"Closing the Holes in the Swiss Cheese Model" – Maximizing the Reliability of Operator Response to Alarms

White Paper

exida
80 N. Main St.
Sellersville, PA
www.exida.com

exida White Paper Library
http://www.exida.com/Resources/Whitepapers

Keywords: Conduct of Operations, Human Factors & Culture, Compliance with Standards, alarm management. Safety alarms, alarm rationalization, ISA-18.2, situation awareness, IPL Alarms, operator response to alarms, classification

Abstract

Layers of protection for abnormal event management can be modeled as slices of swiss cheese according to James Reason [1]. An operator's response to an alarm is one of the first layers of protection to prevent a hazard from escalating to an incident. This paper will present best practices for maximizing the operator's reliability for understanding and responding to abnormal situations as adapted from the alarm management standards ANSI/ISA-18.2-2016 and IEC 62682. Examples include alarm rationalization to ensure all alarms are meaningful and to capture "tribal knowledge", prioritization to help operators determine which alarms are most critical, and creation of alarm response procedures. The treatment of safety alarms, which are those that are deemed critical to process safety or to the protection of human life or the environment, will be specifically highlighted.

The paper will also discuss key human factors considerations for maximizing operator situation awareness (SA) by preventing SA "demons"; such as developing an errant mental model of the process, attention tunneling, data overload, and misplaced salience. As such the resolution of issues which inhibit operator performance, such as nuisance alarms and alarm floods, will also be discussed.

Conclusion

This paper has discussed alarm management and human factors techniques for improving an operator's response to an alarm; thus reducing the size and area of the holes in its "swiss cheese" layer and making it more reliable. These techniques can be applied to all alarms, but are particularly important for safety alarms. Some of the key takeaways are summarized below.

- Performing a thorough rationalization is important to ensure that each alarm is actionable and has a purpose.
- Prioritizing alarms based on consequences and time to respond helps the operator know which alarm to respond to first.
- Classification in (conjunction with prioritization) identifies safety alarms so that they can be managed and maintained appropriately to their risk reduction and can be displayed on dedicated HMIs or annunciated uniquely.
- Alarm system performance is monitored and assessed to ensure that operators are not being overloaded or flooded with too many alarms. Performance of individual safety alarms should be examined in more detail to ensure that they are functioning acceptably.
- Alarm response procedures, created from the results of rationalization, can be presented to the operator as a real-time decision aid.
- Alarms should be designed for reliability from the sensor to the HMI with an appropriate level of salience for quick and easy detection by the operator.
- A nuisance alarm rate of more than 25% can cause operators to ignore alarms or delay their response. Thus significant effort should be put into eliminating nuisance alarms and creating an environment where the operator can quickly and easily confirm the validity of an alarm.

- Training – first believe what the indication is, look for confirmation to prove your mental model…Be careful discounting an alarm without corroborating evidence. Instincts can lead you awry…Challenge when closing of possibilities…When I see an alarm I look for a confirmatory set of actions to confirm that it is real….Factors that correlate with making an error
- Operators should work on developing more and better mental models to provide alternate scenarios for dealing with plant upsets and should be wary of slipping into an attention tunneling episode.

While it is important to follow the alarm management best practices in this document and other references (e.g., ISA-18.2 / IEC 62682), the importance of changing the operator's mindset (behavior) cannot be overemphasized. As discussed, the majority of failures occur in the diagnosis step of the operator response model; thus addressing this failure becomes one of the key actions to improving reliability. No amount of alarm management can make up for operators who mistakenly think that information is not "real", eliminate potential causes too quickly, or exhibit confirmation bias when responding to prevent the escalation of a hazard.

Introduction

The purpose of an alarm is to notify the operator of an equipment malfunction, process deviation or abnormal condition that requires a timely response [2]. Alarms help the operator keep the process within normal operating conditions and play a significant role in maintaining plant safety. 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 IPLs such as relief valves, dikes, and safety instrumented systems (SIS), as shown in Figure 1 [3].

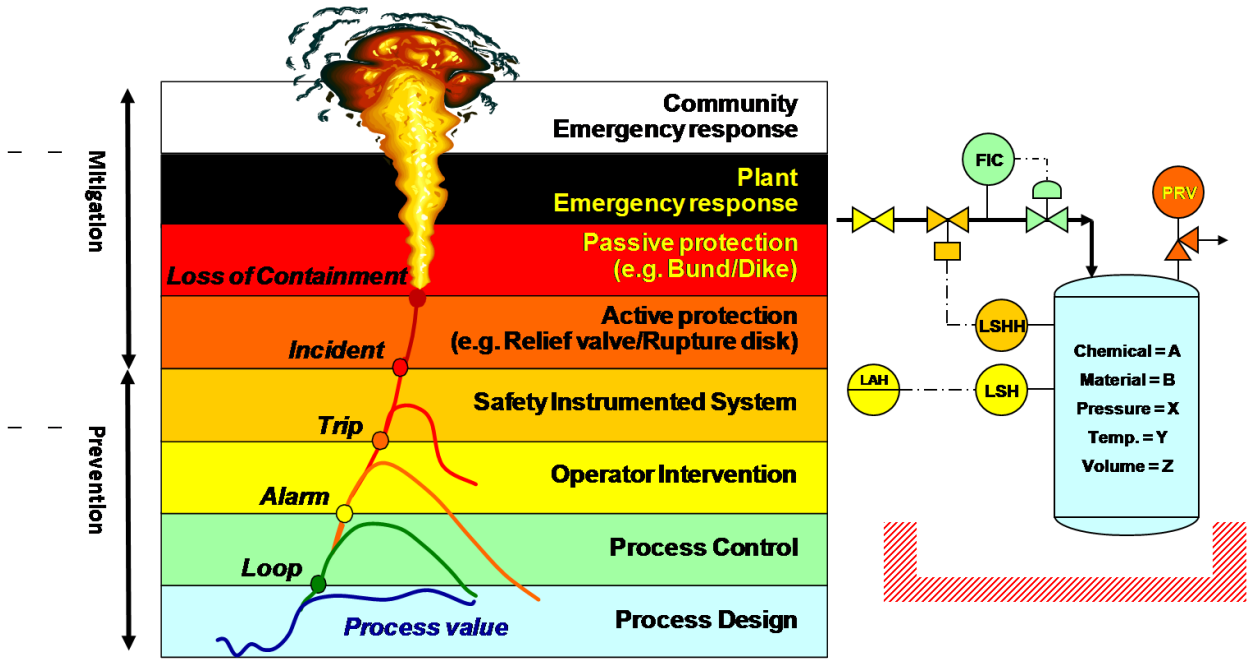

Figure 1. Layers of Protection and Their Impact on the Process [3]

Unlike other safeguards or layers of protection, such as a relief valve or safety instrumented system (SIS), the operator's response to an alarm relies on human intervention. There are numerous potential failure modes for operator response to an alarm including hardware, software, and human behavior. Failures in human behavior become more likely with poor alarm system design and performance (nuisance alarms, stale alarms, redundant alarms, and alarm floods). These failures are often improperly labeled as "operator error"; but are often more appropriately characterized as alarm management failures.

