

INTUIT



GenOS

Accelerating GenAI-powered Experiences at Scale

Merrin Kurian, Distinguished Engineer, Intuit

The AI Conference 2024

Sep 11, 2024



Agenda

AI/ML at Intuit

Intuit Assist

Experiences powered by GenOS

GenOS(Generative AI OS)

Motivation to develop one

Deep Dive into GenOS

Architecture, Components

What does it take to build a GenOS?

and other learnings

Q & A

AI/ML at Intuit

65B

Machine learning predictions per day

500K

Customer/financial attributes per SMB

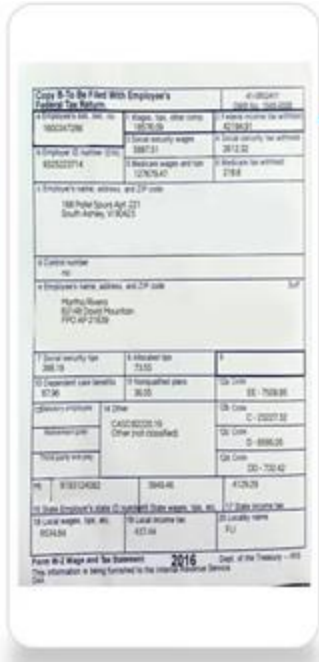
60K

Tax and financial attributes per consumer

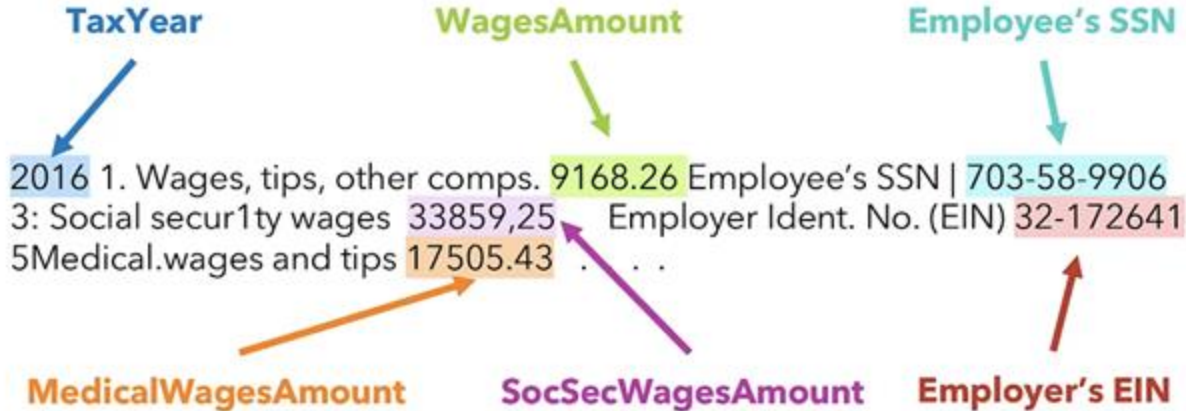


AI-Driven Expert Platform

TurboTax: Automatic Extraction from forms



2016 1. Wages, tips, other comps. 9168.26 Employee's SSN | 703-58-9906
 3: Social secur1ty wages 33859,25 Employer Ident. No. (EIN) 32-172641
 5' Medical.wages and tips 17505.43



