Session Five:

Are your Alarms Safety Related? The Role of Alarms in Functional Safety

Bob Weiss
Principal Consultant, Honeywell Process Solutions

Abstract

Successful management of process alarms allows operators to respond more promptly and more effectively to process disturbances, thus reducing the demand on Safety Instrumented Systems (SIS). In addition, some alarms are so critical, that they deserve special treatment. Drawing on several real-world examples, this paper presents techniques for identifying those alarms that are "safety related" as defined by AS IEC61508. Guidance is also provided on how to implement safety related alarms and how to assess and manage human reliability when the operator is part of the safety function.

Introduction

The failure of so-called "safety related" alarms has been implicated in a number of high profile process industry accidents. Perhaps the most pertinent of these occurred at Buncefield, north of London, in December 2005 when a gasoline tank overflowed in the early hours of the morning. The resulting mist found a source of ignition resulting in Europe's largest peace-time explosion. The principal direct cause was failure of a high level switch to activate a "High High" level alarm because it had been inadvertently bypassed (ref 1). A contributing factor was that the operators had come to rely on alarms to signal when they should stop filling the tank, so they were unaware that their tank gauging system had "flat lined" and their independent protective function had effectively been neutralised. One of the main recommendations of the report was that alarms such as these should be managed in accordance with the functional safety standard IEC61511 (ref 2).

Such alarms are termed "safety related" using the terminology in the basic functional safety standard IEC61508 (ref 3). The distinction between these alarms and normal process alarms is important, but is often not well-understood. This paper explains how such alarms can be indentified during risk assessment and SIL determination and presents requirements for their implementation. When alarms are safety related, the operator effectively becomes part of a "Safety Instrumented Function" (SIF), so quantifying the operator's reliability becomes important. Some approaches to doing this, both simple and more rigorous, are introduced.

Key Concepts

The term "safety related" in functional safety has a specific meaning. To understand this, familiarity with some basic concepts is required. The three acronyms "SIS", "SIF" and "SIL" are explained in Figure 1. A "Safety Instrumented System" (SIS) comprises a number of "Safety Instrumented Functions" (SIF) each of which serves to reduce the risk of a specific hazardous event. The “reliability” with which the SIF achieves this is termed the “Safety Integrity Level” or SIL. Only a SIF can have a SIL. The SIL is directly related to the risk reduction achieved by that SIF, as shown in Figure 2[1].

