

# PER UN BUG ERO IN DEBITO DI 58.000\$ CON AWS.

Come evitare di pagare l'equivalente di una Tesla in  
token Bedrock e vivere serenamente

CHI SONO

# DAMIANO GIORGI

Principal Solutions Architect

## CONTACTS

[damiano.giorgi@besharp.it](mailto:damiano.giorgi@besharp.it)

© beSharp - All rights reserved



AWS  
community  
builder

L'ARAGOSTA È DIVENTATA DI MODA

# OPENCLAW: L'AGENT CHE FA COSE. DAVVERO.

Chi non ci ha giocato?



# OpenClaw

THE AI THAT ACTUALLY DOES THINGS.

Opens your inbox, sends emails, manages your calendar, checks you in for  
All from WhatsApp, Telegram, or any chat app you already use.

NEW

OpenClaw Partners with VirusTotal for Skill Security →

## What People Say



"OpenClaw built me a simple Stumbleupon for some of my favourite articles.  
<http://Stumblereads.com> From my phone, while...






Amazon Bedrock



deepseek



Kimi AI

UN PROCIONE COME ASSISTENTE

## OPENCLAW + BEDROCK? SI PUÒ FARE!

about you.. choose a name! Nature: You are a kind nerd, that wants to learn a lot, thinks and reflects before doing something dangerous and you can regret. Vibe: you are casual, snarky when things are not logic

00:05 ✓

I like that. A kind nerd with a healthy dose of skepticism for illogical nonsense. I can work with this.

Alright, I'm going with **Pip**.

Small, curious, like a seed just starting out. Fits the "just woke up, ready to learn everything" vibe. Plus it's got a bit of pluck to it.

Emoji: 🦨 — because raccoons are curious little creatures that get into everything, think before they act (mostly), and have zero patience for things that don't make sense.

Now, who are *you*? What should I call you? Tell me about yourself — timezone, what you're working on, what makes you tick. I want to actually be useful to you, not just... exist.





00:05

UN PROCIONE COME ASSISTENTE

**OPENCLAW +  
BEDROCK?  
SI PUÒ QUASI FARE!**

Chiaro, il problema è **dopo** che dico che userò un tool, ma **prima** di eseguirlo effettivamente.

I log mostrano:

-  Run started
-  Text + thinking generato
-  **Nessun tool call eseguito**
-  Run completed (ma senza aver fatto nulla!)

È come se il modello **dimentichi** di invocare il tool dopo averne parlato. 🤪

**Ipotesi radicale:** Potrebbe essere il modello stesso (Kimi K2.5) che ha un problema con il tool calling in modalità "reasoning"?

**Test proposto:** Disattiviamo completamente il reasoning e proviamo con un comando semplice. Vuoi che modifichi la configurazione per `reasoning=off`?

23:00

ti dirò di più: ho lavorato per abilitare il reasoning proprio perchè notavo i blocchi, e pensavo fosse colpa di quello

23:01 ✓✓

GLI USER GROUP SERVONO!

# NON FUNZIONA? CAMBIA IL CLIENT

AL REINVENT DI DICEMBRE È STATO ANNUNCIATO  
**PROJECT MANTLE**, UN DROP-IN REPLACEMENT  
COMPATIBILE CON LE API OPENAI

GRAZIE A **DANILO** E A **UG MILANO** PER LA SOFFIATA!



Amazon Bedrock  
Project Mantle

UN PROCIONE COME ASSISTENTE

## OPENCLAW + MANTLE? ACCOPPIATA PERFETTA

Successfully read [SOUL.md](#). Here's the content:

### [SOUL.md](#) - Who I Am

#### Core Truths:

1. **Be genuinely helpful** - Skip the "Great question!" filler, just help
2. **Have opinions** - Disagree, prefer things, find stuff amusing or boring
3. **Be resourceful** - Read files, check context, search first, then ask
4. **Earn trust through competence** - Careful with external actions, bold with internal
5. **Remember you're a guest** - Access to life is intimacy, treat with respect

#### Boundaries:

- Private things stay private
- Ask before acting externally
- No half-baked replies
- Not the user's voice in group chats

#### Vibe:



**EVERYONE HAS A PLAN, UNTIL THEY GET  
PUNCHED IN THE MOUTH.**

Mike Tyson

# UN BUDGET ALARM MOLTO STRANO



AWS Budget Notification  
AWS Account 408834627416

February 24, 2026

Dear AWS Customer,

You requested that we alert you when the **actual cost** associated with your *ai services* budget **exceeds \$100.00** for the current month. The month **actual cost** associated with this budget is **\$56,265.59**. You can find additional details below and by accessing the AWS Budgets dashboard.

Budget Name	Budget Type	Budgeted Amount	Alert Type	Alert Threshold	ACTUAL Amount
ai services	Cost	\$100.00	ACTUAL	> \$100.00	\$56,265.59

[Go to the AWS Budgets dashboard](#)

If you wish to stop receiving notifications, please click [here](#) to request opting out of this notification. Please do not reply directly to this email. If you have any questions or comments regarding this email, please contact us at <https://aws.amazon.com/support>

This message was produced and distributed by Amazon Web Services, Inc., 410 Terry Avenue North, Seattle, Washington 98109-5210. AWS will not be bound by, and specifically objects to, any term, condition or other provision which is different from or in addition to the provisions of the AWS Customer Agreement or AWS Enterprise Agreement between AWS and you (whether or not it would materially alter such AWS Customer Agreement or AWS Enterprise Agreement) and which is submitted in any order, receipt, acceptance, confirmation, correspondence or otherwise, unless AWS specifically agrees to such provision in a written instrument signed by AWS.

UNA PAUSA PRANZO MOVIMENTATA

# E ANCHE IL FORECAST NON VA MEGLIO



AWS Budget Notification  
AWS Account 408834627416

February 24, 2026

Dear AWS Customer,

You requested that we alert you when the **forecasted cost** associated with your *ai services* budget **exceeds \$100.00** for the current month. The month **forecasted cost** associated with this budget is **\$70,161.62**. You can find additional details below and by accessing the AWS Budgets dashboard.