All valid alarms provide a measure of risk reduction to prevent an unwanted consequence. Risk reduction can be qualitative, such as for a safeguard in a Hazard and Operability Study (HAZOP), or semi-quantitative as for an independent protection layer (IPL) in a Layer of Protection Analysis (LOPA). This paper discusses techniques that can be applied to maximize the risk reduction from operator response to alarms and to minimize the chance of "operator error".

Notation

BPCS	Basic Process Control System
HAZOP	Hazard and Operability Study
HMI	Human Machine Interface
IEC	International Electrotechnic Committee
IPL	Independent Protection Layer
ISA	International Society of Automation
KPI	Key Performance Indicator
LOPA	Layer of Protection Analysis
MOC	Management of Change
OSHA	Occupational Safety & Health Administration
PFD	Probability of Failure on Demand
PHA	Process Hazard Analysis
PSM	Process Safety Management
RAGAGEP	Recommended and Generally Accepted Good Engineering Practice
SIS	Safety Instrumented System

Purpose / Problem Description

What is the Swiss Cheese Model?

Investigation of industrial incidents has shown that most include multiple independent failures. Thus it can be useful to invoke the metaphor of slices of swiss cheese to represent the layers of protection as proposed by James Reason [1]. As shown in Figure 2, each slice of swiss cheese represents an opportunity to prevent the hazard from escalating to an incident. No layer of protection is 100% reliable, so the holes in the swiss cheese represent failures. For an incident to occur, the holes in the swiss cheese must be aligned. The area of the holes in the swiss cheese would represent the un-reliability of the layer of protection (the more holes that exist and the larger the holes, the higher the probability of failure on demand).

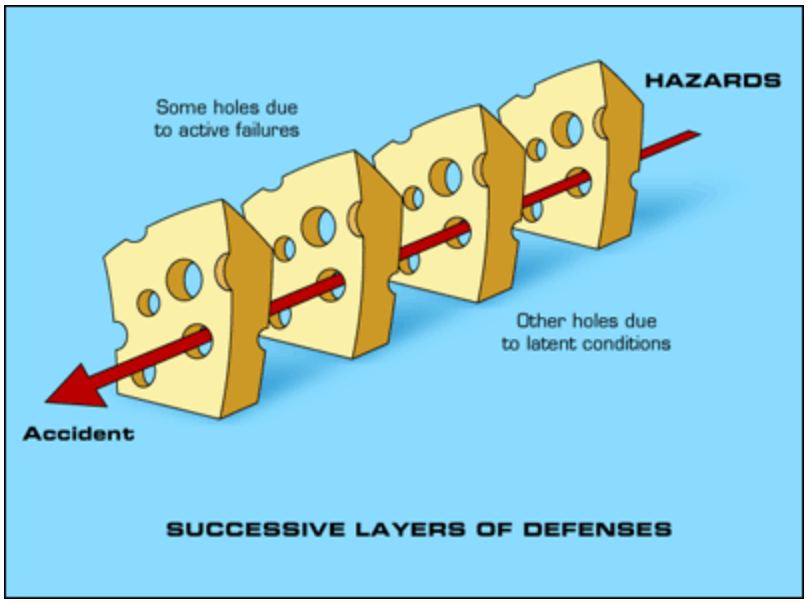

Figure 2. Layers of Protection – Swiss Cheese Model [4]

Typical reliability used in a Process Hazard Analysis (PHA) for an operator response to alarm is 0.9 (PFD = 0.1) assuming the action is simple and well-documented, with clear and reliable indications that the action is required [5]. Applying the swiss cheese model, would indicate that the area of the holes in the slice of swiss cheese is 10%. Keeping with the analogy, to improve performance of the operator response to alarm layer, techniques should be applied to reduce the area of the holes and to make sure the holes don't line up from one layer to the next. For an alarm system with poor performance, the size of the holes would be greater than 10%.

Operator Response Model

To analyze the failure modes for operator response to an alarm it is helpful to define a model for evaluation. For this study the operator response model is defined to consist of the following three components, as shown in Figure 3.

- **Detect** – the operator becomes aware of the deviation from the desired condition
- **Diagnose** – the operator uses knowledge and skills to interpret the information, diagnose the situation and determine the corrective action to take in response
- **Respond** –the operator starts and completes corrective action in response to the deviation [2]

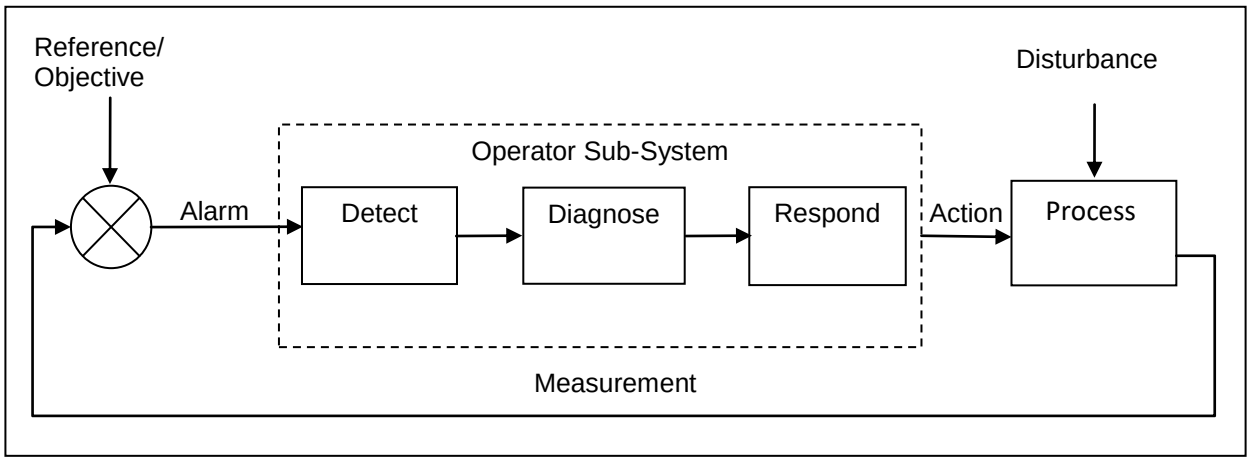

Figure 3. Feedback Model of Operator Process Interaction [2]