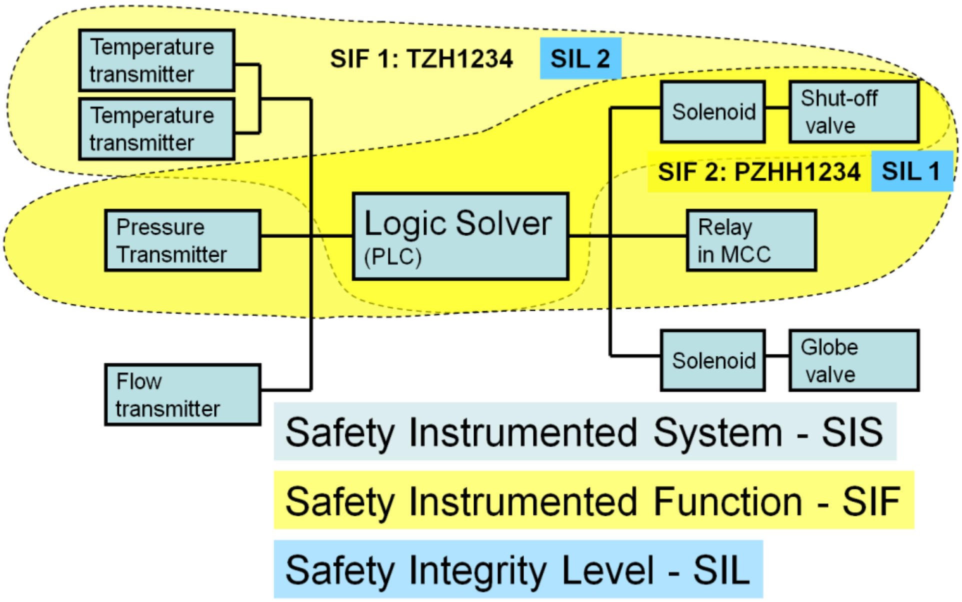

Figure 1 - The relationship between SIS, SIF and SIL

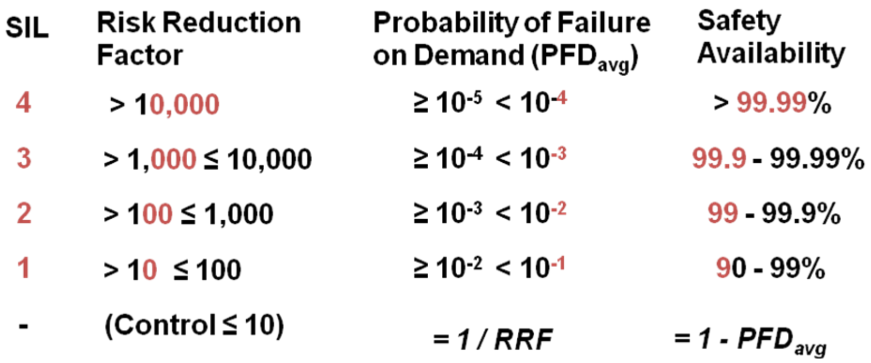

SIL	Risk Reduction Factor	Probability of Failure on Demand (PFD_{avg})	Safety Availability
4	> 10,000	$\geq 10^{-5} < 10^{-4}$	> 99.99%
3	> 1,000 ≤ 10,000	$\geq 10^{-4} < 10^{-3}$	99.9 - 99.99%
2	> 100 ≤ 1,000	$\geq 10^{-3} < 10^{-2}$	99 - 99.9%
1	> 10 ≤ 100	$\geq 10^{-2} < 10^{-1}$	90 - 99%
-	(Control ≤ 10)	= 1 / RRF	= 1 - PFD_{avg}

Figure 2 - SIL versus Risk Reduction Factor

One common means of SIL determination is “Layer of Protection Analysis” or “LOPA” (see ref 2 part 3). This may be represented diagrammatically as shown in Fig. 3. The example is for high pressure protection. A process disturbance or control system fault could result in an uncontrolled increase in pressure. Experience may show that this occurs on average once per year (say[2]). If no protective systems were provided, this would result in a loss of containment. However, the tolerable frequency for such serious events may only be once per

[1] This paper is only concerned with so-called “demand mode” SIFs. The failure measures for “continuous mode” SIFs are different.

[2] The numbers used in all examples in this paper are illustrative only, and will vary depending on the specific application.

10,000 years[3]. So the frequency, and hence the risk, needs to be reduced by a factor of 10,000. Assuming a pressure safety valve (PSV) is provided, this can typically reduce risk by a factor of 100. A further factor of 100 is required to achieve the target risk reduction. For purposes of this discussion, it is assumed that this will be shared between the high pressure alarm (PAH in the figure) and the trip (PZHH). How this is done determines whether or not the alarm is safety related, as is discussed below.

