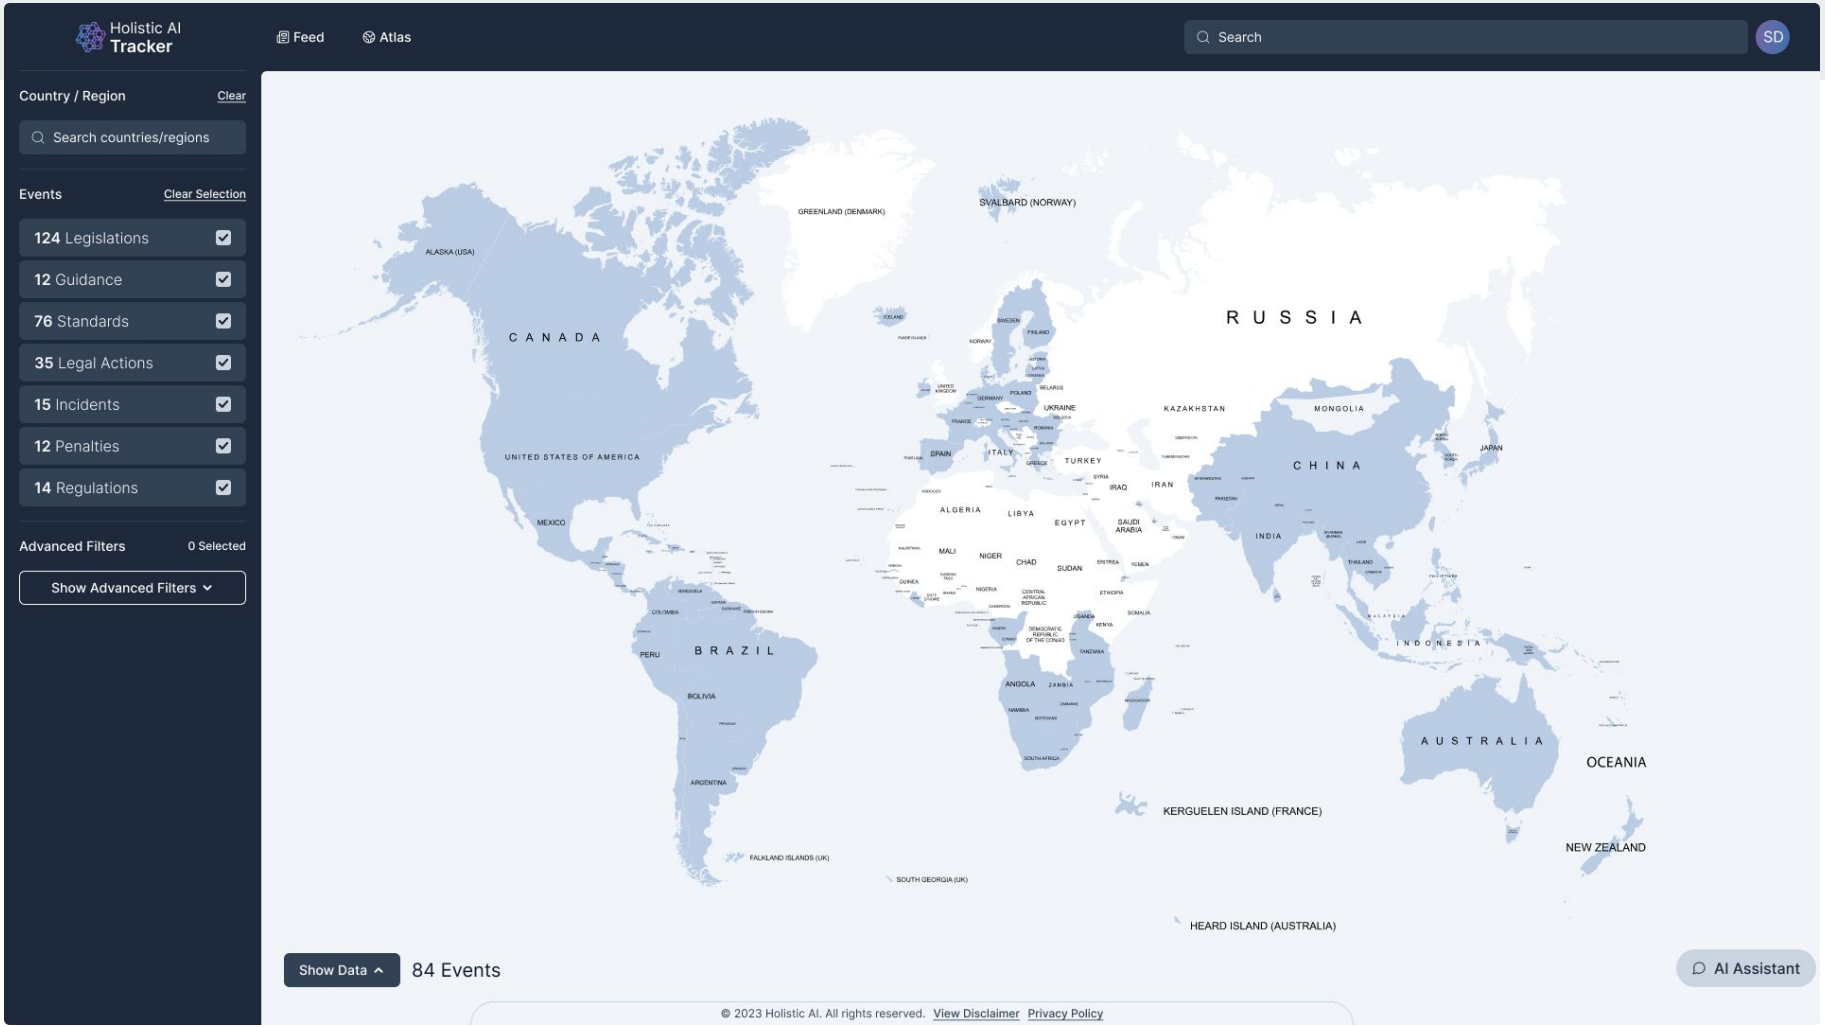
Holistic AI Tracker
Feed
Atlas
Search
SD
Country / Region
Clear
Search countries/regions
Events
Clear Selection
124 Legislations
12 Guidance
76 Standards
35 Legal Actions
15 Incidents
12 Penalties
14 Regulations
Advanced Filters
0 Selected
Show Advanced Filters
GREENLAND (DENMARK)
SVALBARD (NORWAY)
ALASKA (USA)
CANADA
RUSSIA
UNITED STATES OF AMERICA
KAZAKHSTAN
MONGOLIA
CHINA
JAPAN
MEXICO
INDIA
BRAZIL
INDONESIA
AUSTRALIA
OCEANIA
KERGUELEN ISLAND (FRANCE)
NEW ZEALAND
FALKLAND ISLANDS (UK)
SOUTH GEORGIA (UK)
HEARD ISLAND (AUSTRALIA)
Show Data
84 Events
AI Assistant
© 2023 Holistic AI. All rights reserved. View Disclaimer Privacy Policy

