

Don't wait for an outage to improve your reliability

Leveraging game days



Hi, I'm Leo

I'm a product engineer at incident.io

I work on our on-call product





Incidents are stressful





You've been paged

21:06

Pod CPU high (role=web)



Leo 21:07

ok it doesn't actually look too bad

image.png ▼



I can load the dashboard

 Attached 3 alerts to this incident...

Pod CPU high (role=web)

[View in incident.io](#)

Pod CPU high (role=web)

[View in incident.io](#)

HTTP Probe Failed: core-production-api-incident

[View in incident.io](#)

 Attached 3 alerts to this incident...

Pod CPU high (role=web)

[View in incident.io](#)

Pod CPU high (role=web)

[View in incident.io](#)

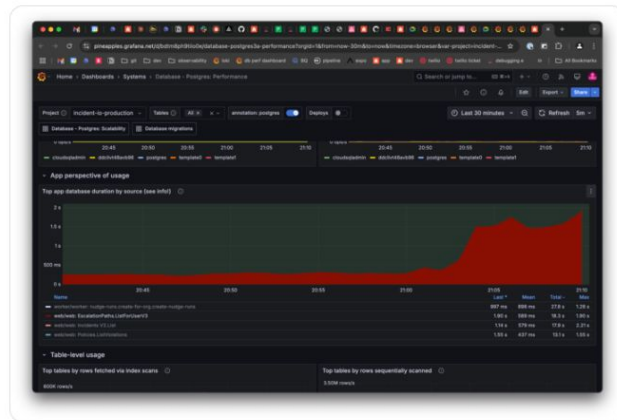
HTTP Probe Failed: core-production-api-incident

[View in incident.io](#)



Leo 21:11

image.png ▼



📎 Attached 3 alerts to this incident...

Pod CPU high (role=web)
[View in incident.io](#)


Pod CPU high (role=web)
[View in incident.io](#)

HTTP Probe
[View in incident.io](#)



Leo 21:11

image.png ▼



Stephen





Just saw this, unsure if related: <https://incident-io.slack.com/archives/C073ZDQGZ4J/p1734470559363529>



Hey team, are we seeing any issues with connectivity? I'm seeing this pop-up on a lot of actions I'm trying to perform currently from slack.

2 files ▼ | 📄 Download all



Thread in # cust- [redacted] 17 Dec 2024 | [View message](#)

Attached 3 alerts to this incident...

Pod CPU high (role=web)

[View in incident.io](#)

Pod CPU high (role=web)

[View in incident.io](#)

HTTP Probe

[View in incident.io](#)

 **Rory**
This

 **kels**
I ma




[Merge It into another](#) [Decline It](#)

Thread in # cust- [redacted] 17 Dec 2024 | [View message](#)

 **Leo** 21:11

image.png



529



Incidents



Fixes

Incidents improve your reliability

Incidents improve your reliability

How do you skip the incident?

How do you skip the incident?

Running a game day

01

Set the stage

What do you want this game day to achieve?

02

Prepare your game day

How do you build scenarios that deliver on that?

03

Game time

How do you make it feel real without breaking anything in production?

04

Learn and celebrate

How do you walk out better and wanting to do it again?

Running a game day

Set the stage

Prepare your game day

Game time

Learn and celebrate

Set the stage

Define your goal

Are you testing your...

