Jump to it! - A New Model for Safety Alarm Operator Response Time Requirements That Avoids Misplaced Conservatism.

Harvey T. Dearden BSc CEng FIET FIMechE FInstMC FIChemE (SISSuite Ltd., www.sissuite.com)

Dr Andy Brazier AIChemE MCIEHF (ABRisk Ltd., www.abrisk,co.uk)

Alarms are a ubiquitous feature of modern process systems (and beyond). They can have a role in control, prevention and mitigation. Many are configured within the Basic Process Control System (BPCS). Although instinctively we may feel that they contribute positively to safety, credit is not usually taken for them in safety studies. However, credit may be taken for operator response for a subset known as 'safety alarms,' which creates a requirement to take account of reliability of the alarm systems hardware and software; and the operator response.

The question often arises about what Risk Reduction Factor (RRF) might be claimed for operator action in response to an alarm. A risk reduction factor of 10, reducing the overall risk by one order of magnitude, is often claimed and may well be considered reasonable provided that:

The alarm is clear and prioritised such that it is unlikely to be overlooked.

AND

The required action is simple to execute.

AND

There is sufficient time for the alarm to be detected, diagnosed, and acted upon effectively to prevent the hazard being realised.

AND

The credibility of the alarm is not compromised by previous experience or operator perception (e.g., frequent spurious initiation).

It seems to be commonly received 'wisdom' that the minimum allowable operator response time for a safety alarm to act as an Independent Protection Layer (IPL) is of the order of 20-30 minutes.

This may strike users as an extraordinarily long time – particularly if the alarm presentation and action requirements meet the other stipulations (clear and simple) - and if they cannot, then taking credit for an alarm would be inappropriate. The users might well feel that they are being denied a legitimate claim for operator response because of this 20-30 minute stipulation. We believe most experienced engineers would expect an operator response to a suitably presented alarm to be much quicker; and Bridges appears to agree [Ref. 1] suggesting that "for actions that require no or very little diagnosis or in simple process units, this value can reasonably be set to five minutes". If truly 20-30 minutes, this would suggest there is something wrong with the alarm configuration or the operator competence.

The UK's Health and Safety Executive (HSE) distinguishes between Safety Instrumented Alarm Functions (SIAF) that are claimed to offer a risk reduction of more than 10 (and therefore correspond with SIL1), and those that do not, which they designate as 'Low Integrity Safety Instrumented Alarm Functions' (LISIAF). Their guidance says "A risk reduction factor of greater than 100 (i.e., SIL 2) should not be claimed for a SIAF as this would require human reliability better than normally achievable". The guidance also stipulates that for any SIAF, "the required operator response should be simple, obvious and invariant". [Ref. 2] The absence of any diagnostic burden associated with a 'simple, obvious and invariant' response argues for a minimal demand on the operator's time.

The CCPS [Ref.3] guidelines are routinely cited, although when their provenance is considered more closely their appropriateness becomes questionable. They identify (from the NUREG Handbook, [Ref. 4]) a successful operator alarm diagnosis of 90% after 10 minutes, and 99% after 40 minutes. (See section 'The Timeline of an IPL Response' within Appendix A.) The actual figure reported in the NUREG Handbook [Ref. 4] is 99% after 20 minutes. (Ref 4, Nominal model table 12-4). These are the times for the response of a control room team (not an individual operator) after 'a compelling signal of an abnormal condition'. Note also that these figures are for 'diagnosis': the recognition and evaluation of the alarm and determination of what action is appropriate.

The probability of a diagnosis failure as a function of time between the points in the table is characterised by a series of straight lines on a log-log plot (Ref 4, fig 12-4):

