Improving Incident Response Workflows with an Embedded AI Assistant

Concept study for an AI-powered assistant designed to help responders orient faster during live incidents

Role: UX Researcher –Jessica Wood
Methods: Moderated concept testing, prototype walkthroughs
Participants: Incident responders and incident managers
Focus: Workflow fit, discoverability, privacy, prompt usefulness, and trust

UX Problem

Responders joining live incidents midway struggle to quickly understand what is happening

High-severity incidents move quickly. When responders join midway, critical context is often spread across live discussion, chat, incident records, and diagnostic tools.

Users need to quickly understand

1. **What happened before they joined**
2. **Current status and impact**
3. **Who is involved**
4. **What actions are underway**
5. **Whether their team or service is affected**
6. **What decisions are needed next**

Core UX problem: Responders need fast situational awareness, but the current workflow forces them to piece together information while trying not to interrupt the group chat or add unnecessary noise.

Why This Matters

Poor orientation creates friction during moments when speed and clarity matter most

Not being able to get situated quickly leads users to…

- Asking repetitive questions
- Interrupting responders who are actively mitigating the issue
- Missing important context
- Switching between multiple tools to reconstruct the situation
- Delaying their ability to contribute meaningfully

Opportunity

Explore whether an embedded AI assistant could help responders orient faster, ask questions privately, and selectively share useful information back to the group.

Design challenge

Add AI support without increasing risk, distraction, or public noise during high-pressure response work.

Research Goals

I evaluated whether the assistant fit real incident workflows

1. **Current behavior**

   How do responders currently catch up when joining an incident midway?

2. **Workflow fit**

   Would an embedded assistant reduce workflow friction?

3. **Discoverability**

   Is the assistant discoverable in the live incident experience?

4. **Capability clarity**

   Do users understand what the assistant can do?

5. **Trust**

   What do users need in order to trust and act on AI-generated responses?

Research focus

Because this was an early concept, the goal was not to measure task completion. The goal was to evaluate workflow fit, mental models, adoption barriers, and trust requirements.

Approach

Scenario-based concept testing to evaluate behavior, expectations, and adoption risk

Method

5
moderated sessions

6
participants total

60
minute sessions

What I evaluated

- Where users naturally went first to catch up
- Whether they noticed and understood the assistant
- How they reacted to private chat and suggested prompts
- What made them trust or question AI-generated output

Key Insight

Private-first investigation was the strongest signal

Users wanted to ask questions privately before sharing anything with the broader group. They were concerned about creating unnecessary noise, especially when their question only mattered to them or when they were still trying to understand the situation.

> *"I wouldn't want to ask everything on the bridge... that just creates noise."*

Why this matters

Private chat lowers the risk of engagement and makes users more likely to try the assistant.

Design implication

Make the assistant private by default and make sharing an intentional action.

> *"I would probably use this in a private way first, before I put anything out on the bridge."*

Additional Findings

The concept had value, but adoption depended on clarity, trust, and discoverability

The assistant was not naturally discoverable

Users defaulted to familiar behaviors, such as scanning chat, listening to the call, or asking others for context.

Prompts did not fully explain capabilities

Users still had questions about what the assistant could access, what data it used, and how prompts differed from one another.

Trust had to be earned

Users were open to AI support, but wanted citations, source visibility, and ways to verify outputs before relying on them.

Static prompts were not enough

Users expected prompts to adapt based on their role, incident phase, and immediate needs.

Solution Direction

The research pointed to a private-first, trust-supported assistant experience

The recommended experience shifted from "an assistant users can find" to an embedded experience that joins them automatically to the AI chat experience.

1. **Private by default**
   Investigate without adding noise to the live incident conversation.

2. **Intentional sharing**
   Users choose what to share back to the group.

3. **Visible at entry**
   Surface the assistant when users join an active incident.

4. **Capability-signaling prompts**
   Prompts communicate timeline, impact, ownership, status, and related incidents.

5. **Verifiable outputs**
   Responses include sources, citations, reasoning, and distinctions between inferred and verified information.

Measuring Success

Success should be measured by adoption, orientation, trust, and reduced workflow friction

Adoption metrics

- Assistant discovery rate
- Prompt engagement rate
- Private chat usage rate
- Share-to-group usage rate

Workflow metrics

- Time to first useful orientation action
- Reduction in repeated catch-up questions
- Reduction in context switching
- Self-reported status clarity

Trust metrics

- Source or citation click-through rate
- Outputs confirmed, dismissed, or refined
- User-rated confidence
- Willingness to use again

Validation plan: Follow-up concept testing should compare entry points, privacy cues, static prompts, and role-aware or phase-aware prompts.

Outcome and Impact

The research clarified what the assistant needed to be useful, usable, and trusted

The study showed that the assistant's value was not just faster answers. Its value was helping responders safely orient, investigate privately, and decide what information was worth sharing during a high-pressure incident.

Research impact

- Prioritized private chat as a core experience
- Reframed sharing as user-controlled, not automatic
- Identified discoverability as a major adoption risk
- Defined prompt clarity as critical to capability understanding
- Established trust requirements for AI-generated outputs
- Recommended role-aware and phase-aware prompts as a future direction

Key takeaway

For AI to work in high-pressure incident response, it needs to fit the workflow, reduce risk, stay private by default, and make outputs easy to verify.