SRECON EMEA

Who Owns the Incident?

The Hidden Costs of Centralized vs. Decentralized Incident Response

Ana A. Lebrón Moreno

Head of TechOps, Enova International

“If everybody owns it, nobody owns it.”

A quick word on who's talking

- **Head of TechOps — Enova International**

  Leading Resiliency, Observability & SRE across a regulated FinTech environment (Chicago).

- **10+ years building & scaling**

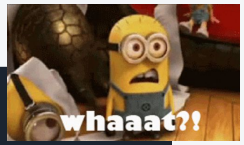
  Incident management, SRE programs, observability platforms, and the teams that run them.

- **Founder, Chicago Women in Tech Conference**

  Running 10+ years; mentoring emerging women tech leaders across the industry.

- **Started in aerospace, not FinTech**

  Systems engineering at EADS/Astrium before making the jump to tech operations.



FUN FACT

My first real engineering job analyzed ranging methods for a mission to Mercury.

These days the only planet I worry about is prod — and it pages a lot more than Mercury did.

HOW I ACTUALLY GOT HERE

A career I did not exactly plan



Spain → Germany→ Spain

Telecommunication Engineering
Ground systems for military aircraft
Radio systems for satellites headed to Mercury.





On to the US

Moved to the US to join Enova as a Project Manager.



5 years in…

“Want to try SRE?”

New skillset, zero notice



Incident Management

ROADMAP

Five stops

00 What actually counts as an incident?

01 Centralized: the expertise trap

02 Decentralized: the ownership mirage

03 Hybrid: design vs. accident

04 A decision framework — and the takeaways

[redacted]
Ana: not sure if it should have been a PI or not. [redacted]

[redacted]
Is this a PI?

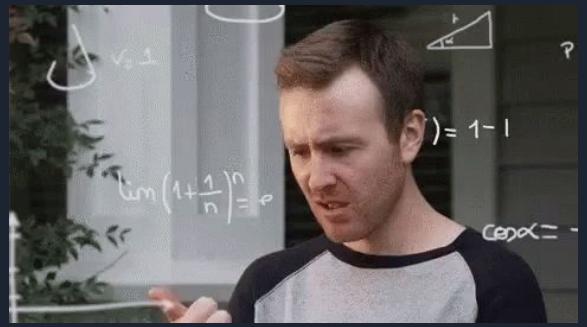
[redacted]
[redacted] Should we call a PI?

[redacted]
Should we call a PI for not being able to call a PI?
7

[redacted]
But all of that was corrected so quickly that I would be fine not calling this a PI



Ana Lebron pi-and-retro-pics
Any reason why we should not treat this as a PI?



What actually counts as an incident?

Even the question “should we call an incident” needed an incident of its own.

What counts as an incident?

!

Under-counting

Only total outages make the list. Everything else is “just a blip.”

!

Definitional chaos

Everything alert counts as an incident

Your model choice is meaningless until you answer this one — incident volume decides which staffing model is even viable.

00 — INCIDENT VOLUME

Incident volume decides everything downstream

Choose your fighter.



The Standing Army

6 RESPONDERS · 4 INCIDENTS A YEAR

Staffing	■■■■■■■■■■	*Six people. Four incidents.*
Rotation	■■■■■■■■■■	*On call all year. Paged twice.*
Sharpness	■	*Last real one: "before the reorg."*
Learning	■	*Annual Postmortem*

Special ability Exceptionally well rested.
Weakness Half the team has LinkedIn open right now.



The Superhero

1 RESPONDER · 400 INCIDENTS A YEAR

Staffing	■	*One person. Also the backup.*
Rotation	■	*Off-call is a rumour.*
Sharpness	■■■■■■■■■■	*Handle incidents in PJs. Did, tonight.*
Learning	■	*Postmortems? After the next one.*

Special ability Types with eyes closed.
Weakness Eyes are closed.

Same job title. Same budget line. One number decided everything else about their year.



Centralized Incident Response Team

Deep expertise. Concentrated pressure.

The appeal

- **Clear ownership**
  One team everyone knows to call.
- **Deep, consistent expertise**
  The same responders, every time — the muscle stays sharp.
- **One team, one process**
  No guessing who's in charge or how it runs.



🎯

But this model only works if incident volume justifies the headcount it requires.

The headcount trap

5–6

The minimum viable on-call rotation for sustainable coverage

Below that minimum: burnout.
Above your incident volume: idle experts drifting off their game.

Amazon's Two-Pizza Rule



Optimal size of a team is the number of people who can be fed by two pizzas

01 — CENTRALIZED

Responder ≠ analyst



The Responder

"Comfortable with ambiguity. For about twenty minutes."

About me Decisive. Fast. Best self at 2 AM with the pager going off.

Green flags Makes the call before the dashboard finishes loading.

Red flags Believes "restarted it, it's fine now" is a root cause.

Looking for The next page. Not the previous one.



The Analyst

"Comfortable sitting with uncertainty. For days."

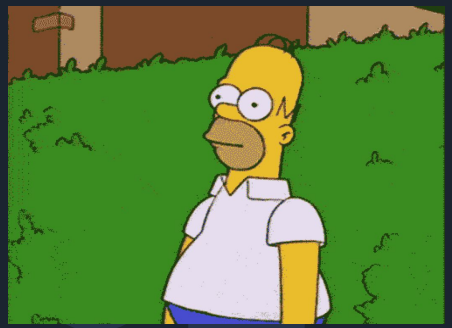
About me Patient. Curious. Traces one factor through six systems — for fun.