D.G. Dunn, et al. performed a study of eleven vessel overflow incidents to analyze the mechanisms where operational discipline broke down and the alarm protection layer failed. The mechanisms studied in the review are summarized in Figure 4 [6].

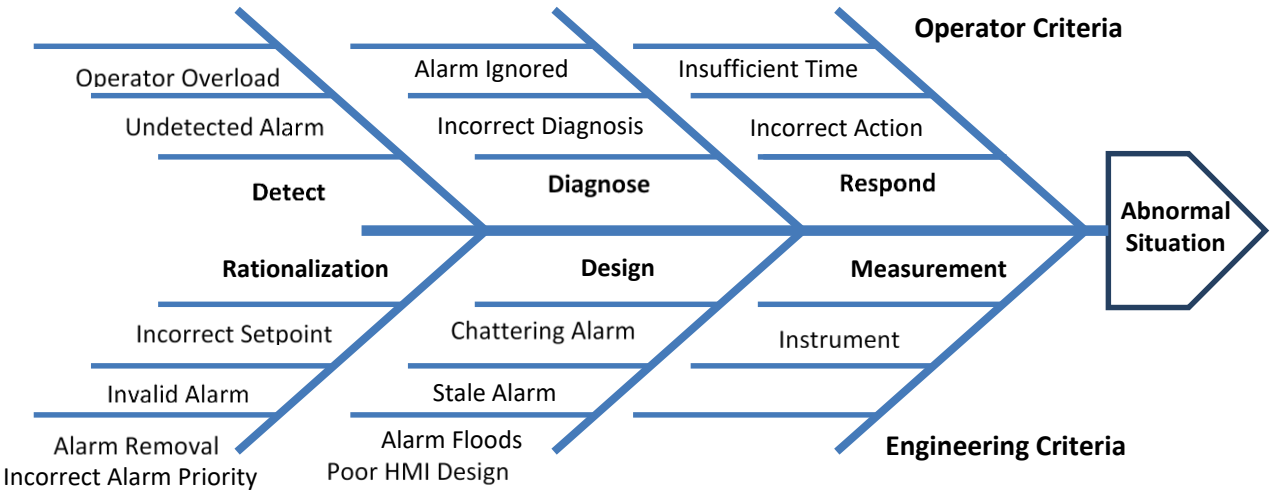

Figure 4. Failure Mechanisms Fishbone [6]

The failure mechanisms in the **"Detect"** component of operator response, include:

- **Undetected Alarm** - the alarm is annunciated but the operator does not notice the alarm
  - Stale alarms can indicate an alarm condition existed previously
  - Alarm floods can hide an alarm in a large number of alarms
  - Poor alarm human machine interface (HMI) design may make alarm annunciation difficult to detect
- **Operator Overload** - Rate of alarm actuation exceeds human capacity for signal detection

The failure mechanisms in the **"Diagnose"** component of operator response include:

- **Alarm ignored** – the operator receives the alarm indication and fails to take any action.
    - Chattering alarms can cause an alarm to be ignored because of frequent annunciation
    - Stale alarms can cause an alarm to be ignored because it is accepted as normal
- **Insufficient training** – the operator does not have enough knowledge to take the corrective action.
- **Incorrect Diagnosis** – the operator does have sufficient knowledge but fails to identify the corrective action.
- **Poor Descriptor** – Alarm information does not convey the real problem.

The failure mechanisms in the **"Respond**" component of operator response include:

- **Incorrect Action** – the operator determines the correct response but fails to take the correct action
- **No Action** – the operator determines the correct response but fails to take the corrective action.
- **Untimely Action** – The operator takes action but not quick enough to prevent the consequence from occurring.

An improper output from the "Respond" component of the operator subsystem, either the wrong or no output, is often referred to as "operator error". Operator error can be thought of as:

- not doing something that should be done (errors of omission),
- doing something in the wrong sequence,
- not doing the action in time, or
- doing something that shouldn't be done (errors of commission).

While failures can occur in all three components (Detect – Diagnose – Respond), most operator response failures in process plants are caused by failures in Detection or Diagnosis. Either the operator failed to notice the problem (Detection) or they incorrectly identified the cause and applied the associated (incorrect) action (Diagnosis). Instances where the operator knew what to do, but performed the incorrect action (such as turning the wrong valve), are much less frequent [7].

Situation Awareness

The concept of Situation Awareness will be applied to help evaluate and categorize failures in the operator subsystem (Detect-Diagnose-Respond). Situation Awareness (SA), which comes from the study of human factors and is more widely known in the airline industry, can be defined as "being aware of what is happening around you and understanding what that information means to you now and in the future [8]." As such, SA drives effective decision making and performance.

As a framework for categorizing SA failures, eight (8) factors, called SA Demons, have been identified which undermine effective situation awareness [8]. Several of these factors, as described below, will be referenced in this paper [6]:

- **Attentional Tunneling** – Focusing on one area or issue to the extent that alarms from another area or issue are excluded.

- **Misplaced Salience** – Incorrect alarm priority or HMI representation of alarm importance and other status information.

- **Errant Mental Models** – Thought process that incorrectly interprets alarms or mistakenly discounts relevant alarms.

The Impact of Human Factors

To improve upon operator performance requires an understanding of how humans process information. The discipline of human factors has become increasingly important with the evolution of technology. Numerous authors have commented on how the span of control (responsibility) of the operator has grown with the adoption of distributed control systems – as indicated by the number of control loops and alarms / operator [7, 9].

