

Welcome



Hosted by SBD Automotive

The Profitability Path of Automotive AI

From “Hype” to “Sustainable Profit”

April 15, 2026



SoundHound AI





Today's agenda

Industry perspectives

AI expert takes

Speaker roundtable

Audience Q&A



Robert Fisher
Moderator





29
YEARS
1997 - 2026

A global team of research and consulting experts helping our clients deliver **Safe, Secure, Sustainable & Seamless** mobility

Our Areas of Expertise



Connected
Vehicles



Autonomous
Vehicles



SW-Defined
Vehicles



Electric
Vehicles



Cyber
Security

USA

UK

Germany

India

China

Japan

Headquarters

Automotive AI is no longer a new concept...

- Most automakers have already adopted an AI VPA (Virtual Personal Assistant)
- Chinese OEMs are very active in AI implementation
- Alibaba offers 3 separate AI models, with cloud and hybrid solutions

Recent Active AI VPA Implementations

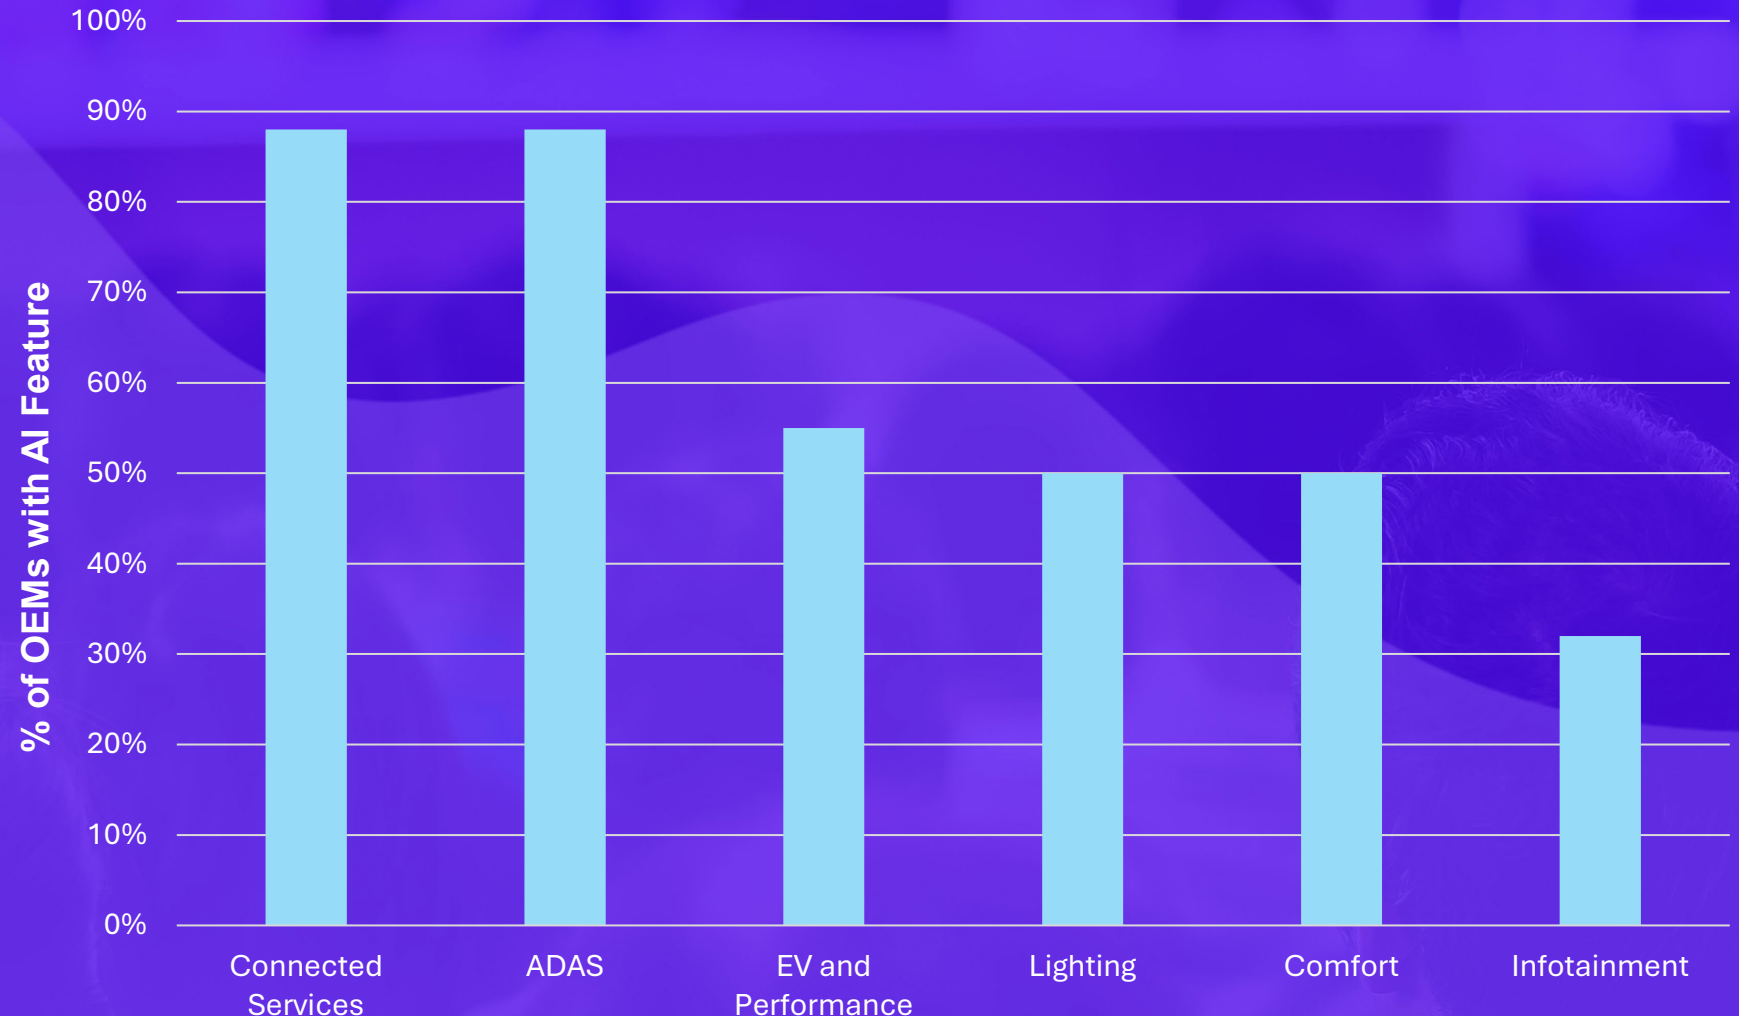
AI Model Name	Supplier	OEMs	Hosting Environment
Mistral	Mistral AI	STELLANTIS	Cloud
Doubao	ByteDance		Cloud
ERNIE 4.5	Baidu		Cloud
Flyme AIOS	MEIZU		Cloud
Gemini AI	Google		Cloud
Grok-4	xAI		Cloud
Grok-3			
GPT-4	OpenAI		Cloud
GPT-4o			
Qwen2.5	Alibaba		Cloud
Qwen1.5			Hybrid
Qwen2			Hybrid
SenseNova 5.0	SenseTime		Cloud
SenseNova 5.5			
Xinghuo 3.5 (Spark)	iFlyTek		Cloud
Xinghuo 4.0 (Spark)			
DeepSeek-R1	DeepSeek AI		Hybrid
StepAI	StepFun		Hybrid

