Plug the Holes in the Swiss Cheese Model

Todd Stauffer, P.E.
exida
Nicholas P. Sands, P.E.
DuPont Protection Solutions
David Strobhar, P.E.
Beville Engineering, Inc.

Stop using operator error as an excuse. Apply human factors considerations to improve your alarm system and help operators respond to alarms effectively.

Alarms play a significant role in maintaining plant safety by notifying operators of an equipment malfunction, process deviation, or abnormal condition that requires a timely response *(1)*. Alarms are one of the first layers of protection for preventing a hazard from escalating to an incident or accident. They work in conjunction with other independent protection layers (IPLs) such as relief valves, dikes, and safety instrumented systems (SIS) (Figure 1) *(2)*.

Operator response to an alarm can have numerous failure modes related to hardware, software, or human behavior. Failures in human behavior are more likely when alarm system design and performance is poor (*e.g.,* nuisance alarms, stale alarms, redundant alarms, and alarm floods). Failures arising from the design of the alarm system are often incorrectly labeled as operator error; this type of failure can be more appropriately characterized as alarm management failure.

This article covers two techniques for maximizing the benefit of operator response to alarms and minimizing the risk of failure:

- follow the best practices of the alarm management standards from the International Society of Automation (ISA) and the International Electrotechnical Commission (IEC), ISA-18.2 and IEC 62682, and guidelines from the Engineering Equipment and Materials Users Association (EEMUA), EEMUA 191
- apply human factors best practices to change operator behavior and improve operator response.

Understanding the problem

The Swiss cheese model. Investigations have revealed that most industrial incidents include multiple independent failures. Imagine each layer of protection as a slice of Swiss cheese *(3)*, with the holes representing vulnerabilities to failure (Figure 2). For an incident to occur, the holes in the slices of cheese must align. The size (area) of the holes in the cheese is proportional to the reliability of the layer of protection. A slice with a large hole area, comprised of many holes and/or big holes, has a high probability of failure. To improve performance of the protection layer, reduce the area of the holes and ensure the holes in adjacent slices do not align.

The typical reliability used in a process hazard analysis (PHA) for an operator response to an alarm is 0.9 (probability of failure on demand [PFD] = 0.1), which assumes that the action is simple and well-documented and there are clear and reliable indications that the action is required *(4)*. Applying the Swiss cheese model, the area of the holes would be 10% of the total area of the slice. A poorly performing layer of protection (unreliable) would have a hole area greater than 10%, easily approaching upwards of 50%, in which case the layer would no longer be considered an IPL.

Operator response model. Failure modes for operator response to an alarm can be evaluated using a simple operator response model, which consists of three steps:

- the operator *detects* the deviation
- the operator uses knowledge and skills to interpret the alarm, *diagnose* the problem, and determine a corrective action
- the operator *responds* with the corrective action necessary to rectify the situation *(1)*.

This article is based on a paper presented at the AIChE 2017 Spring Meeting and 13th Global Congress on Process Safety, San Antonio, TX, Mar. 26–29, 2017.

A study of 11 vessel overflow incidents identified common failure mechanisms in the detect, diagnose, and respond steps (Table 1) *(5)*. Improper outputs in the respond step — either no output or the incorrect output — is often referred to as operator error. Failures can occur in all three steps, but most failures are caused by errors in detection and diagnosis. In many cases, the operator failed to notice the problem (detection) or the operator incorrectly identified the cause and thus applied an incorrect action (diagnosis). Instances in which the operator knew what to do, but performed the incorrect action, such as turning the wrong valve, are much less common *(6)*. Issues such as nuisance alarms, alarm floods, poor human-machine interface (HMI) design, and insufficient training set up operators for failure.

Situation awareness. Outputs in the operator response model are impacted by situation awareness (SA), which is the level of awareness operators have to the events occurring in their environment and their understanding of the meaning and impact of those events now and in the future *(7)*. An operator's SA is impacted by their mental model, which is the cognitive tool that helps a person make sense of a situation by combining disparate pieces of information, interpreting significance, and developing a reasonable projection of the future *(7)*.

Good SA drives effective decision-making and performance, but it can be undermined by various factors, dubbed SA demons *(5, 7)*:

- attention tunneling — focusing on one area or issue to an extent that alarms from another area or issue are ignored
- misplaced salience — incorrect alarm priority or HMI representation of alarm importance and other status information
- errant mental models — incorrect interpretation of alarms or mistakenly ignoring relevant alarms.

