EACG
Enterprise Architecture Consulting Group

SECURITY DIALOGUES
No. 11 · 16 October 2026
Marcus Hale (PSIRT) · Tess Calder (red team) · Jessy Lang

The Workshop That Always Happens

Eight reasons threat models slip, and how to keep one current

A threat model made in a workshop is a photo of a whiteboard: rarely made, built from memory and out of date by the next sprint. The episode looks at what changes when the thorough version simply happens every time, built from what is actually there, with the judgment staying exactly where it belongs.

DEFAULT DILIGENCE
The thorough version happens unless somebody decides to skip it.

1. Why is this threat on the list?

2. What changed since the last run?

3. Who accepted this risk?

WHAT GOES WRONG AND WHAT TO DO ABOUT IT

1	**Done rarely** Eight people for four hours is expensive, so it happens once, if at all.	Make building the model cost very little each time. Then the thorough version becomes what happens by default, and skipping it is the thing somebody has to decide.
2	**Done from memory** The room only knows what the room remembers.	Build the model from four sources and merge them before the first STRIDE question: the SBOM (what we brought in, with its CVEs), SAST results (what we wrote, as CWEs), the OBOM (what really runs) and the deployment structure (what can reach what).
3	**It decays** The product keeps changing and the photo does not.	Keep the four inputs current: SBOM with each build, SAST in the pipeline, OBOM from what is running, the deployment description whenever it changes. Rebuild the model when they change. The question becomes what changed since the last run.
4	**A list is not a model** A list says what is wrong with the parts, not how the whole could be abused.	Model the wiring with the parts. The same CVE weighs differently in an internet-facing service than in a build-time container. Then ask the six STRIDE questions about every part and every connection.
5	**Threats without evidence** A threat that has to be believed instead of checked.	Ask of any tool: why is this threat on the list? The answer must point to a component, a finding or a connection. A countermeasure names what to change, where, and why it reduces that threat.
6	**Judgment handed to the tool** The agent can propose, but it cannot know what you can afford to lose.	The agent does the exhaustive part. You decide how much each risk matters and whether to accept, fix or work around it, based on asset value, risk appetite and budget. A regulator wants a name, not a model version.
7	**What tools cannot see** Business logic, clever abuse of a feature, people being talked into things.	Expect blanks and want them visible: a good model shows no SAST or OBOM for a service instead of pretending it is clean. Repudiation needs a person to confirm that records exist, so a good model flags it as a question.
8	**Missing inputs** Nobody has a complete set of inputs on day one.	Treat the first run as an inventory audit and a list of blind spots. The second run is better than the first. The Cyber Resilience Act expects a documented risk assessment, kept current through the support life.

REVIEWING THE RESULT IN AN HOUR

Check the picture
Is the exposure right? Does the model call something internal that is reachable from outside? A five-minute check by whoever knows the deployment.

Find blanks, pick what matters
Services without static analysis or OBOM show where the model is weakest. Put your own asset values on it: which threats touch your tier-one list?

Decide and keep
Each countermeasure gets an owner, a size and a decision: do it now, schedule it, or accept the risk. Keep the result, the next run is compared against it.

Listen to the episode
The Workshop That Always Happens and the rest of the series.
eacg.de
/en/podcast/s11-default-diligence

Digital Product Design
From the first idea to a viable first release.
eacg.de
/en/consulting/digital-products

PSIRT Services
Intake, assess and disclose vulnerabilities, built or run as a service.
eacg.de
/en/consulting/psirt

© 2026 EACG GmbH