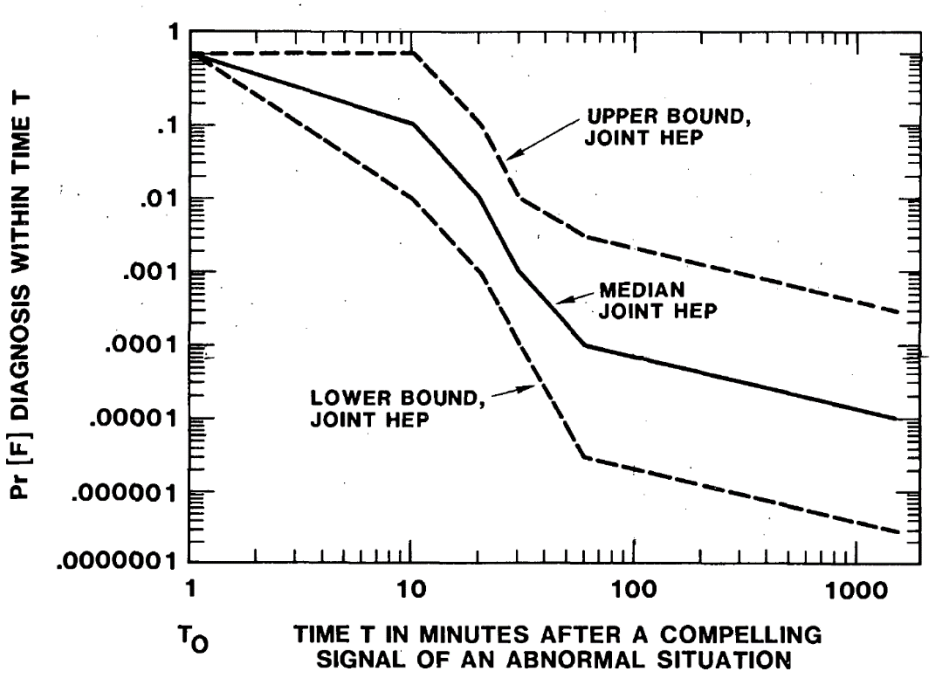

Fig 1. Probability of diagnosis failure vs time [Ref 4]

The straight line between 10 and 20 minutes represents a power law relationship:

$$P_{diag.fail} = 210.t^{-3.32} \qquad (1)$$

Time (t) in minutes. Where t ≥ 10 minutes.

For times from 1 to 10 minutes, the relationship is another straight line in which:

$$P_{diag.fail} = t^{-1} \qquad (2)$$

By subtracting $P_{diag\ fail}$ from 1, we identify the probability of a successful diagnosis. ($P_{diag.success}$)

The use of a log-log plot makes for some difficulty in 'seeing' the relationship. Here it is presented on a conventional 'lin-lin' plot for up to 20 minutes:

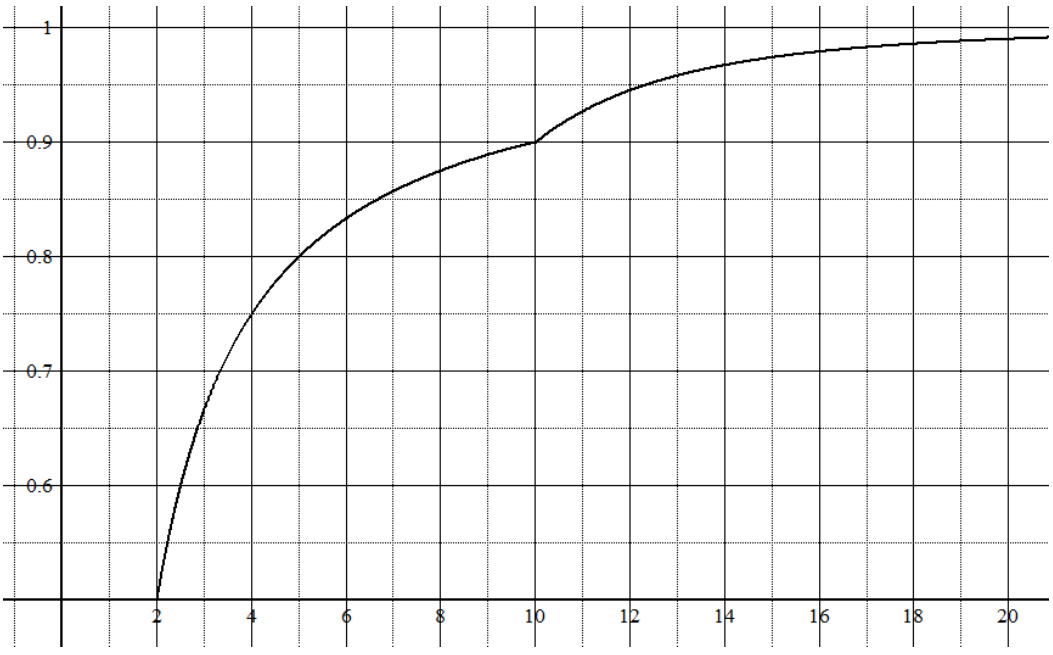

Fig 2. Probability of successful diagnosis (after first detection) against time in minutes.