One of the challenges to applying human factors is that how well a person performs a task cannot be attributed to a single factor. Instead human performance is the product of several variables and their interaction. As shown by Strobhar this multidimensional aspect of human performance presents a couple of significant challenges when trying to understand it. First, deficiencies are often not addressed in the most direct and appropriate manor. In some cases the resolution contributes further to the problem. The response to an operator ignoring a nuisance alarm might be disciplinary action, additional training, more alarms, or longer procedures, instead of determination and resolution of the true root cause. Second, human performance issues rarely have simple and absolute answers [7].

To help understand the causes and resolutions, the above problems will be looked at from the point of view of signal detection theory. Chattering alarms, standing alarms, and alarm floods can all be thought of as visual “noise” to the operator obscuring their ability to correctly detect the alarm signal. According to signal detection theory, as the level of “noise” increases, the ability to discriminate a true alarm from a false alarm is reduced.

Typical Alarm Failure Example

The continuing review of vessel overflow incidents by Dunn et al and Chidambaram et al, attempts to identify the factors contributing to the failure of alarms using both the failure mechanisms shown in Figure 4 and the SA demons described above. In a typical incident where the alarm response layer of protection failed to prevent the overflow incident, the detect step of the measurement and control system worked as designed and presented the alarm to the operator. The failure occurs in the diagnose step.

In some cases, the operator was overloaded with alarms (alarm flood) or with actions related to a process upset. The most common case was that the operator dismissed the alarm as not requiring action either because it has not required action in the past, or because they do not believe the condition indicated by the alarm.

In one of the incidents reviewed, a product was to be transferred to a dedicated storage tank. A high-level alarm annunciated for a second tank because a valving error sent the product to the wrong tank. No action was taken by the operator in response to the alarm because he mistakenly believed no material was going to the second tank [6, 10].

Failure Example – Three Mile Island

One of the most famous and costly cases of operator error was the meltdown of the reactor at the Three Mile Island Nuclear Generation Station. The details below come from documented information and first-hand experience by one of the authors. Numerous factors contributed to the operators misdiagnosing the event and taking the incorrect action (it was correct for what they thought was occurring, but totally inappropriate for the actual event). The alarm flood quickly resulted in the operators abandoning the alarm system. This was one of two nuclear reactors on site. The previously constructed reactor was identical in design, but had about one-half as many alarms as its sister plant that melted down. In addition, an alarm critical to diagnosing the event was located on a remote panel, not given enough salience, and was not viewed until after a rupture disk had blown and returned the alarm to normal.

Failure Example – Petrochemical Plant in Sarnia

As witnessed first-hand by one of the authors, the impact of an alarm cannot be viewed in isolation, it is part of a system that the operator is using to control the plant. A process plant in Canada had a level controller freeze on its treated water tank. While the frozen controller alarmed high, the redundant level indicator alarmed low a few minutes later. The operator took no action for eight hours. When the treated water pumps began to cavitate, the operator realized there was no more water available to feed the plant's boilers. Steam was lost to the plant in January in Canada. The event had been alarmed, but the operator failed to react due to poor practices (did not compare the original high level alarm to the redundant indicator to identify the frozen transmitter), poor alarm management (the low level alarm was one of 15 on a page of standing alarms), and poor salience (the low level alarm differed from the high in only two letters, PVHI vs PVLO).

State of Alarm Management

The discipline of alarm management has evolved significantly over the last ten years. In 2009, the standard ANSI/ISA-18.2, "Management of Alarm Systems for the Process Industries" (ISA-18.2) was released. It provides guidance that can help users design, implement and maintain an alarm system in order to optimize performance for an operator response to alarm [2]. ISA-18.2 was used as the starting point for the creation of an international standard, IEC 62682, which was released in 2014 [11]. The ISA-18.2 standard was updated in 2016 based on lessons learned during the adoption phase and based on input from the IEC committee. The ISA-18.2 standard is considered a recommended and generally accepted good engineering practice (RAGAGEP) by insurance and regulatory agencies.

These standards provide a framework for the successful design, implementation, operation and management of alarm systems. They contain guidance to help prevent and eliminate the most common alarm management problems, as well as a methodology for measuring and analyzing performance of the alarm system. As shown in Figure 5, alarm management activities are structured to follow a lifecycle approach wherein the key activities are executed in the different stages of the lifecycle [2, 11, 12]. The products of each stage are the inputs for the activities of the next stage.

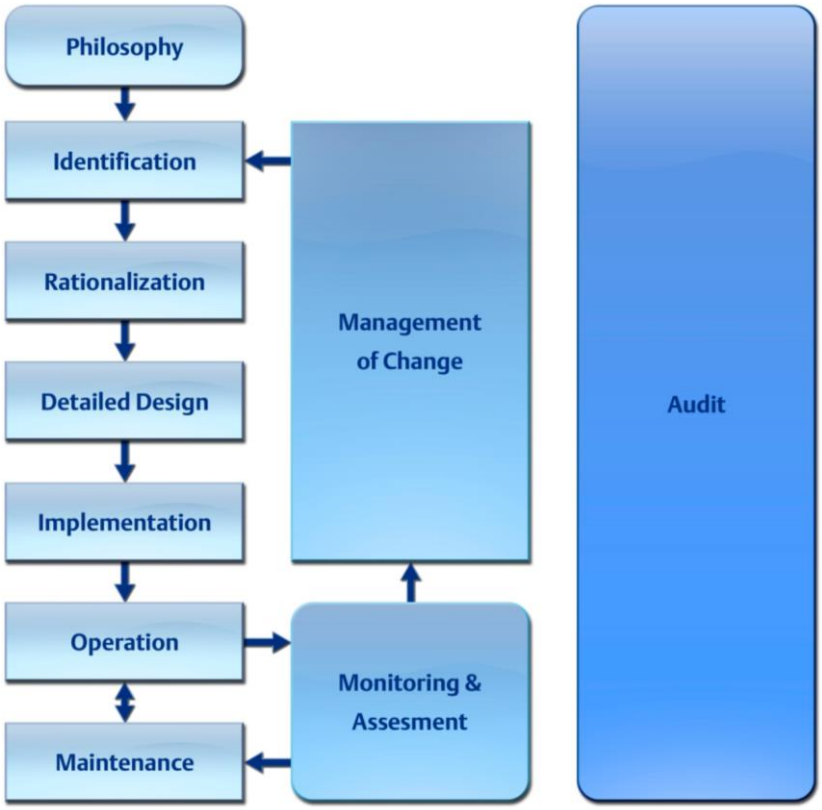

Figure 5. Alarm Management Lifecycle

For the purposes of this paper, the following activities will be highlighted:

- Philosophy
- Rationalization
- Detailed Design
- Operation
- Monitoring & Assessment

Addressing the Problems by following Alarm Management Principles from ISA-18.2

Implementation of an alarm management program is a journey that requires an ongoing commitment. A detailed discussion on how to implement a program is available in the July 2012 issue of Chemical Engineering Progress (CEP) [13]. Key elements of an alarm management program that have a significant impact on operator performance are summarized here below.