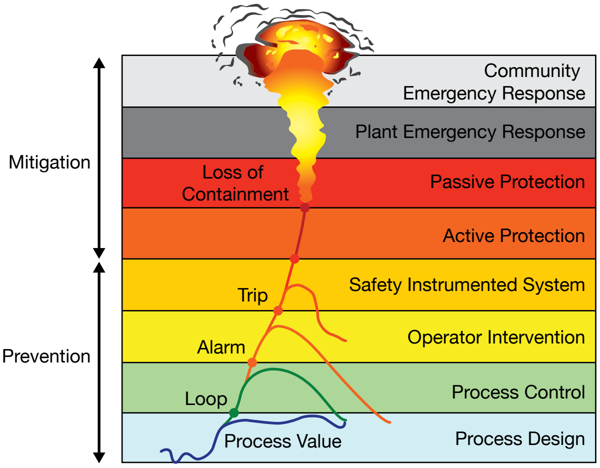

▲ **Figure 1.** Alarms indicate that operator intervention is necessary to prevent an incident from propagating through the various layers of protection designed to prevent accidents *(2)*.

Addressing the problem: Alarm management principles

Operator response can be improved by following the guidelines and alarm management lifecycle defined in ISA-18.2 and IEC 62682. The July 2012 *CEP* article "Implement an Effective Alarm Management Program" *(8)* provides a detailed discussion of these standards, but a few key elements are worth summarizing here in the context of human factors.

Alarm rationalization. To maximize dependability, the operator must believe that every alarm is valid and requires a response. Alarm rationalization is the process of reviewing, validating, and justifying alarms to ensure every alarm meets a set of criteria. This helps to improve operator trust in the alarm system and documents the cause, consequence, corrective action, and time to respond for each alarm.

Alarm prioritization. Priority indicates criticality and helps the operator understand the relative importance of each alarm. Alarms should be prioritized based on the severity of the potential consequences and the time available for the operator to respond.

Alarm classification. Classification is the process of categorizing alarms based on common requirements (*e.g.*, testing, training, monitoring, and auditing). Alarms that need higher reliability, for example, require more management.

Alarm response procedures. It is critical that operators know what to do in the event of an alarm to ensure that their response is dependable. The alarm response

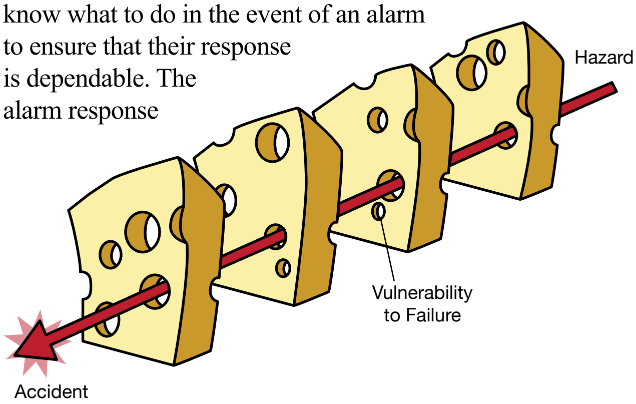

▲ **Figure 2.** The Swiss cheese model depicts layers of protection as slices of cheese and vulnerabilities to failure as holes *(9)*.

Table 1. Each step in the operator response has vulnerabilities to failure *(5)*.

Detect	Diagnose	Respond
Nuisance alarms	Insufficient training	No action (error of omission)
Alarm floods	Errant mental models	Untimely action
Alarm overload	Attention tunneling	Incorrect sequence of actions
Poor human-machine interface design	Ignoring alarms	Incorrect action (error of commission)
	Incorrect diagnosis	

procedure, which contains key information documented during alarm rationalization, can reduce the time to diagnose the problem and determine the appropriate corrective action.

Alarm design. The detailed design of alarms has a significant impact on the reliability of an IPL (*i.e.,* minimizing the area of the holes in the Swiss cheese). Effective alarm design includes:

- configuring alarm attributes that impact alarm behavior (*e.g.,* deadband and on/off delay to prevent alarms from repeating excessively in a short time interval)
- suppressing the alarm when the alarm is not relevant (*e.g.,* prevent an alarm flood after a compressor trip)
- annunciating the alarm (*e.g.,* field horns and lights to indicate the need to evacuate an area due to a severe hazard).