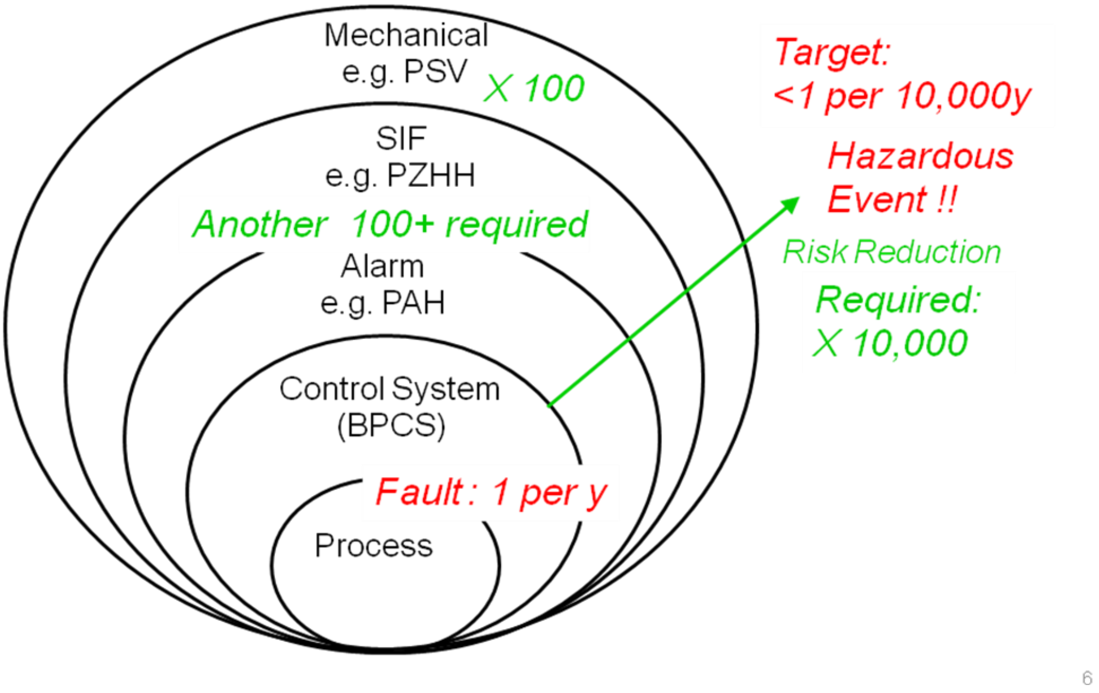

Figure 3 - "Onion Diagram" showing simplified layers of protection analysis.

Types of Alarms

Alarms can perform several different roles related to safety functions. An alarm could:

1. Reduce the demand on a SIF - this is the role of the common process alarm.
2. Prompt the operator to take an action to avoid a hazardous event - this may or may not be a “safety related” alarm.
3. Advise that a SIF has operated
4. Advise that a SIF is partly degraded (e.g. one channel of a voted subsystem is faulty).
5. Mitigate the consequence of a hazardous event (e.g. fire and gas alarms; not discussed further in this paper)

Each of these (except case 5) is discussed in more detail below.

[3] Strictly a probability of 1E-4 in any year.

Process Alarms

Referring to Fig. 3, the necessary risk reduction could be allocated entirely to the SIF, entirely to the alarm or shared. Let's look first at two cases.

1. The risk reduction is entirely allocated to the SIF. The SIF then requires a Risk Reduction Factor (RRF) of greater than 100, so must be implemented as SIL 2. In this case, no risk reduction is claimed for the alarm. Instead, the alarm's role is to reduce the demand on the SIF and keep it below the nominal once per year assumed in the analysis.
2. The alarm is allocated a risk reduction factor < 10 and the SIF takes the rest. For example. the alarm has a RRF of 5, and the SIF then requires > 20, equivalent to SIL 1. In this case, the alarm may still be implemented as a process alarm, but its role in risk reduction has been explicitly stated, but is limited.

In either of these cases, the alarm is not safety related, and may be implemented in the control system HMI as a normal process alarm. Approach 1 is generally preferred when analysing existing facilities. It can usually be readily determined how often a trip (the potential SIF) operates i.e. its demand rate. Part of the contribution to this demand rate will be failure of the process alarms to prompt adequate response. Teasing out the contribution of alarm response failure from other failures is usually difficult.