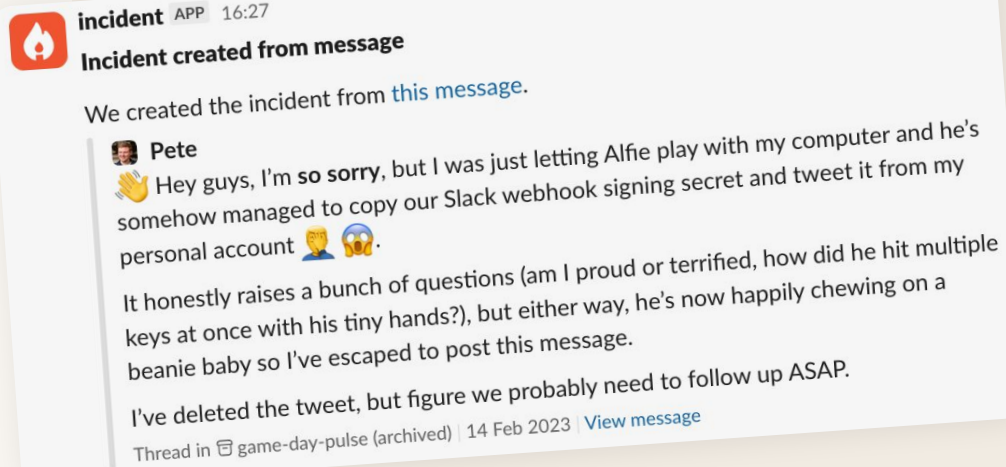
- Systems?
- Processes?
- People?



Define your goal

Test response to unlikely scenarios

“How will my responders react when faced with this scenario”



incident APP 16:27
Incident created from message

We created the incident from [this message](#).

Pete
👉 Hey guys, I'm **so sorry**, but I was just letting Alfie play with my computer and he's somehow managed to copy our Slack webhook signing secret and tweet it from my personal account 🙄🤖.

It honestly raises a bunch of questions (am I proud or terrified, how did he hit multiple keys at once with his tiny hands?), but either way, he's now happily chewing on a beanie baby so I've escaped to post this message.


I've deleted the tweet, but figure we probably need to follow up ASAP.


Thread in game-day-pulse (archived) | 14 Feb 2023 | [View message](#)

Define your goal

Test response to failure modes

“How do I build more confidence in the reliability of our system”


 **incident** APP 10:35


 **High web pod memory use**


- Web pods are pinned at 100% CPU
- Looks like the app is still usable but it's **likely we'll start dropping traffic for all customers**
- Customers may see slowness via the app and public API
- Other pods look ok, this means that statuspages, alerts should still be working
- Marking this as major for now

 **Type:** Game day


 **Severity:** Major

 **Status:** Closed

 **Incident Lead:** Henry Course


 **Channel:** [#game-day-inc-high-web-pod-memory-use](#)

 **Features:** [API](#)

 [Incident homepage](#)

 [Subscribe](#)


Last updated at 12:31, Friday, 11 October 2024 (*Game day*)


 **5 replies** Last reply 2 years ago


Define your goal


Test fixes from a prior outage


“If the same failures happened, would we hold strong”


 **incident** APP 11:22


 **AWS is down**


 Type: Game day


 Severity: Major

 Status: Closed

 Reporter: @oli

 Incident Lead: @oli

 Channel: [#game-day-inc-aws-is-down](#)