The discontinuity at 10 minutes on the NUREG plot has no justification in terms of operator behaviours. It seems likely it was found convenient to draw a straight line on the log-log plot from there to the point where the probability of failure was 1 ($P_{diag.success} = 0$). A continuous function would appear a sensible proposition and the use of the upper curve function for values below the break point is conservative.

This plot is of the median probability. The limitations of the model are acknowledged in the handbook; "…at present, there is no sufficient body of data that would enable us to derive a data-based model. Hence, our nominal model is merely a gross approximation that is adequate for most PRA purposes." The actual probabilities will of course vary with the specifics of any particular alarm and its context (the handbook identifies an error factor of 10 about the median values reported), but these influences should be constrained if standards and guidance related to alarm system design including BS EN IEC 62682 [Ref. 5] EEMUA 191 [Ref. 6] are adhered to.

The report proposed a factor 10 improvement in performance ("use lower (uncertainty) bound") if:

"(a) the event is a well-recognized classic, and the operators have practiced the event in the simulator requalification exercises,

and

(b) the talk-through and interviews indicate that all the operators have a good verbal recognition of the relevant stimulus pattern and know what to do or which written procedures to follow".

The lower bound gives a 99% probability of successful diagnosis after 10 minutes (rather than the median plot value of 20 minutes).

The NUREG handbook was essentially based on loss of (reactor) coolant accidents and identified 'diagnosis' (in table 12-1) as "the attributing of the most likely cause(s) of the abnormal event to the level required to identify those systems or components whose status can be changed to reduce or eliminate the problem; diagnosis includes interpretation and (when necessary) decision-making said to offer a conservative and proportionate basis for the allocation of alarm response probabilities as a function of the time available".

This is clearly associated with a more nuanced operator process than the 'simple, obvious and invariant' response stipulated in the HSE guidance. The median probability model from NUREG may be considered correspondingly conservative. We suggest that we might reasonably equate the above 'lower bound' improvement qualification to the HSE stipulation that the response be 'simple, obvious and invariant'.

Although the extract above talks of probability as a function of 'time available', at other points in the handbook the probability plots are discussed in terms of 'elapsed time' after first detection. A subtle and potentially confusing difference. Since we are concerned here with identifying the time needed to be sufficiently confident of the correct response, our interest is in the variation of probability with the available time.

The NUREG handbook considers the possibility of the correction of erroneous diagnoses with elapsed time but given the stipulated simple diagnosis requirement we may expect no significant contribution from increasingly refined diagnostic effort which would anyway be retarded by confirmation bias. If an operator cannot be relied upon to make a successful diagnosis to a simple alarm after a few minutes, it becomes questionable whether further time will enhance performance. A correct diagnosis probability limit of 99% seems sensible.

In terms of correct response, it may be that the alarm may be superseded by other stimuli: the operator may be distracted from the required action. Rather than approaching perfection with time, performance may begin to deteriorate as the significance of the now silenced and acknowledged alarm is displaced in the operator's consciousness.

If the required response is truly 'simple, obvious and invariant' in accordance with the HSE guidance, the question arises – why not automate it? The fact it is an alarm (rather than a trip) implies some requirement for diagnosis. This is a recognition that, whilst we may often focus on human failure in regard to safety, plant operators are able to use judgment, reasoning, and decision-making skills when determining how to respond appropriately to events and circumstances. Increased automation tends to lead to more plant disturbances. Whilst the most immediate benefit may be viewed as being production (i.e., fewer disturbances), plant trips and subsequent restarts are in themselves risky.

It may be that the operator is asked to evaluate the alarm in the context of the prevailing circumstances – validating the requirement for the specified response, or to identify which of a limited choice of specified actions is the appropriate one. (Although there may be alternative actions, e.g., stop pump A or stop pump B, the response 'stop the running pump', might be said to remain 'invariant'. If there is more than one unit that could give rise to a given alarm condition, the operators will have to identify the appropriate action in the context of the specific alarm – the same action e.g., 'initiate deluge', but in relation to different units.)