Monitoring and assessing alarm system performance. Monitoring and assessing the performance of an alarm system helps to determine whether the reliability of the associated IPL is improving or degrading. An evaluation should look at the overall alarm system and individual alarms that are used as protection layers.

Addressing the problem: Human factors

Human factors have become increasingly important because evolving technology and automated control systems have expanded operator responsibility. How well a person performs a task cannot be attributed to a single factor, which makes applying human factors principles challenging. Deficiencies are often not addressed in the most direct or appropriate manner, and in some cases, the resolution makes the problem worse. For example, if an operator ignores a nuisance alarm, the consequences may include disciplinary action, additional training, more alarms, or longer procedures, instead of resolving the root cause of the nuisance alarm. The resolution to a performance issue is rarely simple or absolute *(6)*.

Operator response to alarms can be analyzed by signal detection theory, which quantifies an operator's ability to discern between useful patterns that provide information and random patterns that distract from necessary information (*i.e.,* noise). Chattering alarms, standing alarms, and alarm floods are examples of visual noise that inhibit an operator's ability to detect the alarm signal. According to signal detection theory, as noise increases, the operator's ability to discriminate a true alarm from a false alarm decreases.

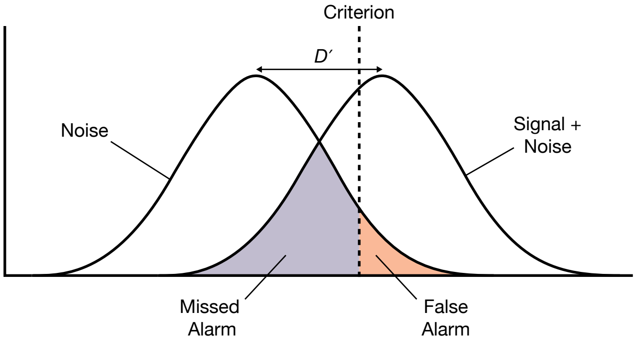

▲ **Figure 3.** As the operator becomes less tolerant of false alarms, the criterion for action moves to the right. The area under the curve representing missed alarms increases, while the area under the curve for false alarms decreases.

All judgments must be made in an environment of uncertainty. All signals indicating an event have some degree of noise, and that signal is then conveyed to the operator in an environment with its own degree of noise (Figure 3). The larger the difference (*D′*) between the nature of the noise in the environment and the characteristics of the signal, the more likely that the signal will be detected. As the environmental noise (audible and/or visual) more closely matches the signal, *D′* approaches zero and the more likely it becomes that the signal will be missed.

The difference between the noise and the signal is not the only factor that affects the detection of a signal. Operators, serving as signal detectors, create their own criteria for what signal to accept as true or valid, typically based on past history. And, these criteria vary from operator to operator.

A tradeoff exists between the probability of getting a false alarm and missing an alarm. If the system designer prioritizes reducing the number of false alarms, then the likelihood of missing a valid alarm increases. Ensuring large differences between alarm signals and noise is important not only for signal detection, but also for operators to establish their criteria for detection.

Nuisance alarms

A nuisance alarm (or false alarm) is an alarm that annunciates excessively or unnecessarily, or that does not return to normal after the correct response is taken (*e.g.,* chattering, fleeting, or stale alarms) *(1)*. A false alarm can occur because the condition is not true or when no action is needed on the part of the operator. It does not take a high nuisance alarm rate for the operator to doubt the veracity of the alarm system. A 25% false-alarm rate is enough for operators to stop relying on the system for detecting an abnormal event.

Reluctance to respond immediately to a system that produces many false alarms is a rational behavior; responding takes time and attention away from other important tasks *(7)*. Thus, ignoring nuisance alarms is not a behavior that can be changed by training or disciplinary measures. The best way to eliminate this unwanted behavior is to eliminate the nuisance alarms.

Alarm rationalization helps to increase operators' reliance on the alarm system. As redundant alarms are eliminated, operators begin to realize that alarms are true and that they must pay attention when an alarm actuates.

Well-designed HMI displays help to support situation awareness and verify true alarms. Alarm summary displays offer only minimal assistance because they do not include the variables related to the alarm condition. For example, a high overhead temperature alarm on a distillation column requires the operator to understand the other column temperatures, reflux flow, and reboil temperature to determine whether the alarm condition is true. Visual displays should do more than just indicate an active alarm; they should help the operator confirm that the alarm is true *(7)*.