Green flags Says "contributing factors." Plural. Reads the timeline twice first.

Red flags Asks "but why?" a sixth time. At 2 AM. Prod is still down.

Looking for A root cause that fixes something. Happy to wait.

Ask a responder to write a 5-page root-cause doc and watch "comfortable with ambiguity" evaporate in real time.



Decentralized Incident Response

Scales with the org. Diffuses the ownership.

The appeal

- **No dedicated headcount problem**
  Nobody has to be hired just to hold the pager.
- **Scales naturally as the org grows**
  More teams, more responders — automatically.
- **Responders are close to the systems they own**
  The person paged already knows the code.



🥺

But this is not the job I was hired for

Everyone's responsible. Is anyone responsible?

THE THEORY

Diffusion of Responsibility · Ringelmann Effect

1913: people pulled a rope alone, then in groups.
Alone, each pulled ~**63** kg.
In a group of three, per-person force dropped to ~**53** kg — less coordination, less felt accountability.



The tax — and the warning signs

The context-switching tax

Engineers whose actual job is different from incident response get pulled into firefighting instead — every incident is a context switch away from the roadmap

Signals it's failing

- Inconsistent communication and response process
- Long-running incidents, incomplete documentation
- The same incident type recurs across teams with no shared learning



Hybrid

Most orgs already are one. Few designed it.

Most organizations are already hybrid

Few designed it that way. Most inherited it — a pattern solidified before anyone asked whether it was the right one.



Someone went on call

→



Something broke

→



A pattern solidified

Three questions every hybrid model must answer

1. **Which team member is in charge of Incident Response?**
2. **Who runs the incident vs who actually resolved it?**
3. **Who does the analysis afterwards?**

Answer these on purpose — or answer them at 3 AM, under pressure.

Train the muscle, whichever model you pick

🧪 Chaos engineering and tabletop exercises work across all three models — they build the same response muscle, whoever's holding the pager.

💣 *Nothing builds team spirit like simulating a database meltdown at 2pm on a random Tuesday.*



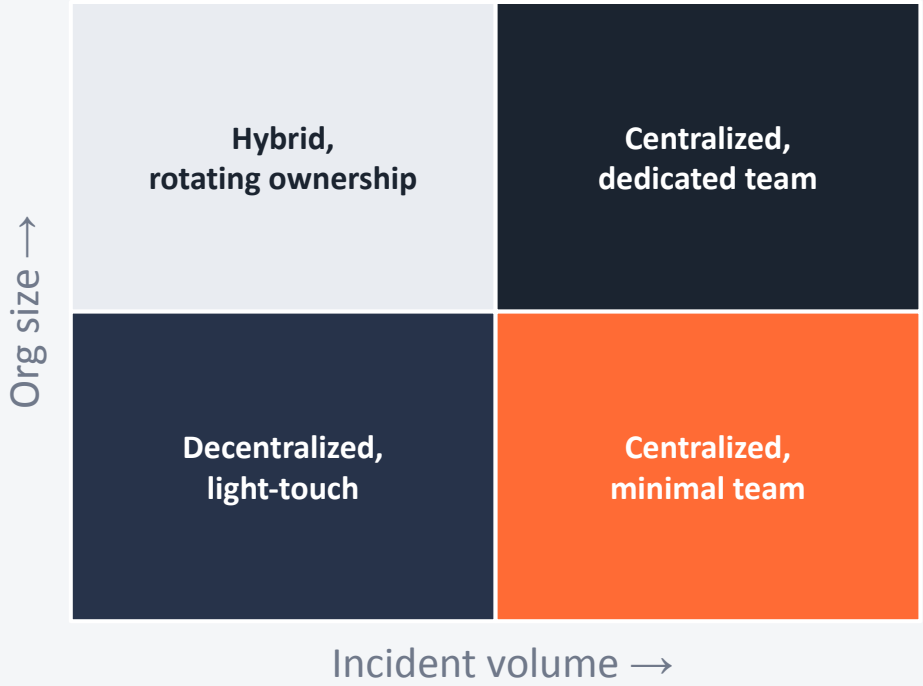


The Framework

Four variables. One 2×2. Choose deliberately.

Four variables decide the model that fits

1. Organization size
2. Incident volume & cadence
3. Observability & tooling maturity
4. Is analysis a distinct discipline? (yes/no)



Three signals your model needs to change

1

Rotation unsustainable

Too few people, or too little load to keep skills current.

2

Shallow postmortems

Consistently produce no durable, lasting change.

3

Recurring incidents

Same incident type across teams — with no cross-team learning.

04 — CLOSE

This was never really about org charts

Your incident response model encodes who owns reliability, what learning is worth funding, and how you value engineers' time. Choose it — don't inherit it.

- A practical definition of “incident” — the prerequisite to any model choice
- A structured way to evaluate your current model
- Concrete signals that tell you a change is due
- The responder / analyst distinction — and what to do about it

Thank you!



Ana A. Lebrón Moreno · Head of TechOps, Enova International
alebronmoreno@gmail.com