Time Scaling

The expectation is of a rising probability of successful operator diagnosis that shows an initial rapid rise which then slows progressively. We may take the power law relationship identified for the probability of diagnosis failure in the NUREG handbook for times beyond 10 minutes and introduce a time scaling factor (k) to accommodate less conservative models:

$$\log P_{diag.fail} = -3.32 \log kt + \log 210 \qquad (3)$$

We may subtract this probability from one to identify the probability of successful diagnosis.

We might adopt a curve reaching 99% after 10 minutes, (corresponding with the lower band value) where the alarm management practices are considered exemplary, and the required response meets the 'simple, obvious and invariant' stipulation. This may be modelled using the same relationship but scaling time by a factor k = 2.

An intermediate model, reaching 99% after 15 minutes, may be identified with time scaling factor k = 1.33.

The corresponding curves (without deadtime – see below) are shown as dashed on the plot.

Note that the logarithmic curve model implicit in the NUREG model has been accepted as a reasonable basis for modelling the growth of probability with time. The actual nature of the probability-time relationship will likely depend on a variety of factors relating to individual alarms and the context in which they arise. That said, our expectation is of a relatively fast initial rise with a progressive slowing, which the log relationship provides with a relatively distinct 'knee' compared with say, an exponential curve, which might have been considered as an alternative candidate.

The user might identify an appropriate value for k by considering what period must be available (T99), following recognition of the alarm, for there to be essentially complete confidence (>=99%) that the operator would identify the right action.

$$k = 20/T_{99} \qquad (4)$$

It would seem prudent to impose a low limit of 2 minutes for T99, (k=10) which might only be employed if the alarm was of the most unequivocal nature, in a control room where exemplary alarm management practices are adopted.

With k=1 we have the NUREG median probability, and we may regard this as a highly conservative model.

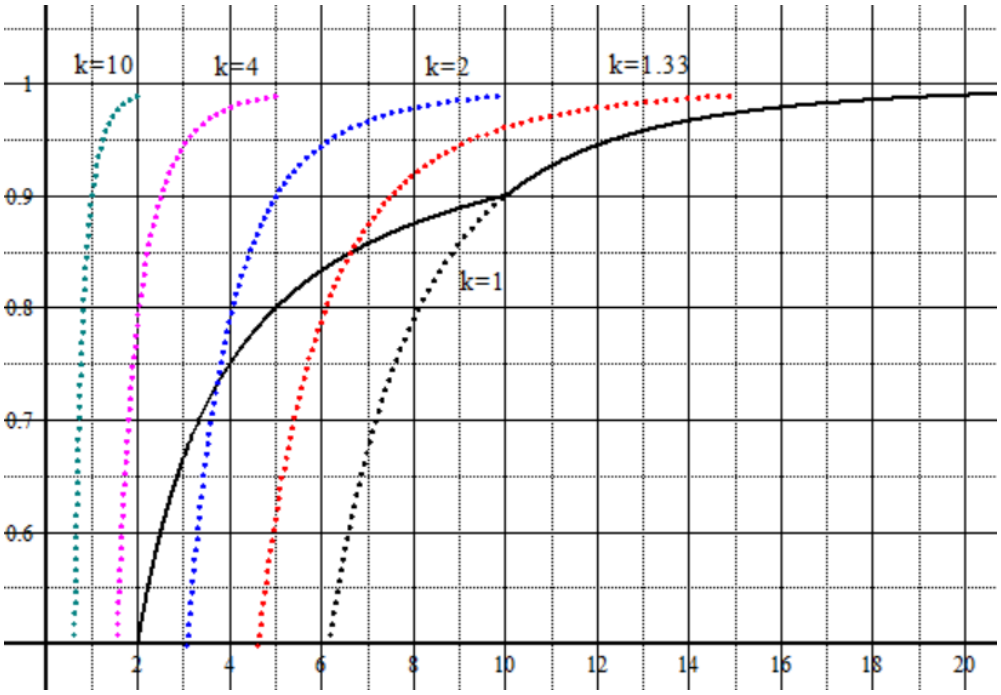

Fig 3. Probability of successful diagnosis (after first detection) against time in minutes, for different values of k.

