Process Safety and Environmental Protection, Volume 87 - No. 1 - January 2009

Hazards Equal Trips or Alarms or Both

Clive Timms

C&C Technical Support Services

www.silsupport.com

Abstract

Anyone who has been involved in the application of IEC 61508and IEC 61511by undertaking the Safety Integrity Level (SIL) determination for Safety Instrumented Systems (SIS) will appreciate the amount of effort and tenacity that is required to undertake the task. SIL determination of Safety Instrumented Systems requires considerable commitment and tenacity to get the job done, but it is like climbing to the top of a hill only to be faced with a mountain when we come to consider what is involved in reviewing or configuring a typical alarm system.

A medium sized process facility may have a few hundred or so primary Safety Instrumented Functions (SIF) or trips configured into a Safety Instrumented System, but the number of alarms configured into a process control system (PCS), that need to be assessed and prioritised, can often run into the thousands.

There is synergy between safety instrumented functions and alarms because they both make a contribution to reducing the risk of having unwanted events, and both need an assigned their appropriate criticality.

This paper details various methods of criticality assessment which have been successfully applied to set the appropriate priority, identify the critical alarms that need to be upgraded to trips and to rationalise those of no value. It will also cover the use of software tools which can significantly reduce the effort involved in this process.

Keywords

Alarms, trips, prioritisation, rationalisation, SIF, SIL, risk

Introduction

Anyone who has been involved in the application of IEC 61508 (IEC, 1998-2000) and IEC 61511 (IEC, 2003) and the Safety Integrity Level (SIL) determination for Safety Instrumented Systems (SIS) will appreciate the amount of effort and tenacity that is required to undertake the task. SIL determination of Safety Instrumented Systems, or shut down systems as they are traditionally called, requires considerable commitment and tenacity to get the job done, but it is like climbing to the top of a hill only to be faced with a mountain when we come to consider what is involved in reviewing or configuring a typical alarm system.

A medium sized process facility may have a few hundred or so primary Safety Instrumented Functions (SIF) or trips configured into a Safety Instrumented System. These need to be assessed and assigned an appropriate SIL, but the number of alarms configured into a process control system (PCS) that need to be assessed and prioritised can often run into the thousands. The requirements for alarms usually involve

different disciplines such as instruments, process, maintenance and the operators themselves. The latter often have the misconception that their life will be easier if they have alarms on everything. Thus the demand for more alarms, along with the ease of configuration afforded by PCS's, regularly leads to a proliferation of alarms. In other words, alarm configuration can all too easily get out of hand.

There is synergy between safety instrumented functions and alarms because they both make a contribution to reducing the risk of having unwanted events, and both need an assigned criticality. It is also important to be able to determine when an alarm should be upgraded to a trip to provide automatic protection, and conversely, when a trip can be downgraded to alarm status.

A SIF is engineered to provide protection against a hazard caused by some kind of failure, and has a concise and automatic role to play when a process moves out of its normal operating envelope. Using good practice to comply with the IEC 61508/61511 standards, a risk assessment can be undertaken to determine its criticality or Safety Integrity Level (SIL).

This risk assessment is related to the consequences that would occur if the SIF were to fail on demand and the frequency of a demand. The consequences can be any combination of safety, societal, financial and environmental impact.

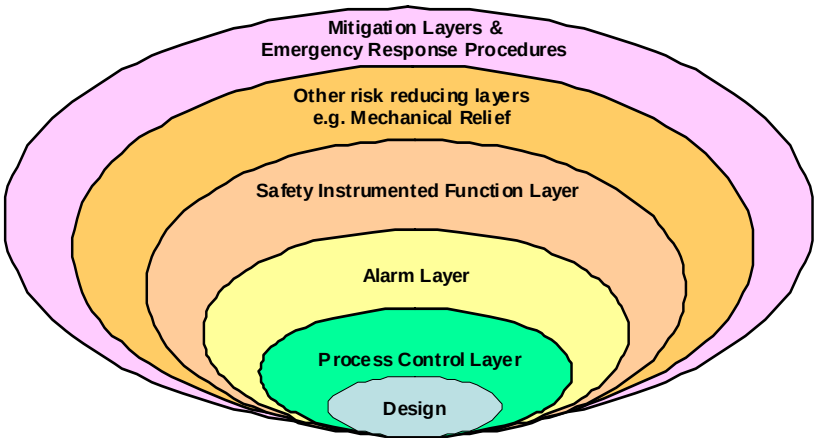

Figure 1: Typical Risk Reduction Layers

An alarm function works through the human interface – 'A Methodology for Alarm Classification and Prioritisation' - (Timms, 1999) to provide an early warning that the process has moved away from the normal operating envelope to:

- Alert the operator to disturbed plant conditions,
- Provide indication of further developments that may need attention,
- Trigger a trained operator response.