QuickBooks: Extraction & Categorization

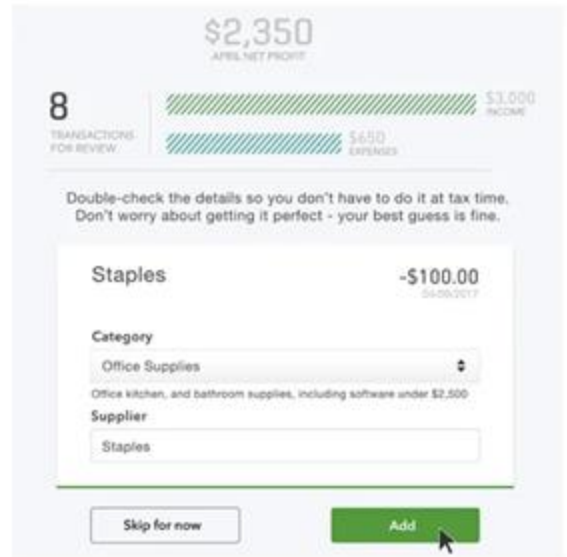


Vendor: Aroma Cafe

Date: 2016-03-12

Credit Card: 8281

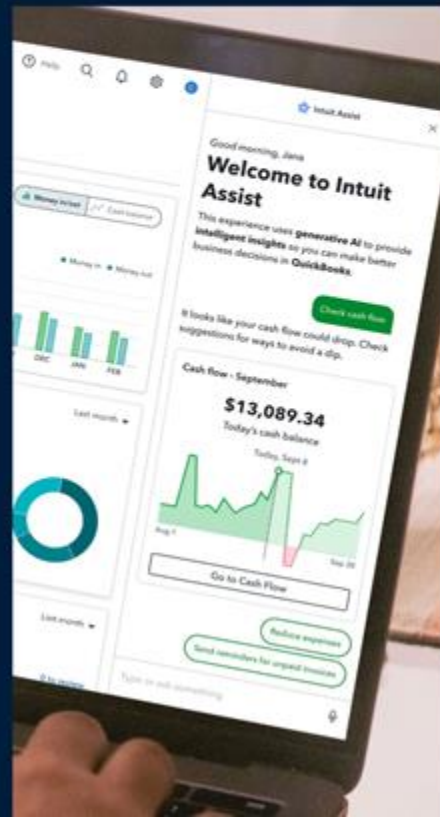
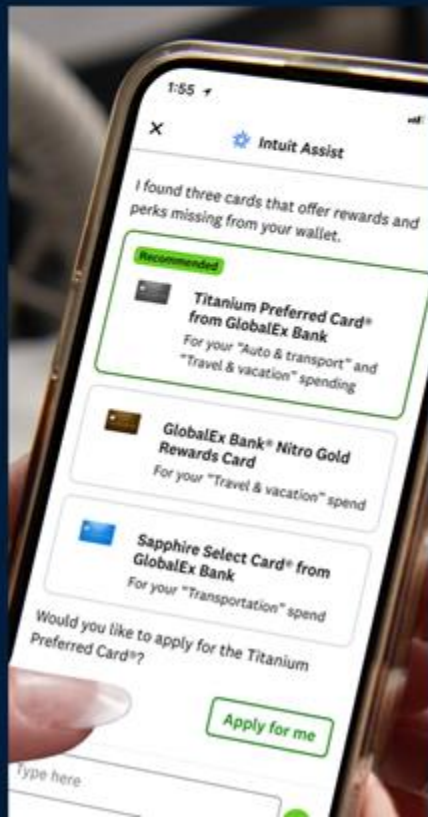
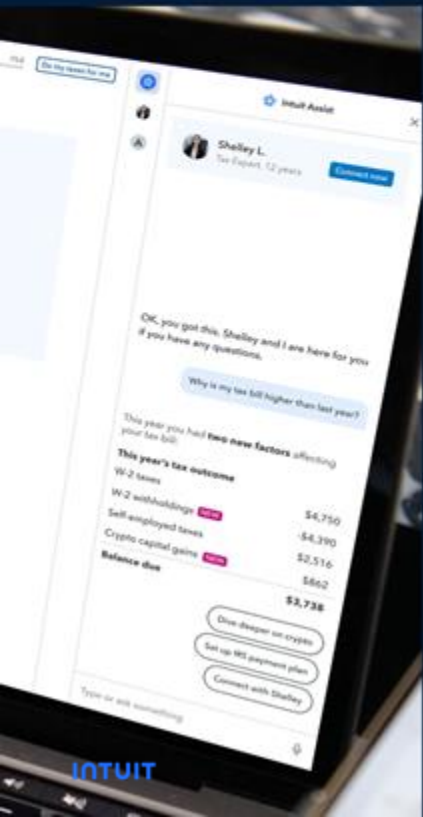
Total: \$90.52



Intuit Use Case:
Extract information from documents

Match to QB transactions

Intuit Assist





Connecting People to Tax Experts




Intuit Assist Q&A

We introduced Intuit Assist Q&A to iEP and Salesforce for experts. This powerful tool leverages generative AI, utilizing our internally authored content, Elevate community contributions, and additional sources to provide precise answers to expert inquiries. This launch has significantly enhanced our mission to empower experts with advanced tools that boost productivity and effectiveness.