Alarm Recognition-Detection Deadtime

Diagnosis is only one component of the time required to deal effectively with a condition creating an alarm. We may add a recognition-detection deadtime on the basis that there is a minimum period within which no diagnosis may be expected: the operator may be distracted by some other activity (e.g., phone call) and there may be a delay before his/her attention is properly directed to the alarm. (Its effect is to shift the curve to the right by the deadtime value.

$$\log P_{diag.fail} = -3.32 \log k(t - T_D) + \log 210 \qquad (5)$$

$$P_{diag.success} = 1 - 10^{[-3.32 \log k(t-T_D) + \log 210]} \qquad (6)$$

Deadtime is not identified in the NUREG model; the response times are from the time at which the operators "…notice that an abnormal condition exists".

We may couple estimates for the required diagnosis time with estimates for dead time:

1 minute: Best possible. Continuous attendance by more than one operator able and authorised to respond.

2 minutes: Continuous attendance by one operator. More than one may be in attendance, but not guaranteed to be continuously more than one.

5 minutes: Continual (as distinct from continuous) attendance by one operator but with the possibility of intermittent and brief distancing from the control panel e.g., to an adjacent mess room.

10+ minutes: Continual attendance by one operator with occasional excursions to the field whilst carrying a pager alert prompting immediate return. Time to be adjusted with anticipated potential remoteness from the control room.

The human factors to be considered in nomination of k and the dead time are identified in the appendix.

There is a separate question of the possibility of error in the action taken after a correct diagnosis ($P_{corr.act.}$). Given that the response for a safety alarm should be simple, we might assign a probability of successful action of 0.99. The probability of the correct response then becomes:

$$P_{corr.resp.} = P_{corr.act.} \times P_{diag.success} \qquad (7)$$

Contributions to Overall Function Probability of Failure on Demand (PFD)

If an overall alarm function PFD average of 0.1 is claimed and the failure rate for the alarm system, without the operator, (sensor-logic-annunciator-final element) is 0.1/year, (typical for a BPCS provision) then, if tested annually, the PFD contribution from the alarm system would be 0.05, leaving a corresponding 'headroom' of 0.05 for the operator contribution. The operator would need to respond correctly 19 times out of 20. Using the conservative model above, this would require a minimum allowable operator diagnosis time of 13 minutes without any deadtime allowance. If we assume 1 minute deadtime, the operator must not need to act before 14 minutes have elapsed. This may debar an alarm IPL claim for many hazard scenarios.

If the alarm system was engineered separately from the BPCS and was 'SIL1' rated, we may expect an improved system PFD contribution, say 0.01 - this would then leave a 'headroom' of 0.09 if an overall PFD of 0.1 is to be achieved.

If a higher risk reduction overall is to be claimed, for example a mid-SIL1 band PFD figure of say 0.03 (RRF 33), then using the same SIL1 alarm system 0.01 PFD, the operator contribution 'headroom' would be 0.02.

If a modest risk reduction factor 4 is to be claimed with a BPCS alarm system designed to contribute PFD of 0.05, then the operator PFD contribution must not exceed 0.2.

Typical System Values

In summary, for a range of values for k:

Overall PFD Claim	RRF	Alarm System PFD	'Headroom' PFD	$P_{diag.success}$ Required**	Minimum Allowable Operator Response Time (min) AFTER dead time has elapsed.*		
					K=4	K=2	K=1
0.1	10	0.05 (BPCS annual test)	0.05	0.96	3	7	13
0.25	4	0.05 (BPCS annual test)	0.2	0.81	2	4	8
0.1	10	0.01 (SIL1)	0.09	0.92	3	5	11
0.03 (mid-SIL1)	33	0.01 (SIL1)	0.02	0.99	5	9	18

Table 1. Typical System Values

*Note that the appropriate detection deadtime must be added to these values to identify the overall minimum time requirement.

**Assuming probability of correct response of 0.99

General Equation

In general, we may identify the Minimum Allowable Operator Response Time ($AORT_{min}$) in minutes needed to support a given target PFD claim as:

$$AORT_{min} = T_D + \frac{10^{\{[\log(P_{diag.fail}) - \log 210]/-3.32\}}}{k} \quad (8)$$

Where:

$$P_{diag.fail} = 1 - P_{diag.success.} \quad (9)$$

$$P_{diag.success.} = \frac{P_{corr.resp.}}{P_{correct.act}} \quad (10)$$

$$P_{corr.resp.} = 1 - \left(PFD_{claim\ target} - PFD_{alarm\ system}\right) \quad (11)$$

$P_{corr.resp.}$ being the minimum required to meet the target.

