Operational Security Rules That Actually Work

A Practical Security Program for Mid-Sized Companies

Derek Carson · carson.de · derek@carson.de · November 2026

What If Your Employees Do Not Recognize the Attack?

A perfectly-polished phishing email. A phone call using the CEO’s cloned voice. An urgent request to change a supplier’s bank account details.

Such attacks routinely succeed. According to the 2026 Bitkom Economic Protection Report, German companies estimate their losses from theft, espionage, and sabotage at 211 to 271 billion euros; 76 percent of this is attributed to cyberattacks. One quarter of the 1,003 surveyed companies suffered losses from ransomware within the last twelve months, and 16 percent suffered losses from phishing.

Even well-trained employees can fall for such attacks. This is not an individual failure, but rather a limitation of any training that relies solely on recognizing threats.

Therefore, a different question is crucial: What happens when someone is deceived? Are there rules and procedures in place to prevent an error from resulting in damage? Can employees easily report a suspicion? And can it be verified whether these rules are actually followed in everyday operations?

Using a manufacturing company as an example, this paper demonstrates how selected information security risks can be translated into concrete operational rules, appropriate measures, and suitable training—and how their implementation can be verified.

Why a Sign-in Sheet Is Not Enough

The German BSI Act (BSIG) obliges the management bodies of essential and important entities to participate in regular training sessions (Section 38(3) BSIG). Furthermore, ISO/IEC 27001 (Clauses 7.2 and 7.3) requires organizations to ensure that employees are competent, aware of the information security policy, and understand their contribution to ISMS effectiveness.

The simplest proof of both is an attendance list. It proves who was in the room and nothing else; it does not state whether even a single person acts differently afterward. However, damage is caused by actions. A phishing email only causes harm when someone clicks a link or types in credentials; payment fraud only succeeds when someone changes bank details without verification or approves a transfer.

What Studies Show About Standard Training

At UC San Diego Health, a large US clinical network with approximately 19,500 employees, an eight-month study investigated whether standard training changes this behavior (Ho et al., IEEE Symposium on Security and Privacy 2025). Participation in annual mandatory information security training did not lead to safer handling of simulated phishing emails. Furthermore, displaying short training modules immediately after an erroneous click reduced future error rates by an average of only 1.7 percentage points; more than half of these remedial modules were dismissed in under ten seconds.

A second study ran for 15 months in an unnamed large company with more than 14,000 employees (Lain, Kostiainen, and Čapkun, ETH Zurich, IEEE Symposium on Security and Privacy 2022). This study also showed that training notices displayed following an erroneous click did not make employees more resilient, and in some cases even made them more vulnerable.

Both studies demonstrate why certain training measures have little effect: they assume that employees recognize a threat and react correctly at the critical moment. The annual mandatory training is intended to impart the necessary knowledge beforehand. The notice following an erroneous click is intended to correct behavior immediately afterward. However, neither succeeds reliably.

In the second study, a different measure proved more successful: adding a reporting button to the mail client led to a sharply increased probability that simulated phishing emails were reported to IT security. To do this, employees did not have to judge for themselves whether a message was actually a phishing attempt. They could simply report a suspicious email; IT security checked the case and was able to identify new attacks early. Reporting thus became a simple action that did not require a reliable assessment of the threat.

An important consequence for the design of a security program emerges from this: it is not enough to impart knowledge and hope that employees will apply it correctly in an emergency. More effective are clearly defined rules and procedures that function even when someone does not recognize a threat or is uncertain. This leads to three principles:

- **Default secure pathways:** Procedures where secure behavior is the normal and simplest path.
- **Low-friction reporting:** Mechanisms that enable simple, constructive action in case of doubt, thereby reducing errors in judgment.
- **Operational metrics:** Measurements that show what is actually happening in operations rather than merely documenting training attendance.

The cited studies examine phishing. However, the insights gained can be applied as design principles to other areas where employees must make decisions under uncertainty or time pressure—such as payment fraud, phone-based social engineering, or the use of AI tools. There, too, clear rules and established procedures can prevent security from depending solely on the correct assessment of individual employees.

What This Concept Achieves—and What It Does Not

Corporate information security research is extensive and shows reproducible findings: certain interventions work measurably, while others do not. Management and information security officers can rely on this evidence base to separate achievable program goals from unattainable ones. This paper draws this boundary explicitly and describes a security program focused on levers built upon documented evidence and best practices.

It differentiates between various levers that are often conflated in daily operations: awareness (knowing about the threat), training (practicing a specific action), organizational rules, operational procedures, technical mechanisms, and management decisions.

What it does not achieve:

- Reducing the error rate on suspicious emails. No training achieves this reliably; technical filtering, however, never reduces it to zero.
- Improving employees' individual recognition capabilities. A well-crafted attack fools even trained eyes.
- Replacing technical measures (mail gateways, phishing-resistant MFA, EDR, isolated browser environments). These remain the responsibility of IT.
- Providing protection against all attacks. The "resistant core" of repeat clickers requires operational access controls, not additional training.