Alarm Philosophy

A necessary and often first step is creating an alarm philosophy document, which is the cornerstone of an effective alarm management program. It establishes the guidelines for how to address all aspects of alarm management, including the criteria for determining what should be alarmed, roles and responsibilities, human machine interface (HMI) design, management of change (MOC), and key

performance indicators (KPIs). This document is critical for helping plant staff maintain an alarm system over time and for driving consistency.

Establishing the methodology for alarm prioritization and classification is particularly important before beginning rationalization. Priority is used to indicate criticality and to help the operator understand the relative importance of each alarm. To ensure consistency, alarms should be prioritized based on the severity of the potential consequences and the time available for the operator to respond. Alarm classification organizes alarms based on common requirements (e.g., testing, training, MOC, reporting). Certainly an alarm that is identified as safeguard in a hazard and operability study (HAZOP) or as an independent protection layer (IPL) will have more stringent requirements for testing and operator training than the "average" process alarm. A good philosophy will provide a listing of relevant alarm classes (e.g., personnel safety, quality critical, environmental critical; safety critical, OSHA PSM critical), and their requirements.

Alarm Rationalization

To maximize dependability the operator must believe that every alarm is valid and requires their response. Alarm rationalization is the process for ensuring that every alarm configured in the system is valid and justified. Rationalizing the alarms in the system helps to improve the operator's trust in the alarm system, and serves to document the cause, consequence, corrective action, and time to respond. It also defines the priority of the alarm, which is a measure of the alarm's criticality. Priority, which is typically assigned based on the severity of the potential consequences and the time available for the operator to respond, tells the operator which alarm they should respond to first. While rationalization only requires that a single operator action and consequence for each alarm be identified, the process can become a rich source of training material and decision aids if a thorough documentation of all causes, responses, and consequences is conducted.

Classification of Alarms

Along with prioritization, the activity of classification is important to the effectiveness of alarms. Classification is process of separating alarms into alarm classes based on common requirements (e.g., testing, training, monitoring, and auditing requirements). Not all alarms are created equal and for alarms that need higher reliability, a higher level of management is needed. In fact, ISA-18.2 uses the term Highly Managed Alarms (HMA) for the classes of alarms that need higher reliability. Most of the requirements for highly managed alarms are derived from OSHA PSM. These requirements target:

- Training and documentation of training, both initial training and refresher training for both operations personnel and maintenance personnel
- Testing and documentation of testing
- Procedural control of shelving, including authorization and documentation
- Procedural control of out-of-service suppression, including authorization, documentation, and alternate risk reduction.

The HMA requirements apply to alarms classified as safety alarms; alarms critical to process safety for the protection of human life or the environment.

Alarm classification is usually based on the consequence the alarm is designed to prevent and the method used to identify the alarm (e.g., PHA, LOPA). It is worth noting that no risk reduction needs to be explicitly assigned to a safety alarm. If the alarm is specifically designed to detect a life-threatening

condition, it is a safety alarm. All alarms on chlorine detectors, for example, might be assigned to a safety alarm class. The testing, calibration, training and suppression would all be documented.

Operation – Alarm Response Procedures

For an operator response to alarm to be dependable, it is critical that the operator know what to do in the event of the alarm. This is best achieved through training and by making alarm response procedures available. The alarm response procedure, which contains key information documented during rationalization, can be provided in context to the operator from within the HMI. Use of alarm response procedures can reduce the time it takes the operator to diagnose the problem and determine the appropriate corrective action, as well as promote consistency between operators.

According to ISA-18.2 [2], alarm response procedures "shall be readily accessible to the operator as specified in the alarm philosophy." ISA-18.2 recommends that the following content be included in an alarm response procedure:

- Tag name for alarm
- Tag / alarm description
- Alarm Type
- Alarm Setpoint
- Potential Causes
- Consequence of Inaction
- Operator Action
- Allowable Response Time
- Alarm Class.

Training on how to respond to safety alarms is especially important as these alarms will not occur frequently and because they are most likely to occur during stressful situations such as a major plant upset. Providing operators with alarm response procedures is a recommended practice that should be considered mandatory for safety alarms according to some practitioners [3].

Well trained operators should possess knowledge necessary to diagnose an alarm; however, as the Deepwater Horizon incident illustrates, in times of stress, operator decision processes can be slow and/or incorrect. Making this type procedure available to the operator will support operational discipline by providing information that can improve the speed and accuracy of diagnosis, especially in times of stress [14].

Monitoring & Assessment

Monitoring the performance of an alarm system provides an indication as to whether the holes in the swiss cheese are getting smaller or larger. In a perfect plant, the alarm summary is a blank screen in the control room. The only time a message appears on the screen is when the operator needs to take an action to prevent an undesired consequence. The message disappears after the corrective action is taken and the screen is blank again. In such a plant, the alarm system does not negatively impact the ability of the operator to respond to an alarm, though other factors may. There are some plants that have achieved this level of alarm system performance.

In some plants the alarm system may overwhelm the operator, significantly increasing the probability the operator will fail to take the corrective action in time to prevent the consequence. Metrics allow categorization and trending of the alarm system performance.

The most commonly used performance metrics are average alarm rate and time-in-flood.

- Average alarm rate is the average number of alarms to an operator position in a unit of time (e.g., alarms per operator per hour).
- Time-in-flood is the percentage of 10-minute time intervals with greater than 10 alarms (per operator).

