

# Empowering SRE Teams & Incident Management with AI

Transforming Chaos into “some” Control with GenAI

# Being On-Call is Stressful & Exhausting

- Being on-call disrupts sleep, mental focus, and productivity
- Incidents lead to high-stress levels and kill innovation
- Incidents need multitasking



# A High-Stakes Challenge

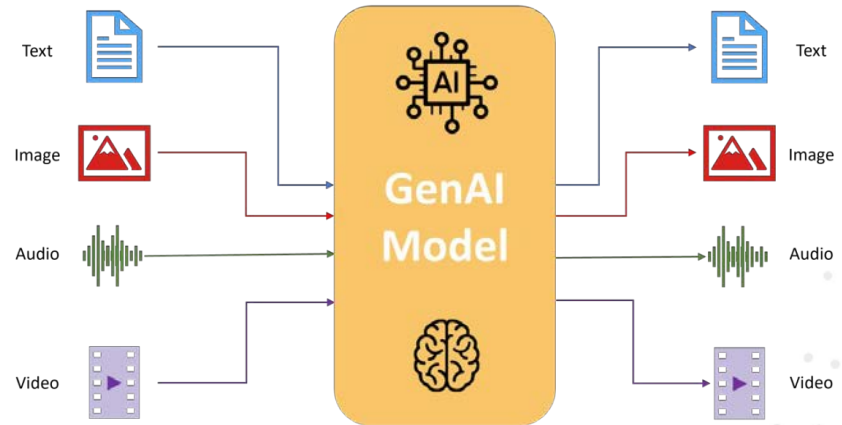
- You have to find the problem while communicating updates.
- Give updates to multiple teams and customer with timely and accurate information.
- War rooms are chaotic with logs, graphs, and endless threads.



# GenAI capabilities

# GenAI: More Than Just Text Analysis

- **Multimodal Capabilities:** AI can analyze text, logs, graphs, images, and more.
- **Contextual Understanding:** By integrating various data types, GenAI provides context-aware insights, offering a broader understanding of situations.