Holistic AI Tracker

Feed

Atlas

Search

SD

News

Blogs

World

Asia

Africa

Asia

Asia

Sectors

Subheading

Subheading

Technical

Subheading

Subheading

Risk Management

Subheading

Subheading

Explainer

Subheading

Subheading

Re-usable components built using Radix UI and Tailwind CSS

News Anti-discrimination

Nanda Smith, Sim Deol
07 Jun 2023

Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud

Re-usable components built using Radix UI and Tailwind CSS

News Anti-discrimination

Nanda Smith, Sim Deol
07 Jun 2023

Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud

Re-usable components built using Radix UI and Tailwind CSS

Blogs Tag

Nanda Smith, Sim Deol
07 Jun 2023

Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud

Re-usable components built using Radix UI and Tailwind CSS

Blogs Tag

Nanda Smith, Sim Deol
07 Jun 2023

Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud

Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud exercitation ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute...

Hot Topics

OpenAI
14 Articles

ChatGPT
14 Articles

Google
14 Articles

Holistic AI Expert Community

Tracked Experts

Open Source Resources

Github

Holistic AI Library

Upcoming Events

24th Feb 2023
Event Title
Event detail copy to describe the event taking place

24th Feb 2023
Event Title
Event detail copy to describe the event taking place

24th Feb 2023
Event Title
Event detail copy to describe the event taking place

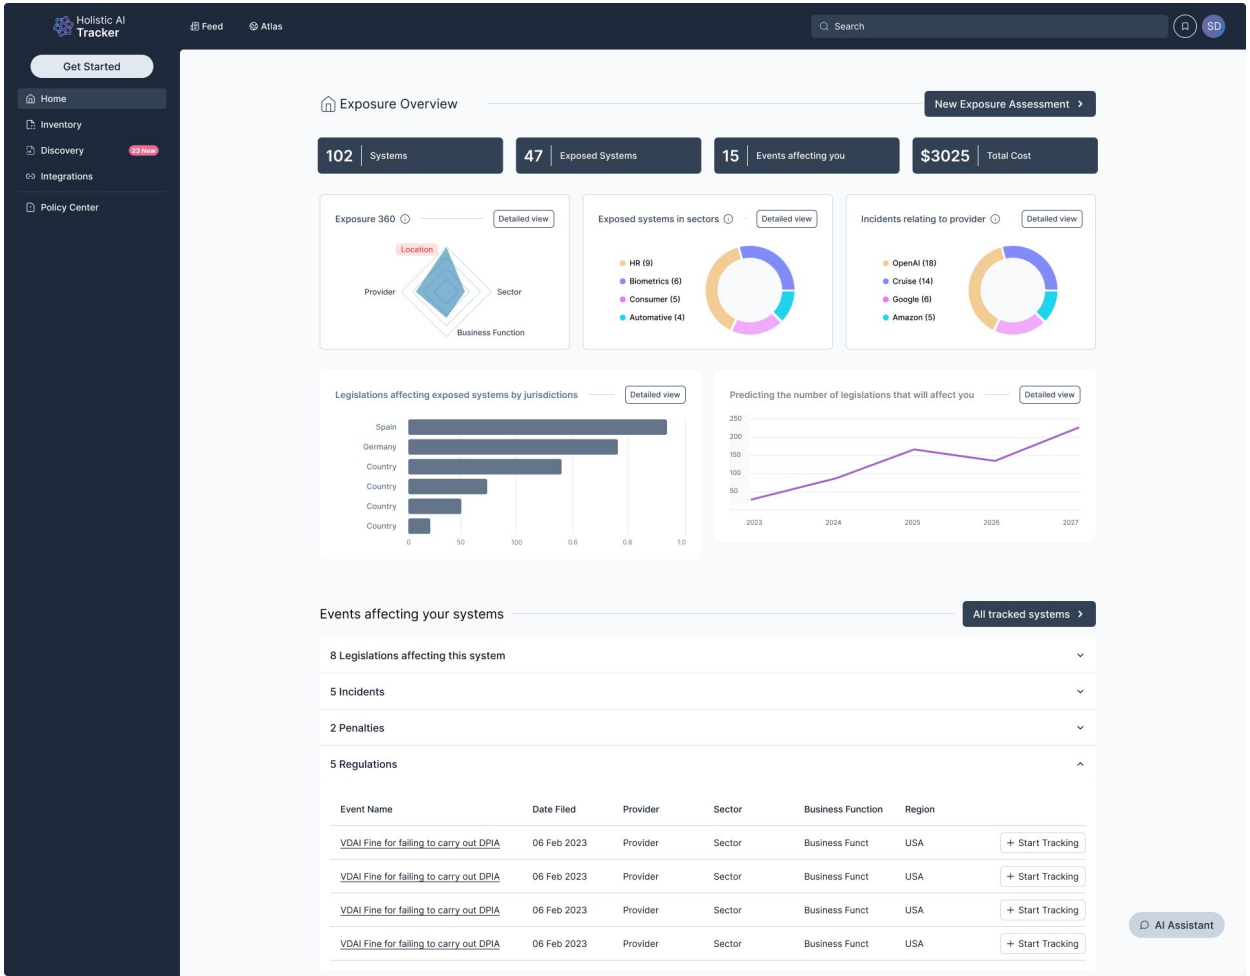
Holistic AI Tracker
Feed
Atlas
Search
SD
Get Started
Home
Inventory
Discovery
23 New
Integrations
Policy Center
Exposure Overview
New Exposure Assessment
102 Systems
47 Exposed Systems
15 Events affecting you
$3025 Total Cost
Exposure 360
Detailed view
Location
Provider
Sector
Business Function
Exposed systems in sectors
Detailed view
HR (9)
Biometrics (6)
Consumer (5)
Automative (4)
Incidents relating to provider
Detailed view
OpenAI (18)
Cruise (14)
Google (6)
Amazon (5)
Legislations affecting exposed systems by jurisdictions
Detailed view
Spain
Germany
Country
Country
Country
Country
0
50
100
0.6
0.8
1.0
Predicting the number of legislations that will affect you
Detailed view
250
200
150
100
50
2023
2024
2025
2026
2027
Events affecting your systems
All tracked systems
8 Legislations affecting this system
5 Incidents
2 Penalties
5 Regulations
Event Name
Date Filed
Provider
Sector
Business Function
Region
VDAI Fine for failing to carry out DPIA
06 Feb 2023
Provider
Sector
Business Funct
USA
+ Start Tracking
VDAI Fine for failing to carry out DPIA
06 Feb 2023
Provider
Sector
Business Funct
USA
+ Start Tracking
VDAI Fine for failing to carry out DPIA
06 Feb 2023
Provider
Sector
Business Funct
USA
+ Start Tracking
VDAI Fine for failing to carry out DPIA
06 Feb 2023
Provider
Sector
Business Funct
USA
+ Start Tracking
AI Assistant