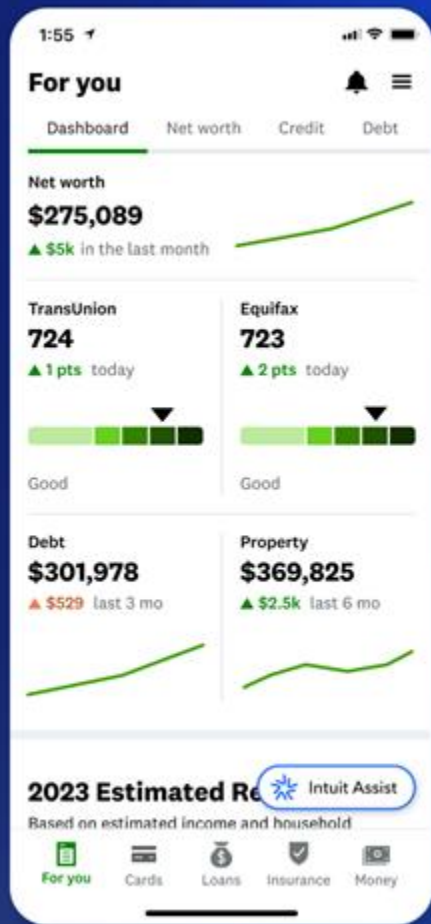
Call Summary

We launched Call Summary to iEP for experts. This capability provides real-time conversation summaries, allowing experts to focus on customer interactions. It shows our commitment to leveraging generative AI for enhanced productivity. We also added the function to edit and verify summaries, catering to user needs and improving the overall experience.



Generative Email

We rolled out Generative Email to iEP for experts. This feature reduces the time spent on writing follow-up emails by automatically generating concise summaries of conversations. This feature underscores our dedication to accelerating expert productivity through innovative AI solutions.



1:55

X Intuit Assist

Help me find a credit card

Sure thing!

First, let's look at your spending.

Top spending categories

Food & dining	\$2,459.30
Transportation	\$1,074.24
Travel & vacation	\$642.68

Type here

1:55

X Intuit Assist

Based on your spending, here are the best card options.

Recommended

Titanium Preferred Card® from GlobalEx Bank
 For your "Food & dining" and "Transportation" spend

GlobalEx Bank® Nitro Gold Rewards Card
 For your "Travel & vacation" spend

Sapphire Select Card® from GlobalEx Bank
 For your "Transportation" spend

Would you like to apply for the Titanium Preferred Card®?

Apply for me

Type here

Email builder

Layouts Strategy Build email Settings

Intuit Assist

Get more customers with email builder

Create an email campaign to announce a new product - the Cavendish Comfort chair

Email purpose Furniture Sort: Newest

Create an email campaign to announce a new product - the Cavendish Comfort chair

Email purpose Furniture

The screenshot displays the Mailchimp Email Builder interface. At the top, there's a navigation bar with 'Email builder' and tabs for 'Layouts', 'Strategy', 'Build email', and 'Settings'. A 'Start from scratch' button and a 'Continue' button are also visible. Below the navigation, the 'Intuit Assist' feature is active, showing a campaign title: 'Create an email campaign to announce a new product - the Cavendish Comfort chair'. There are dropdown menus for 'Email purpose' (set to 'Furniture') and 'Sort' (set to 'Newest'). A search bar below the navigation contains the same campaign title. The main area shows a grid of five email templates, each featuring a yellow armchair in a modern living room setting. The templates are labeled 'Option #1' through 'Option #5'. Option #1 is titled 'Cavendish Comfort Chair'. Option #2 is titled 'Framework'. Option #3 is titled 'Be the first to secure your new Cavendish Comfort Chair and preorder today!'. Option #4 is titled 'Be the first to secure your new Cavendish Comfort Chair and preorder today!' and includes a 'Preorder' button. Option #5 is titled 'Be the first to secure your new Cavendish Comfort Chair and preorder today!' and includes a 'Preorder' button. Each template has a 'Preview' icon in the bottom right corner.