# A real Incident

## High Database CPU > 90%- Paged at 2:00 AM

**Me:** *It's 00:40 AM, and the first thought is, "Is this real? Is something going to break?"*

**Wife:** *And I'm over here wondering, 'What's going to crash first: the servers or you?'*

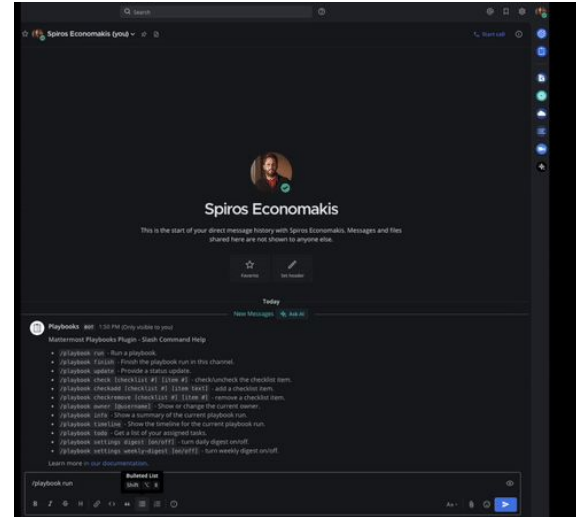


# Incident response



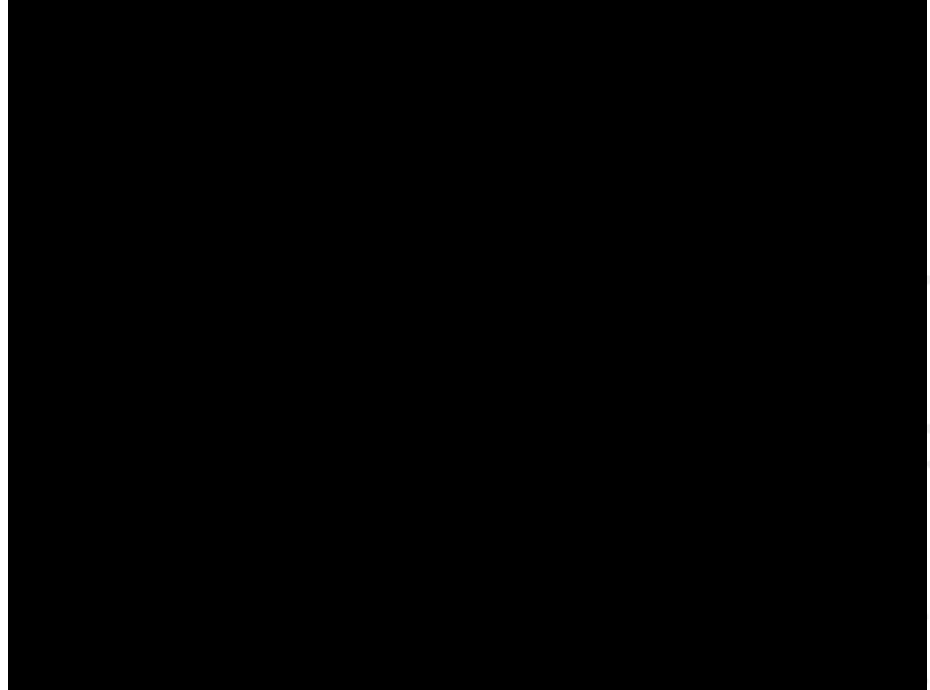
# Incident Response Playbook

- Start the predefined incident response playbook
- Automatically creates a dedicated channel as a "war room" to centralize communication and updates.



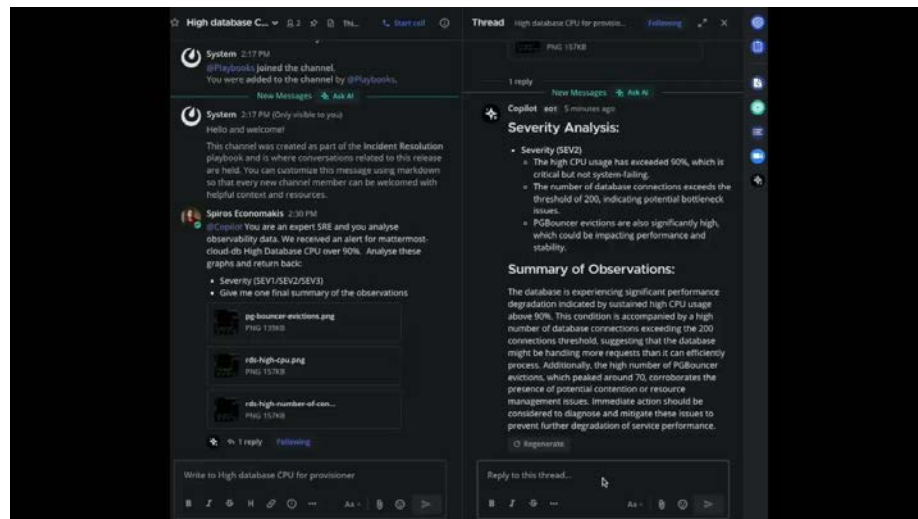
# Identification with AI

- Identify the severity based on the data we gathered from observability graphs



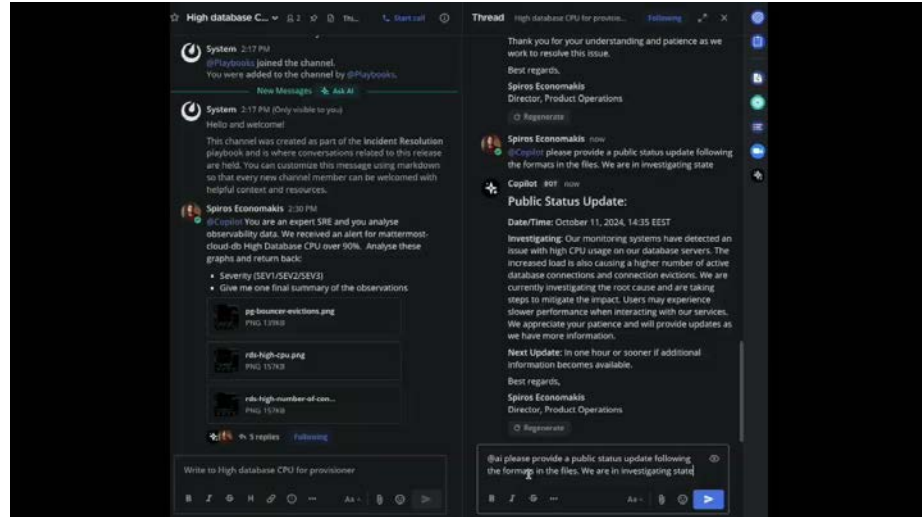
# Internal status updates with AI

- Identify the severity based on the data we gathered from observability graphs
- Internal status update