These metrics only roughly measure system performance. Generally, if the alarm rate is less than 6 per hour and the time in flood is less than 1 percent, the alarm system is performing well. Achieving this performance may also indicate that nuisance alarms are not a significant problem.

Another useful performance metric is the alarm priority distribution. When alarms have been prioritized based on the urgency for the operator to respond, the priority distribution can be an indication of the risk the operator is managing. The general guidance is ~80% of alarm annunciations should be in the lowest priority. Better guidance might be the higher the better. In a 3-priority alarm system, the general guidance is ~5% of alarm annunciations should be in the highest priority. Better guidance might be the lower the better. These numbers vary on the method of prioritization and type of process.

More focused metrics measure the performance of a subset of alarms, like the alarms in a safety alarm class. These are not system performance metrics, but can still be safety KPIs. Some class performance metrics could be alarm rate, time in alarm, and the number of stale alarm occurrences.

- Alarm rate could be used to infer how often the operator response to alarm layer of protection (the slice of swiss cheese) is challenged by an escalating hazard.
- Time-in-alarm is the measure of the average time from alarm annunciation to alarm return to normal.
- Number of stale alarm occurrences is the count of the number of times the stale alarm threshold, typically 24 hours, is exceeded.

These metrics indicate the amount of time alarms in the measured class, for example a safety alarm class, are annunciated to the operator. If safety alarms remain active for an extended period of time, the alarm is not effective and the alarm state will be normalized, making future alarms less effective. Again, the holes in the swiss cheese grow larger. With poor alarm system performance or poor alarm class performance, a safety alarm cannot be considered an effective layer of protection at all.

Design

The detailed design of alarms also has a significant impact on reliability, or minimizing the area of the holes in the cheese. Alarm design includes several aspects, as reflected in Figure 3 including:

- Selection and design of the process measurement
- Selection of alarm attributes that impact alarm behavior
- Designed suppression of the alarm when the alarm is not relevant
- Annunciation of the alarm to the operator (HMI)

It is a given that the process measurement must be reliable for the alarm to be reliable. Yet in some plants, measurements become faulty over time and the alarms become ineffective. With the monitoring described above, this can be detected and corrected.

Certain alarm attributes greatly impact alarm behavior, especially alarm deadband, on-delays, and off-delays. These attributes can be adjusted to reduce nuisance alarm behaviors like chattering or fleeting.

Designed suppression is used to hide alarms from the operator when the alarms are not relevant to the current mode of operation. An example is suppressing a low pump discharge pressure alarm if the pump is not running.

The alarm annunciation is also critical to reliability. If the alarm is hidden on a graphic in the HMI it will be less effective. Use of alarm priorities increase the salience of the most urgent alarms. For true safety alarms, an indication independent of the basic process control system (BPCS) may be needed to achieve the desired reliability. Often field horns and lights are used to indicate the need to evacuate an area due to a severe hazard. There is a significant difference in reliability between an alarm passing through the BPCS HMI for the operator to take action, and an alarm in the field that indicates the action is immediate evacuation. This was demonstrated in the Deepwater Horizon accident where the general evacuation alarm was NOT automated, but instead had to be initiated by the control room operator. The delay in was a contributing factor to the fatality and injury count [14]. Independent field alarms may have a reliability of 0.99 (PFD = 0.01) as long as the system performance is good as described above.

Addressing the Problems by Applying Human Factors Principles

Decision Theory applied to Operator Response to Alarms

Operator detection of an abnormal event, like the detection of any event or signal, has been studied by the human factors community since WWII. They have found that performance is based upon two key variables: (1) difference in the signal from the surrounding noise and (2) the willingness of the operator to accept false alarms versus missing a true event. The latter can be highlighted by torpedo lookouts on Navy vessels; they were willing to have more false alarms in order to ensure that they did not miss a true event.

All judgements are made in an environment of uncertainty. The signal indicating the event has some degree of noise associated with it. That signal then is conveyed to the operator in an environment with its own degree of “signal” noise (Figure 6). The greater the difference in the nature of the noise in the environment and the characteristics of the signal (D’), the more likely is that the signal will be detected. As the environmental noise (audible or visual) more closely matches the signal (e.g., red is used for both warning and pump status), the more likely it is that the signal will be missed. As D’ approaches zero (e.g,, large number of standing alarms or false alarms), the ability of the operator to detect a valid alarm is little better than 50%, a coin toss.

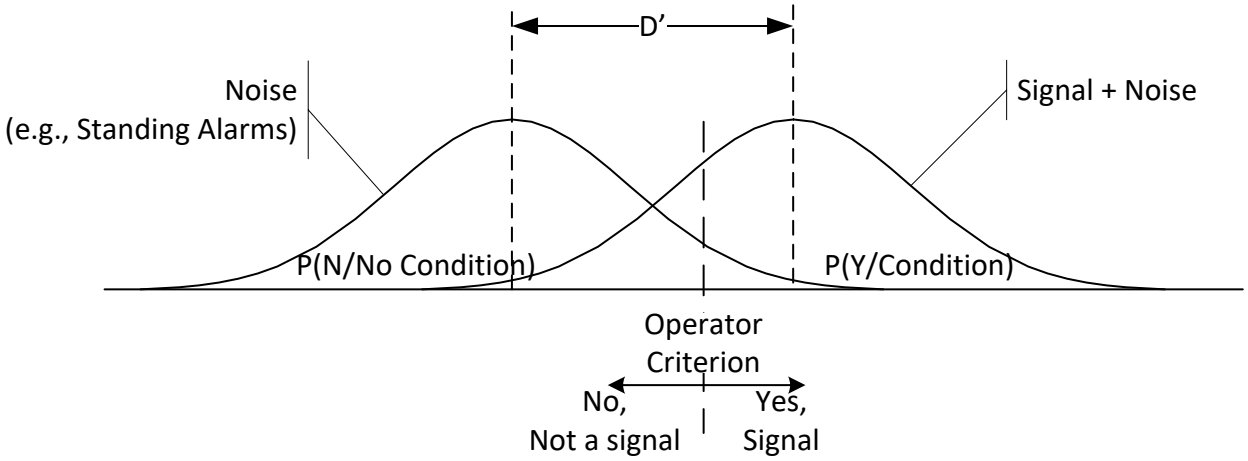

Figure 6. Signal and Noise Conditions

While the difference between the noise and signal affects the detection of the signal, it is not the only factor. The person, the operator, creates their own criterion for what signal to accept as true or valid. As the signal and noise become more alike, the probability of mistaking a false condition for a true event increases (e.g., a poorly designed alarm system). If experience has placed little importance on missing the actual event, then the user will generally opt to reduce false alarms at the expense of missing a true event – the "cry-wolf" phenomena. In reducing the number of false alarms, the operator increases the probability of missing a valid alarm. Ensuring a large difference in alarm signal/noise is important not only for the ability to detect the signal, but the operator's criterion for doing so.

