

Introduction:

Recent industrial accidents at Texas City, Buncefield (UK) and Institute, WV have highlighted the connection between poor alarm management and process safety incidents. At Texas City key level alarms failed to notify the operator of the unsafe and abnormal conditions that existed within the tower and blowdown drum. The resulting explosion and fire killed 15 people and injured 180 more.¹ The tank overflow and resultant fire at the Buncefield Oil Depot resulted in a £1 billion (1.6 billion USD) loss. It



Figure 1: Fire & Explosion at Texas City Refinery¹

could have been prevented if the tank's high level safety switch, per design, had notified the operator of the high level condition or had automatically shut off the incoming flow.² At the Bayer facility (Institute, WV) improper procedures, worker fatigue, and lack of operator training on a new control system caused the residue treater to be overcharged with Methomyl - leading to an explosion and chemical release.³



Figure 2: Aftermath of Explosion at Buncefield Oil Depot²

Accidents like these demonstrate what can happen when an alarm system and operator response fail as a layer of protection in a hazardous process. They also provided the motivation for the new ISA-18.2 standard "Management of Alarm Systems for the Process Industries", which provides a framework for the successful design, implementation, operation and management of alarm systems in a process plant. It offers guidance on how alarm management can be used to help a plant operate more safely. ISA-18.2 can also be used to bring together the disciplines of alarm management and safety system design, which must work more closely to prevent future accidents.

The Alarm System and the Operator are one of the first layers of protection

The operator's response to alarms is crucial in preventing a process upset from escalating into a more serious event. As shown in Figure 3, there are multiple layers of protection that can prevent an incident from occurring and to mitigate its impact if it does occur. Operator intervention is one of the first layers of protection. Next comes the Safety Instrumented System (SIS) whose job is to drive the process to a safe state, as needed, to protect people, the environment, and equipment. When a safety system trips

it typically results in lost production, which can be very significant – for an oil refinery it can easily exceed \$1M / hour. Therefore implementing proper alarm management to improve the operator’s performance can help your plant run more efficient and also reduce the likelihood that a process upset reaches the SIS layer of protection.

