STORY: AI Threat to Critical Dam Infrastructure

10 a.m., Stony Brook

At the Stony Brook Cancer Center, surgeon Brice Radclyffe was scrubbed and ready for the operation. He knew his patient was the wife of a state senator, but he'd operated on many important people before, and as long as there were no unexpected hiccups, it all should be fine. He had a brief moment of gratitude that his own wife, Glenda, was safe and sound at home in White Plains babysitting their grandson, Caleb.

But he had to focus on this patient and not on family matters—though when he spoke with his son Zach last night, Zach was stressed over the new AI control system they were installing at the Kensico Dam where he worked as an engineer. Well, it would all work out. It always did.

11 a.m., Kensico Reservoir*

Zach Radclyffe stared at the computer screen in front of him. They'd had the new AI software, HydroAssist, up and running for a month now, but Zach still had some lingering doubts. The process had seemed rushed and staff training time too skimpy—what staff were left, that is, after management fired half the team. The software was from an Indian company, but Zach wasn't as much worried about the company's place of origin as he was potential conflicts, like when that 300-million Mars lander crashed

because one team was working in metric and one in imperial units.

Still, everything seemed to be running smoothly at first, and since most everything was automated, engineers like Zach were mostly glorified computer babysitters. But just yesterday, Zach had noticed something odd: the overflow gauge alarm kept winking off and on. They'd had Short Duration High-Intensity Rainfall event two days ago, leading to a larger inflow of water into the reservoir, which threatened to overtop the dam, so an alarm like that should be taken seriously. But his supervisors determined it was just an alert-settings glitch, which HydroAssist would catch and fix. But if it didn't? The dam gates couldn't be opened in time to release the excess water, leading to overtopping and possible structural damage. Zach looked around at the mostly empty control room and turned back just in time to see the alarm flash once more.

12:15, White Plains

In White Plains, Zach's mother, Glenda Radclyffe, was standing in line at the grocery store checkout with her grandson Caleb when the lights flickered off. No illumination, except for what was streaming through the front doors, and no working cash registers. Glenda joked with the cashier, "Do you have an old-school adding machine in the back somewhere? Or a calculator?" The enterprising cashier used the calculator app on her otherwise-useless cellphone, adding, "I'm glad this thing is handy for something."

Everyone breathed a sigh of relief when the power flicked back on again. Glenda thought of her son at the dam's power plant, wondering if he'd noticed anything there, or

if it was just another transformer blowing out. They'd had a lot of rain lately, so that was probably it. Zach worked too hard, and was such a serious boy. He'd been so worried about some new program at work, and he'd tried to explain it to her, but she didn't under a word of it. Something about a lot of "ops"—Machine Learning Ops, DevOps, DataOps. She preferred "oops," anyway. Trusting everything to robots seemed just silly. But she trusted Zach. He'd take care of everything.

1 p.m., Stony Brook

Brice Radclyffe finished the operation and congratulated his team on a job well done. He headed for the break room for a snack, grateful the administration had nixed those "intelligent" vending machines. As if he needed a robot telling him what it thought he wanted. But that made him think of his son again, as he tried not to worry about Zach and the new AI system at the dam.

But did it really matter since AI was gunning for everyone's jobs, anyway? The hospital administration may have passed on those vending machines, but they were still very much interested in AI system for operating management—predicting surgical case durations, optimizing post-anesthesia care unit resource allocation, and the like.

That was tomorrow's problem, and it probably wouldn't come for the surgeons' jobs anytime soon. But Zach was having to deal with it today, forced to retrain on a new system that he feared hadn't been wrung out properly, and left without a lot of human backup after the recent post-AI engineer "cleanse." Hopefully, everything was fine, and

Zach could console himself with the big Thanksgiving dinner Glenda had planned at home.

2 p.m., Kensico Reservoir

North of White Plains, Zach Radclyffe was frantic because the alarms were constant now, and it wasn't just a phantom malfunction—the valve mechanisms kept opening and shutting after getting conflicting commands. If that kept up, they might fail altogether and lead to a spillway cascade, sending signals to the dam to release too much water at once. He'd worried about having the systems connected to the internet. Had they been hacked? Or was that much-vaunted HydroAssist hallucinating? The software's manual hadn't mentioned anything about that, but Zach had read about cases where an AI perceived nonexistent patterns and created inaccurate outputs.

There were several alarms now, and Zach frantically tried to get a supervisor on the phone. He finally got Jim Whitmer, although Whitmer was skeptical at first—until he heard all the sirens in the background. He said he'd hurry on down to the site, but just as he said those words, the power winked off in the control room.

If the backup power didn't hold, the water from the Kensico Reservoir could become a torrent heading downstream, its thirty billion gallons of water causing floods up to seventy feet in parts of White Plains—let alone disrupting New York City's water supply since ninety percent of it ran through the reservoir.

He checked the status monitor. The software seemed to be setting up a floodgate

release. That meant the AI algorithm was picking the wrong option—it should be allowing a controlled water rate that would minimize any loss of human life, but instead, it was signaling a full left-gate principal spillway release only. That would likely push more water away from a country club and college, due to the area topography and bends in the terrain, but make it more likely to flood the more heavily populated area downriver like White Plains. Why would it do that? The only reason Zach could think of was HydroAssist had decided to minimize damage to economic interests, not people.

Zach wasn't waiting for his boss, he was sending out a warning to local governments via the disaster protocols in place for just such an occasion. But would it get there in time? Would the warning sirens work if there was a power outage due to another AI algorithm glitch or it battery backups failed? He was worried about his mother and son in the potential path of a flood. He tried to call her, but there was no answer. "Come on, Mom, come on. Please pick up." But his efforts to reach her failed.

2:05 p.m., White Plains

Glenda Radclyffe piled all the groceries into the back of her car and bundled little Caleb into his child seat, strapping him in tight. She picked up her cellphone to check on the traffic warnings, but then she had to laugh. She'd forgotten her cellphone at home to charge after the battery ran down. Oh well, she'd just have to hope there weren't any accidents blocking her path.

Finally, safe and sound, she returned home with Caleb and settled in for the

evening, enjoying the extra time with her grandson. Maybe Brice would even be home early from the hospital? She glanced over at her phone in the charger and noticed for the first time the light was flashing, indicating a message. Well, it would have to wait, since she was in the middle of something more important.

As Glenda and Caleb played “Go Fish,” a sudden uncontrolled release of impounded water began to cascade over the Kensico Dam spillway. The waters from the reservoir would soon be an unstoppable tsunami heading directly for White Plains.

*The names and situations in this story are fictitious and present one possible scenario in a representative composite facility, but are not meant to impugn any one specific individual, utility, or organization.