Human Factors Consideration for Nuisance Alarms

It does not take a high false alarm rate to result in the operator essentially abandoning their alarm system. A 25% false alarm rate, one alarm out of every four, is enough to have the operator no longer rely on the alarm system to detect an abnormal event. A false alarm is does not necessarily indicate that the condition is not true (e.g., level is high), but simply that no action is needed on the part of the operator ("It's not high enough to warrant my doing something about it").

Ignoring nuisance alarms is not a behavior that can be changed by training or disciplining the operator. According to Endsley "A person's reluctance to respond immediately to a system that is known to have many false alarms is actually quite rational. Responding takes time and attention away from other ongoing tasks perceived as important [8]".

Shifting the operator's reliance on the alarm system is often a byproduct of the alarm rationalization process. As redundant alarms are eliminated, the operators begin to realize that they must pay attention when an alarm actuates. They will not get a second, third, and fourth alarm prior to their need to take action. Several operators have commented during rationalization, "No more ignoring alarms, we're going to have to act when they come in".

Additionally design of HMI displays should support situation awareness. As stated by Endsley "Visual displays need to do more than simply provide a visual indication of the alarm; they should help people rapidly confirm that the alarm is a real indication [8]".

Human Factors Consideration for Training

Helping operators identify valid alarm conditions from invalid ones should be part of the basic operator training program. Expert operators know what to look for in order to determine if the alarm condition is real. Rarely does a process variable move in isolation. Something else will have changed as well, either another variable or the control system. As noted in the earlier example, had the operator in the Canadian refinery compared the redundant level transmitters, one would have been identified as failed and the facility would not have lost steam in January. A key part of operator training needs to be how to cross-check alarms with other process variables to determine their validity. For example, if a low flow alarm occurs, is there a change in level in the up or down stream vessels.

Human Factors Consideration for Alarm Response Procedures

A general guideline in alarm rationalization is DO NOT alarm the normal or expected. That provides no real information to the operator. So it is with alarm response procedures. Providing the operator with information from the rationalization that is obvious for a trained operator just creates noise, not value. While the entire output of the alarm rationalization is useful for training, only a subset of that information should be provided to the operators as a real-time decision aid. Alarm response procedures should contain or highlight that which is unique about this particular alarm, either in cause, response, or consequence of inaction. Alarm response procedures are an aid for the operator, not a replacement for good operating procedures.

Human Factors Consideration to Best Deal with Alarm Floods

Even well designed alarm systems can generate a large number of alarms as a result of a major process upset, such as a loss of power. Humans have a relatively limited capacity for processing information, so the potential to exceed that capacity is high. However, there are several techniques that can be employed to maximize the potential that alarms will be processed.

Two of the techniques reduce the number of alarms that must be processed. The easiest is to automatically suppress low priority alarms for some period of time. If the alarm prioritization has been done correctly, this will reduce the number of alarms by 80% with limited risk in the short term, given that low priority alarms generally do not need to be responded to with the same urgency as high priority. A more complicated technique is to utilize state based alarming such that those alarms that are expected to occur during such upsets do not actuate. For example, once a unit has shutdown, alarms due to low energy (e.g., low temperature, pressure, flow) will likely be “normal”. Alarm for these conditions should therefore be suppressed when entering the low-energy state, prior to annunciating.

The third technique is to aggregate the alarm information into higher order alarms for entire systems, thereby reducing what needs to be processed by the operator. For example, all alarms for a tower or set of towers operating in series can be combined so the operator need not process them individually, but qualitatively as the tower is either “okay” or “upset”.

Human Factors - How to Prevent Attention Tunneling

Due to the limit of human processing capabilities, it is possible for the operator to become overly focused on the task at hand and miss events in other parts of the system. This is a loss of the “big picture” or situation awareness. A single display for the operator’s entire span of control is needed to ensure that they can, at all times, assess the status or health of all the equipment for which they are

responsible. Included in this display should be the status of the alarm system – what units / areas have alarms and what is their priority/importance.

Human Factors - How to Prevent Misplaced Salience

People are very poor at detecting changes, resulting in a phenomenon known as change blindness. Overcoming change blindness requires that changes be highlighted when they happen. This is the purpose of having new alarms flash until acknowledged directing the operator's attention to the variables that have changed. It would be very difficult for any individual to notice that some values on a screen had changed color if that were the only indication that an alarm had occurred, particularly if the display was very colorful.

Human Factors – How to Improve Use of Mental Models

Mental models are an important mechanism for interpreting new information. They are thinking tools that provide a construct for a person to combine disparate pieces of information, interpret the significance of that information, and to develop reasonable projections of what will happen in the future [8]. Operators work with mental models of how the process works (**if** the reactor feed flow rate is increased, **then** the reactor temperature will increase without additional coolant).

Application of mental models run into problems when the operator uses an incomplete or incorrect mental model. One of the keys to using mental models is to be able to realize when you are using the wrong one. Operators may misinterpret alarms or events as fitting into their current mental model without realizing these cues indicate that they should be thinking differently (using a different mental model). Unfortunately people tend to explain away conflicting cues to their current mental models (confirmation bias) and can be slow to realize this mistake.

One way to improve operator response is to have them develop multiple mental models so that they can pick the right one for the situation. To help in the development of more and better mental models one recommendation is to apply a pre-mortem strategy [15]. This involves analyzing / brainstorming how a process or operation could fail through the creation of if – then scenarios and discussion of how to rectify the situation. Experienced operators have more and more varied models of plant operation, developed after years of experience plant upsets. These models need to be transferred to new operators without them having to experience the upsets for themselves. This could be especially useful before starting up or shutting down equipment or changing its mode of operation. Rationalization can be effective in identifying where multiple models exist, and training is how we share with the entire staff.