What it achieves:

- **Standardized safeguards:** Embedding procedures where secure behavior is the standard path—such as calling back before changing bank details, requiring dual-person approval for payments, or prohibiting payment approval solely based on a phone call. Measurable via spot checks in operations.
- **Actionable reporting:** Establishing a functioning reporting channel for suspicious emails so the company recognizes new attack campaigns early. Measurable via the reporting rate.
- **Minimizing damage per click:** If the callback rule and the rule "no payment approval by phone alone" are followed, a BEC attack via email will not lead to an erroneous transfer, and a deepfake call will hit a dead end.
- **NIS-2 documentation:** Providing documentation that can be used to demonstrate compliance and support liability release arguments for management (Section 38(2) BSIG).

The Working Example: Muster Antriebstechnik GmbH

Muster Antriebstechnik GmbH in the working example is a fictional mechanical engineering company in East Hesse with 250 employees. That the company falls under NIS-2 is an assumption of the example and not a legal classification. The advantage of a fictional company is that all documentation can be shown in full depth without disclosing a customer's confidential data.

From Risk to Rule

Every rule in the program stems from a risk assessed by the company, not from a catalog of topics. For example, if risk assessment classifies payment fraud via forged supplier emails as high, a rule for accounting follows: sensitive data such as bank details and salary records are only modified when confirmed via already stored contact details, documented, and approved by a second person. Training practices this rule through role-playing, and an operational spot check verifies compliance every quarter.

Four Target Groups, Four Formats

A machine operator does not have a PC, but has a brief moment before a shift begins. Management must make strategic decisions and therefore needs practice in decision-making, not e-learning. The concept divides the program into four tailored modules:

- **Management (M1):** A 90-minute briefing in Q1, and a 60-minute decision-making exercise in Q4. Focus: weighing risks under pressure and enforcing the rule that payment instructions are never approved based solely on a phone call.
- **Office Workstations (M2a):** A 60-minute training session and four short modules distributed throughout the year. Focus: reporting suspicious emails and using only officially approved AI tools.
- **Production and Logistics (M2b):** 15-minute shift briefings led by supervisors, supported by concise reference sheets. Focus: zero unauthorized USB devices on production machinery.
- **Finance, HR & Purchasing (M3):** A 150-minute workshop with role-playing exercises in groups of up to five people. Sensitive data such as bank details and salary records are only modified when confirmed via stored contact details, documented, and approved by a second person.

Rules Instead of Recognition

With AI, a voice on the phone can be synthesized, and the displayed phone number can be spoofed. Anyone tasked with spotting a forgery is at a disadvantage. Instead, the working example relies on a rule that protects even if no one recognizes the forgery: a payment instruction from management is never valid solely via phone call, voice message, messenger, or video conference. Employees confirm it using the number in the directory that they dial themselves.

The rule only works if employees feel confident asking management for clarification. Therefore, management announces it directly: “We hold ourselves to this rule as well. When we instruct a payment, we expect your callback.” No employee suffers a disadvantage for verifying, even if the instruction was genuine.

Reporting Without Disadvantage

Employees will only report what they notice if they do not fear negative consequences. Therefore, all modules contain the same pledge: “Anyone who reports something, including their own mistake, faces no disadvantage.” Management must explicitly confirm and honor this promise in everyday operations.

Testing Under Operational Pressure

In the fourth quarter, management, IT leadership, production management, and the information security officer simulate a cyberattack where malware halts production. The scenario reflects real organizational decisions—if a planned measure was implemented, the attacker tries another path; if skipped, the vulnerability remains.

32 minutes in, a caller claiming to be a managing director calls and demands an urgent payment. This tests whether the rules hold up under stress.

Measuring What Matters

Each module is checked through two distinct mechanisms:

- **Learning Objective Verification:** The trainer checks against predefined criteria whether guidelines are applied (e.g., is the bank account updated only after a callback via the stored number?). Failed learning objectives require repetition with new case scenarios.
- **Operational KPIs:** The company draws a random quarterly sample to verify operational compliance—confirming that verification, approval, and modification occur in that exact sequence.

If learning checks pass but the operational KPI remains low, the issue is structural or technical—such as a software workflow that permits changes without secondary approval—rather than a lack of employee knowledge.

Data Protection and Co-Determination

Because employee performance is measured, data protection and works council requirements are integrated from the start: store results solely as pass/fail, evaluate data only in groups of ten or more, and ensure supervisors do not grade their direct reports. Phishing exercise clicks have no employment law consequences.

The Full Working Example

The full working example spans 165 pages, comprising the training concept and all documents for modules M1, M2a, M2b, and M3. A 23-page excerpt showing the core concept and sample module documents is available upon request: Derek Carson, derek@carson.de.