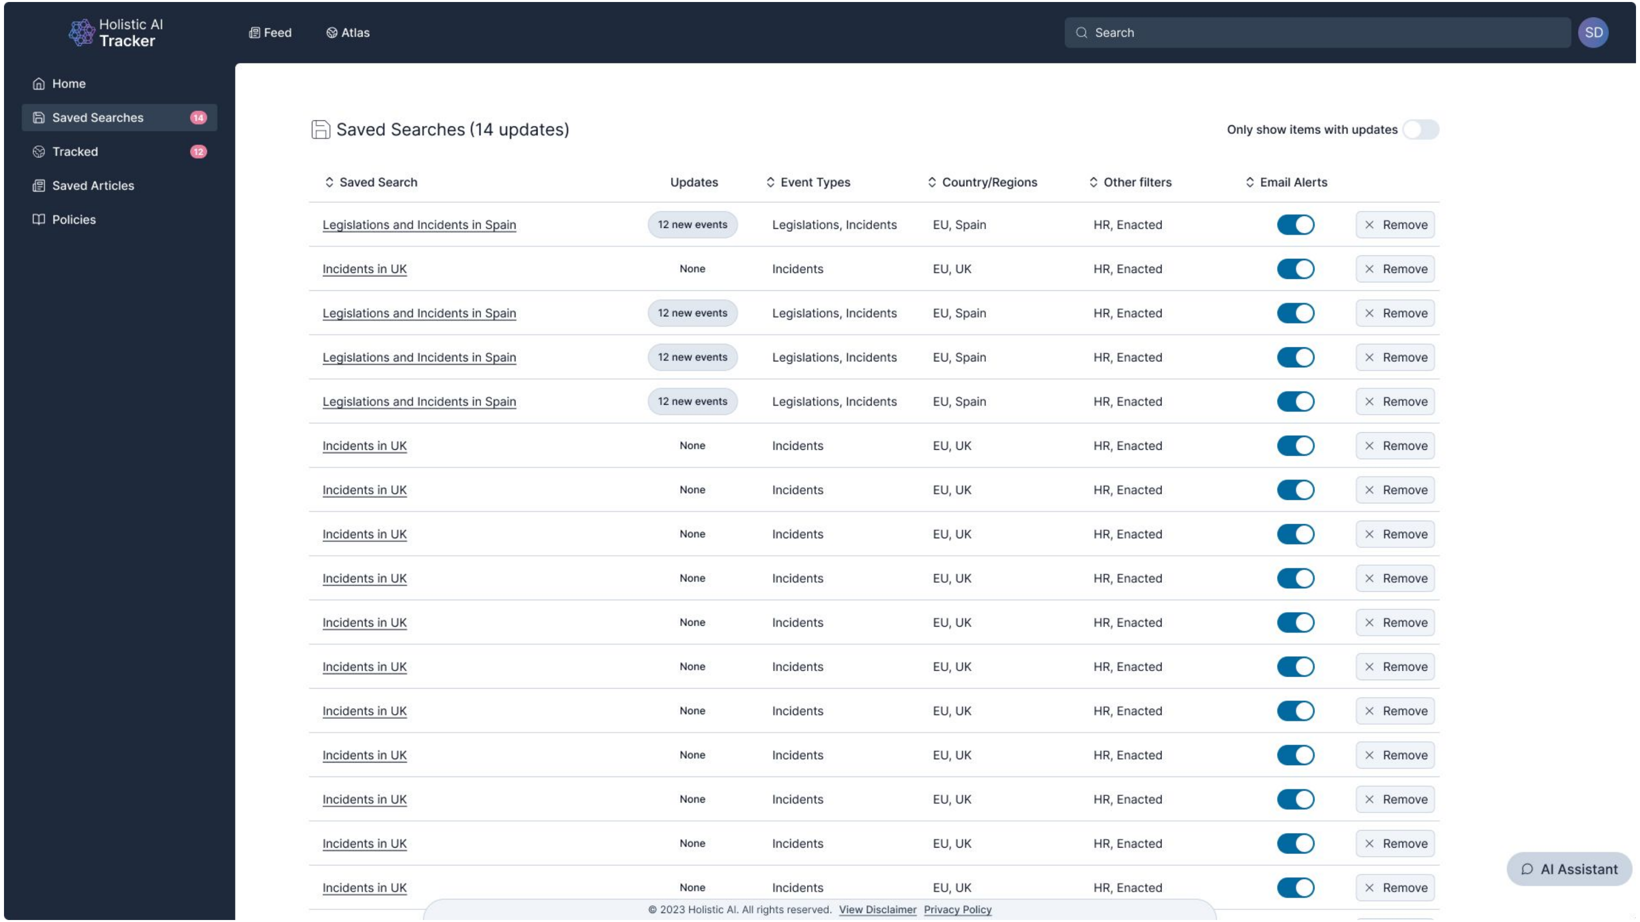
Holistic AI
Tracker
Feed
Atlas
Search
SD
Home
Saved Searches 14
Tracked 12
Saved Articles
Policies
Saved Searches (14 updates)
Only show items with updates