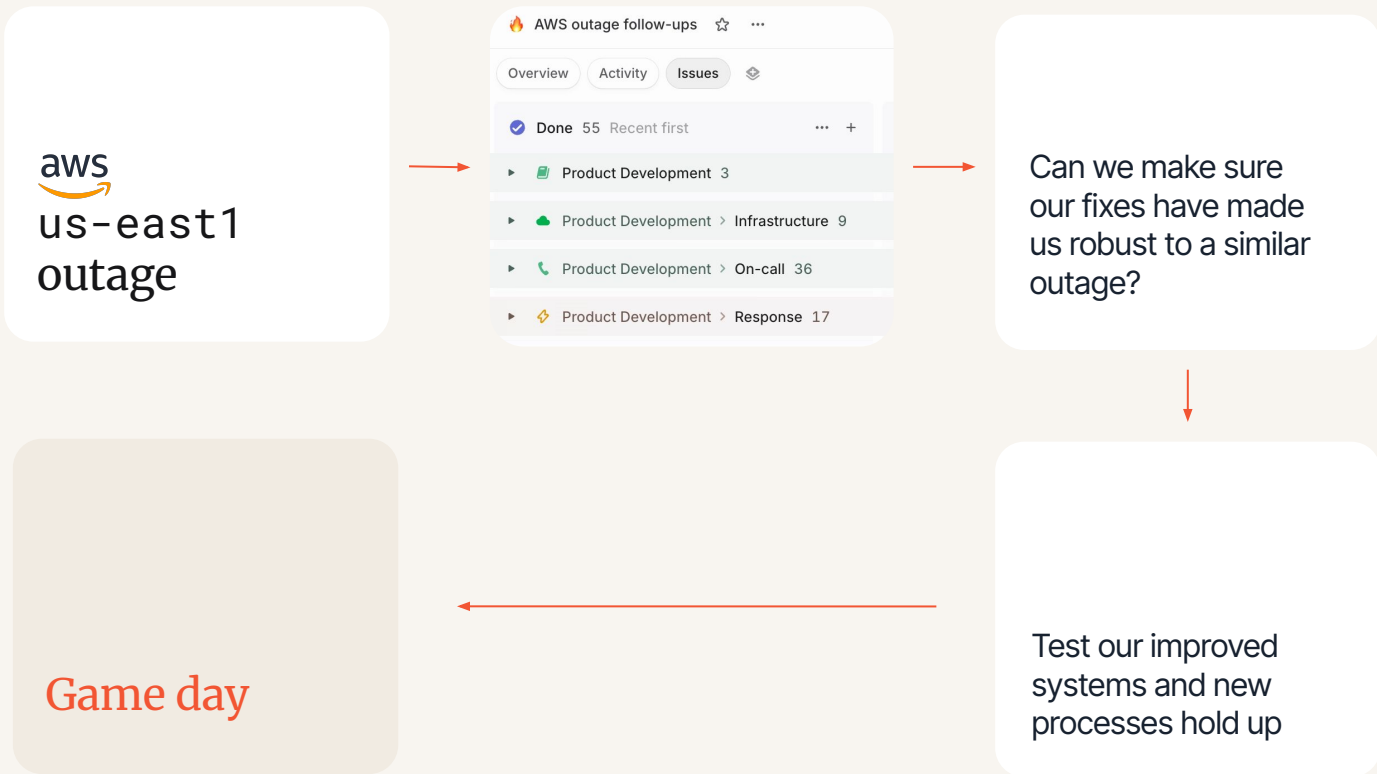
 [Incident homepage](#)

 [Subscribe](#)

Last updated at 11:42, Thursday, 11 December (*Game day*)

 **1 reply** 6 months ago

Define your goal



Set the stage

Assemble your champions

- 1x Engineering Manager
- 1x Engineer
- 1x Platform Engineer



Running a game day

Set the stage

Prepare your game day

Game time

Learn and celebrate

Prepare your game day

Coming up with *what* to break

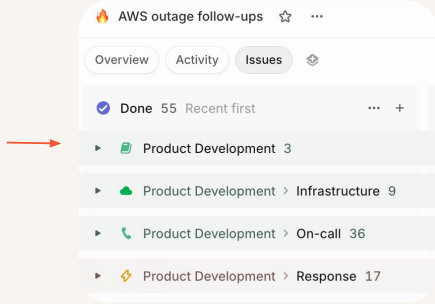
- Which part of your architecture have you always wanted an excuse to break?
- What parts of your system make you nervous?
- What failure modes have you seen before, and don't want to see again?

Prepare your game day

How will you break it?

- Can you plausibly break it?
- Do you have an environment to break things on?
- How will your breakages manifest?
- What mitigation steps do you expect?


aws
us-east1
outage



Can we make sure
our fixes have made
us robust to a similar
outage?