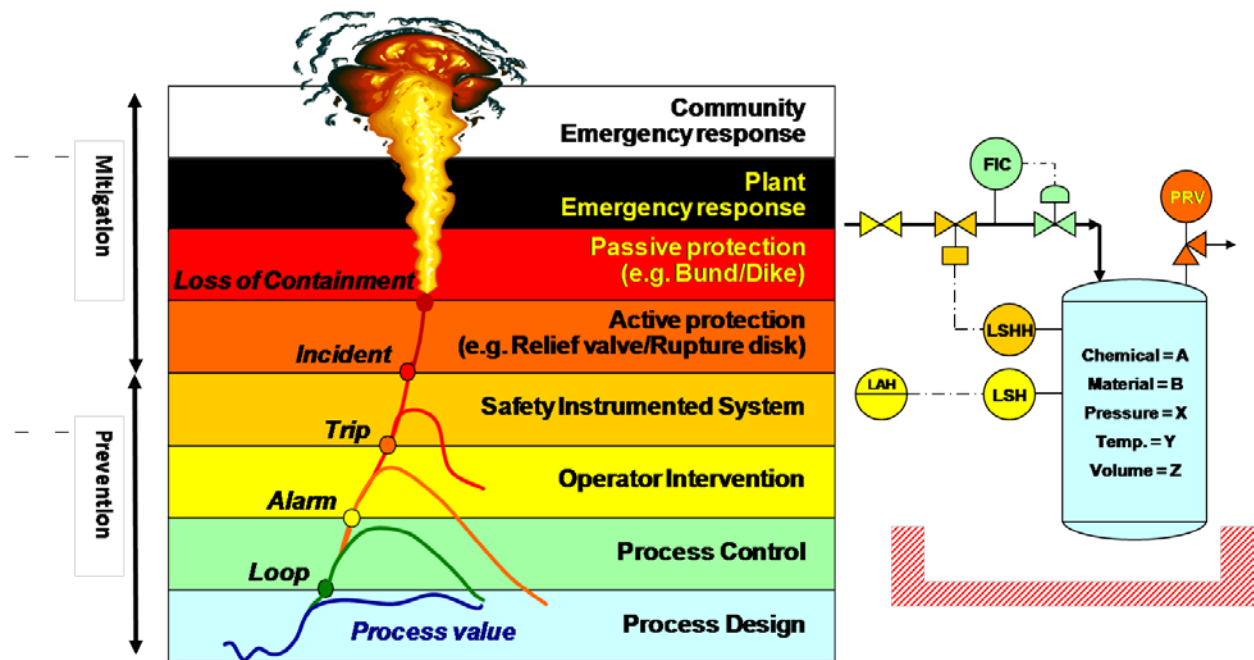


Figure 3. Layers of Protection and Their Impact on the Process

How much Risk Reduction can the Alarm System and the Operator Provide

According to the IEC 61511/ISA 84 process safety standards, process risk must be reduced to a tolerable level as set by the process owner. This is done using multiple layers of protection including the basic process control system (BPCS), alarms, operator intervention, mechanical relief systems, and (if necessary) an SIS. As shown in Figure 4, the more risk that can be reduced by the alarm system and the operator, the less risk reduction (Safety Integrity Level – SIL) which must be provided by the SIS. The higher the SIL level, the more complicated and expensive is the SIS. Additionally, a higher SIL will require more frequent proof testing, which adds cost and can be burdensome in many plants. Unfortunately, human performance factors provide constraints on the level of risk reduction that an operator can actually provide. By “getting the most” from the operator, the demands on the SIS are reduced, which in turn reduces its chance of failure.

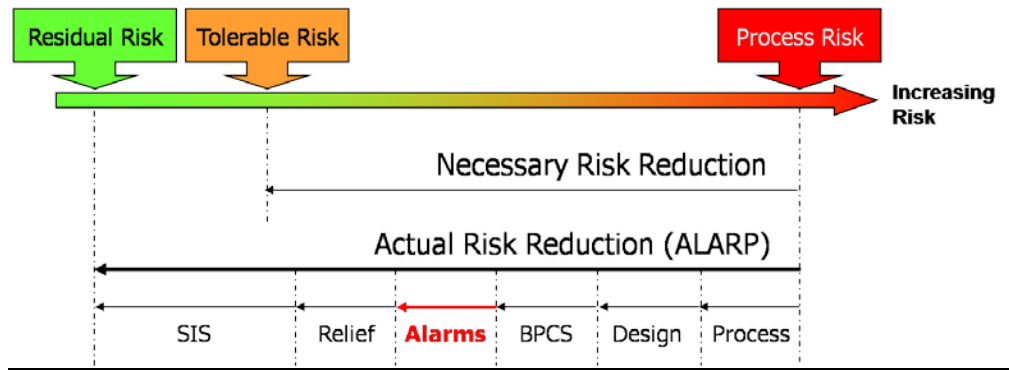


Figure 4 – Risk Reduction is achieved through use of multiple protection layers⁴

The reliability of the alarm system and operator are considered when performing a Layer of Protection Analysis (LOPA), which is one of several methods for calculating the required SIL target. In a LOPA the frequency of a potentially dangerous event is calculated by multiplying the probability of failure on demand (PFD) of each individual layer of protection times the frequency of the initiating event. In the example LOPA of Figure 5, the likelihood of a fire occurring after the release of flammable materials is calculated assuming that the initiating event (the loss of jacket cooling water) occurs once every two years.

Initiating Event	Protection Layer #1	Protection Layer #2	Protection Layer #3	Protection Layer #4	Outcome
Loss of Cooling Water	Process Design	Operator Response (to Alarm)	Pressure Relief Valve	No Ignition	Fire
				0.3	2.10E-05
				0.07	Fire
				0.2	
0.5 / yr	0.01				
					No Event

Figure 5. Example Layer of Protection Analysis (LOPA) Calculation⁵

How Reliable is the Operator and the Alarm System

The example LOPA calculation assumes that each protection layer, including the operator, is specific, auditable, independent, and dependable. The calculation uses a 20% chance that the operator will fail to respond correctly and in time to prevent the outcome (PFD = 0.2). Assuming an 80% success rate might seem conservative, but studies have shown that human error is one of the leading causes of industrial accidents. A review of the top 100 plant accidents determined that operator failure was the second leading cause (after equipment mechanical fatigue).⁶

On the other hand, an 80% success rate might be generous. Consider that safety-critical alarms are most likely to occur during major plant upsets. Throw in operator fatigue, lack of proper training, increasing operator workload, physical condition (age, amount of rest), along with alarm overload – and one can



see the challenge to improving the operator’s response. Table 1 presents representative values for estimating operator response as part of a safety calculation.