Operators working with a well-designed alarm system will still require some training to identify valid alarm conditions. Rarely do process variables fluctuate in isolation — thus, confirming a true alarm is a key operator skill. Training should teach operators how to cross-check alarms with other process variables to determine their validity. For example, if a low-flow alarm occurs, the operator could check the level in the upstream or downstream vessels.

Training that addresses human factors limitations can help operators develop a different way of thinking about and responding to alarms. Training should teach operators to:

- believe that the alarm (indication) is real, and look for confirmation to validate their mental model of the situation
- cross-check the alarm with other process variables to confirm that it is true
- be careful when discounting an alarm without corroborating evidence
- challenge themselves when ruling out possible explanations to ensure that they are applying the best mental model
- beware of factors that increase errors, such as lack of sleep, high stress, and long work hours.

Alarm response procedures

A general guideline for alarm rationalization is to not alarm the normal or expected situation because it does not provide useful information to the operator. A similar principle can be applied to alarm response procedures. The operator should not receive information that would be obvious to a trained operator because that just creates noise.

The entire output of the alarm rationalization is useful for training purposes, but only a subset of that information should be provided to operators as a real-time decision aid. Alarm response procedures should highlight what is unique about the particular alarm, either in cause, response, or consequence of inaction. The procedure should also detail how to confirm that the alarm is real.

Response procedures are an aid for the operator, not a replacement for good operating procedures. Effective response procedures:

- include objective actions
- use action verbs (*e.g.,* start, stop)
- use simple and precise language (no technical jargon).

Operators should not receive information that is obvious, because that creates noise.

Alarm floods

Even a well-designed alarm system can generate a large number of alarms in response to a major process upset, such as a loss of power. Humans have a relatively limited capacity for processing information, so the potential for overload in this type of situation is high. However, several techniques can be employed to maximize the potential that alarms will be processed.

The easiest technique for handling alarm floods is to automatically suppress low-priority alarms for a set period of time. If the alarm prioritization has been done correctly, this will reduce the number of alarms by about 80%. This method has limited risk because low-priority alarms generally do not need to be responded to with the same urgency as high-priority alarms.

A more-complicated technique is state-based alarming, which prevents alarms that are expected to occur during an upset. For example, once a unit has shut down, alarms due to low energy (*e.g.,* low temperature, pressure, and flow) will likely return to normal. Alarms for these conditions are automatically suppressed when entering the low-energy state.

Alarms can also be aggregated into higher-order alarms for the entire system to reduce the number of alarms. For example, all alarms for a tower or set of towers operating in series can be combined so each does not have to be processed individually. Instead, the collective alarm can be evaluated qualitatively to determine the tower condition.

Attention tunneling

Because humans have a limited processing capability, operators can become overly focused on a task and miss other important events, *i.e.,* a loss of situation awareness and attention tunneling. A single display for the operator's entire span of responsibility can help prevent attention tunneling. The display should include the status of the alarm system, the units and areas with alarms, and the alarm priority and importance. This style of display will help alert operators when they need to switch attention to other areas or equipment.

Misplaced salience

Humans have trouble detecting changes — a phenomenon known as change blindness. Highlighting changes when they happen can help overcome change blindness. For example, when a new alarm is active it flashes until it is acknowledged to direct the operator's attention to the changed variable. When looking at a color display, it can be difficult to detect a value change indicated by a color variation, unless it is highlighted in an additional way.

The salience of information on an HMI should be related to its operational importance. Background information should be given low visibility, normal plant measurements medium visibility, and abnormal conditions (*e.g.,* alarms, values, and states) high visibility.

Normalization

Standing alarms create visual noise. Alarms can go unrecognized for extended periods of time if the summary display is clogged and alarms blend together visually. Standing alarms also create visual noise that makes it increasingly difficult to detect the occurrence of a signal.

Standing alarms can be addressed via rationalization and by varying the alarm's setpoint or suppression status based on state (*i.e.,* dynamic alarming).

Errant mental models

Mental models are an important mechanism for interpreting new information. Operators use mental models of the process, such as *if* the reactor feed flowrate is increased, *then* the reactor temperature will increase without more coolant.