Game day

Test our improved
systems and new
processes hold up


 **Scenario 1:** third-party vendors are down
Scenario 2: dashboard usage spikes
Scenario 3: we can't deploy

Prepare your game day


A third-party vendor we use for a feature is down



Increased web dashboard usage




We can't deploy because building images in CI is broken



Prepare your game day

A third-party vendor we use for a feature is down




Block traffic with a DNS rule 

Increased web dashboard usage



We can't deploy because building images in CI is broken



Prepare your game day


A third-party vendor we use for a feature is down




Block traffic with a DNS rule 

Increased web dashboard usage




Trigger a malicious endpoint that increases CPU by calculating pi 


We can't deploy because building images in CI is broken




Prepare your game day


A third-party vendor we use for a feature is down




Block traffic with a DNS rule 



Increased web dashboard usage



Trigger a malicious endpoint that increases CPU by calculating pi 

We can't deploy because building images in CI is broken



Break image building in CI 
Declare manually 

Prepare your game day

Put it together

1 Group 1

▼ Issue 1: Scribe

How to break it:

- Add a fake DNS record for the us-east-1 Recall API
 - ▶ Instructions
- Prove that Scribe is broken
- [Nudge] Post a message as an engineer "AWS region is down"
 - ▶ **Expected impact**
 - ▶ **Expected response actions**

Expected response time: ~20m

- ▶ **Prerequisites**


Cleanup:

- Remove fake DNS entries

Prepare your game day


Who are your responders?

- ✓ 3-4 people per 2-3 incidents
- ✓ Mix of experience
- ✓ Small organisation → mix product teams
- ✓ Bigger organisation → one session = 1 product team
- ✓ Involve customer success and other stakeholders



Prepare your game day

Final logistics

- ✓ 1 session = half day
- ✓ Book rooms
- ✓ Organise food/snacks
- ✓  Private villain channel
- ✓ #game-day-chat
- ✓ Incident tool

Running a game day

Set the stage

Prepare your game day

Game time

Learn and celebrate

Game time

Welcome your responders

- ✓ Give timings
- ✓ Establish rules of engagement
- ✓ Icebreaker/Warm-up

Timings

- 10:00-10:15 : [All] Chat / Rules of engagement / Questions
- 10:15-10:30 : [Villains] Scenario prep
- 10:30-11:30 : [All] Incidents begin
- 11:30-12:00 : [Villains] Scenario cleanup, debrief doc prep
- 12:00-12:30 : [All] Debrief with the group

Game time



oli 10:20

it's going to be a very quiet day today looking forward to getting lots of work done.



3



1



1



1



Game time

Time to break things

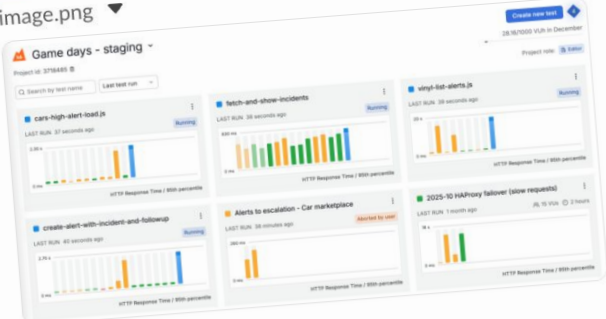
☆ game-day-private-2025-12-10 Used for villain coordination

Messages Files

Leo 10:00
Do you want to begin background load now?

Leo 10:11
✓ Tom is now on [Game day schedule](#)
✓ Background noise kicked off

image.png ▾




The screenshot shows a Grafana dashboard with the following panels:

- cars-high-alert-load.js**: LAST RUN: 37 seconds ago. HTTF Response Time / 80th percentile.
- create-alert-with-incident-and-followup**: LAST RUN: 40 seconds ago. HTTF Response Time / 80th percentile.
- Alerts to escalation - Car marketplace**: LAST RUN: 38 seconds ago. Alerted by Leo.
- 2025-10 HAProxy fallover (low request)**: LAST RUN: 1 month ago. JEL TO: 10 min. 3 hours.
- Batch-and-show-incidents**: LAST RUN: 38 seconds ago. HTTF Response Time / 80th percentile.
- why4-fst-alerts.js**: LAST RUN: 38 seconds ago. HTTF Response Time / 80th percentile.