Category	Description	Probability that Operator responds successfully	PFD	SIL
1	<p>Normal Operator Response – In order for an operator to respond normally to a dangerous situation, the following criteria should be true:</p> <ul style="list-style-type: none"> • Ample indications exist that there is a condition requiring a shutdown • Operator has been trained in proper response • Operator has ample time (> 20 minutes) to perform the shutdown • Operator is ALWAYS monitoring the process (relieved for breaks) 	90%	0.1	1
2	<p>Drilled Response – All of the conditions for a normal operator intervention are satisfied and a “drilled response” program is in place at the facility.</p> <ul style="list-style-type: none"> • Drilled response exists when written procedures, which are strictly followed, are drilled or repeatedly trained by the operations staff. • The drilled set of shutdowns forms a small fraction of all alarms where response is so highly practiced that its implementation is automatic • This condition is RARELY achieved in most process plants 	99%	0.01	2
3	<p>Response Unlikely / Unreliable – ALL of the conditions for a normal operator intervention probability have NOT been satisfied</p>	0%	1.0	0

Table 1 – Simplified Technique for Estimating Operator Response ⁵

It's just a matter of time

So how can we improve the operator's performance to keep our plants safer? One way is to think about what constitutes a successful operator response. As described in ISA-18.2, the operator must be able to **detect**, **diagnose**, and **respond** within the appropriate timeframe, called the Maximum Time to Respond. The operator's response must be quick enough that the process has time to react to the corrections that have been made before reaching the consequence threshold. If the total time elapsed exceeds the Process Safety Time, which is the time between the initiating event and occurrence of the hazardous event, then the upset will escalate to create a demand on the SIS, initiate a trip, or cause an accident.