Saved Search	Updates	Event Types	Country/Regions	Other filters	Email Alerts	
Legislations and Incidents in Spain	12 new events	Legislations, Incidents	EU, Spain	HR, Enacted		Remove
Incidents in UK	None	Incidents	EU, UK	HR, Enacted		Remove
Legislations and Incidents in Spain	12 new events	Legislations, Incidents	EU, Spain	HR, Enacted		Remove
Legislations and Incidents in Spain	12 new events	Legislations, Incidents	EU, Spain	HR, Enacted		Remove
Legislations and Incidents in Spain	12 new events	Legislations, Incidents	EU, Spain	HR, Enacted		Remove
Incidents in UK	None	Incidents	EU, UK	HR, Enacted		Remove
Incidents in UK	None	Incidents	EU, UK	HR, Enacted		Remove
Incidents in UK	None	Incidents	EU, UK	HR, Enacted		Remove
Incidents in UK	None	Incidents	EU, UK	HR, Enacted		Remove
Incidents in UK	None	Incidents	EU, UK	HR, Enacted		Remove
Incidents in UK	None	Incidents	EU, UK	HR, Enacted		Remove
Incidents in UK	None	Incidents	EU, UK	HR, Enacted		Remove
Incidents in UK	None	Incidents	EU, UK	HR, Enacted		Remove
Incidents in UK	None	Incidents	EU, UK	HR, Enacted		Remove
Incidents in UK	None	Incidents	EU, UK	HR, Enacted		Remove
Incidents in UK	None	Incidents	EU, UK	HR, Enacted		Remove

AI Assistant
© 2023 Holistic AI. All rights reserved. View Disclaimer Privacy Policy

The AI Safeguard

Oversight for LLMs. A platform for monitoring and auditing the usage of Large Language Models (LLMs) within enterprises.

- Created real-time interface for AI prompt testing (toxicity, PII, hallucinations)
- Designed dashboards to visualise model outputs and flag risks
- Built flows for automated test creation, result reporting, and compliance tracking

Result: successful launch of a new product, with 2 major clients adopting the product.

Holistic AI
Safeguard

Get Started

- Dashboard
- Audits
- ROI
- Monitoring
- Chatbot
- Governance

MONITORING

- Settings
- User Management

ADMIN

- Settings
- User Management

Dashboard

Overview | Risks | Usage | Leaderboard

33	15	36	$7k
Audits	APIs	Users	Cost Savings (ROI)