✓ Game day alert route updated


Ben 10:28
💣 Recall DNS records now added
Scribe is definitely dead

Game time






Game time


Managing stress



 **Chris Class** 9 Oct 2024 at 18:29
Gameday coffee break order ☕

7 replies

 **Chris Class** 9 Oct 2024 at 18:30
[@Leo](#) [@Liz](#) [@milhouse](#) [@Leo Sjöberg](#) [@Martha](#) [@Ben](#) [@Alicia](#)

 **Leo** 10:32
[ Lucy] Hey guys, my fave customer Ben's Used Car Marketplace is reporting they're not receiving incident updates for their major incident - but they've got subscriptions set up. Help please? (edited)

 **Harry Austin** 🌴 11:11
Can we move back to to a previous version?

 **Ben** 11:12
[ Matthew] I'm at a rave in the woods right now, but maybe you can try rollback?

Game time

Finding weaknesses as you go



Ben 🏠 11:22



"Can't find just one dashboard for web"



Leo 11:14



I still don't know why we got 422s from 

Game time

We created the incident from [this message](#).



Leo

[[#engineering](#)]

Hey guys, I'm in an incident and Scribe isn't joining in. Can you take a look please?

Thread in [# game-day-chat](#) | 10 Dec 2025 | [View message](#)

Game time

We created the incident from [this message](#).



Leo

[\[#engineering\]](#)

Hey guys, I'm in an incident and Scribe isn't joining in. Can you take a look please?

Thread in [# game-day-chat](#) | 10 Dec 2025 | [View message](#)



Gabriel de Olim Gaul 10:38

[@Tom DW](#) do we have any dashboards for scribe?



Tom DW 10:39

I'm taking a look. There's one for recall

[https://pineapples.grafana.net/d/adiw4fbize2o0b/integrations?orgId=1&refresh=1m&var-proj\[...\]tegration=recall&var-organisation=&var-organisation_name=latency](https://pineapples.grafana.net/d/adiw4fbize2o0b/integrations?orgId=1&refresh=1m&var-proj[...]tegration=recall&var-organisation=&var-organisation_name=latency) seems up

Game time

We created the incident from [this message](#).



Leo

[[#engineering](#)]

Hey guys, I'm in an incident and Scribe isn't joining in. Can you take a look please?

Thread in [# game-day-chat](#) | 10 Dec 2025 | [View message](#)



Gabriel de Olim Gaul 10:38

[@Tom DW](#) do we have any dashboards for scribe?




Tom DW 10:39

I'm taking a look. There's one for recall

[https://pineapples.grafana.net/d/adiw4fbize2o0b/integrations?orgId=1&refresh=1m&var-proj\[...\]tegration=recall&var-organisation=&var-organisation_name=latency](https://pineapples.grafana.net/d/adiw4fbize2o0b/integrations?orgId=1&refresh=1m&var-proj[...]tegration=recall&var-organisation=&var-organisation_name=latency) seems up



Leo 10:40

[ Lucy] Hey guys, just joining the channel because we've got a couple more customers reporting issues with Scribe - maybe it's not just us

Game time

We created the incident from [this message](#).



Leo

[#engineering]

Hey guys, I'm in an incident and Scribe isn't joining in. Can you take a look please?

Thread in # game-day-chat | 10 Dec 2025 | [View message](#)



Gabriel de Olim Gaul 10:38

@Tom DW do we have any dashboards for scribe?