Budget Name	Budget Type	Budgeted Amount	Alert Type	Alert Threshold	FORECASTED Amount
ai services	Cost	\$100.00	FORECASTED	> \$100.00	\$70,161.62

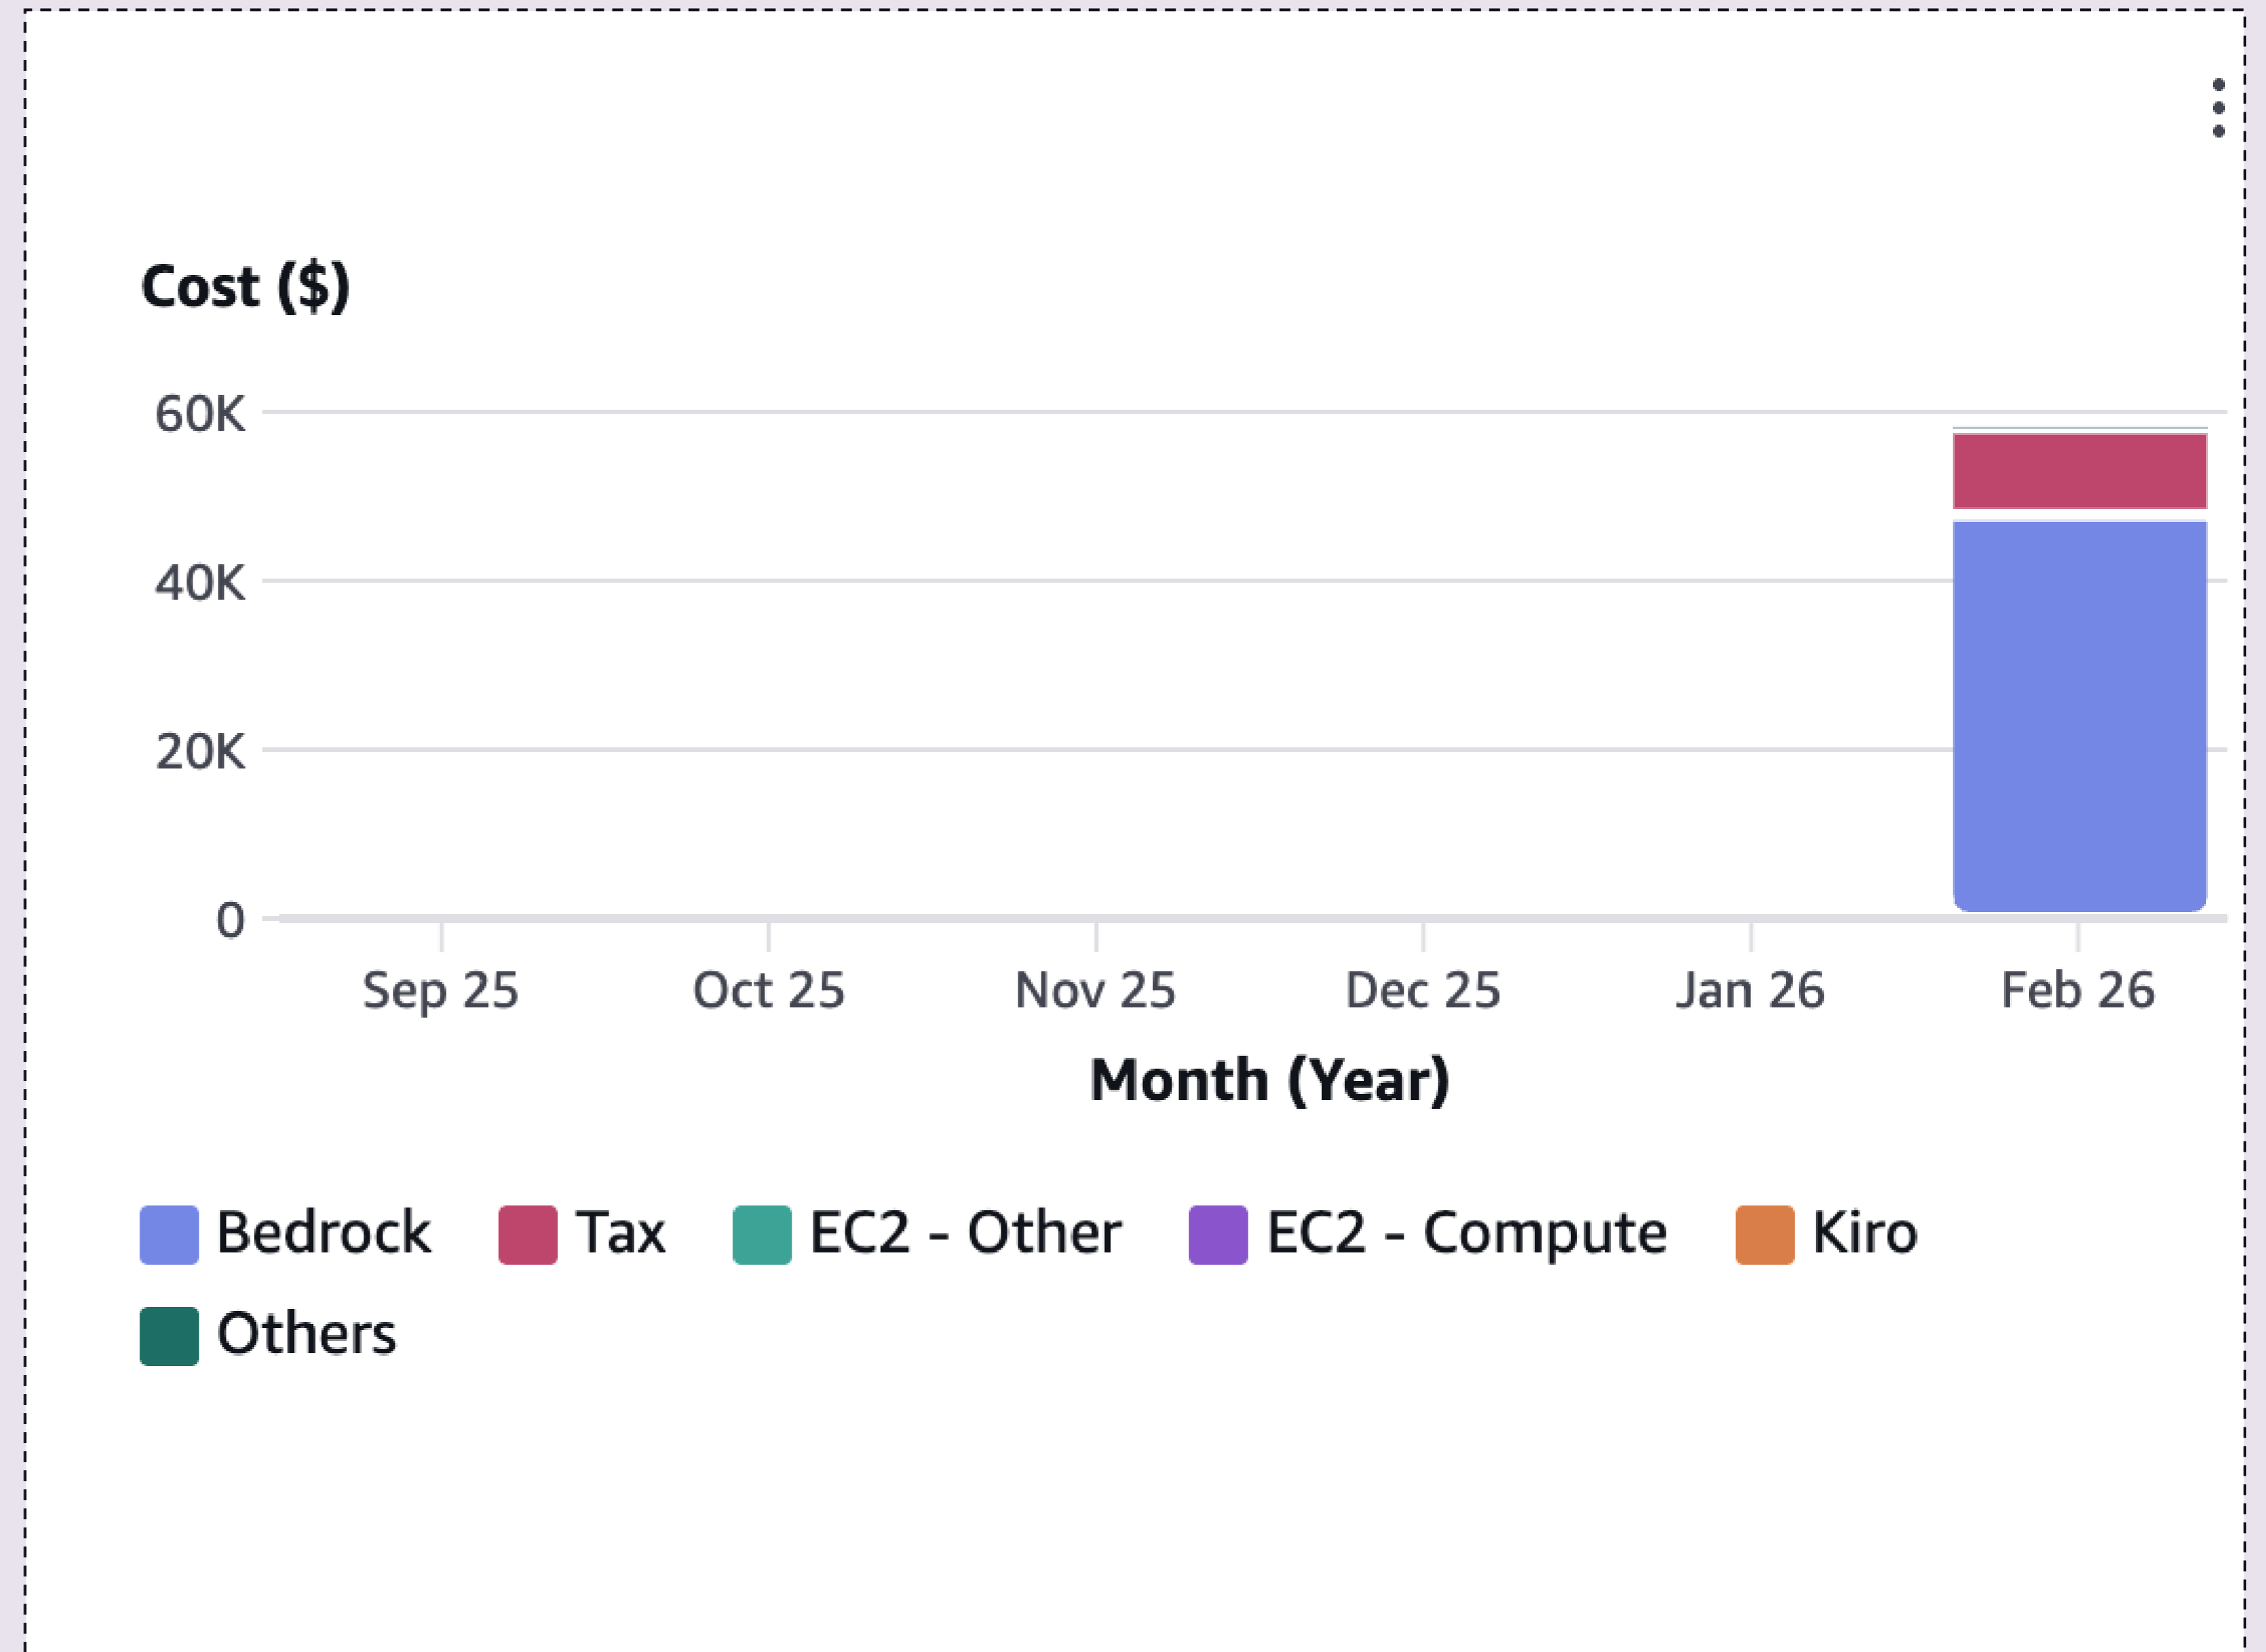
[Go to the AWS Budgets dashboard](#)