...but making AI pay for itself is still a challenge.

- Most OEMs have some form of AI in their vehicles
- BMW tops our charts with 22 separate AI features
- CapEx is just the tip of the iceberg. OpEx for AI features can be very high

How should automakers turn this AI expense into a profitable service?

Global Automotive AI Feature Adoption



Meet today's experts!



Stefano Marzani

Head of Emerging Technologies



Dani Cherkassky

Chief Executive Officer



Stas Matviyenko

VP of Voice AI Monetization



Andy Qiu

Senior Manager



Stefano Marzani



Head of Emerging Technologies





The Automotive Flywheel Runs on Tokens

Stefano Marzani

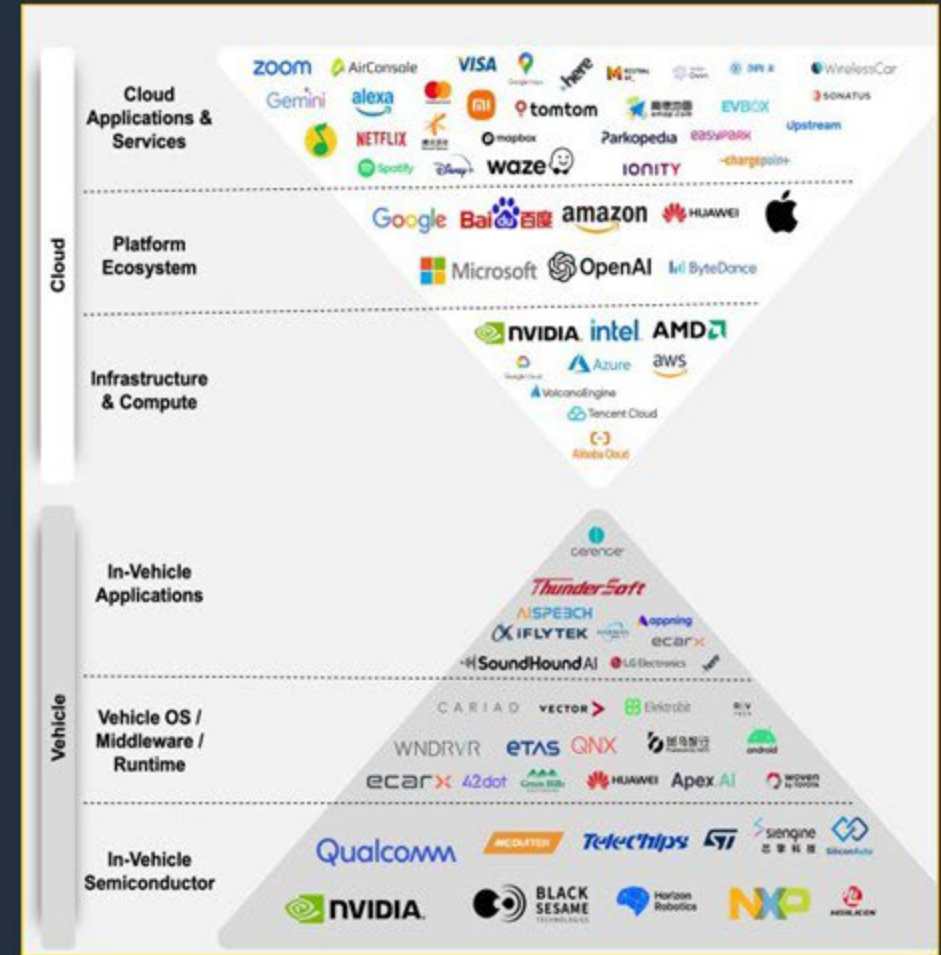
WW Head of Emerging Tech, Auto & Manufacturing
Amazon Web Services, Inc.