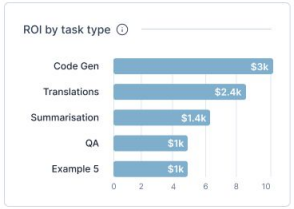

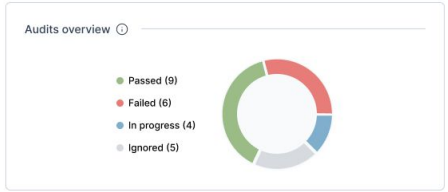

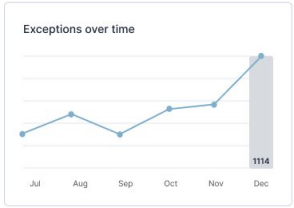

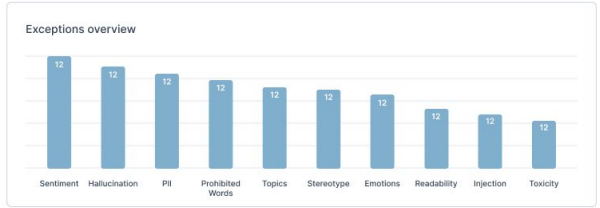

Top scoring LLMs (Global benchmarks)

LLM	Overall	Efficacy	Privacy	Robustness	Bias
GPT-4-turbo	0.1453	0.453	0.453	0.453	0.453

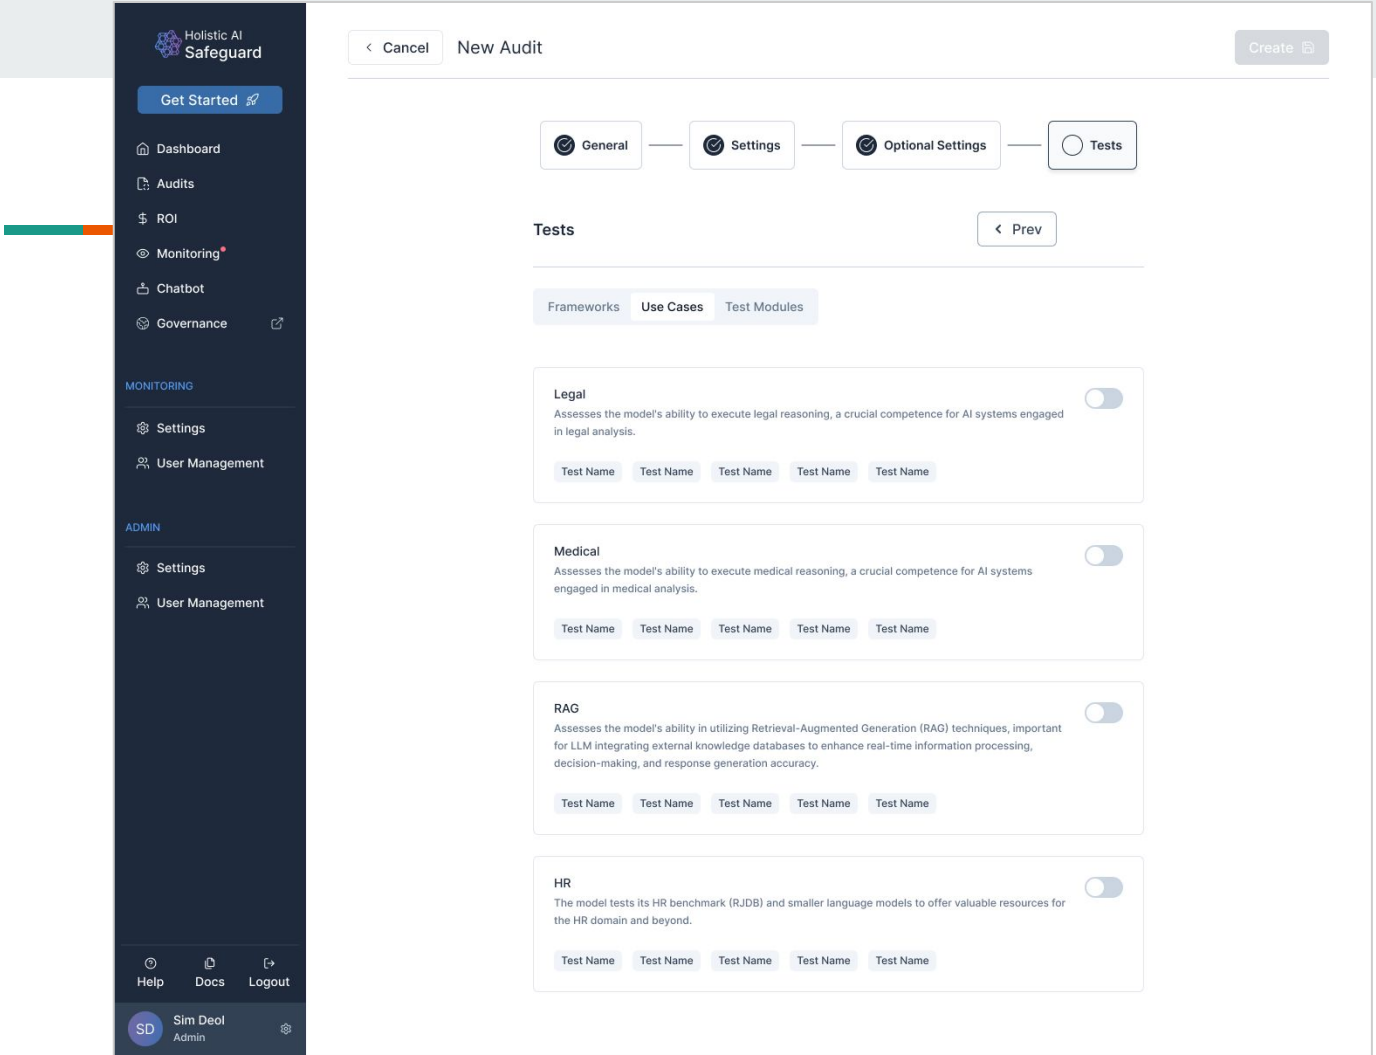
Holistic AI
Safeguard
Get Started
Dashboard
Audits
ROI
Monitoring
Chatbot
Governance
MONITORING
Settings
User Management
ADMIN
Settings
User Management
Help
Docs
Logout
SD
Sim Deol
Admin
Cancel
New Audit
Create
General
Settings
Optional Settings
Tests
Tests
Prev
Frameworks
Use Cases
Test Modules
Legal
Assesses the model's ability to execute legal reasoning, a crucial competence for AI systems engaged in legal analysis.
Test Name
Test Name
Test Name
Test Name
Test Name
Medical
Assesses the model's ability to execute medical reasoning, a crucial competence for AI systems engaged in medical analysis.
Test Name
Test Name
Test Name
Test Name
Test Name
RAG
Assesses the model's ability in utilizing Retrieval-Augmented Generation (RAG) techniques, important for LLM integrating external knowledge databases to enhance real-time information processing, decision-making, and response generation accuracy.
Test Name
Test Name
Test Name
Test Name
Test Name
HR
The model tests its HR benchmark (RJDB) and smaller language models to offer valuable resources for the HR domain and beyond.
Test Name
Test Name
Test Name
Test Name
Test Name