(This equation looks complicated, but it isn't really. It looks complicated because it represents a chain of calculations and uses a power law.)

Available Operator Response Time

We must distinguish between the time for the operator to recognise and diagnose the alarm (Operator Response Time) and the time taken for the intervention to be effective. If the time taken for the required action, AND for the process to respond, is a significant proportion of the available process safety time (PST), the time available in which the operator must respond (AORT) is reduced correspondingly.

There are three time components:

Available Operator Response Time (AORT)

System Response Time (SRT): Sensor delay (usually negligible) + time for required action execution.

Process Safety Time (PST): Time from alarm trigger point to realisation of the hazard if no intervention.

$$AORT = PST - SRT \quad (12)$$

Given the requirement that any action required should be suitably simple, the action execution time will typically be short, but in some circumstances, even after action is initiated it may take a significant time to execute e.g., stroke time on large remotely operated valves, operation action in the field, evacuation of personnel.

If the process safety time (PST) from the alarm point to the latest time at which action could prevent the hazard was 30 minutes and the action (SRT) takes 15 minutes, then there would only be 15 minutes available for the operator to respond (AORT). If the action took 10 minutes, the available operator response time would be 20 minutes.

For a safety alarm IPL claim to be permissible, the available operator response time must be greater than, or equal to $AORT_{min}$, the minimum allowable operator response time (which is needed to give the required probability of the correct alarm response). If you want a healthy alarm action to be guaranteed to prevent a subsequent trip operation, in assessing the process safety time, the 'hazard' point should be set to the trip point.

In terms of hazard prevention (as distinct from trip prevention), an independent alarm may be claimed as long as there is sufficient time for effective action to prevent the hazard being realised. The alarm and trip may be regarded as two separate layers: if the trip should be in a dangerous failed state, the alarm would still offer a defence. However, this assumes that the operator would respond correctly to the alarm prompt regardless of the trip activation. If there is any expectation that the operator would not take action if s/he anticipated or saw the trip initiation, then the available operator response time would be restricted to a value that would prevent the trip point being reached.

Note that IPL alarms that are triggered by Basic Process Control System (BPCS) failure should NOT be claimed in defence against such failures unless both the alarm and the means of corrective action are implemented with sufficient independence from the failed control system.

Many excursions that are triggered by control failure are likely to prompt operators to switch the affected control to manual, but if the control valve or the associated indications are themselves compromised, the manual action may well fail.

Concluding Remarks

It appears that the reference that underpins the widely cited figures for operator safety alarm response times is not well aligned with the circumstances prevailing in modern process operations with properly managed safety alarm provisions. The routinely cited requirements of 10-30 minutes do not match many users' expectations. That said, the underlying power law model does meet expectations of a progressively slowing rise with available time of the probability of successful diagnosis in response to an alarm. The original (1983) NUREG model was developed by consensus judgement of what was representative of the control room team response to abnormal events in nuclear power plants. On closer examination, it becomes clear that these judgements were in respect of a fairly sophisticated diagnosis that some (including HSE [Ref. 2] may not accept as appropriate for a safety alarm. However, humans can be very capable at applying judgement under appropriate conditions, and so allowing them to use this can have a positive effect on overall risk of an operation.

The NUREG model was acknowledged at the time by the authors to be 'speculative'. It would be specious to speak of the model now being 'calibrated': the model remains speculative but by accepting the median probability curve as a conservative boundary case (being constituted for more sophisticated diagnoses) and adjusting it for 'simple, obvious, invariant' responses, it may, with the appropriate nomination for scaling factor k and the addition of deadtime, be considered to offer a practicable and proportionate basis for evaluation purposes.

Unwarranted conservatism in the claims permitted for alarms will place a higher burden on other protection layers and may result in more complicated requirements in respect of automated Safety Instrumented Functions. Although 'on paper', a defence avoiding a claim for an alarm and relying on a SIL3 SIF say, may be shown to offer risk equivalence to an alarm + SIL2 SIF, that is not to say it should be preferred. This paper has been focussed on safety alarms, but applying similar logic more widely should build the case for better alarm system design and management.