For new facilities, approach 2 may be adopted, where some risk reduction is allocated to the alarm and other failures are assessed to come up with a demand rate. Whichever approach is adopted, it is important that the alarm's contribution is not counted twice by assuming its role in demand reduction and giving it credit for independent risk reduction.

Care should be taken in assigning an alarm a risk reduction factor of 10. Although the standards equate SIL 1 to risk reduction factors greater than 10, the aim is always to achieve a risk reduction greater than the target. By setting requirements for each layer right on the upper limit of a SIL RRF category, one does not achieve this requirement. So in this example, the alarm could potentially be assigned a RRF = 10, but the SIF would then require a RRF of >10 to achieve an overall RRRF >10,000, so would still require to be SIL 1. In cases such as this, the author prefers to avoid subtle distinctions as far as possible, by assigning risk reduction factors that are not on the boundary between categories. This becomes particularly important for alarms, as we will see.

Safety Related Alarms

Consider the case where the RRF assigned to the alarm is > 10. In this case, the alarm becomes "safety related". (ISA S18.2 (ref. 4) terms this a "Manual Safety Function Alarm", a type of "Highly Managed Alarm"). In practice, this means that the alarm is part of a SIL 1 SIF. In this case the operator is the "logic solver" subsystem of the SIF and the actuation subsystem is undefined. This situation should be designed out as far as possible, as the operator is the least reliable component of the SIF. Will they see the alarm? Will they respond effectively and quickly enough?

However, although rare, there are circumstances where such alarms cannot be avoided. A common example is a flare header low temperature alarm. Under most foreseeable circumstances of depressurisation or emergency relief, the temperature will remain above the metal's minimum safe operating temperature to avoid brittle fracture, but it is conceivable that this may not always be the case. If it is not practicable to automate the response to the alarm[4], as the events leading to it or the response are ill defined, an alarm will be provided. This alarm is probably safety related, as the consequences of failure will be severe, and its RRF is likely to be > 10.

"Safety Critical" Alarms

The normal risk reduction approach does not take account of the sequence in which the different layers of protection occur. Consider two cases based on the high pressure protection example:

1. The alarm setpoint is less than that of the SIF, the usual case, often called a "pre-alarm".
2. In rare situations, the alarm setpoint may be set greater than that of the PSV. In this case, the purpose of the alarm is to warn that the PSV has failed to function.

From a LOPA point of view the RRF of the alarm is exactly the same in the two cases. However, from the operators' point of view, the response to the second case will be very different. In case 2, they are now the last line of defence in avoiding a major incident. Although the case 2 alarm will operate much less frequently than the case 1 alarm, when it does operate the operator response is much more critical.

The author prefers the term "Safety Critical Alarm" for this special type of safety related alarm. They should be avoided, but where this is not possible, their implementation requires special care, with particular emphasis required on operator training and periodic testing of their response.

Implementation of Safety Related Alarms

Given their importance as part of risk reduction, and the critical importance of effective operator response, it is essential that safety related alarms are implemented differently to process alarms. They need to be treated like any other SIF that provides an independent protective layer (ref 5). In practice, this generally means that:

1. The number of safety related alarms should be minimised through inherently safe process design and provision of automated SIFs as far as is practicable.
2. The sensors should be separate and independent from those used for process alarms and control.
3. The alarm logic should be independent of the control system. This typically means that alarm detection will be part of the SIS logic solver. An independent alarm annunciator or display could also be used.

[4] In cases where automated depressurisation can lead to low temperatures, the SIS should limit the blowdown rate to avoid the need for operator manual intervention. However, this is not always possible.

4. Safety related alarms should be displayed independently from the control system HMI, typically on a separate lamp box driven by the SIS logic solver, or via a dedicated alarm annunciator. The alarm state should be repeated in the control system for logging and actuation of clean-up functions as appropriate.
5. Because the operator is the logic solver part of the SIF, it is essential their response is "programmed" by providing specific response procedures.
6. Similarly, the "logic solver" needs to be tested periodically by ensuring that every operator knows the significance of, and how to respond to each safety related alarm. The testing should follow a documented test procedure and the results be recorded.
7. The safety related alarms should be clearly identifiable through colour coding or by other means.