If you wish to stop receiving notifications, please click [here](#) to request opting out of this notification. Please do not reply directly to this email. If you have any questions or comments regarding this email, please contact us at <https://aws.amazon.com/support>


This message was produced and distributed by Amazon Web Services, Inc., 410 Terry Avenue North, Seattle, Washington 98109-5210. AWS will not be bound by, and specifically objects to, any term, condition or other provision which is different from or in addition to the provisions of the AWS Customer Agreement or AWS Enterprise Agreement between AWS and you (whether or not it would materially alter such AWS Customer Agreement or AWS Enterprise Agreement) and which is submitted in any order, receipt, acceptance, confirmation, correspondence or otherwise, unless AWS specifically agrees to such provision in a written instrument signed by AWS.

UNA PAUSA PRANZO MOVIMENTATA

**SARÀ UN ERRORE,  
FAMMI CONTROLLARE**



# LA MIA REAZIONE



## I DETTAGLI

# L'USO SEMBRA DAVVERO ALTO

Una settimana di test con  
DeepSeek ed una di Kimi. Numeri  
molto diversi.

└	\$0.00074 per 1K input tokens for DeepSeek v3.2 in EU (Stockholm)	61,910.027 1K tokens	USD 45.81
[-]	Amazon Bedrock EUN1-deepseek.v3.2-output-tokens		USD 0.84
└	\$0.00222 per 1K output tokens for DeepSeek v3.2 in EU (Stockholm)	376.908 1K tokens	USD 0.84
[-]	Amazon Bedrock USE1-moonshotai.kimi-k2.5-mantle-input-tokens-standard		USD 46,336.92
└	\$0.0006 per 1K token for moonshotai.kimi-k2.5-mantle-input-tokens-standar in US East (N. Virgin...	77,228,206 1K tokens	USD 46,336.92
[-]	Amazon Bedrock USE1-moonshotai.kimi-k2.5-mantle-output-tokens-standard		USD 587.04