GenOS: A Paved Road for GenAI App Development at Intuit

Motivation for Building a Proprietary GenOS

- Lack of enterprise-grade GenAI app dev solutions
- Critical importance of responsible AI & data governance
- Enabling product development teams to move with velocity

Build a Paved Road for GenAI App Development

Leverage our investments in platform, data and AI
Enable rapid experimentation as technology evolves
Streamline end-to-end app development

What does GenOS offer?

Responsible AI Development

- Secure, private and reliable access to LLMs
- Out-of-the-box guardrails on security, safety, privacy, compliance

Rapid Experimentation at Scale

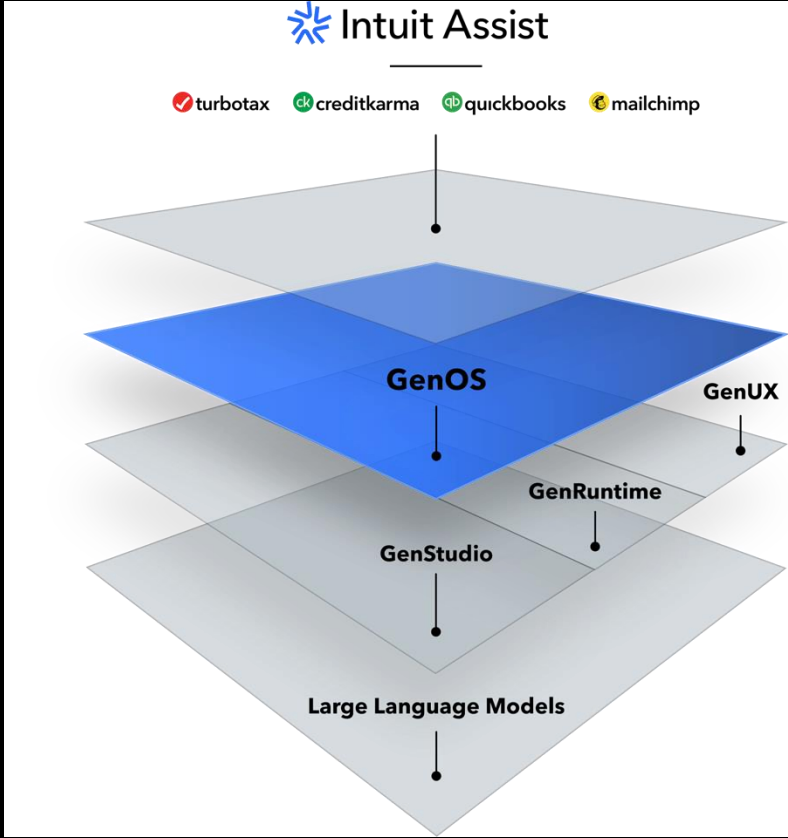
- Data handling to meet governance needs
- End-to-end observability & analytics via instrumentation

Extensible Platform to overcome limitations of LLMs

- Plug-in domain capabilities & knowledge to ground LLMs
- Extensibility for guardrails, evaluation metrics & more

Use case teams focus on solving customer problems.