The Squeeze is Real, but it Has a Hidden Dimension

OEMs ARE BEING SQUEEZED — AND THE SUPPLY CHAIN MAKES IT WORSE

- Cloud platforms & apps: \$71B → \$161B by 2035, Semiconductors & E/E: \$103B combined
- OEMs need to own the layer that talks to the driver, and that layer is becoming an *agentic* layer
- Missing dimension: **integration** difficulty
 - **Surface content** [high-level APIs, can ship today]: music, video, gaming, podcasts, etc.
 - **Deep content** [ECU data, OEM data silos, systems outside the car]: very difficult to monetize through the supply chain [1]
- OEMs can't afford legacy architectures: SDV transition is not completed and the product still needs to be simplified and opened: EE consolidation, virtualization, platform engineering
- The business case breaks not at the idea layer, it breaks at the integration layer




"The Squeeze", Augustin Friedel [LinkedIn Post](#)

Agents Are the Revenue Engine – Holistically

AGENTIC AI IS THE ABSTRACTION LAYER IOT PROMISED AND NEVER DELIVERED

- IoT promised: connect your car, home, devices, digital life — seamlessly. It never happened.
- Agents are the ultimate abstraction layer, acting as connective tissue:
 - **cloud ↔ vehicle edge ↔ user context ↔ other environments**
- Tokens are the units of that new *agentic intelligence* layer, holistically speaking:
 - tokens consumed to **generate revenues from end-users**: for example, hyper-personalized experiences
 - *"hey, I noticed you seem a bit congested, want me to raise the temperature up a notch at home?"*
 - tokens used to **produce engineering efficiency**
 - *Low hanging fruit: Engineering and Software Development Lifecycle agentic automation; then code gen.*



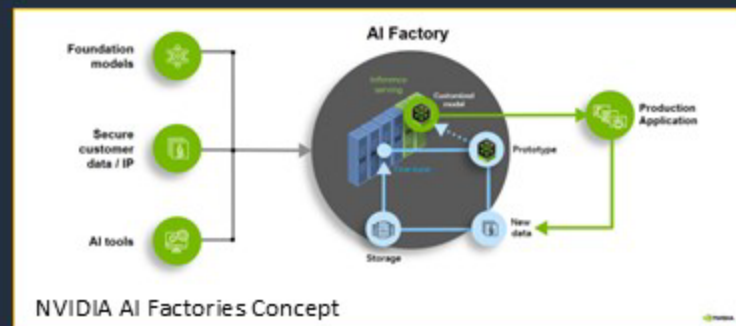
A New Industrial Revolution, Nvidia GTC 2026



It's a Flywheel. But the Game Just Got Harder.

TOKEN ROI AND EFFICIENCY IS THE NEW HORSEPOWER — AND THE ATTACK SURFACE JUST EXPANDED

- **The Flywheel:** Tokens consumed → costs reduced (engineering, ops) + revenue generated (experiences, subscriptions) → more data → better models → better experiences → more engagement → more data
- Who owns the **agentic control plane** owns the cockpit in 2035
- **Early signals:**
 - CES Demo (Alexa Custom Assistant with Alexa+ and Strands agents with vehicle integration), BMW iX3 Alexa Custom Assistant with Alexa+
 - Customers reducing feature development from weeks to days, thanks to KIRO
- It's still early days, the window is open. The imperative:
experiment, experiment, experiment
- But this new game comes with new risks.
Autonomous agents = autonomous attack surfaces. This isn't a feature choice — it's a *structural architecture choice*. Mindset shift: from "AI as a tool" to "AI as an actor" with trust, governance, and guardrails





Thank you!

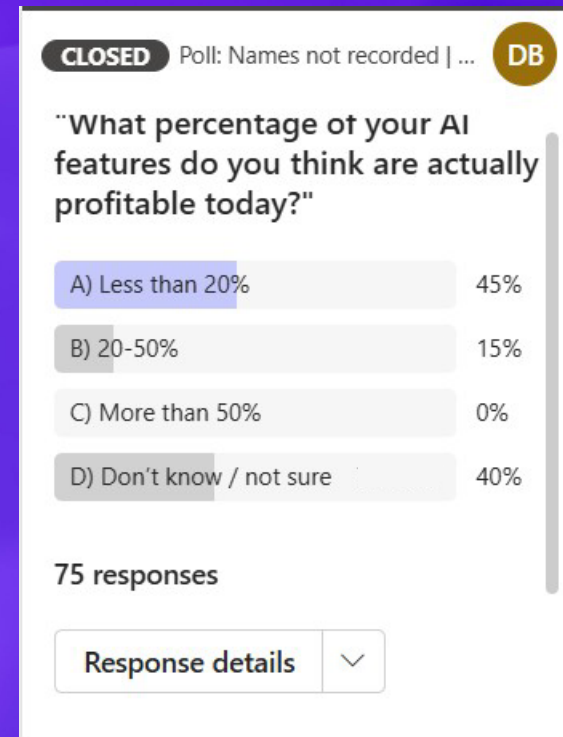
Stefano Marzani

marzanis@amazon.com

Audience poll question

"What percentage of your AI features do you think are actually profitable today?"