# External status updates with AI

- Identify the severity based on the data we gathered from observability graphs
- Internal status update
- Customer facing status update



# Investigate further with AI

## Contextual analysis:

- With extra input from service logs
- With extra input by teams

The screenshot shows a Microsoft Teams chat window with a thread titled "High database CPU for provisioner". The thread includes the following content:

- System** (2:17 PM): "Hello and welcome! This channel was created as part of the Incident Resolution playbook and is where conversations related to this release are held. You can customize this message using markdown so that every new channel member can be welcomed with helpful context and resources."
- Spiros Economakis** (2:30 PM): "@Copilot You are an expert SRE and you analyse observability data. We received an alert for mattermost-cloud-db High Database CPU over 90%. Analyse these graphs and return back:
  - Severity (SEV1/SEV2/SEV3)
  - Give me one final summary of the observations"
- Three screenshots: "pg-bouncer-evictions.png" (139KB), "rds-high-cpu.png" (157KB), and "rds-high-number-of-con..." (157KB).
- Playbooks bot** (3:00 PM): "@Spiros Economakis posted an update for High database CPU for provisioner"
- Copilot bot** (38 minutes): "Major Incident  
Incident with Database  
Investigating #1  
We are investigating reports of degraded performance for the database. We're investigating the cause and will provide more information as soon as it's available."  
"@ai Please further analyse the attached logs. Also people reporting:
  - "Support: Customers report that they cannot create trials."
  - "CSMs: A customer reported that tries to upgrade from Professional to Enterprise SKU and they get an error.""



# Post-incident with AI

- Generate a post-mortem and a timeline of events and evidences.
- Suggests action items to prevent the issue in the future
- Capturing metrics which can be used for insights

The image shows a Slack interface with a channel on the left and a thread on the right.

**Channel: High database C...**

System 2:17 PM (Only visible to you)  
Hello and welcome!  
This channel was created as part of the Incident Resolution playbook and is where conversations related to this release are held. You can customize this message using markdown so that every new channel member can be welcomed with helpful context and resources.

Spiros Economakis 2:30 PM  
@Copilot: You are an expert SRE and you analyse observability data. We received an alert for mattermost-cloud-db High Database CPU over 90%. Analyse these graphs and return back:

- Severity (SEV1/SEV2/SEV3)
- Give me one final summary of the observations

Attachments:  
pg-bouncer-evictions.png (PNG 139KB)  
rds-high-cpu.png (PNG 157KB)  
rds-high-number-of-con... (PNG 157KB)

9 replies Following

Playbooks bot 3:00 PM  
@Spiros Economakis posted an update for High database CPU for provisioner

**Thread: High database CPU for provisioner**

1. Immediate Actions:
  - **Scale Resources:** Consider increasing the database server capacity (CPU and Memory) to handle the current load.
  - **Database Tuning:** Optimize queries and indexes to reduce the load.
  - **Connection Management:** Evaluate and adjust the connection pool settings to better handle peak loads.
2. Incident Communication:
  - **Public Status Update:** Inform customers that high database load is impacting certain operations (trial creation, SKU upgrades), and the team is actively working on resolving it.
  - Maintain updates at regular intervals or when significant progress is made.

**Updated Public Status Update**

Date/Time: October 11, 2024, 15:35 EEST

**Investigating #1**

We are investigating reports of degraded performance for the database. Customers may experience issues while creating trials or upgrading from Professional to Enterprise SKU. We're investigating the cause and will provide more information as soon as it's available.

**Next Steps**

Please monitor the database performance and observe if scaling resources and other immediate actions help mitigate the issues. Continue to update stakeholders as we make progress and have new information.

Regenerate

Reply to this thread...



**AI isn't perfect yet**

# AI Isn't Perfect, But It's Improving

- **Not an artificial SRE**—AI automates repetitive tasks during stressful incidents.
- AI can give key insights fast, so you **focus** on fixing problems.
- Spend more time **improving**, less time **managing**.





# Key takeaways

# Takeaways

- AI helps understand the problems faster by giving you key insights right away.
- AI reduces stress by decreasing multitasking with automating updates and summaries.
- You spend less time managing the incident and more time fixing it.



# Thanks & follow

 @spirosoik

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