GenOS Deep Dive: Architecture & Components



LLMs available in GenOS

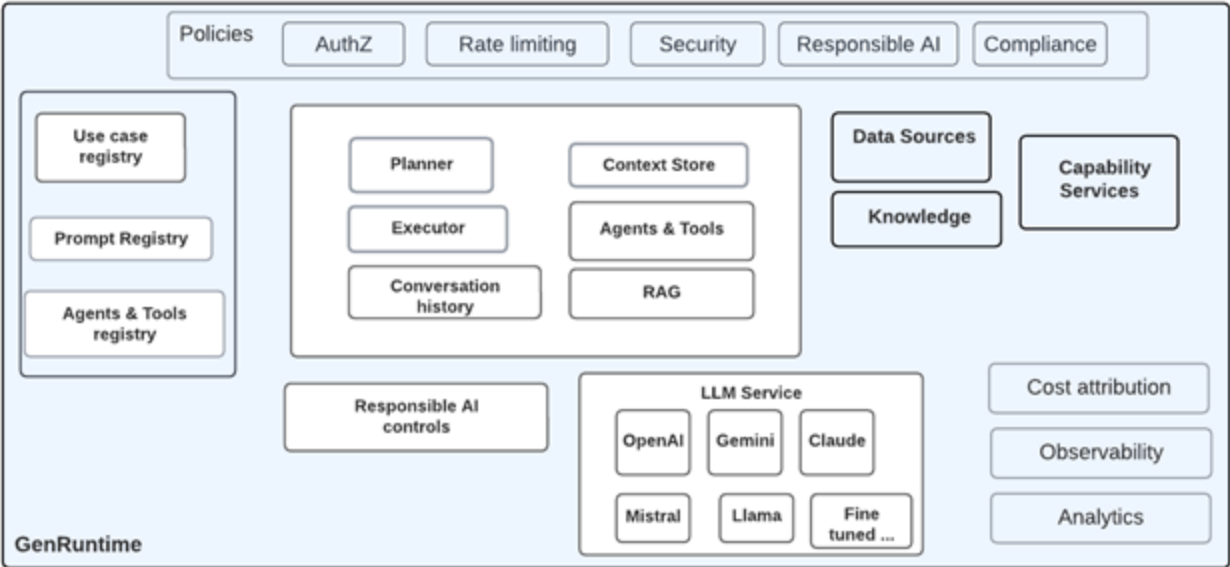



...and their fine-tuned derivatives

Fine tuning & inference

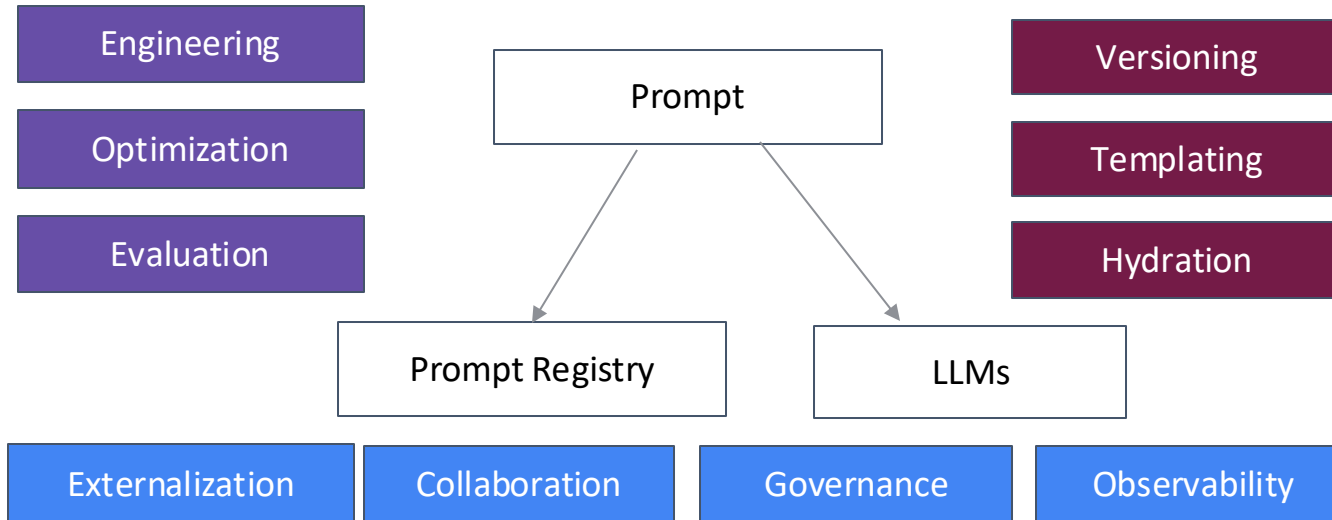
- Leverages Intuit's standard MLOps practice and infrastructure
- Specialized training/inference pipelines with GPUs
- Configurable pipelines with Intuit-specific benchmarks and evaluations
- Fine-grained access control for secure hosting & rapid experimentation