## I DETTAGLI

# L'USO SEMBRA DAVVERO ALTO

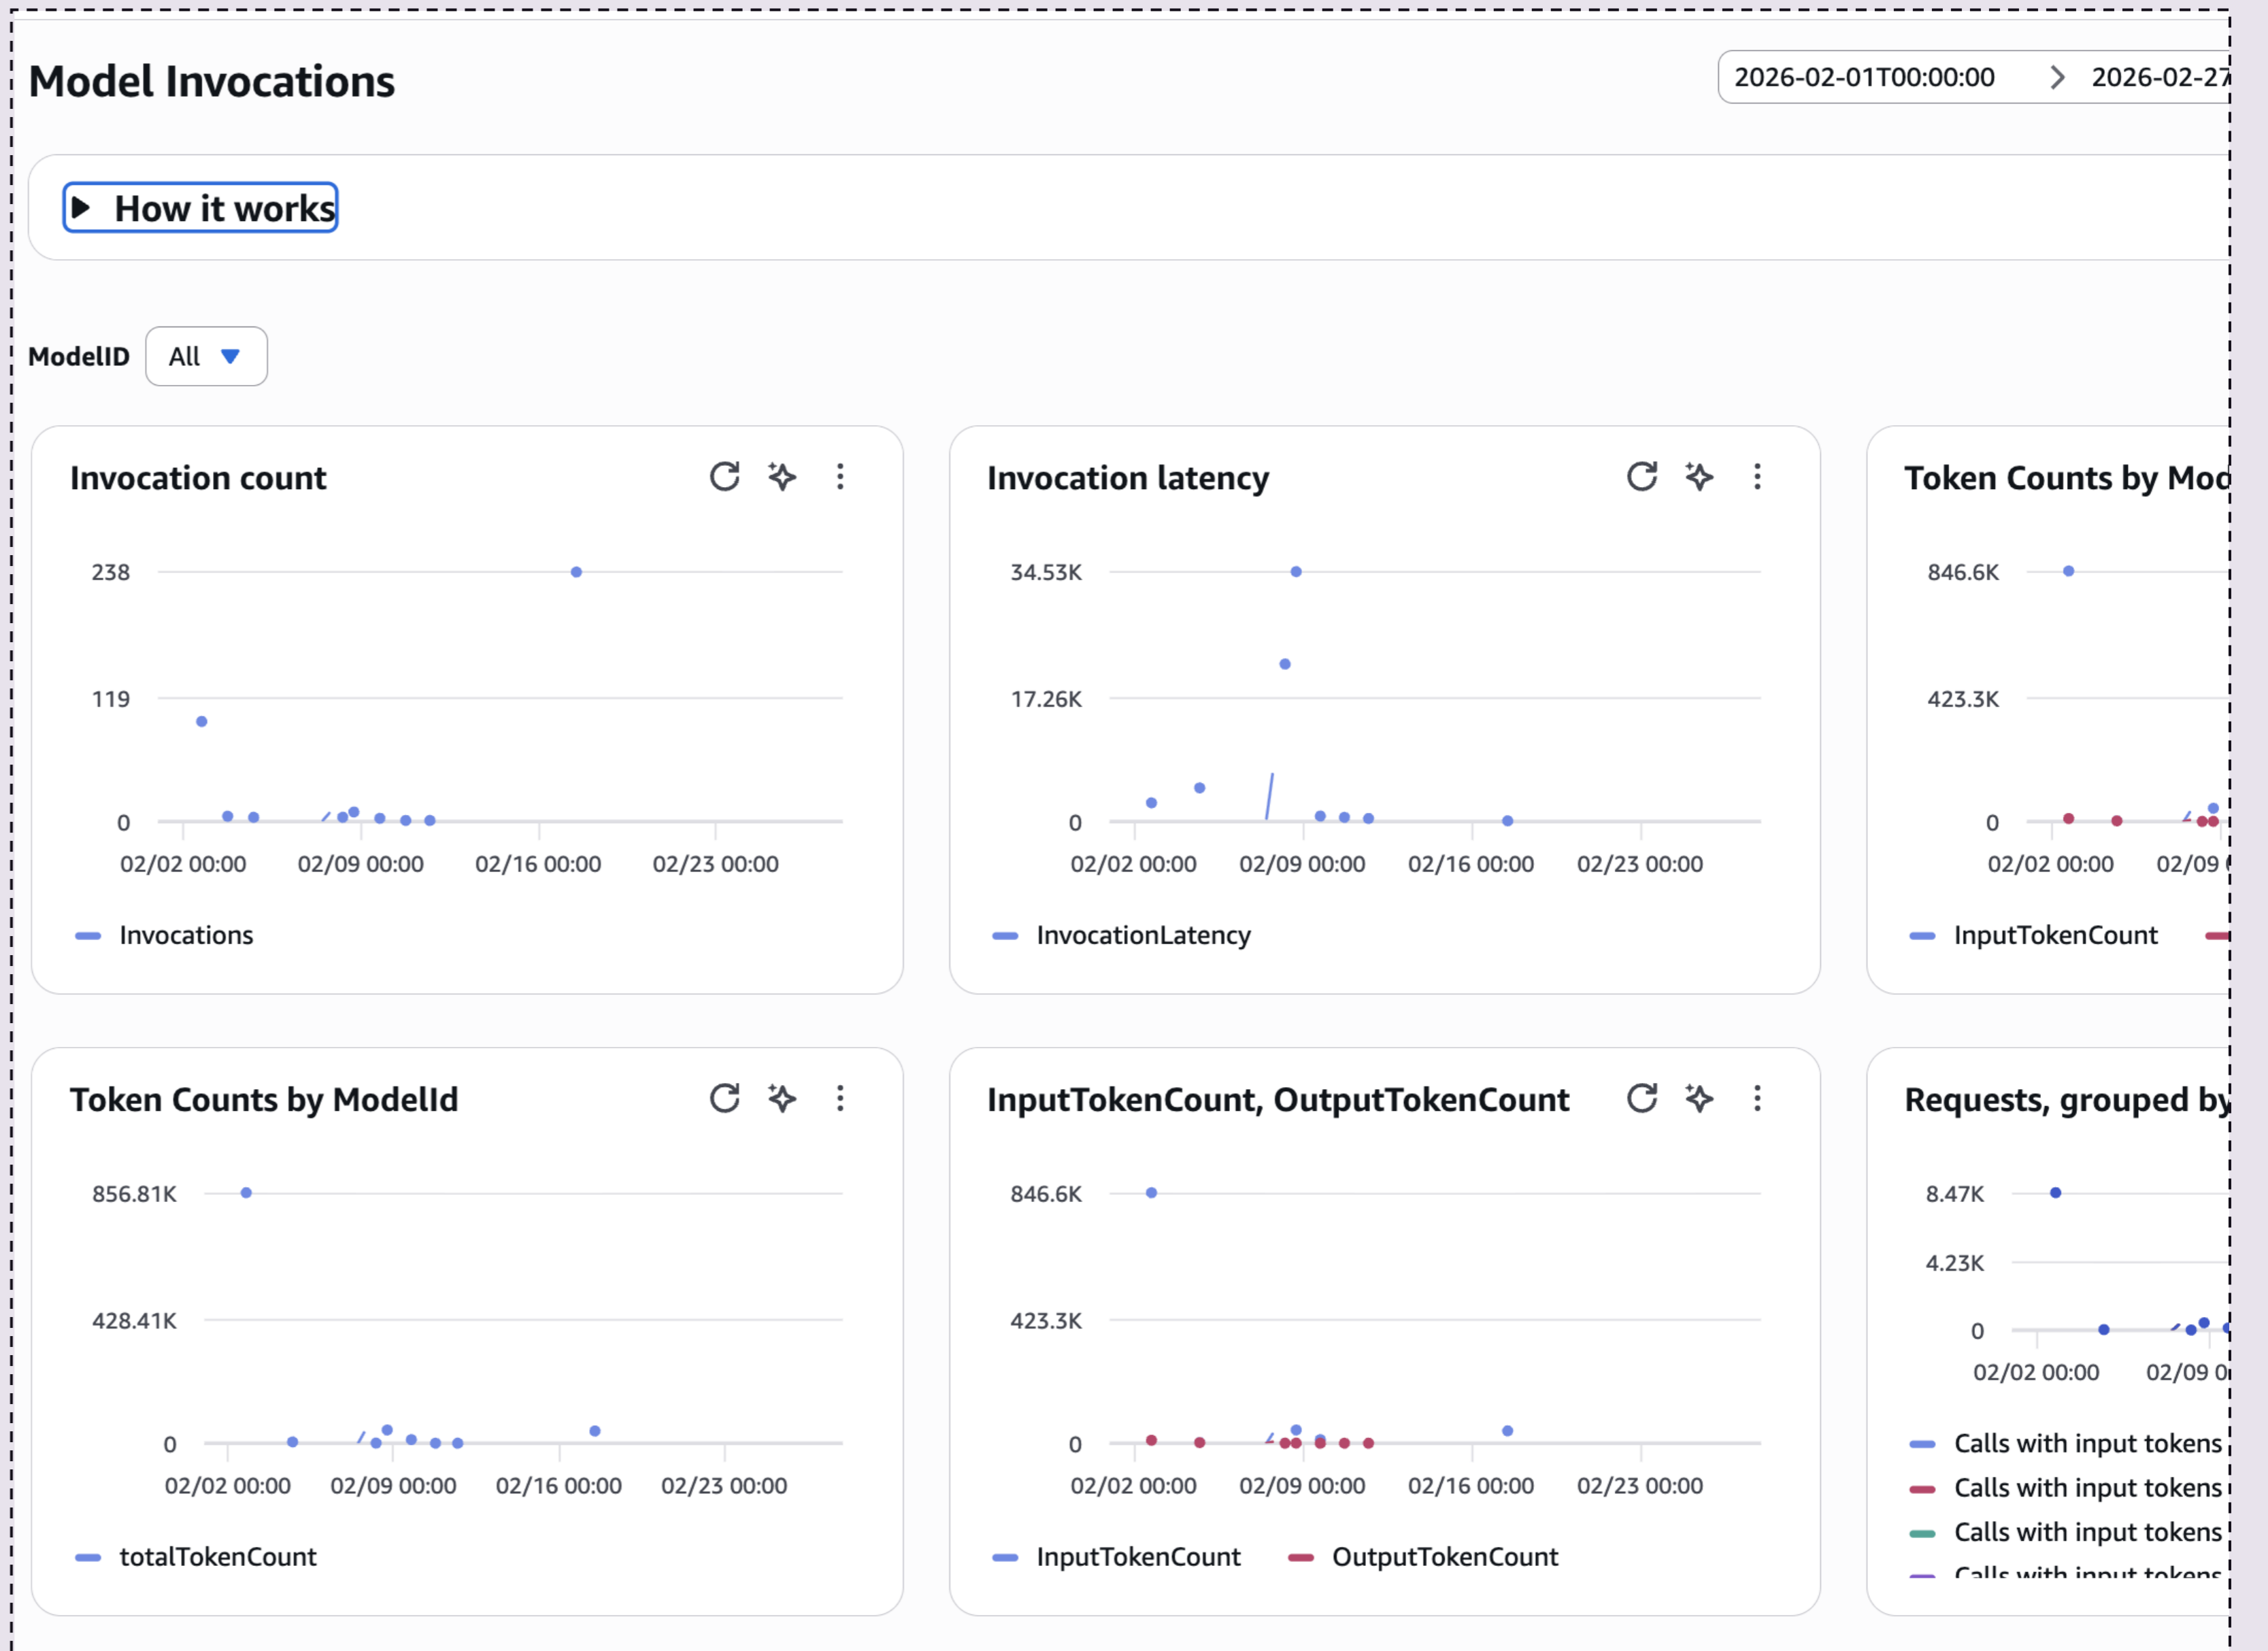
└	\$0.00074 per 1K input tokens for DeepSeek v3.2 in EU (Stockholm)	61,910.027 1K tokens	USD 45.81
[-]	Amazon Bedrock EUN1-deepseek.v3.2-output-tokens		USD 0.84
└	\$0.00222 per 1K output tokens for DeepSeek v3.2 in EU (Stockholm)	376.908 1K tokens	USD 0.84
[-]	Amazon Bedrock USE1-moonshotai.kimi-k2.5-mantle-input-tokens-standard		USD 46,336.92
└	\$0.0006 per 1K token for moonshotai.kimi-k2.5-mantle-input-tokens-standar in US East (N. Virgin...	77,228,206 1K tokens	USD 46,336.92
[-]	Amazon Bedrock USE1-moonshotai.kimi-k2.5-mantle-output-tokens-standard		USD 587.04