Human Factors – Conditioning / Normalization

Standing alarms are a negative in alarm metrics as they create noise in the visual environment, in this case the alarm display. In the case cited earlier of the Canadian Refinery, the alarms for the faulty and valid level indications were present over eight hours before water was lost to the boilers. How does an operator go eight hours with alarms in their visual field and not notice them? It happens because some 15 or more alarms already existed on the alarm summary display. The two alarms of the anomalous condition simply blended into the other alarms. More standing alarms results in greater visual noise in the environment and increasing difficulty to detect a signal when it occurs.

References

1. J. Reason, *Managing the Risks of Organizational Accidents*. Burlington, VT: Ashgate Publishing, 1997.
2. ANSI/ISA–18.2–2016. Management of Alarm Systems for the Process Industries. June 2016.
3. Stauffer, T., Clarke, P., Using Alarms as a Layer of Protection – AICHE 8th Global Congress on Process Safety, April 2012.
4. J. Reason, (1990) *Human Error*. Cambridge, Cambridge: University Press, 1990.
5. CCPS. *Guidelines for Safe and Reliable Instrumented Protective Systems.* Center for Chemical Process Safety, American Institute of Chemical Engineers, Hoboken, NJ, 2007.
6. D.G. Dunn, N.P. Sands, and T. Stauffer, When Good Alarms go Bad: Learning from Incidents, 70th Annual Instrumentation and Automation Symposium – Texas A&M University, January 2015.
7. D. Strobhar, P.E., *Human Factors in Process Plant Operation*, Momentum Press, New York, NY, 2013.
8. M. Endsley, *Designing for Situation Awareness: An Approach to User-Centered Design*, CRC Press, Boca Raton, Fl, 2012.
9. B. Hollifield, E. Habibi, *Alarm Management: A Comprehensive Guide, Second Edition*, International Society of Automation, Research Triangle Park, NC, 2011.
10. Chidambaram, P., Sands, N., Analysis of Overflow Incidents: Searching for Situational Awareness Demons – AICHE 12th Global Congress on Process Safety, March 2016.
11. IEC62682-2014 "Management of Alarm Systems for the Process Industries". October 2014.
12. Stauffer, T., Sands, N., and Dunn, D., "Alarm Management and ISA-18 – A Journey, Not a Destination", Texas A&M Instrumentation Symposium (2010).
13. T. Stauffer, P.E., Implement an Effective Alarm Management Program, Chemical Engineering Progress, pp. 19-27, July 2012.
14. Chastain-Knight, D., Stauffer, T., Managing Alarms to Support Operational Discipline – AICHE 12th Global Congress on Process Safety, March 2016
15. G. Klein, *Sources of Power: How People Make Decisions*, The MIT Press, Cambridge, MA, 1998.

Revision History:

Authors: Todd Stauffer, PE, Nicholas P. Sands, CAP, PE, David Strobhar, PE

Prepared for Presentation at
American Institute of Chemical Engineers
2017 Spring Meeting and 13th Global Congress on Process Safety
San Antonio, TX
March 26 – 29, 2017

exida – Who we are.

exida is one of the world's leading accredited certification and knowledge companies specializing in automation system cybersecurity, safety, and availability. Founded in 2000 by several of the world's top reliability and safety experts, exida is a global company with offices around the world. exida offers training, coaching, project-oriented consulting services, standalone and internet-based safety and cybersecurity engineering tools, detailed product assurance and certification analysis, and a collection of online safety, reliability, and cybersecurity resources. exida maintains a comprehensive failure rate and failure mode database on electrical and mechanical components, as well as automation equipment based on hundreds of field failure data sets representing over 350 billion unit operating hours.

exida Certification is an ANSI (American National Standards Institute) accredited independent certification organization that performs functional safety (IEC 61508 family of standards) and cybersecurity (IEC 62443 family of standards) certification assessments.

exida Engineering provides the users of automation systems with the knowledge to cost-effectively implement automation system cybersecurity, safety, and high availability solutions. The exida team will solve complex issues in the fields of functional safety, cybersecurity, and alarm management, like unique voting arrangement analysis, quantitative consequence analysis, or rare event likelihood analysis, and stands ready to assist when needed.

Training

exida believes that safety, high availability, and cybersecurity are achieved when more people understand the topics. Therefore, exida has developed a successful training suite of online, on-demand, and web-based instructor-led courses and on-site training provided either as part of a project or by standard courses. The course content and subjects range from introductory to advanced. The exida website lists the continuous range of courses offered around the world.

Knowledge Products

exida Innovation has made the process of designing, installing, and maintaining a safety and high availability automation system easier, as well as providing a practical methodology for managing cybersecurity across the entire lifecycle. Years of experience in the industry have allowed a crystallization of the combined knowledge that is converted into useful tools and documents, called knowledge products. Knowledge products include procedures for implementing cybersecurity, the Safety Lifecycle tasks, software tools, and templates for all phases of design.

Tools and Products for End User Support

- exSILentia® – Integrated Safety Lifecycle Tool

- PHAx™ (Process Hazard Analysis)
- LOPAx™ (Layer of Protection Analysis)
- SILAlarm™ (Alarm Management and Rationalization)
- SILect™ (SIL Selection and Layer of Protection Analysis)
- Process SRS (PHA based Safety Requirements Specification definition)
- SILver™ (SIL verification)
- Design SRS (Conceptual Design based Safety Requirements Specification definition)
- Cost (Lifecycle Cost Estimator and Cost Benefit Analysis)
- PTG (Proof Test Generator)
- SILstat™ (Life Event Recording and Monitoring)

- exSILentia® Cyber- Integrated Cybersecurity Lifecycle Tool
  - CyberPHAx™ (Cybersecurity Vulnerability and Risk Assessment)
  - CyberSL™ (Cyber Security Level Verification)

Tools and Products for Manufacturer Support

- FMEDAx (FMEDA tool including the exida EMCRH database)
- ARCHx (System Analysis tool; Hardware and Software Failure, Dependent Failure, and Cyber Threat Analysis)

For any questions and/or remarks regarding this White Paper or any of the services mentioned, please contact exida:

exida.com LLC

80 N. Main Street

Sellersville, PA, 18960

USA

+1 215 453 1720

+1 215 257 1657 FAX

info@exida.com