Tom DW 10:39

I'm taking a look. There's one for recall

[https://pineapples.grafana.net/d/adiw4fbize2o0b/integrations?orgId=1&refresh=1m&var-proj\[...\]tegration=recall&var-organisation=&var-](https://pineapples.grafana.net/d/adiw4fbize2o0b/integrations?orgId=1&refresh=1m&var-proj[...]tegration=recall&var-organisation=&var-)



Leo 10:40

[Lucy] Hey guys, just joining the channel because customers reporting issues with Scribe - maybe it's no



incident APP 10:49



Gabriel de Olim Gaul became the **Communications Lead**



Gabriel de Olim Gaul 10:51

@Leo [norberto] I'm going to make a status page update if you can review please



1



Game time

We created the incident from [this message](#).



Leo

[#engineering]

Hey guys, I'm in an incident and Scribe isn't joining in. Can you take a look please?

Thread in # game-day-chat | 10 Dec 2025 | [View message](#)



Gabriel de Olim Gaul 10:38

@Tom DW do we have any dashboards for scribe?



Tom DW 10:39

I'm taking a look. There's one for recall

[https://pineapples.grafana.net/d/adiw4fbize2o0b/integrations?orgId=1&refresh=1m&var-proj\[...\]tegration=recall&var-organisation=&var-](https://pineapples.grafana.net/d/adiw4fbize2o0b/integrations?orgId=1&refresh=1m&var-proj[...]tegration=recall&var-organisation=&var-)



Leo 10:40

[Lucy] Hey guys, just joining the channel because customers reporting issues with Scribe - maybe it's no



incident APP 10:49



Gabriel de Olim Gaul became the **Communications Lead**



Gabriel de Olim Gaul 10:51

@Leo [norberto] I'm going to make a status page update if you can review please



Tom DW 11:04

I'm going to follow the runbook to switch over the Scribe region to us-west-2

Game time

We created the incident from [this message](#).



Leo

[#engineering]

Hey guys, I'm in an incident and Scribe isn't joining in. Can you take a look please?

Thread in # game-day-chat | 10 Dec 2025 | [View message](#)



Gabriel de Olim Gaul 10:38

@Tom DW do we have any dashboards for scribe?



Tom DW 10:39

I'm taking a look. There's one for recall

[https://pineapples.grafana.net/d/adiw4fbize2o0b/integrations?orgId=1&refresh=1m&var-proj\[...\]tegration=recall&var-organisation=&var-](https://pineapples.grafana.net/d/adiw4fbize2o0b/integrations?orgId=1&refresh=1m&var-proj[...]tegration=recall&var-organisation=&var-)



Leo 10:40

[Lucy] Hey guys, just joining the channel because customers reporting issues with Scribe - maybe it's no



incident APP 10:49



Gabriel de Olim Gaul became the **Communications Lead**



Gabriel de Olim Gaul 10:51

@Leo [norberto] I'm going to make a status page update if you can review please



Tom DW 11:04

I'm going to follow the runbook to switch over the Scribe region to us-west-2



Gabriel de Olim Gaul 11:08

Cool Scribe is joining now



Game time

We created the incident from [this message](#).



Leo

[#engineering]

Hey guys, I'm in an incident and Scribe isn't joining in. Can you take a look please?

Thread in # game-day-chat | 10 Dec 2025 | [View message](#)



Gabriel de Olim Gaul 10:38

@Tom DW do we have any dashboards for scribe?



Tom DW 10:39

I'm taking a look. There's one for recall

[https://pineapples.grafana.net/d/adiw4fbize2o0b/integrations?orgId=1&refresh=1m&var-proj\[...\]tegration=recall&var-organisation=&var-](https://pineapples.grafana.net/d/adiw4fbize2o0b/integrations?orgId=1&refresh=1m&var-proj[...]tegration=recall&var-organisation=&var-)