Ref 1. Bridges, William G., 2017, Proven Approaches to Ensuring Operators Can Respond to Critical Process Deviations in Time (Human Response IPL), 13th Global Congress on Process Safety, San Antonio, Texas, USA.

Ref 2. Operator Response within Instrumented Safety Functions in the Chemical, Oil & Gas, and Specialist Industries, Appendix 1, HSE, 2018.

Ref 3. Guidelines for Initiating Events and Independent Protection Layers of Protection Analysis, Appendix A, AIChE, Wiley, 2015.

Ref 4. Swain and Guttman, 1983, NUREG/CR- 1278 Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications.

Ref 5. BS EN IEC 62682:2022, Management of alarm systems for the process industries, 2023.

Ref 6. EEMUA 191, Alarm systems - a guide to design, management and procurement, 2014.

APPENDIX: Human Factors to Consider.

The following provides a narrative that could be used to support expert judgement when determining operator alarm response deadtime and diagnosis. Whilst the factors described can be used to assess existing circumstances, they will be most powerful if used to improve alarm system design and management.

Noticing.

Time to realise an alarm requires attention. (Dead time contribution)

Operator needs to hear the audible signal (if there is one). Reliability depends on:

Factor – time to realise	Optimal arrangement	Things to consider
Volume of the alarm's audible output. Other sounds masking the alarm's audible output.	Low background noise at all times so that alarm audible volume can be set to always be heard.	Intermittent sources of noise (e.g., people talking, equipment start-up, operator prompts, maintenance work).
Distance of operator from audible output	Operator always present within short distance of the audible output. Control room always manned by same number of fully competent operators.	Size of control room. Operator leaving the control room to get a drink, reset a switch. Operator breaks covered by reducing manning in control room or by people with lower competence.
Items insulating sound between alarm's audible output and operator	No physical barriers. Repeater audible outputs in all locations operator may be present.	Items in the control room creating a barrier (e.g., control panel, DCS screens, room dividers, acoustic panels, mobile notice boards). Wall/door if operator ever leaves the control room.
Discrimination between the audible output and other sounds	Unique sound that cannot be confused. Unique location of the output giving directional attention.	People can discriminate between sound creating devices (e.g., bell, buzzer, air horn) but less able to discriminate different tones from the same source. Computer generated alarms from the same location are difficult to reliably differentiate and may all be treated as equal.
Distress caused by the sound	Not too loud, which would startle the operator. Not high pitched, which can make it difficult for the operator to think.	Difficult to choose a suitable sound when background noise levels vary, especially if the operator may be at different locations (nearer or further away). Consider the effect if the operator is concentrating on a task and the alarm sounds unexpectedly.
Operator's hearing health	All of the above should consider this.	Hearing health will vary between individuals, with age being a significant factor.

Table 2. Human Factors (Noticing)

Rule of thumb for time to realise an alarm requires attention: allow 10 seconds for the operator to realise an alarm requires attention if optimal arrangements can be guaranteed all of the time. Add the worst-case time delays for any sub-optimal arrangements, particularly if the operator leaves the control room, in which case the maximum time away should be added.

Identification.

Time to identify the alarm. (Dead time contribution)

Operator needs to accurately identify the item effected (especially if there are duplicates/similar items) and understand its meaning. Reliability depends on:

Factor – time to identify	Optimal arrangement	Things to consider
Clarity of the visual indicator.	Item can be identified from anywhere in the control room. For a very small number of alarms a spatial indication may be enough, provided the operator knows the identification of each. In most cases text will be required so the size of the characters has to be legible from any place in the control room.	Reliance on spatial indication may not be reliable (operator thinks they recognise the location but confuse which indicator is active or identify it incorrectly). Text has to be very large to be read from any significant distance with reliability.
Distance of operator from visual indicator	Operator is always close enough to reliably identify the alarm	If operator may be some distance from the indicator they may be inclined to assume or guess the alarm identification.
Light levels	Spatial indicators (assuming lights are used for this) are easily visible above background light. Optimum light levels mean all text is easy to read.	Light levels and colour may vary according to time of day (sunlight). Operators may adjust light levels (may turn off lights). Low light levels may make fixed labels difficult to read. High light levels cause glare on screens.
Operator's eyesight	All of the above should address this.	Eyesight will vary between individuals, with age being a significant factor.
Operator attention	Indication is clear that immediate attention is required, no matter what else the operator is doing.	Operator may be attending to something they perceive to be important (e.g. high priority alarm, critical communications).