Alarms normally contribute to the overall risk reduction as they represent one of the many typical independent risk reduction layers as shown in Figure 1.

The criticality of the alarm should also be assessed in order to set an appropriate priority. It makes sense to assess the criticality of an alarm in a similar way to a SIF but based on the consequences that would follow if the alarm fails or is missed by the operator.

However, the contribution that an alarm makes to risk reduction can become clouded if the operator cannot identify the important alerts against a background of alarm problems. The three main problem areas that can potentially compromise safety, production and the environment are:

- Nuisance alarms.
- Standing alarms.
- Alarm avalanches or floods.

Nuisance alarms and standing alarms are usually caused by instrument faults, out-of-service equipment or inappropriate limit and/or dead-band settings. They can be relatively easily identified and rectified by maintenance or adjusting the configuration parameters. However, alarm avalanches or floods are usually the result of consequential or secondary events following a primary event, and the more alarms that are configured; the more there are to appear before an operator in a plant upset condition. The problem for the operator is how to distinguish between the primary initiating event and the secondary consequence events.

The primary objective must therefore be to rationalise the alarm system to a configuration which alerts the operator to alarms in order of importance, so as to give him/her the best chance to take corrective action. Inability to take corrective action can have significant safety, economic and environmental consequences. We must also eradicate those alarms that serve no purpose as this will significantly reduce the alarm overload. This can be achieved by a well defined methodology, and the effort can be significantly reduced by engaging specialist software tools as discussed later.

The UK HSE Position

The UK HSE often uses the Texaco Refinery explosion in 1994 as the prime example of how the poor application alarms and human factors can result in serious incidents – Health & Safety Executive 'The explosion and fires at Texaco Refinery, Milford haven, 24 July 1994, HSE, 1997'. This paper is not going to regenerate the UK HSE findings, but their position on alarm handling has been made very clear. They have produced an UK HSE information sheet 'Better alarm handling' (U.K. HSE, 2000) to provide some basic guidance, and the Hazardous Installations Directorate (HID) have outlined their strategy with respect to inspection and enforcement, and their expectations with respect to users and designers, in an article entitled 'Better Alarm Handling – a Practical Application of Human Factors' (Wilkinson and Lucas, 2002).

In both publications the UK HSE guidance provides a simple 3-stage approach:

- Find out if you have a problem
- Decide what to do and take action
- Manage and check what has been done

The HSE also reference the Engineering Equipment and Materials Users Association (EEMUA) publication 191, 'A Guide to Design Management and Procurement' (EEMUA, 1999) as 'the nearest thing to a standard currently available'.

Undertaking an Alarm Rationalisation Exercise

Avoid the common pitfalls

The initial reaction when faced with alarm problems can often be to look for ways of using technology to suppress unwanted alarms. PCS vendors are eager to demonstrate how sophisticated their technology can be and commit their customers to using these techniques. There may be possible scenarios where suppression of alarms is simple (e.g. main and standby equipment with auto changeover) but as a rule, the more complex the plant then the more complex the suppression scenarios, leading to very time consuming and complicated solutions. It is also all too easy to loose the focus due to the complexity, and this could result in flawed logic for the scenarios and hence compromise the alarm integrity.

An alarm flood reduction will almost certainly require a rationalisation exercise to challenge each alarm and reduce the number of configured alarms. In essence, an alarm review following the methodology outlined below in conjunction with software tools to aid the process will achieve the most significant benefits. Channelling efforts into this type of activity should be the first priority.

Software Tools will help

The quantity of data to be manipulated, sorted and rationalised will often be considerable, and can amount to thousands of alarms on a modest size process facility. It makes little sense to undertake an alarm review as a paper exercise, since dealing with large numbers of alarms will simply overwhelm those involved in the process, and the final paper report will be hard to manage and update as it will only represent a snapshot in time It can be tempting to use spreadsheets to manipulate the data, but they are not the most appropriate solution since they do not have the integrity or sorting power afforded by database based approaches.