Failure to appreciate the significance of a safety related alarm can have disastrous consequences. For example, in 1998 a heat exchanger at a gas plant at Longford, near Melbourne, was allowed to cool below its minimum safe operating temperature when a pump circulating warm lean oil stopped unexpectedly. The alarm signifying that the pump had stopped was displayed in red on an annunciator window along with other process alarms. (At this time, the plant was operated primarily from its original 1969 control panel.) Red was used to signify that some equipment had stopped. Most of these events resulted in, at most, lost production. A few, such as failure of this pump, were "safety related". The operators were not made aware of the distinction, and treated the interruption primarily as a production loss incident. Closure of a manual isolation valve would have averted the subsequent explosion that killed two people and injured eleven (ref 6).

Non-safety-related Alarms

There are several other types of alarms that may be erroneously regarded as "safety related":

1. An alarm that indicates that a SIF has operated should usually be treated like a process alarm, or have its own special category, unless operation of the SIF may create a specific hazardous event that requires operator intervention to prevent.
2. Diagnostic alarms that signal failure of a SIS component are important, but are not safety related in the sense used here. For example, an alarm indicating that one channel of a redundant system requires maintenance must be responded to, but has a lower priority than safety related alarms.

Quantifying Operator Error

Given that the operator becomes the "logic solver" in a SIF involving a safety related alarm, it becomes important to quantify the probability that the operator will respond correctly. This is also relevant when assessing operator performance as part of determining the demand frequency for a SIF.

The simplest approach is to assume that the overall SIF that includes a safety related alarm is SIL 1, and has a risk reduction factor of between 10 and 100.

EEMUA 191 (ref 5) supports this approach provided that the implementation guidelines described above are followed. However, in some circumstances a more quantitative approach may be warranted. The primary purpose of this section is to alert the functional safety practitioner to the existence of methods for quantifying human reliability, which are part of the Human Reliability Assessment discipline. Further details are readily available in the references provided.

A number of different techniques for quantifying human error are available. The UK Health and Safety Executive has published a useful survey of different techniques, together with assessment of their effectiveness (ref 7). Smith (ref 8 section 8.6) provides further details of the most used of these. Two that are in widespread use are TESEO and HEART, discussed briefly below. Refer to references 7 and 8 and the references contained therein for further details.

HEART

Human Error Assessment and Reduction Technique (HEART) was developed by J.C. Williams in the early 1980s. The user first assigns a probability of error to one of nine basic error task types. This ranges from 0.55 for a totally unfamiliar task performed at speed with no idea of the outcome to as low as 0.00002 for a correct response when there is an “augmented supervisory system” providing interpretation. One of 38 “error-producing conditions” is each assigned a maximum multiplier that can then be derated by the analyst dependent on the perceived applicability of the condition to the task under review. For example, a “shortage of time for error detection” could attract a multiplier of up to 11. As many of these multipliers may be applied as are considered relevant.

For example, consider an operator response to a flare header low temperature alarm. The task characteristics are:

1. The alarm occurs rarely.
2. Rapid diagnosis of the cause of the problem is required. This involves checking several well-defined possible relief sources.
3. The operator is well trained with twice yearly checks of response to the alarm.
4. The response is to reduce the flaring rate by manipulating controller setpoints.
5. The flow measurements indicating flare source flowrates are unreliable.

The basic error task type that most accurately fits this case could be:

Complex task requiring high level of comprehension and skill.

Probability of error: 0.16

or possibly

Fairly simple task performed rapidly or given scant attention

Probability of error: 0.09

It is then necessary to assess which of the 38 “Error Producing Conditions” (EPCs) are relevant, and to what extent. In this case two are chosen, and the maximum factor is derated based on the analyst’s assessment of the task.