GenOS Architecture

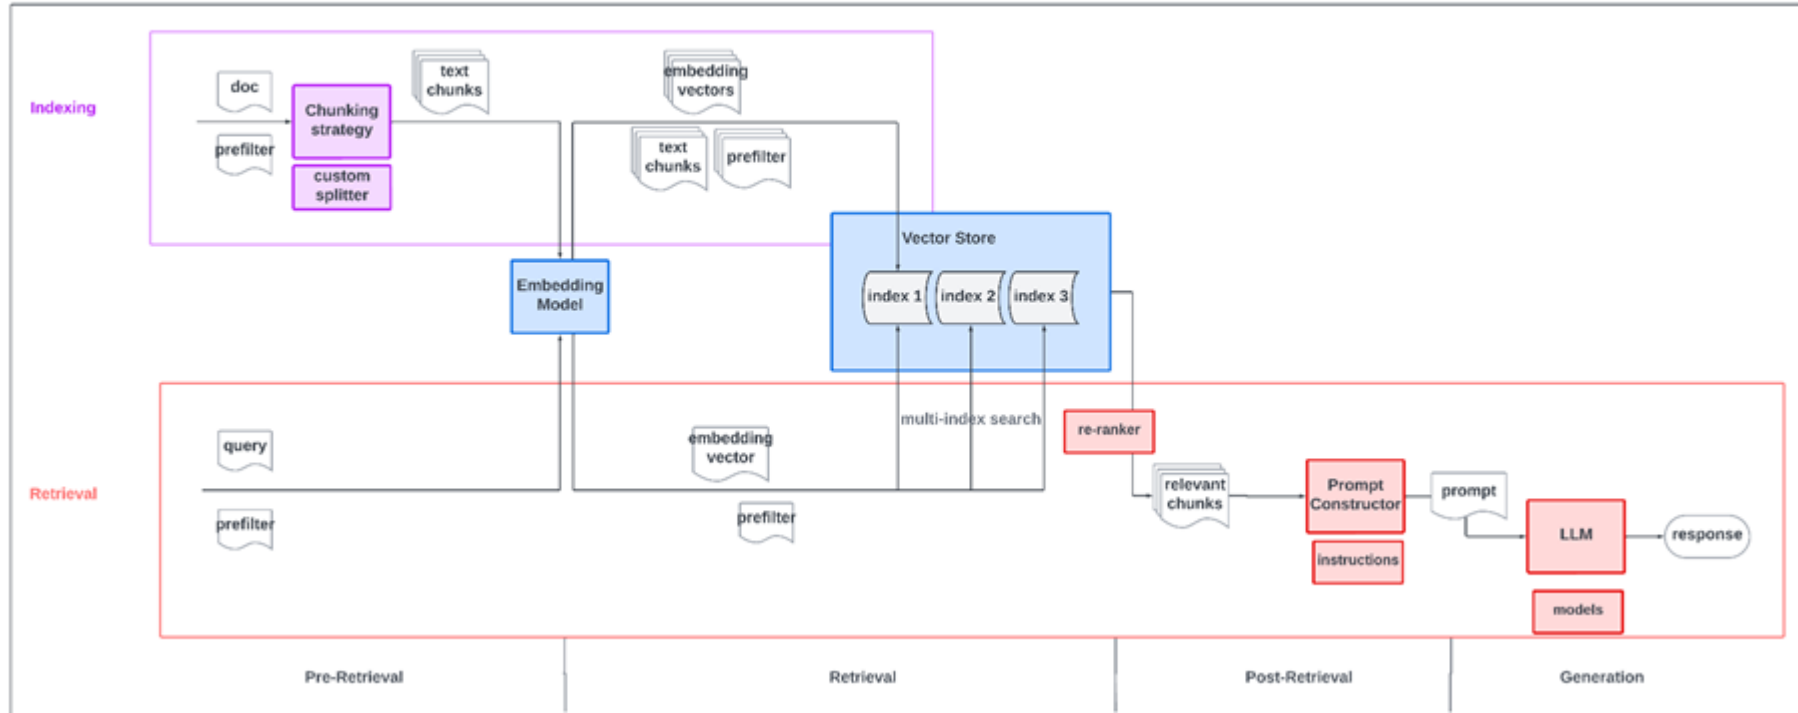


Prompt management in GenOS

Prompts are first-class entities with their own lifecycle

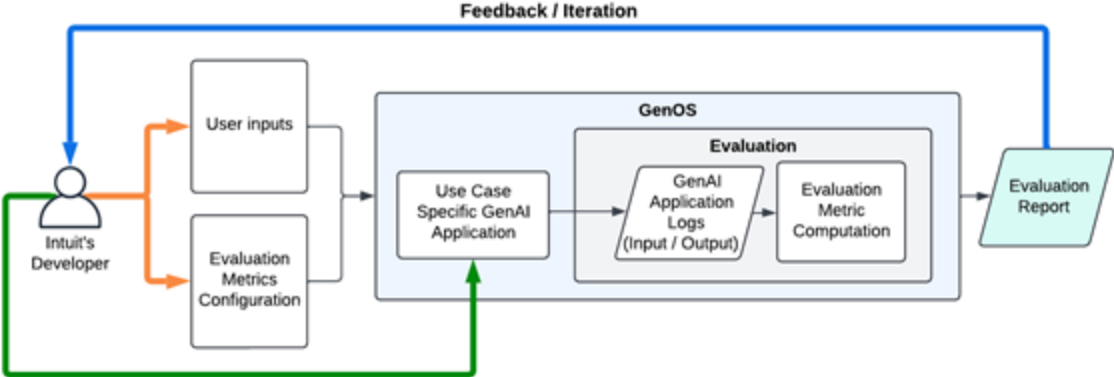


RAG pipelines

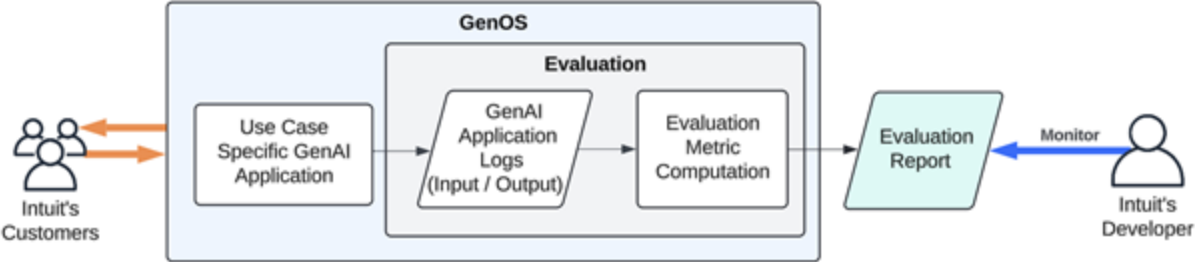


Evaluations

Development: Offline Evaluation



Production: Monitoring



What it Takes to Build a GenOS and Other Learnings...

Working with LLMs

- Prompt portability is a challenge across LLMs
- Prompt engineering is collaborative effort
- Working with LLMs is science; learn evaluation in addition to prompt engg
- Widely available benchmarks are more academic; develop domain/task-specific ones
- Real world problems ground us on cost, latency, accuracy of LLMs

GenOS: A mission like no other!

What does it take to build a GenOS?

- **Commitment** from the highest levels of leadership
- Clear decision making **processes** & discussion forums to move fast
- Scrum teams from various organizations in a **single mission**
- **Unify** Existing capabilities, Enhance as needed, Build net new
- Extend with **contributions** across the organization
- **Experimentation** outside of the platform
- **Adaptable** teams quickly pivoting as technology/customer needs evolve

GenOS accelerates GenAI app development

Enable every technologist to leverage GenAI to solve problems

Collective learning is accessible to everyone

Rapid experimentation at scale requires enforceable guardrails

Self-serve onboarding and guided flows accelerate experimentation at scale

288 projects(Spring GED 2023) -> 900 projects (Spring GED 2024)

1000s of Intuit technologists

100s of GenAI experiments with

customers

Game changing new experiences with

Intuit Assist

Thank you!

Q&A