



Coming next

The session will begin shortly · September 3, 2026



IT G E H G I T I G I T H I I

The Rise of Agentic AI

Claude in Copilot Studio & Microsoft Foundry



Mitul Tailor

IT G E H G I T I G I T H I I



Sarang Bhrame

Principal Architect · LVS

SEPT 3, 2026 · VANCOUVER

Our Sponsors

PLATINUM



GOLD



SILVER



BRONZE



COMMUNITY PARTNER



MEDIA PARTNER

DevIT:>_

About Me

Mitul Tailor

Sr. Data Engineer & Consultant

Long View Systems — Toronto, Canada

Microsoft Certified Trainer (MCT)

Delivering official Microsoft training in data & AI

Passionate Technologist

Data, generative AI, and enterprise Copilots

Agentic AI Builder

Autonomous agents with Copilot Studio & Microsoft Foundry

Community Contributor

Sharing on Claude, Copilots, and applied AI

Let's Connect


LinkedIn: [@mitultailor15](https://www.linkedin.com/in/mitultailor15)



About Me

Sarang Brahme

[linkedin.com/in/sbrahme](https://www.linkedin.com/in/sbrahme)



Almost 2 Decades in the Tech Trenches

A Principal AI Architect, Technical Product Owner/Leader who turns big ideas into real, working AI roadmaps.

C-Suite and VC Translator

I speak fluent "Business Value," "Data Science," and "Software Engineer"—and help them all get along.

Current Obsession

Agentic AI, MLOps, building/advising startups. If it makes life easier intelligently, I'm building it.

Tech Stack of Choice

Azure at the moment, but also - Databricks, Fabric, GCP, Open Source

Speaker & Practitioner

I do talk about AI on stages; and I actually build the tech when I get off stage.

Warning

Once started, I can't stop talking about agentic AI, quantum computing, heavy metal, and philosophy.

Agenda

01



02



03



04

The Background

How claims are triaged today

Claim form vs. policy history

What adjusters need

Data & Architecture

Bronze, silver, gold in Fabric

Reference architecture

Extraction & Reasoning

AI Builder extraction

Reasoning over claim history

Rules vs. judgement

Claude & Copilot Studio

Claude in Foundry

The adjuster front end

DEMO!!

A claim that should take days takes weeks

CAUSE 1

Intake is unstructured

Photos, a police report PDF, two repair estimates in different formats, an FNOL call, free-text notes.

CAUSE 2

Policy reading is judgment

Coverage, deductible, exclusions, endorsements. Forty pages that must be read whole, not in chunks.

CAUSE 3

No API on the system of record

In most carriers the adjuster is the integration layer — reading one screen, typing into another.

CAUSE 4

Exceptions escalate to a queue

And the queue is a person. This is where the days actually disappear, undocumented.

11 days expected vs 23.9 days actual (J.D. Power 2024)— none of it a model problem, all of it context and orchestration.

From Collision to Clarity

STEP 1

Claim arrives

A form lands in the queue with photos and free text attached.

STEP 2

Adjuster opens two systems

The document in one window, policy and history in another.

STEP 3

Manual cross check

Coverage, limits, deductible, and past claims are compared by eye.

STEP 4

Judgement call

Approve, deny, or escalate, with the reasoning undocumented.


Every claim repeats the same manual comparison, and no two adjusters weigh it the same way.

A claim that should take days takes weeks

What policyholders expect



Industry average cycle time



Source: J.D. Power 2024 U.S. Auto Claims Study

≈ 13-day gap — what you get measured on



Intake is unstructured

Photos, a police report PDF, two repair estimates in different formats, a recorded FNOL call, free-text notes.



Policy reading is judgment

Coverage, deductible, exclusions, endorsements. Forty pages that must be read whole, not in chunks.



System of record has no API

In most carriers the adjuster is the integration layer — reading one screen, typing into another.



Exceptions escalate to a queue

And the queue is a person. This is where the days actually disappear.

None of these is a model problem. Every one is a context problem or an orchestration problem — and that distinction is the whole talk.