Error-producing conditions:

	EPC	Proportion	Calculated Factor
Shortage of time for error detection	11	0.2	(11-1) * 0.2 +1 = 3.0
Unreliable instrumentation	1.6	0.5	(1.6-1) * 0.5 + 1 = 1.3

Overall probability of error = 0.16 * 3.0 * 1.3 = 0.624 (or 0.234 for the 2nd case)

Clearly this is significantly greater than would be acceptable for SIL 1.

TESEO

Whereas HEART applies to more general human tasks, TESEO was developed specifically for the plant control operator. The method requires selection of five factors that are then multiplied together. The factors are shown in the table below.

Table 1- TESEO error shaping factors

Activity	
Simple	0.001
Requires attention	0.01
Non-routine	0.1
Time stress (in seconds available)	
2 (routine), 3 (non-routine)	10
10 (routine), **30** (non-routine)	1
20 (routine)	0.5
45 (non-routine)	0.3
60 (non-routine)	0.1
Operator	
Expert	0.5
Average	1
Poorly trained	3
Anxiety	
Emergency	3
Potential emergency	2
Normal	1
Ergonomics (i.e. ***Plant interface)***	
Excellent	0.7
Good	1
Average	3-7
Very poor	10

Applying TESO to the above example could result in the following factors:

Activity:	Requires attention	0.01
Time Stress:	30 (non-routine)	1
Operator:	Average	1
Anxiety:	Potential Emergency	2
Ergonomics:	Average	4

Overall: 0.01 x 1 x 1 x 2 x 4 = 0.08

This value is significantly lower than the error rate produced using HEART. Note that the factors used are different, and in particular, the time factors considered by TESEO are very short, the longest being 60 seconds. Clearly the judgement used to select the value of each factor is crucial. In any case, it is strongly recommended that the guidance in EEMUA 191 (ref 5) be respected, and that a probability of error of smaller than 0.01 be selected. It should also be noted that the basis for these methods is primarily expert judgement, and there is generally no empirical basis for the factors chosen. However, use in practice has shown that they can provide reasonable agreement with actual human behaviour, provided that sufficient care is taken in understanding the task being evaluated.

Conclusion

Correctly identifying and implementing safety related alarms is critical to the safety of many process plants. As the human operator is effectively the “logic solver” of safety instrumented functions that rely on safety related alarms, it is essential that sufficient care is given to the human factors aspects of alarm implementation and use, and that the operator is clearly aware of which alarms are truly safety related and knows how to respond effectively. It is often necessary to quantify the human’s reliability as part of a safety instrumented function, either as one contributor to the demand on a SIF or as the “logic solver” responding to a safety related alarm. As well as simple default values, a range of techniques exist to evaluate human reliability and two of these techniques were introduced.

It is hoped that application of the techniques mentioned in this paper will reduce the occurrence of incidents due to the incorrect response to safety related alarms.

References

1. Buncefield - Why did it happen? - the Competent Authority; 2011 (available from www.hse.gov.uk)
2. AS IEC61511-2004 Functional safety—Safety instrumented systems for the process industry sector Part 1: Framework, definitions, systems, hardware and software requirements.
3. AS IEC61508-2010 Functional safety of electrical/electronic/ programmable electronic systems Parts 1 to 7 Standards Australia 2011
4. ANSI/ISA-18.2-2009 Management of Alarm Systems for the Process Industries ISA 2009
5. EEMUA publication 191 - Alarm Systems- A Guide to Design, Management and Procurement Ed. 2 Engineering Equipment Manufacturers and users Association 2007
6. The Esso Longford Gas Plant Accident - Report of the Longford Royal Commission Parliament of Victoria June 1999
7. Health and Safety laboratory; HSE guidance on human error Review of human reliability assessment methods. HSE 2009
8. Smith, David J.; Reliability Maintainability and Risk 5th ed. 2000