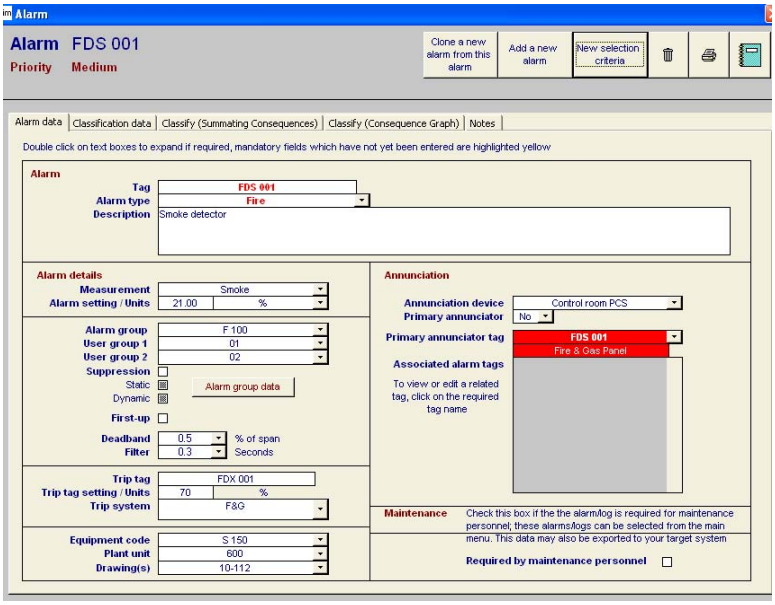

Figure 2: Alarm Data Form

In this context, it is sensible to invest in a good software application tool as this will be a significant aid to managing and maintaining your alarm configuration as well as the integrity of a 'master' alarm database. Experience shows that the best tools are database based with a good user interface to facilitate the manipulation and sorting of the large quantities of data involved in an alarm review. They will save considerable time and effort in the execution of the process. Figure 2 shows an example of a typical alarm data form which could be manually populated, but the advantage of a database structure, is that it can be pre populated from existing listings, or by bulk data exchange vehicles from the current PCS configuration, such as simple spreadsheets.

Each alarm will then have an entry in the master database with the typical parameters shown in Figure 2. These parameters are the 'keys' for the sorting and grouping of alarms to aid the all important prioritisation and rationalisation process that then follows. It is also very important to be able to handle alarms that appear on more than one annunciation device e.g. a Fire and gas alarm may have primary annunciation on the Fire and gas systems and secondary annunciation on a PCS.

Some tools can be PCS system specific, but the best choice will be an open system tool with the capability of importing an alarm configuration from any PCS, via a suitable data exchange or conversion.

An appropriate tool should then be capable of performing the following:

- Importing the PCS alarm configuration;
- Handling multiple annunciation;
- Sorting on various alarm parameters such as type, group, measurement;
- Selection by parameter or type;
- Cloning new alarms from existing alarm template;
- Cloning a selection from an existing template;
- Performing alarm prioritisation;
- Producing alarm metrics;
- Maintaining the master database configuration;
- Exporting alarm configuration to the PCS;
- Producing reports and statistics.

Methodology

The methodology for undertaking an alarm review detailed in this paper embraces both the (EEMUA) publication 191, 'A Guide to Design Management and Procurement' (EEMUA, 1999) and the U.K. Health & Safety Executive's own guidance, 'Better Alarm Handling' (U.K. Health & Safety Executive, 2000), but it has been enhanced to offer a very practical and pragmatic step-by-step approach to achieving a high degree of success on established brown field facilities or proposed green field developments. It will also describe how software based tools can significantly improve the process.

The team

It is essential to set up an optimal alarm review team to provide productivity and quality of output. A small team is recommended for alarm reviews:

- a process engineer, preferably with operational experience, to interpret process design information and to ensure that the design intent is not compromised;
- a control room operator, preferably from the facility under review, to provide information on operator requirements and likely plant dynamics and to ensure that the alarms provided meet the needs of operational personnel;
- an automation and control engineer, preferably with experience of the relevant type of PCS, to advise on aspects of implementation and to ensure that the proposed alarm changes are correctly described for implementation personnel.