Governance and rollout strategy

- Created documentation within Figma with usage rules and naming logic
- Ran handoff walkthroughs and design-dev-pm syncs to promote adoption and prioritise which components needed developing first
- conducted system demos and shared libraries to support scale and quality
- Created a contribution model (proposal → review → build → release).
- Created clear and consistent guides, usage rules, and rationale for components and patterns.

Leadership

I lead by amplifying strengths, removing roadblocks, and fostering true collaboration between design, product, and engineering for a shared goal

Measurable impact

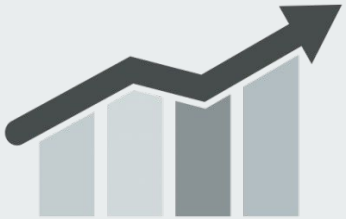

- 60% faster onboarding via redesigned flows and automation
- Retention of Unilever as a key client
- Signed Mapfre as a major new customer
- System adopted by engineers across three frameworks
- Launched AI Tracker to over 1,000 sign-ups in the first month
- 2 major clients signing-up for the AI Safeguard

Reflection

- Wins: scaled UX, created shared language between teams, accelerated roadmap
- Challenges: legacy code, aligning stakeholders, no full-time researcher
- Next time: introduce design system governance earlier, invest in onboarding documentation sooner
- When the time is right: Introduce a better, more customised UI

Further exploration

- ➔ A singular platform with modules that allowed each organisation to tailor the platform to their needs
- ➔ AI-driven onboarding tips (e.g. GPT-powered assistant for workflows)
- ➔ Persistent shared project activity stream as a separate screen (cross-team collaboration view)
- ➔ Smart templates - creating project flows based on the project type, not just the user type

How about it? Let’s work together and build the future together :)