Incomplete, multimodal, and slightly inconsistent by design.

WHAT ARRIVED

- Claim submission
- Claimant Information
- Loss Details
- Description of Loss
- Policy and vehicle details
- Prior claims summary

The narrative describes a rear impact. The submitted photos primarily show front-left damage.

Define the outcome before you pick the platform

TARGET OUTCOMES, WRITTEN AS DELTAS



FNOL triage: hours → minutes

with severity, coverage determination and a fraud signal already attached



Straight-through on the simple tail

humans keep the contested claims, which is where they add value anyway



Adjuster time reallocated

from assembling documents to exercising judgment



Every decision carries a citation

back to the policy clause it relied on — untraceable answers are unusable in a regulated line

40%+

of agentic AI projects will be cancelled by the end of 2027

GARTNER NAMES THREE CAUSES

Escalating cost

Unclear business value

Inadequate risk controls

All three are architecture decisions you make in week one and discover in month nine.

Cost per claim, cycle time and audit trail are design inputs — not phase two.

A recommendation is not enough. Show the evidence.

RECOMMENDED ROUTE

CONFIDENCE · MEDIUM


Complex damage review

- 1 Coverage appears active for the incident date.
- 2 Front-left damage is visible across submitted photographs.
- 3 The narrative describes a rear impact.
- 4 Suspension work in the estimate is not clearly supported by available images.
- 5 A required wide-angle scene photograph is missing.

HUMAN REVIEW REQUIRED

What this workload needs from the solution?

Model selection: six reasoning demands this investigation places on the model.




The model should be chosen based on the reasoning workload – not because one model is always best.

Reasoning over claim history

Claude reads the extracted claim next to the customer's full history, and the pattern is what matters.

Frequency	Three windshield claims in 18 months Short interval, same damage type
Consistency	Amounts cluster under the review threshold Worth a second look, not an accusation
Coverage	Comprehensive, deductible met Nothing disqualifying in the policy
Context	No prior escalations on this policy History is otherwise clean


From conversation to governed action.



Enterprise APIs remain the systems of record.

WORKFLOW

Operational claim interaction is separated from analytics, historical insight, and model lifecycle services.



One model should not do every job.

LIGHTWEIGHT MODEL

Classification and routing

Speed + cost

EXTRACTION SERVICE

Forms and document fields

Deterministic structure

CLAUDE

Multimodal evidence synthesis

Deep reasoning

BUSINESS RULES

Coverage and threshold checks

Consistency


HUMAN ADJUSTER

Material customer-impacting decision

Accountability

Architecture starts with the task, risk, and evidence requirement, not the model name.

Agent Flow



What Agents Look Like in a Demo

Copilot Studio

Dataverse/CSV

What Agents Look Like in Production

Copilot Studio

Orchestrator

Hotel Agent

Logistics
Agent

Quotation
Agent

Static Web App / API

DirectLine
SDK

Speech

Foundry
Summarize

Dataverse
Proxy

Microsoft Foundry

Destination Expert
(gpt-5.4)

AI Search

Code Interpreter

Power Automate Flows

Search Hotels

Lookup Reservation

Create Draft Quote

Data & Storage

Dataverse

Blob Storage

AI Search

External APIs

Sabre GDS

Expedia OTA

Google Maps
Platform

Demo Time!



Key takeaways

What to remember about bringing Claude into Copilot Studio.

01

Not just a model swap

Claude in Copilot Studio is a step-change in what your agents can understand and deliver.

02

Use it where it counts

Document analysis, long-form summarization, and multi-step problem solving are its sweet spots.

03

Choose per task


Multi-model means picking Claude or GPT per agent or skill — balanced, not absolute.

04

Design decides value

Thoughtful prompts and well-structured knowledge sources matter more than flipping a selector.

Share Your Feedback



Share comments for the presenter,
that may help them in **future sessions**



Q&A

Thanks for listening — let's connect.

J 6ETHIG
yoursite.com

IHI HGGHI
[/your-name](https://your-name)

VANCOUVER MICROSOFT 365 SUMMIT 2026