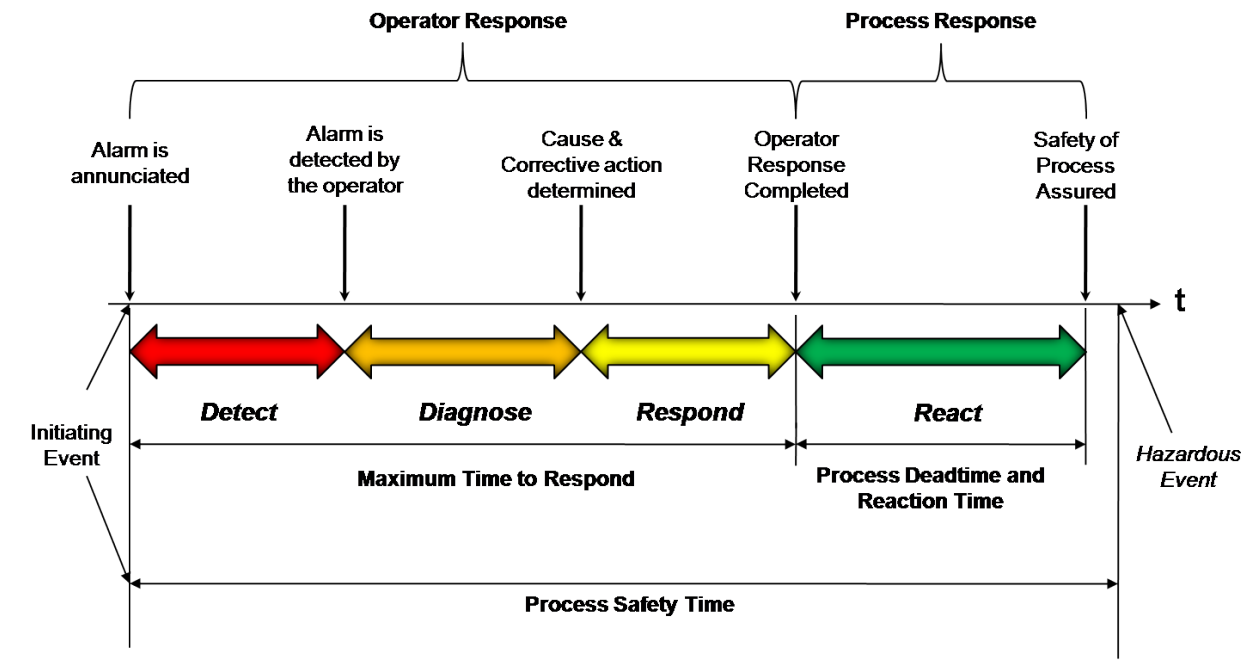


Figure 6. Operator Response Timeline

Operator response time should be considered up-front during design. Creating a situation where an operator has only a few minutes to **detect**, **diagnose**, and **respond** increases the probability for failure and means that they cannot be a significant safety layer. One company has set a threshold requirement of 10 minutes, meaning any alarm which has a process safety time of less than 10 minutes cannot be claimed as a layer of protection (PFD = 1.0).

Applying a Lifecycle Approach to both Safety and Alarm Management

Both the ANSI 61511/ISA-84 standard on process safety and the ISA-18.2 standard on alarm management advocate the use of a lifecycle approach. As shown in Figure 7, the two lifecycles are very similar. There are several phases where they can be “connected”. Results from the Safety Hazard and Risk Assessment are an input to alarm management’s Identification phase. Alarms which are being counted on as a safety protection layer will be assigned a (high) priority during Rationalization.

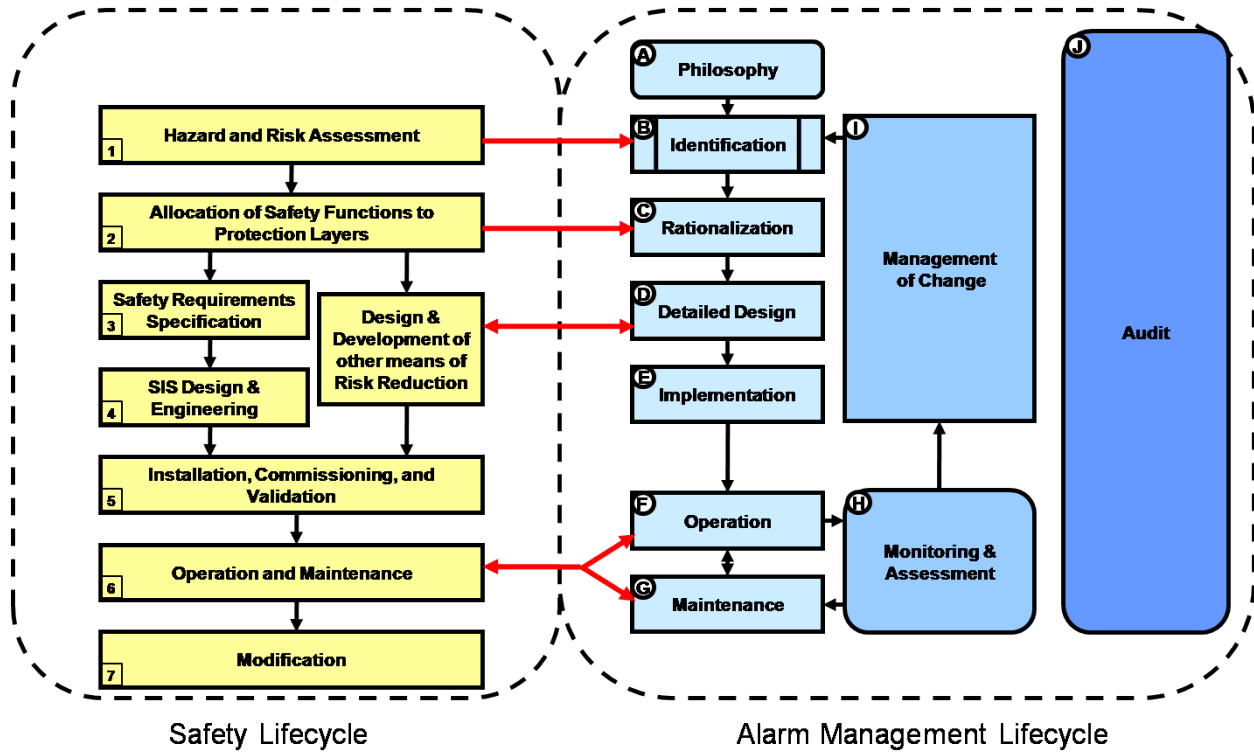


Figure 7. The Alarm Management & Safety Lifecycles (Abbreviated for clarity) ^{7,8}