- A) Less than 20%
- B) 20-50%
- C) More than 50%
- D) Don't know / not sure



Dani Cherkassky



Chief Executive Officer





Crossing the Engagement Threshold

The Profitability Path of Automotive AI

Dani Cherkassky,
CEO & Co-Founder, Kardome



Why current voice AI fails the engagement test



Too expensive to stay "Awake"



Too slow for natural conversation



Wrong tools for trivial tasks

The architectural shift

System 1:
The Reflex Layer
(Edge SLM)

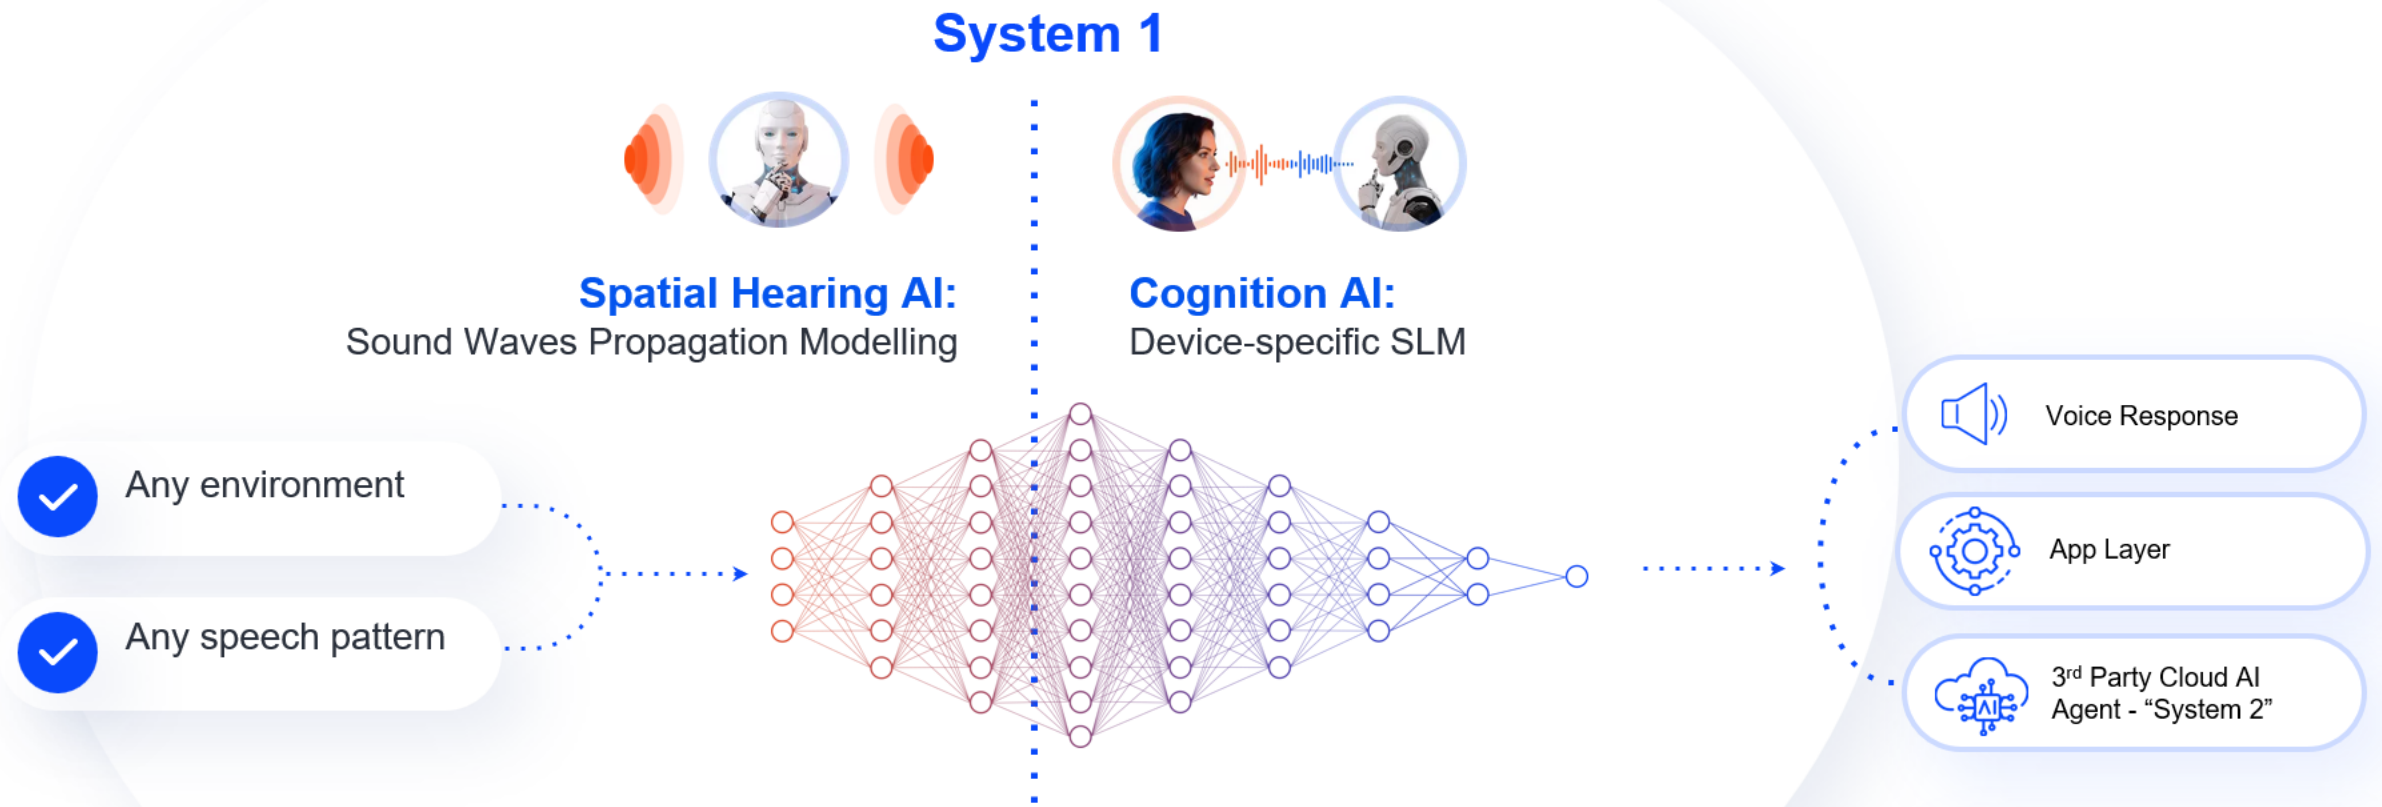
Fast, intuitive, private, and
always-on.