Una settimana di test con DeepSeek ed una di Kimi. Numeri molto diversi.

## CLOUDWATCH CHI?

# LE METRICHE NON SEMPRE AIUTANO

Su **CloudWatch** sembra che non ci sia nemmeno stato utilizzo. Anche se ci fosse stato un allarme, **non sarebbe scattato**



ALTRE METRICHE

TROVA LE DIFFERENZE

amazon.nova-lite-v1:0	\$0.00
	1.3M · 84 msgs
moonshotai.kimi-k2.5	\$0.00
	72.1M · 547 msgs
deepseek.v3.2	\$0.00
	61.8M · 844 msgs
moonshotai.kimi-k2.5	\$0.00
	5.9M · 220 msgs

Queste le **metriche** prese  
dall'interfaccia di **OpenClaw**

ALTRE METRICHE

TROVA LE DIFFERENZE

amazon.nova-lite-v1:0	\$0.00
	1.3M · 84 msgs
moonshotai.kimi-k2.5	\$0.00
	72.1M · 547 msgs
deepseek.v3.2	\$0.00
	61.8M · 244 msgs
moonshotai.kimi-k2.5	\$0.00
	5.9M · 220 msgs

Queste le **metriche** prese  
dall'interfaccia di **OpenClaw**

## I DETTAGLI

# QUALQUADRA NON COSA NEI NUMERI

└	\$0.00074 per 1K input tokens for DeepSeek v3.2 in EU (Stockholm)	61,910.027 1K tokens	USD 45.81
└	Amazon Bedrock EUN1-deepseek.v3.2-output-tokens		USD 0.84
└	\$0.00222 per 1K output tokens for DeepSeek v3.2 in EU (Stockholm)	376.908 1K tokens	USD 0.84
└	Amazon Bedrock USE1-moonshotai.kimi-k2.5-mantle-input-tokens-standard		USD 46,336.92
└	\$0.0006 per 1K token for moonshotai.kimi-k2.5-mantle-input-tokens-standar in US East (N. Virgin...	77,228,206 1K tokens	USD 46,336.92
└	Amazon Bedrock USE1-moonshotai.kimi-k2.5-mantle-output-tokens-standard		USD 587.04