Table 3. Human Factors (Identification)

Rule of thumb for time to identify: allow 10 seconds for the operator to identify an alarm if the number of alarms is very small and optimal arrangements can be guaranteed all of the time. Increase this to 20 seconds where the number of alarms is greater (5+) but optimal arrangements are guaranteed. Add the worst-case time delays for any sub-optimal arrangements, particularly if the operator has to move closer to the visual indicators to identify the alarm and/or they may be engaged in other activities that may be perceived as more important.

Diagnosis.

Time to diagnose the alarm. (T99 contribution)

Operator needs to accurately identify the cause of the alarm. Reliability depends on:

Factor – time to diagnose	Optimal arrangement	Things to consider
Alarm clarity	The cause of the alarm is unambiguous and can only have one cause.	Alarm may have several causes. The most common occurrence will be most familiar to the operator and so they are likely to assume that cause. The less frequent instances may be associated with the more significant consequence. Alarm may be false. Operator action may be different if they know (or suspect) it is false.
Alarm response clarity	Same response to the alarm every time it occurs.	Providing an alarm instead of automated response usually means there is some requirement for human judgement. Response to a false alarm is likely to be different to confirmed genuine in practice, even if procedures say every alarm should be assumed to always be genuine.
Access to supporting information	Minimal information required to diagnose and all is immediately available and visible.	Different causes of alarm may require different supporting information. Human tendency to look for information to support their diagnosis rather than considering all potentially relevant information.
Obtaining information from the field	There is no requirement for obtaining information from anywhere outside of the control room.	Reliability of communication. Possible location of person asked to attend the field. Competence of the person to evaluate what they see, hear, smell etc.
Operator knowledge of alarms and causes	For a very small number of alarms the operator knowledge can be high and immediately recalled.	If there are multiple alarms the operator is unlikely to remember the details of each. They will use a mental model to make sense of the alarm and so direct where to look for supporting information. Operator may have to find and read a procedure to support their diagnosis.

Table 4. Human Factors (Diagnosis)

Rule of thumb for time to diagnose: allow 10 seconds for the operator to diagnose the cause of the alarm if the number of alarms is very small and optimal arrangements can be guaranteed all of the time. Increase this to 20 seconds where the number of alarms is greater (5+) but optimal arrangements are guaranteed. Add the worst-case time delays for actions taken by the operator to access and interpret supporting information. If this includes asking for information from the field the time taken for communication and travel to the location has to be added.

Decision.

Time to decide what to do. (T99 contribution)

Operator needs to decide to act. Reliability depends on:

Factor – time to decide	Optimal arrangement	Things to consider
Clarity and reliability of the alarm	Operator does not need to decide anything. If the alarm sounds they have to implement the defined action.	Operator may have self-doubt, concerned that they may be overreacting. Time taken to check more information.
Authority to respond	Operator decides without any hesitation and with no requirement for validation from others.	Operator is required to inform a superior or obtain permission to act. Or perceives they need to. Less experienced operators more inclined to want reassurance from others, not necessarily a superior.
Operator instinct	Willingness to over-react in the first instance, knowing the response can be scaled back later.	Likelihood of under reacting in the first instance, hoping they can implement the full response later if it proves necessary. Overly optimistic view that the situation can be controlled, or the plant can be shutdown manually without activating an emergency (crash) shutdown.

Table 5. Human Factors (Decision)

Rule of thumb for time to decide: no time required for the operator to decide to act if optimal arrangements can be guaranteed all of the time. Add 20 seconds for operator hesitancy. If permission to act will be asked for there could be a significant time delay, although contact may already have been made earlier in the process, so this may not be the case (i.e., if the operator's early action is to inform the supervisor of an alarm, they may still be in contact by the time a decision has to be made). A tendency to under-react in the first instance could be significant if the effectiveness of the full reaction is reduced if it is started later.