For large review studies, it may be desirable to assign a team leader or facilitator to ensure that the correct balance is struck between effort and results. Such a person must have a good understanding of the alarm review objectives and methodology, and should preferably have previous experience of alarm reviews.

Planning

It is anticipated that there will be eight main review phases:

1. Management Plan.
2. Preparation of documentation and data sourcing.
3. A review of alarm system performance.
4. Categorising and functional grouping of alarms.
5. Prioritisation and rationalisation.
6. Assessing results and findings.
7. Other considerations e.g. alarm suppression.
8. Implementation of the changes.

A planning schedule should be set up for each phase.

Management Plan

There has to be management commitment before an alarm review can proceed, as an alarm review will use significant time and resource. Where process plants have comprehensive automatic shutdown facilities, then process deviation alarms that are missed, or not acted on by the operator, will invariably lead to trips. So it helps to develop a rationale for the alarm review which defines the problem with an analysis of the following:

- Details of the number and frequency of trips
- Estimates of the cost of loss production
- Details of unsafe consequences (e.g. near misses, injuries)
- Details of any environmental impact (e.g. increased flaring)

Provide estimates of the cost and duration of the review and predicted payback from improved performance and present this to management. Once they understand the economics then approval should be forthcoming.

Preparation of documents and data sourcing

The following documents must be available at the start of the review:

- alarm schedule in suitable electronic format;
- P&I diagrams;
- shutdown Cause and Effect drawings;
- details of fire & gas system;
- details of any controlled sequences;

- configuration details for any complex points, such as digital composites;
- a list of standing alarms during typical steady operation;
- alarm journal print-outs following typical upsets.

The following documents would also be useful:

- a definition of the PCS network configuration;
- a definition of any application programs handling data within the PCS.

To minimise delays, all preparatory work should be done prior to the team convening so that no valuable time is lost during the review process.

A review of alarm system performance

Using alarm journal printouts or output from other dedicated logging facilities, review the existing alarm system performance in order to identify alarm frequency, nuisance alarm sources and standing alarms. Whatever type logging facility is used it must be fast enough to provide true sequence of events recording (SER). Logging facilities on PCS systems can fail to have sufficient speed and buffer capacity to capture the real time snapshot of events resulting in SER overload and indeterminate time stamping.

Each nuisance alarm should be subjected to a detailed review to determine if it caused by a fault or inadequacy in a measurement instrument or the actual alarm configuration. If it is the latter then the alarm trigger point and dead-band should be checked to see whether adjustment of these parameters would eliminate the problem.

If possible, obtain logs of alarm performance under a variety of process upset conditions as this will provide an indication of the number of alarms generated and the frequency at which they are generated. Compare these figures with Appendix 11 of the (EEMUA) publication 191, 'A Guide to Design Management and Procurement' (EEMUA, 1999) which sets out guidance on performance metrics.

The majority of standing alarms materialise from spared equipment which is not running spared or that operate intermittently. List any associated alarms that will be active when the equipment is shut down, or put on stand-by, as a source of generating standing alarms when the plant is operating normally. These will be prioritised as outlined in the next section.

Categorising and functional grouping of alarms

The alarm schedule must be broken down into selected or batches or categories to match the major functional groups of alarms that can be considered as single entities during the review e.g. all alarms of a specific type within a certain area such as fire or gas. The time involved in the next phase, the actual prioritisation and rationalisation process, will then be significantly reduced. This categorisation exercise is where the use of software tools can provide a powerful aid to the facilitation and preparation of a comprehensive the alarm schedule.

Providing the selected tool has a comprehensive data selection structure, similar to that shown in Figure 3, it is a simple matter to sort alarms on a wide range a of pre defined parameters such as tags, alarm groups, alarm types e.g. High (the pre trip alarm), high high (the actual trip alarm), low (the pre trip alarm), low low (the actual trip alarm), open, closed, not open, not closed), measurement types, annunciation device, alarm priority, plant unit, equipment code etc. The user can also define specific parameters on which a selection is made. Each grouping is prioritised as a single entity as detailed in the next section.