Leo 10:40

[Lucy] Hey guys, just joining the channel because customers reporting issues with Scribe - maybe it's no



incident APP 10:49



Gabriel de Olim Gaul became the **Communications Lead**



Gabriel de Olim Gaul 10:51

@Leo [norberto] I'm going to make a status page update if you can review please



Tom DW 11:04

I'm going to follow the runbook to switch over the Scribe region to us-west-2



Gabriel de Olim Gaul 11:08

Cool Scribe is joining now



incident APP 11:18



Scribe not joining incident calls was resolved on the Game days status page

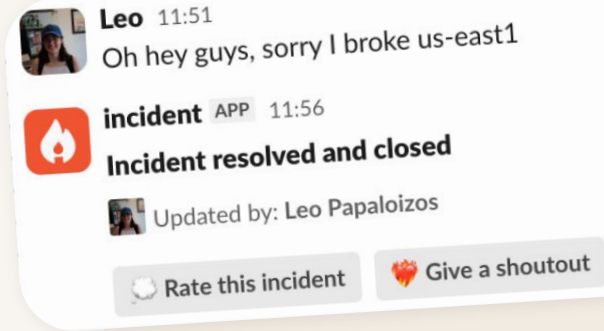


1 reply 6 months ago

Game time


Wrap it up

- ✓ No more pages!
- ✓ Let your responders know game time's over




Game time

Clean up your breakages

 **Joe Hart** 17 Mar at 17:27
📱 Howdy folks, do we have any weird staging egress rules that would prevent staging from talking to Twilio?
✅ 1 😊

11 replies

 **Ben** 17 Mar at 17:27
No
edit: Yes (edited)
😂 1 😊
👁️
🐸

Running a game day

Set the stage

Prepare your game day

Game time

Learn and celebrate

Learn and celebrate

Regroup and reflect

What went well 😊 + 🗨️

Utilizing Scribe for Update: 3 Cards

Using Scribe was a great shout, and produced really accurate updates that allowed people to follow along

Leo Sjöberg

Solid Liz Comms 3 Cards

Incident updates felt good - willing to go back on incorrect assumptions from before, frequent enough that I felt up to date (roleplaying a founder)

Ed Stephinson

What didn't go well 😞 + 🗨️

Role Clarity and Communic 6 Cards

Should the incident lead been doing more "directing"? It felt a bit as though people were going off and pulling on their own threads.

Ben Wheatley

Data Recovery Confidence. 2 Cards

We didn't have the knowledge/confidence to do a restore from DB

Ben Wheatley

Learn and celebrate

Follow-ups

- ✓ Collect notes from what you found as a villain
- ✓ Combine your learnings with your responders'
- ✓ Find owners for the follow-ups

Gameday: 2025-10-20 AWS outag...

Aa Name	Category	Team
Log client timeouts on the "Integrations" Grafana dashboard	Observability	Response
Add data links/drill down for more Splash dashboard components: e.g. Allow jumping from components table to "detailed" view.	Observability	Infrastructure
Workflow to nudge you to fallback to Telnyx if Twilio has a timeout / if we detect a latency error	Alerting/Nudging	On-call
Improve the Recall failover runbook nudge	Alerting/Nudging	Response
Saying "Runbook" should maybe link to our runbooks	Alerting/Nudging	Response
Update our API keys and the runbook for our third Recall region in staging	Code changes	Response
Consider alerting on "retriable" errors, either generally or for Scribe specifically	Code changes	Response
Retry Twilio timeout errors with Telnyx, not just twilio errors	Code changes	On-call

Learn and celebrate

Celebrate

- ✓ End it with a fun activity
- ✓ Make game days collectibles



Summary

Running a successful game day: a template

Set the stage

- ✓ Define your goal
- ✓ Assemble your champions

Prepare your game day

- ✓ What will you break?
- ✓ How will you break it?
- ✓ How will responders fix it?
- ✓ Who are your responders?

Game time

- ✓ Rules of engagement
- ✓ Use a private channel
- ✓ Manage stress
- ✓ Take notes

Learn and celebrate

- ✓ Run a retrospective
- ✓ Find owners for follow-ups
- ✓ Celebrate!

Incidents improve your reliability
Go cause them intentionally