A key deliverable is to create an alarm philosophy document which defines how a company or site will address alarm management throughout all phases of the lifecycle. It should contain information like the criteria for classifying and prioritizing alarms (safety-related alarms are classified as “Highly Managed Alarms”), what colors will be used to indicate an alarm in the HMI, and how changes to the configuration will be managed. It should also establish key performance benchmarks (like the acceptable alarm load for the operator). Special procedures for handling safety-related alarms, such as the frequency of testing and refresher training, would be documented here.

Alarm Detection made Quick and Easy

Operator performance can be optimized by ensuring that alarms are annunciated in a way that makes it easy for them to be detected and so that none are missed. Some helpful techniques are:

Design the HMI to promote operator situational awareness - The operator’s performance is directly linked to the proper use of color, text, and patterns within the HMI, which should be configured to uniquely indicate the state of the alarm (normal, unacknowledged, acknowledged, suppressed). Since 8-12% of the male population is color-blind, it is important to consider what colors are used. Ideally the colors used for alarm indication should be reserved for alarming only and should be different depending upon priority.



Minimize the number of alarms on the Operator – Alarm overload is a key reason why operators “miss” alarms. An operator should be hit with no more than 1-2 alarms every 10 minutes during steady-state operation. In many control rooms, operators are hit with one alarm every minute, which is considered unmanageable.

Alarm Performance Metrics		
Based upon at least 30 days of data		
Metric	Target Value	
Annunciated Alarms per Time:	Target Value: Very Likely to be Acceptable	Target Value: Maximum Manageable
Annunciated Alarms Per Day per Operating Position	~150 alarms per day	~300 alarms per day
Annunciated Alarms Per Hour per Operating Position	~6 (average)	~12 (average)
Annunciated Alarms Per 10 Minutes per Operating Position	~1 (average)	~2 (average)
Metric	Target Value	
Percentage of 10-minute periods containing more than 10 alarms	<1%	

Table 2: Alarm Performance Metrics from ISA-18.2 ⁸

Performing a thorough rationalization ensures that every alarm in the system is necessary, has a purpose, and follows the cardinal rule – that it requires an operator response. In today’s DCS it is really easy to add alarms that aren’t called for – we must all resist the temptation.

Treat High Priority Alarms as “High” Priority – Another way to ensure operators don’t miss important alarms is to ensure that high priority alarms can be differentiated from other alarms. ISA-18.2 recommends using 3-4 different priorities, where no more than 5% of alarms are configured as high priority. Priority is set based on the potential consequences and on the time available to respond. Establishing consistent priorities aids the operator in determining the order of response during upset conditions.

Eliminate “Nuisance” Alarms – The presence of standing alarms (lasting > 24 hours) and chattering alarms (points that go needlessly in and out of alarm on a frequent basis) can obscure the operator’s view and make it more difficult for him to detect a new alarm. Poor configuration practices are one of the leading causes of nuisance alarms. The proper use of alarm deadbands and on / off-delays can go along way to eliminating them. An ASM study found that the use of on-off delays in combination with other configuration changes was able to reduce the 10 min alarm rate by 45 – 90%. ⁹



Diagnose the issue accurately

Make Information on Cause and Corrective Action available – The Information documented during alarm rationalization and hazard and risk assessment can be indispensable for helping the operator diagnose the problem and determine the best response. The cause of the alarm, corrective action, consequence, time to respond, and safety implications should be made available ideally in real-time and in the proper context.

Suppress Unimportant Alarms During a Flood - Plant upsets, which can generate tens to hundreds of alarms, are one of the most challenging times for the operator. At the Milford Haven refinery, operators were inundated with 275 alarms in the 11 minutes leading up to the explosion.¹⁰ Advanced alarming techniques, such as state-based alarming, can be used to temporarily suppress alarms when they are not meaningful. When a distillation column crashes it is best to present only those few alarms which effect the diagnosis and response, rather than all of the temperature and pressure alarms that occur.