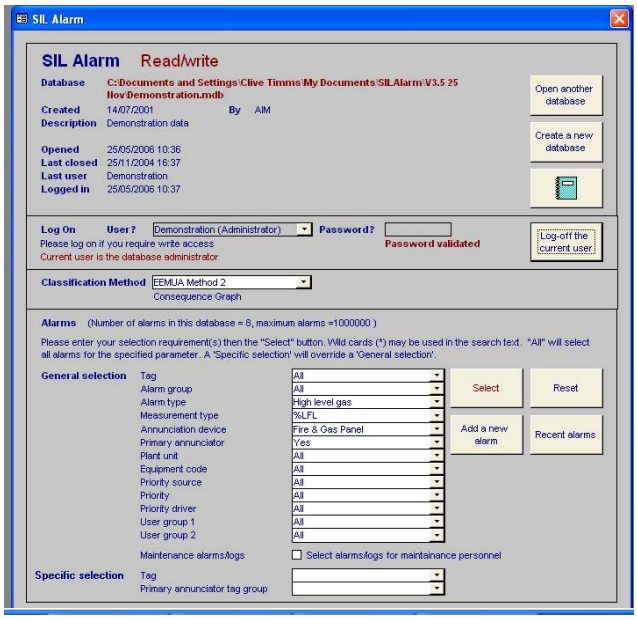

Figure 3: Alarm Categorisation and Selection Form

It is likely that many of the alarms to be considered during an alarm review process will not be covered, by the functional groups as discussed above. However, these will be readily revealed from a database and they will have to be reviewed individually.

Determine primary annunciation

It is also important to determine the primary point of annunciation for each alarm e.g. the PCS, Fire and Gas panel, individual graphic tile etc., as many alarms are often repeated and this is a common source of alarm overload. Alarms may also have different priorities at each annunciation location. For example, a fire alarm may be given a high priority on the fire and gas annunciation but is only required to change graphical status colours and be event recorded within the PCS.

The prioritisation and rationalisation process

Before attempting to prioritise an alarm it is necessary to understand the purpose of the alarm and the message that it conveys to the operator, as this determines the required operator action.

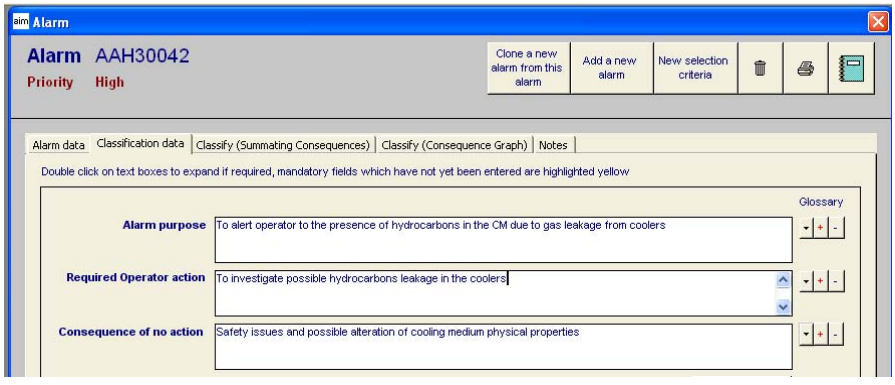

Figure 4: Record of Alarm Purpose and Required Operator Action

The required action must be clearly defined and the consequences of no action determined and described. No action may be as a result of a failure of the alarm itself or from the operator failing to respond but the consequences will be the same. Determining the no action consequences is fundamental to establishing the criticality of the alarm. The alarm purpose, the required

operator action and the consequences of no action are recorded and form part of the audit trail for setting the priority as shown in Figure 4.

The EEMUA publication 191, 'A Guide to Design Management and Procurement' (EEMUA, 1999), Appendix 5 details two methods for prioritising alarms as 'Taking the Maximum Consequence' and 'Summating Consequences'. Both methods are explained below.

1. Taking the Maximum Consequence:

IEC 61508 SIL determination is a risk-based assessment of the consequences that would arise from a SIF failing to operate correctly, where risk is a combination of probability of occurrence and the degree of harm arising from the consequence. The maximum consequence method uses a similar approach for setting alarm priorities but it is focused entirely on consequences, since the regularity of the alarm occurrence has no bearing on its priority i.e. an "Emergency" alarm will indicate an emergency condition whether it occurs once a year or once in ten years. The priority of the alarm should be set purely on the severity of the consequences of an unwanted event occurring. Therefore, if operators are presented with highest priority alarms, they will also be addressing those with the worst potential consequences.