Problems arise when operators use an incomplete or incorrect mental model. A key to using mental models is knowing when you are using the wrong one. Operators may misinterpret alarms or events as fitting into their current mental model, without realizing that cues indicate they should be using a different mental model. People tend to explain away cues that conflict with their current mental models (confirmation bias) and can be slow to notice the mistake.

Operators should be trained to develop multiple mental models for a situation to improve their response. A premortem strategy can help in the development of more and better mental models *(10)*. Premortem involves creating if-then scenarios to analyze how a process or operation could fail and discussing how to rectify the situation.

Experienced operators have more and richer models of plant operation developed from living through process upsets. These models need to be transferred to less-experienced operators. These models can help a new operator effectively start up or shut down equipment or change the equipment's mode of operation. Alarm rationalization can help identify where multiple models exist, and training disseminates these models across staff.

Closing thoughts

While it is important to follow the alarm management best practices in this article and other references (ISA-18.2 and IEC 62682), changing the operator's mindset and behavior is vital. Alarm management cannot make up for an operator who mistakenly thinks information or alarms are not real, eliminates potential causes too quickly, or exhibits confirmation bias when responding to a hazard. CEP

Literature Cited

1. **International Society of Automation,** "Management of Alarm Systems for the Process Industries," ANSI/ISA-18.2-2016 (June 2016).
2. **Stauffer, T., and P. Clarke,** "Using Alarms as a Layer of Protection," CCPS 8th Global Congress on Process Safety, Houston, TX (Apr. 2012).
3. **Reason, J.,** "Managing the Risks of Organizational Accidents," Ashgate Publishing, Burlington, VT (1997).
4. **Center for Chemical Process Safety,** "Guidelines for Safe and Reliable Instrumented Protective Systems," Wiley and CCPS, Hoboken, NJ (2007).
5. **Dunn, D. G., *et al.,*** "When Good Alarms Go Bad: Learning from Incidents," 70th Annual Instrumentation and Automation Symposium, Texas A&M Univ. (Jan. 2015).
6. **Strobhar, D.,** "Human Factors in Process Plant Operation," Momentum Press, New York, NY (2013).
7. **Endsley, M.,** "Designing for Situation Awareness: An Approach to User-Centered Design," CRC Press, Boca Raton, FL (2012).
8. **Stauffer, T.,** "Implement an Effective Alarm Management Program," *Chemical Engineering Progress,* **109** (7), pp. 19–27 (July 2012).
9. **Klein, G.,** "Sources of Power: How People Make Decisions," The MIT Press, Cambridge, MA (1998).
10. **Reason, J.,** "Human Error," Cambridge Univ. Press, Cambridge, U.K. (1990).

TODD STAUFFER, P.E., is responsible for marketing and business development for exida's alarm management products and services (80 N. Main St., Sellersville, PA 18960; Email: tstauffer@exida.com). He is an editor and voting member of the International Society of Automation (ISA)-18.2 standards committee on alarm management, cochair of ISA's WG3-Basic Alarm Design, and cochair of the ISA-84.91.03 standard on safety alarms, controls, and interlocks. Stauffer earned a BS in mechanical engineering from Pennsylvania State Univ. and an MS in mechanical engineering from the Univ. of Pennsylvania. He is a registered professional engineer in Pennsylvania.

NICHOLAS P. SANDS, P.E., is a senior manufacturing technology fellow in DuPont's Kevlar, Nomex, and Tyvek businesses and the global alarm management leader for DuPont (Email: nicholas.p.sands@dupont.com). He is an International Society of Automation (ISA) Fellow and a past ISA vice president of standards and practices. He is a cochair of the ISA-18 committee on alarm management, a director of the ISA-101 committee on human machine interface, and a director of the ISA-84 committee on safety instrumented systems. Sands earned his BS in chemical engineering from Virginia Tech.

DAVID STROBHAR, P.E., is the founder of Beville Engineering, Inc., which conducts human factors engineering analyses of plant modernization, operator workload, and alarm/display systems for major energy and chemical companies (Email: dstrobhar@beville.com). He has evaluated operator workload at more than 1,000 units at more than 80 different sites, and he has participated in more than 150 alarm rationalizations. He is a founder of the Center for Operator Performance — a collaboration between operating companies, distributed control system (DCS) suppliers, and academia that researches human factors issues in process control. Strobhar has a BS in human factors engineering from Wright State Univ. (Dayton, OH), where he now serves on the engineering advisory board. He is a registered professional engineer in Ohio.