System 2:
The Reasoning Layer
(Cloud LLM)

Slow, deliberate, and energy-
intensive.

Voice AI architecture zoom in



The value equation: Sustainable profit is only achieved when the assistant moves from a "heavyweight" cloud inconvenience to a seamless, on-edge reflex.

Key Takeaway

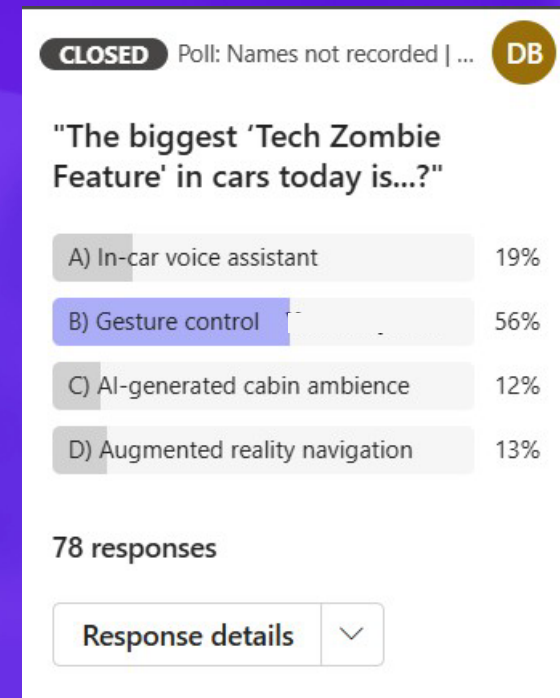
Sustainable profit starts at the edge.

Cloud reasoning extends capability,
but most engagement must happen on-device

Audience poll question

"The biggest 'Tech Zombie Feature' in cars today is...?"

- A) In-car voice assistant
- B) Gesture control
- C) AI-generated cabin ambience
- D) Augmented reality navigation



Stas Matviyenko



VP of Voice AI Monetization

SoundHound AI

SoundHound AI

MONETIZING THE CONNECTED VEHICLE

FROM VOICE ASSISTANTS TO VOICE COMMERCE

Stas Matviyenko
Vice President, Voice AI Monetization
SoundHound AI



VOICE ASSISTANTS DON'T MAKE MONEY.
COMMERCE AGENTS DO.

SoundHound AI

CATEGORIES WE SHOW UP IN:



FOOD ORDERING



TABLE RESERVATIONS



PARKING



EVENTS & TICKETS



TRAVEL BOOKING
AND BEYOND...

The industry invested billions in in-car AI. Most of it generates zero revenue.

The shift: voice assistants answer questions → agentic AI that completes transactions autonomously

Feature cost center → revenue-generating platform through agentic commerce

SoundHound's agentic commerce stack:

Voice UX → Intent → Agent Execution → Payments → Merchant/Partner Network → Fulfillment – all embedded in the vehicle

THREE REVENUE LAYERS THAT WORK TODAY.

SoundHound AI



PER-TRANSACTION COMMERCE & CONVENIENCE FEES

Every transaction an agent completes — food, coffee, tickets, parking — generates revenue. Whether it's a commission, a convenience fee, or both — standard in delivery and mobile ordering — for the added value of hands-free, agentic transactions. No subscriptions. No ads. Revenue tied to real consumer spending.



PARTNER-FUNDED DISCOVERY

Agents surface brands when intent is highest. Example: a driver says "Order coffee on my way" without naming a brand → the agent presents options → partner brands pay a higher commission for that incremental order. The driver gets relevant choices. The brand gets a sale they wouldn't have had otherwise. Everyone wins.



ADS & AI-POWERED UPSELL

Agentic AI makes intelligent, contextual suggestions at the right moment. At checkout, the agent recommends relevant add-ons that match user intent (e.g., "Add a pastry with your coffee?" or "Upgrade to a meal?"). Deals from partner brands surface based on location, time, and behavior — not generic banners. High conversion because the agent understands context.



FROM PILOT TO PROFIT.

SoundHound AI

THE OEM PLAYBOOK

TRY – OEMs can experience the full agentic commerce flow today by installing our app. No integration required. See agents work before committing.

INTEGRATE – SDK + API into existing infotainment. Agents work alongside the native assistant.

PILOT – Employee → controlled public → GA. Measure order volume, completion rate, GMV.

SCALE – Expand agent categories, activate partner-funded discovery, grow revenue share.

TRY

INTEGRATE

PILOT

SCALE

THE OEM PLAYBOOK

TRY – OEMs can experience the full agentic commerce flow today by installing our app. No integration required. See agents work before committing.

INTEGRATE – SDK + API into existing infotainment. Agents work alongside the native assistant.

PILOT – Employee → controlled public → GA. Measure order volume, completion rate, GMV.

SCALE – Expand agent categories, activate partner-funded discovery, grow revenue share.