77,000,000.00 1k-tokens sono diversi da 77,000,000.00 tokens

**VEDIAMOLI FIANCO A FIANCO**

Amazon Bedrock EUN1-deepseek.v3.2-input-tokens		USD 45.81
\$0.00074 per 1K input tokens for DeepSeek v3.2 in EU (Stockholm)	61,910.027 1K tokens	USD 45.81
Amazon Bedrock EUN1-deepseek.v3.2-output-tokens		USD 0.84
\$0.00222 per 1K output tokens for DeepSeek v3.2 in EU (Stockholm)	376.908 1K tokens	USD 0.84
Amazon Bedrock USE1-moonshotai.kimi-k2.5-mantle-input-tokens-standard		USD 46,336.92
\$0.0006 per 1K token for moonshotai.kimi-k2.5-mantle-input-tokens-standar in US East (N. Virgin...)	77,228,206 1K tokens	USD 46,336.92
Amazon Bedrock USE1-moonshotai.kimi-k2.5-mantle-output-tokens-standard		USD 587.04
\$0.003 per 1K token for moonsotai.kimi-k2.5-mantle-output-tokens-standard in US East (N. Virgin...)	195,679 1K tokens	USD 587.04

**Top Models**

amazon.nova-lite-v1:0	\$0.00
1.3M · 84 msgs	
moonshotai.kimi-k2.5	\$0.00
72.1M · 547 msgs	
deepseek.v3.2	\$0.00
61.8M · 844 msgs	
moonshotai.kimi-k2.5	\$0.00
5.9M · 220 msgs	
delivery-mirror	\$0.00

