Podcast Name: *ACM ByteCast*
Episode: 87. Kelly Shortridge

Welcome to the *ACM ByteCast* podcast, a series from the Association for Computing Machinery! The podcast features conversations with researchers, practitioners, and innovators at the intersection of computing research and practice about their experiences, lessons learned, and visions for the future of computing. In this episode, host Scott Hanselmen is joined by Kelly Shortridge, Chief Product Officer at Fastly.

First, Kelly explains that security chaos engineering is fundamentally about resilience engineering—building systems that can recover quickly from inevitable failures rather than assuming every failure can be prevented. She argues that no system is perfectly secure because complex systems are vulnerable to unexpected failures. Instead of focusing solely on prevention, organizations should design for adaptability, redundancy, and recovery. Software has unique advantages over physical systems because engineers can duplicate production environments, simulate failures safely, and test how systems respond before incidents occur.

The conversation also explores how organizations can realistically measure and improve resilience. Kelly criticizes "metrics theater" and argues that teams claiming complete control over their environments are often the least prepared for unexpected events. Instead, security teams should embrace uncertainty by running small, low-risk chaos experiments that test assumptions, such as removing cookies or authentication headers from duplicated requests to see whether applications behave as expected. She encourages greater collaboration between security and platform engineering teams to validate shared assumptions. Finally, discussing large language models, Shortridge says they can be useful for generating tests, configurations, and repetitive tasks, but they should augment rather than replace human judgment, since resilience ultimately depends on thoughtful decision-making rather than automated responses.

Kelly argues that LLMs should augment human expertise rather than replace it, handling repetitive tasks such as reviewing compliance documents, extracting information, and assisting with brainstorming while leaving critical judgment to experienced engineers. She believes AI can make security more accessible, but cautions against relying on it for decisions that require context and expertise, such as evaluating whether a system is truly resilient. They also discuss using LLMs as "rubber ducks" for problem-solving and for challenging assumptions. The conversation then shifts to the shortcomings of compliance and security metrics. Kelly explains that regulations and compliance checklists often become outdated, encouraging organizations to satisfy auditors instead of improving actual security. Resilience requires continually updating mental models rather than merely patching vulnerabilities or following legacy processes.Finally, they discuss the shared responsibility model in cloud security, with Shortridge explaining that vendors should secure the underlying platform while customers remain responsible for the security of their own applications and dependencies, since organizations ultimately must own the risks introduced by the software and services they choose to use.

Links:
Learn more about Kelly Shortridge.

Learn more about the ACM ByteCast podcast at https://learning.acm.org/bytecast.