TRY

INTEGRATE

PILOT

SCALE

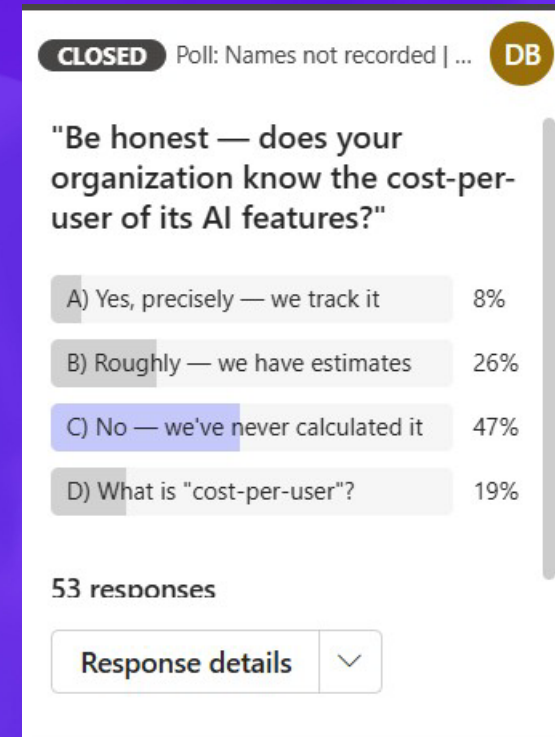
SoundHound AI

THANK YOU

Audience poll question

"Be honest — does your organization know the cost-per-user of its AI features?"

- A) Yes, precisely — we track it
- B) Roughly — we have estimates
- C) No — we've never calculated it
- D) What is "cost-per-user"?



Andy Qiu



Senior Manager



From CapEx to OpEx: The New Battleground



The AI monetization challenge is a **P&L problem**, not a technology problem—shifting automotive software from a predictable one-time CapEx to an unpredictable, perpetual OpEx.

Old Model – Software 1.0

- One-time Development Cost (CapEx)
- Embedded Software, Fixed Margins
- Cost predictable at Start of Production (SOP)
- Development = CapEx

THE TRAP

New Model – Generative AI

- Perpetual Cloud Inference Costs (OpEx)
- Every user interaction "burns money"
- Margins erode as usage scales
- Usage = OpEx — unpredictable & growing

1 The Profitability Gap

OEMs compete with consumer electronics playbooks – but scale amplifies losses, not revenue.

2 The Zombie Feature Risk

Without usage data, cabins fill with high-cost, low engagement AI – hidden liabilities on the P&L.