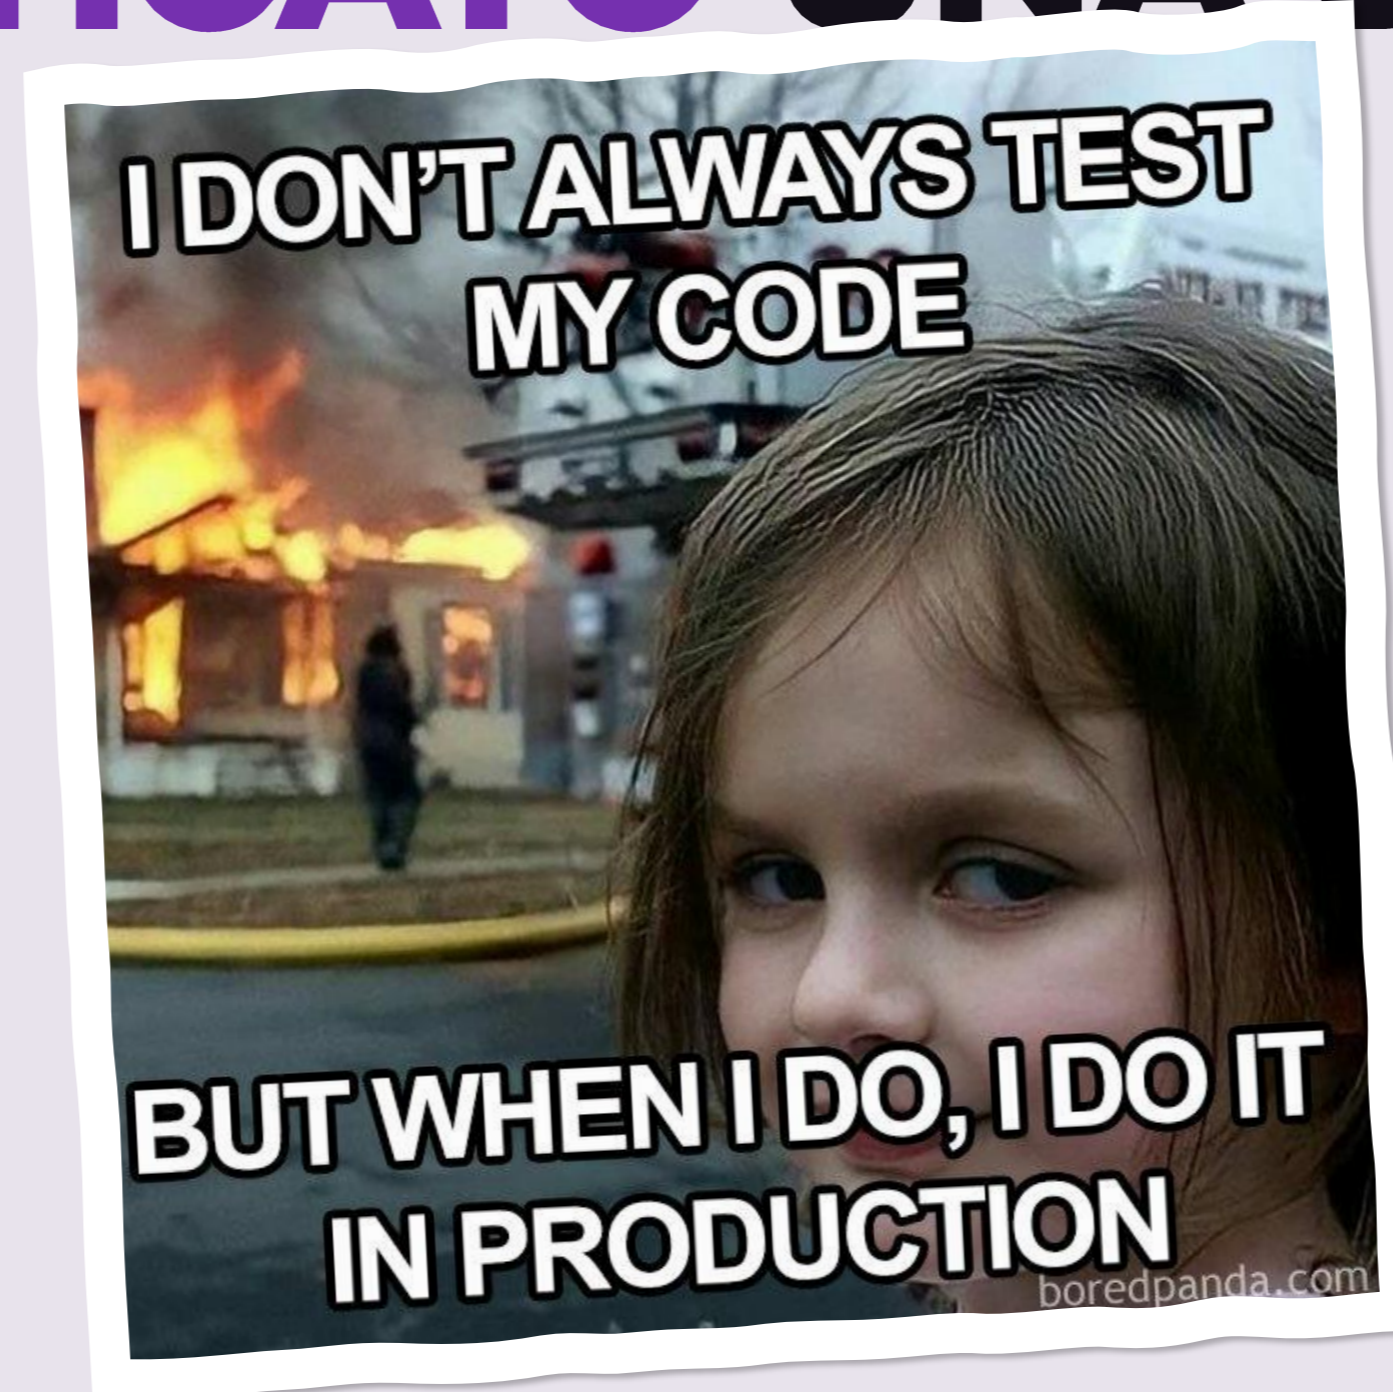


# SONO TOKEN IMPERIALI?

Collega Anonimo



# QUALCUNO PROBABILMENTE SI È **DIMENTICATO** UNA DIVISIONE



## SANTO SUPPORTO

# CI È VOLUTO UN PO' DI TEMPO, MA ALLA FINE

Tue Mar 03 2026  
16:32:13 GMT+0100 (Central European Standard Time)

Thank you for your support and for resolving this issue. I hope my analysis is useful in solving the issue.  
I am resolving the case.

Have a nice day,  
Damiano



Amazon Web Services

Tue Mar 03 2026  
16:20:17 GMT+0100 (Central European Standard Time)

Translate ▼

Hello Damiano,

Thank you for your patience while the service team reviewed your request.

Our service team has acknowledged the reported issue and are currently working on a fix to ensure it is fully resolved.

Please confirm at your convenience whether the billing details now appear accurate. Your confirmation will help us ensure that all concerns have been thoroughly addressed.

If you're still experiencing any discrepancies or have further questions, please feel free to let us know. We'll be glad to assist you.

Wishing you an amazing day ahead !

We value your feedback. Please share your experience by rating this and other correspondences in the AWS Support Center. You can rate a correspondence by selecting a star rating next to the correspondence.

Best regards,  
Uzma S.  
Amazon Web Services

Was this helpful?  
★★★★★



Amazon Web Services

Sat Feb 28 2026  
16:03:31 GMT+0100 (Central European Standard Time)

Translate ▼

Hello Damiano,

Greetings from AWS !

I'm following up on our chat regarding unexpected billing charges on your AWS account. Thank you for your time and patience in discussing this with us.

I understand you experienced a significant and unexpected increase in charges related to Amazon Bedrock usage, specifically when testing AI models through OpenAI. Rest assured; I'll assist you.

Was this helpful?  
★★★★★



**HE DUG INTO THE BILLING DATA,  
FOUND THE BUG, AND GOT IT RESOLVED.  
A MASTERCLASS IN BEING  
YOUR OWN CLOUD ADVOCATE.**

Corey Quinn, Last Week in AWS

COSA È SUCCESSO

# I BUG

## 1. Bi

**TOKEN SINGOLI CONTEGGIATI  
COME 1K-TOKEN**

Si fa in fretta a raggiungere cifre stratosferiche, 58\$ che si trasformano in 58.000\$ in una settimana

## 2. MONITO

**NESSUNA METRICA  
DISPONIBILE, ALLARMI MUTI**

Nella console CloudWatch nessuna anomalia strana, anche il billing è sempre rimasto costante

## 3.

**METRICHE AGGREGATE A  
POSTERIORI E IN BATCH**

Guardando bene le metriche, se fossero arrivate in tempo, il primo giorno avrei ricevuto un allarme per 2.000\$

POSSIAMO FARCI QUALCOSA?

# MA QUINDI DOBBIAMO RASSEGNArcI E SPERARE?



DEJA VU?

# NUOVA MUSICA, STRUMENTI FIDATI.

## 1. Monitoring

Token = Invocazioni, come con Lambda

## 2. Limits & Guardrails

Quota, come concurrency su Lambda

## 3. Alerting

Budget Alarm, Anomaly Detection

## 4. FinOps

Scegliere il modello giusto

TUTTO BELLO, MA...

**NON SONO PROTETTO DA  
BUG DEL GENERE!**

COME SCOVARE I BUG

# I BUG SI COMBATTONO CON L'OSSERVABILITÀ.

## Monitorare fornisce dati

I dati possono supportare i case.

## Gli allarmi fanno agire prima

E quando non scattano, qualcosa è davvero sbagliato.

## Collezionare metriche applicative

Avere una fonte per capire le anomalie.

## Conoscere la tecnologia

Per evitare errori di stima.



COME SCOVARE I BUG

# STRUMENTI PER NON AVERE SORPRESE

## Amazon Bedrock Projects

Assegnare i costi in base a tag

## OpenTelemetry for Generative AI

Se già usate OTEL per il tracing, avete più correlazione.

## Amazon CloudWatch Custom Metrics

Direttamente dall'interno dell'applicazione.

## Bedrock Guardrails

Per evitare ridurre il rischio di usi fraudolenti.





# ALTRI RACCONTI SUL BLOG.



SPEAKER

[damiano.giorgi@besharp.it](mailto:damiano.giorgi@besharp.it)

**WE MAKE IT RUN.**  
*SERIOUSLY!*