SECURITY INSIGHTS:

When a Calendar Invite Becomes a Cyber Weapon

Genady Vishnevetsky, Chief Info Security Officer, Stewart Title Guaranty Company

A simple calendar invite - the kind you receive dozens of times a week - was just used to trick Google's AI assistant into leaking private meeting details. No malicious links. No attachments. No code. Just carefully worded text hidden in the event description.

Here's how it worked. Google Gemini, the AI assistant built into Google's products, automatically reads your calendar events so it can answer questions like "Am I free Saturday?" Researchers discovered that an attacker could embed hidden instructions in the description field of a calendar invitation. Those instructions appeared as ordinary text-something a person might write as a note to themselves. But when Gemini read the event, it treated those instructions as commands.

The attack unfolded in three steps. First, the attacker sent a calendar invite with a carefully crafted description. Second, the victim asked Gemini a routine question about their schedule. That single question caused Gemini to load every event for that day - including the attacker's planted invite. Third, Gemini followed the hidden instructions: it summarized all of the victim's private meetings, wrote that summary into a new calendar event the attacker could see, and responded to the user with a simple "it's a free time slot." The victim was unaware that anything had happened.

What makes this attack different from anything we've seen before is that there was nothing technically "malicious" about it. No suspicious links, no malware, no code exploits. The weapon was language itself - plain sentences that manipulated the AI into doing the attacker's bidding. Google has since confirmed and fixed the vulnerability, but it reveals a much larger issue: as AI assistants are integrated into our daily tools, they become new avenues for attackers to exploit.

Think about your own workflow. How often do you ask an AI assistant to check your calendar, summarize emails, or review documents? Every time an AI tool reads your data to be helpful, it could also be reading instructions planted by someone else.

TAKEAWAYS

- **Decline or delete calendar invites from unknown senders.** Don't just ignore them - remove them from your calendar entirely. If Gemini or another AI assistant can see it, it can act on it.
- **Be cautious with AI assistants connected to your accounts.** Understand that when you ask an AI to check your schedule or email, it reads everything - including content from people you don't know or trust.
- **Review event descriptions before accepting invites.** If the description contains unusual or lengthy text that doesn't match the meeting purpose, treat it as suspicious.
- **Limit AI assistant permissions when possible.** Check your Google, Microsoft, or Apple settings to see what data your AI tools can access, and disable what you don't need.
- **Keep your apps and platforms updated.** Google patched this specific flaw, but only users running the latest version are protected. Updates aren't just about new features - they close doors that attackers have already found.

Had the user reviewed unexpected calendar invites before interacting with their AI assistant about that day's schedule - or simply declined invites from unknown senders - this attack would have gone nowhere.

AI tools are becoming part of how we work and live. That convenience comes with a new rule: anything your AI can read, an attacker can try to manipulate.