Shelving helps the operator stay focused - Alarm shelving allows an operator to temporarily suppress an insignificant alarm, removing it from view. It is a great tool for improving response during a process upset. The alarm will come back later (like after 30 minutes) so that it can be addressed when things have calmed down in the control room. It is important to provide controls on who can shelve an alarm and which alarms can be shelved.

Respond correctly

Practice makes perfect – Not to be underestimated is how important it is to train the operators so they are comfortable with the system, and so they trust it to help them do their job. The last you thing you want is the operator abandoning the control system during an upset, like the operators did prior to the explosion at the Milford Haven refinery. Training the operator as part of process simulation can create a “drilled response” where corrective action is so-well reinforced that it is automatic.

Provide Alarm Response Procedures – Written alarm response procedures should be created which include the potential causes and consequences of the alarm, the recommended corrective action, the alarm limit, and the allowable response time –information that was fleshed out during rationalization and during hazard and risk assessment.

Maintenance and Change Control

Review and Know which Alarms are Out-of-Service - Alarms will periodically be taken out-of-service for maintenance - repair, replacement, or testing. It’s important to document why an alarm was removed from service, the operation of interim alarms, special handling procedures, as well testing required prior to returning to service. For safety reasons the system should be able to produce a list of which alarms are currently out-of-service. This serves as a reminder of what alarms are suppressed. The list can then be reviewed before putting a piece of equipment back into operation to ensure that all critical alarms are functional.



Manage and Control Configuration Changes – Even the most well-designed alarm system can run into problems if there is poor control over who can make changes. A Management of Change procedure should be implemented to ensure that modifications (such as changing an alarm limit, disabling an alarm, or adjusting its priority) are reviewed and approved prior to implementation. Modifications should not be made without proper analysis and justification, particularly if the alarm is a safety layer of protection.

Conclusion

As plants run closer to their performance limits with fewer operators and support staff, the importance of alarm management to maintaining plant safety is becoming paramount. The key to maximizing the safety protection provided by the operator is to create an environment where they are able to **detect**, **diagnose**, and **respond** to alarms properly within time. This can be achieved by adopting the requirements and recommendations of the new ISA-18.2 standard on alarm management and by taking a coordinated approach to both alarm management and SIS design.

References

1. “BP America Refinery Explosion” U.S. CHEMICAL SAFETY BOARD
www.chemsafety.gov/investigations
2. “The Buncefield Investigation” - www.buncefieldinvestigation.gov.uk/reports/index.htm
3. “Bayer CropScience Pesticide Waste Tank Explosion”, U.S. CHEMICAL SAFETY BOARD
www.chemsafety.gov/investigations
4. ANSI/ISA-84.00.01-2004 Part 3 (IEC 61511-3 Mod) “Functional Safety: Safety Instrumented Systems for the Process Industry Sector- Part 3”
5. Marszal, E. and Scharpf, E. “Safety Integrity Level Selection”. ISA (2002)
6. Coco, James, editor, The 100 Largest Losses of 1972-2001, Marsh Risk Consulting Practice, February 2003
7. ANSI/ISA-84.00.01-2004 Part 1 (IEC 61511-1 Mod) “Functional Safety: Safety Instrumented Systems for the Process Industry Sector”
8. ANSI/ISA ISA18.00.02-2009 “Management of Alarm Systems for the Process Industries”.
9. Zapata, R. and P. Andow, “Reducing the Severity of Alarm Floods”, Proceedings, Honeywell Users Group Americas Symposium 2008, Honeywell, Phoenix, Ariz (2008).
10. “The Explosion and Fires at the Texaco Refinery, Milford Haven, 24 July 1994”, HSE Books, Sudbury, U.K. (1995).



About the Authors:

Todd Stauffer is director of alarm management services for exida and an editor / voting member of the ISA-18 committee. He can be reached at tstauffer@exida.com.

David Hatch is a Certified Functional Safety Engineer (CFSE), as well as an active member of the ISA-18 committee and EEMUA 191 working group. He can be reached at david.hatch@exida.com.

Acknowledgements:

Portions of this whitepaper were published in the September 2009 issue of Intech- "Operators on alert: Operator response, alarm standards, protection layers key to safe plants".