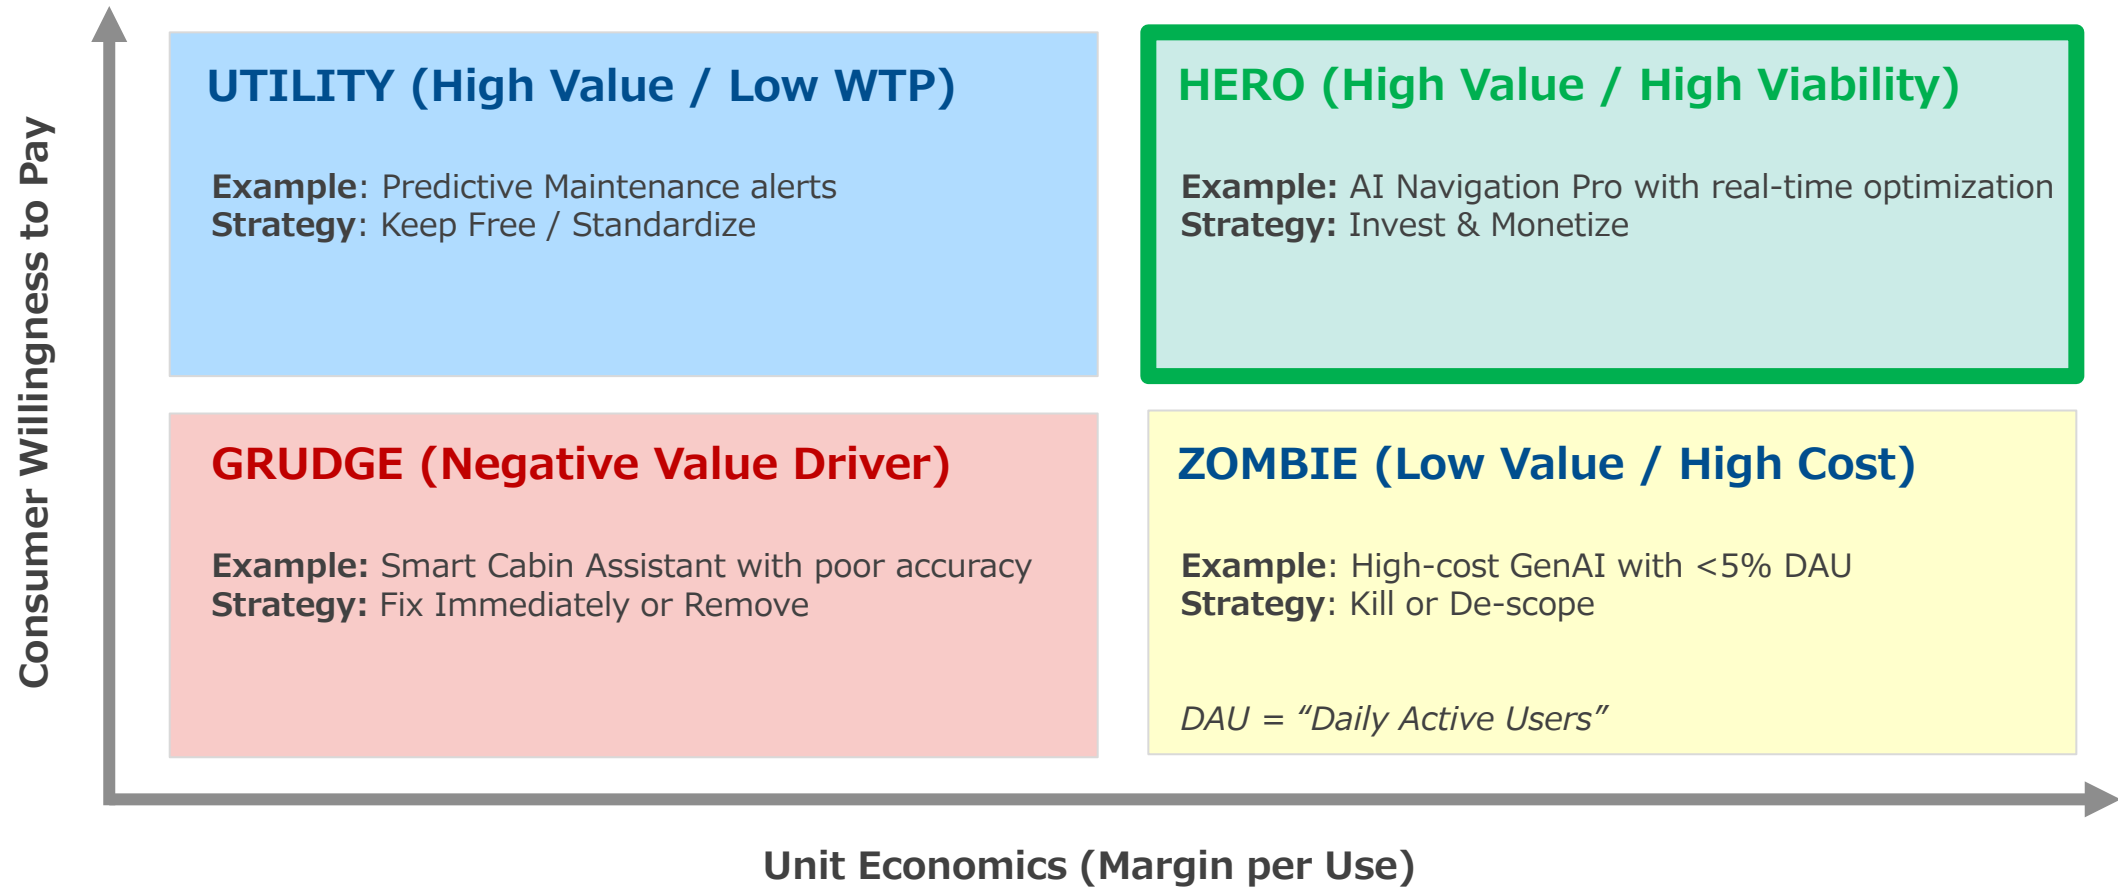
3 The 2026 Multiplier

Agentic AI shift from passive response to proactive execution – inference costs rise 5-10x.

Hero vs. Zombie: Classifying AI Features



To prevent margin erosion, OEMs must ruthlessly categorize their **AI portfolio**—shifting investments away from hidden "Zombie" liabilities to fund high-value "Hero" features.

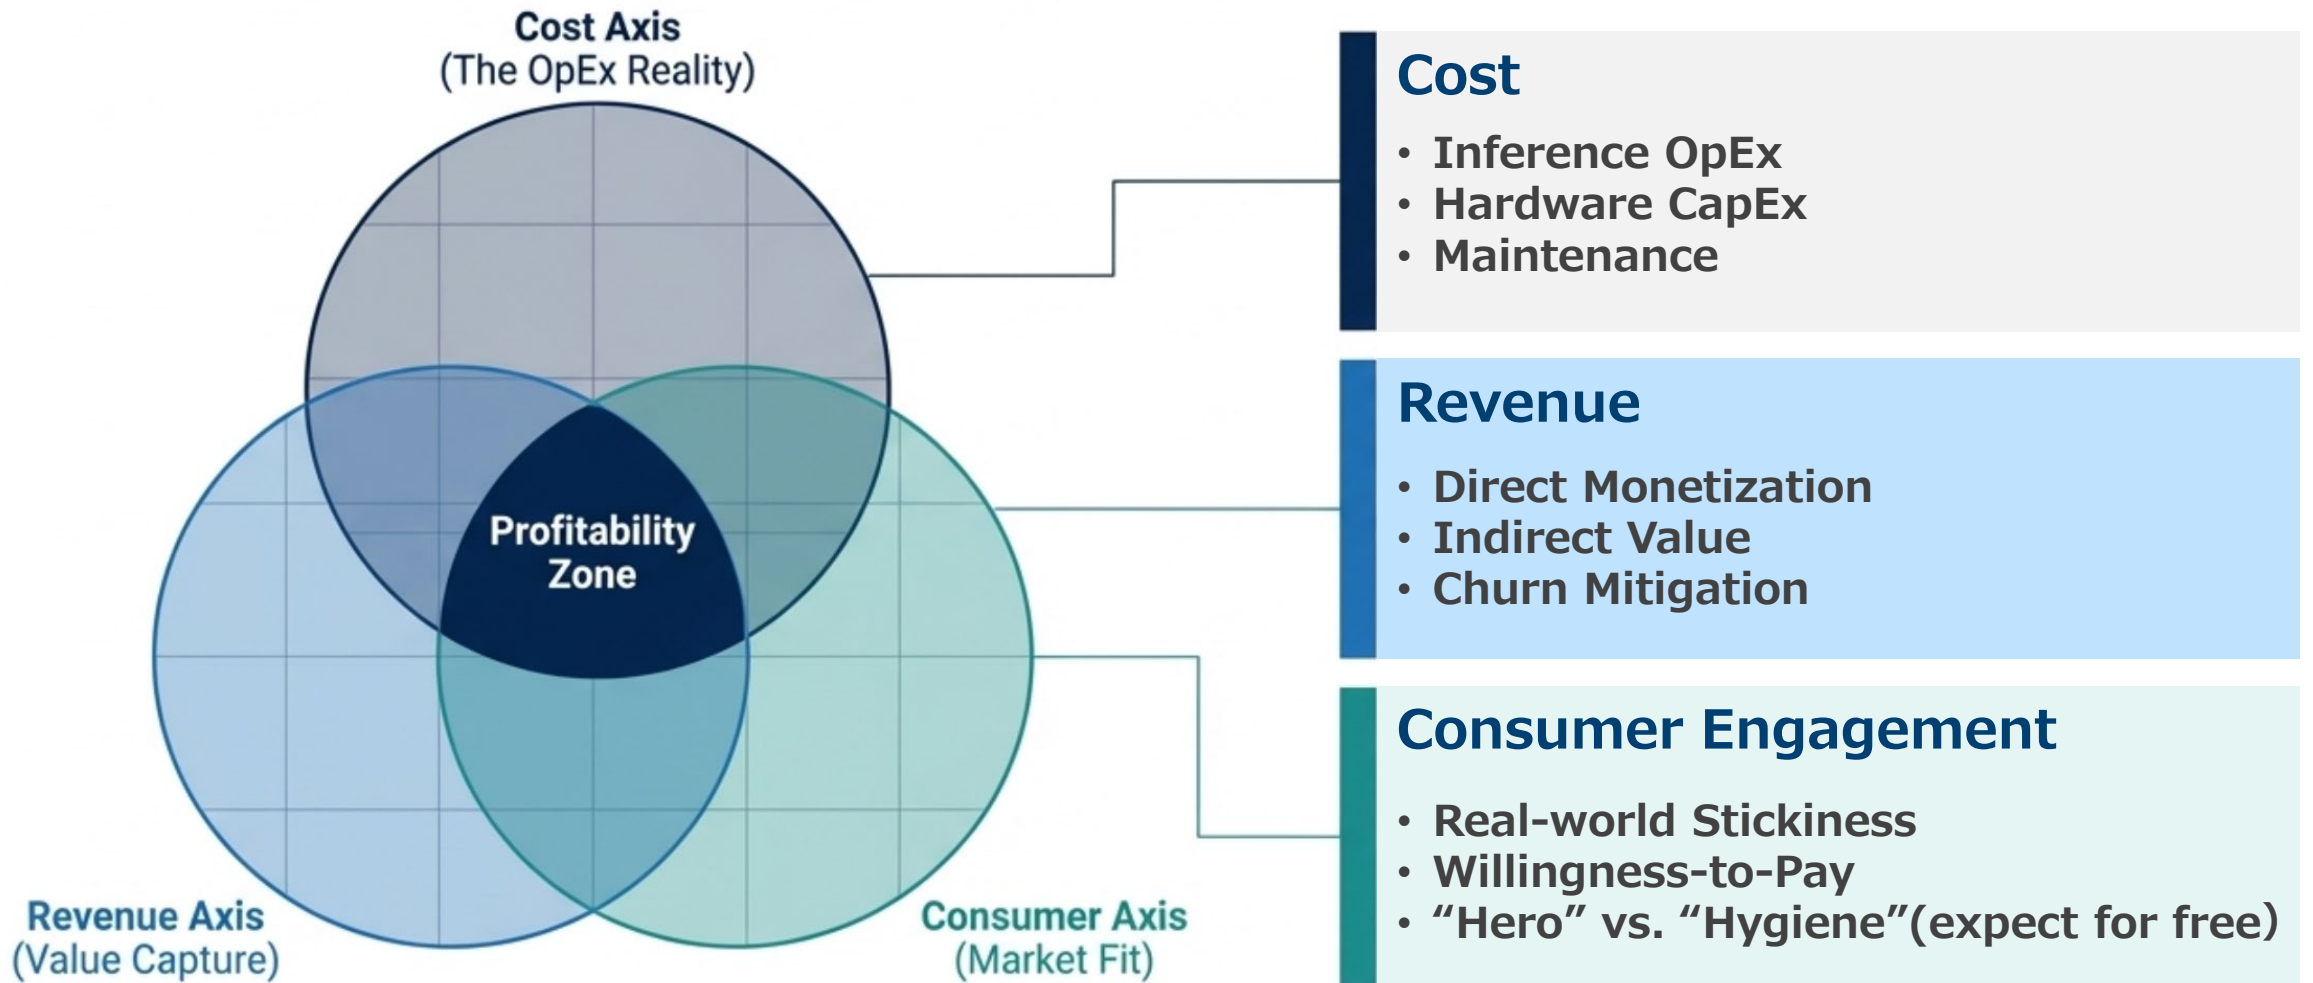


Strategic Goal: Shift resources from Zombie & Grudge quadrants to fund Hero features.

Feature Lifecycle Margin: The Triangulation Model



Moving from “**Vehicle Margin**” to “**Feature Lifecycle Margin**” requires triangulating true cost-per-query, revenue streams, and consumer engagement.



Ask Yourself: Do you know the true **cost-per-user** of every AI feature in your vehicle today? If not, that is exactly where your margin risk lives.

Speaker roundtable



Stefano Marzani

Head of Emerging Technologies



Dani Cherkassky

Chief Executive Officer



Stas Matviyenko

VP of Voice AI Monetization



Andy Qiu

Senior Manager



Robert Fisher

Moderator





Audience Q&A

Please submit *your* questions now!

Our parting thoughts...

Takeaway #1



AI is no longer a one-time investment

Takeaway #2



The biggest opportunity isn't building more AI features

Takeaway #3



Revenue won't come from subscriptions alone