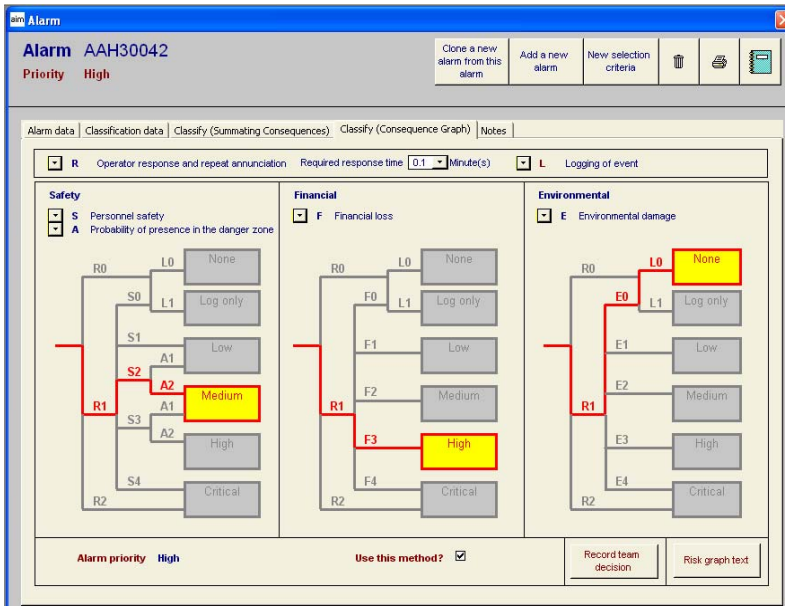

Figure 5: Consequence Graphs

The maximum consequence method has also been enhanced by the author in the SILAlarm™ tool to strengthen its practical application, and provide a consistent and transparent interchange between alarms and automatic SIF protection criticality.

It will highlight those instances where an alarm does not provide sufficient risk reduction, and automatic protection is required.

The prioritisation process assesses the consequences resulting from of an alarm failing, for some technical reason, or is missed by the operator. This assessment considers the following factors:

- the consequences of the resulting event in terms of personnel safety, financial loss and environmental impact;
- in the case of personnel safety, the probability of personnel being present in the danger zone at the time.

As an example, if a plant trip were to occur after an alarm was missed, then this could result in financial consequences from lost production. The priority is then based on the severity of the consequences.

Figure 5 shows an example set of consequence graphs based on the EEMUA Personnel Safety, Financial Loss and Environmental Damage graphs**.** They have been slightly adapted to help with practical application and the parameters for each graph are described below.

Personnel Safety Graph

The selections that are made for operator response (R), safety consequences in terms of the degree of severity (S), journal requirements (J) and presence in the danger zone (T) factors that are selected to establish the safety priority are described in the table of Figure 6. Examples of the safety consequence factors S0 through to S4 are shown in the table of Figure 6 and these represent the scale of injury or fatality that could result if an alarm fails or is missed by an operator.

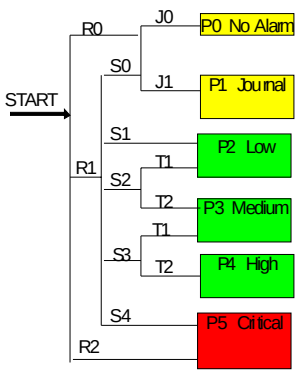

Factor	Qualitative definition
R0	The operator is unable to respond in a timely and effective manner
R1	The operator is able to respond in a timely and effective manner
R2	The operator is unable to respond but consider automatic protection
J0	No requirement to journal, and no activated graphic status
J1	Needs to journal, and/or activates grahpic status change
S0	No possibility of injury
S1	Negligible risk that failure to respond to alarm will result in injury
S2	Remote possibility of injury
S3	Potentially dangerous situation, some possibility of risk of injury
S4	Dangerous situation with a real potential for injury/death
T1	Rare to more often exposure in the hazardous zone
T2	Frequent to permanent exposure in the hazardous zone

Figure 6: Personnel Safety Consequence Graph

An assessment of R2 or R1, S4 resulting in a priority P5 indicates that the alarm is equivalent to or greater than a Safety Integrity Level 1 and serious consideration should be given to provide automated protection.

Financial Loss Graph

The Financial Loss Graph is shown in Figure 7 and operator response capability parameters R0, R1 and R2 and journal requirements J0/J1 are exactly the same as in the Personnel Safety graph. Examples of the financial loss consequence factors F0 through to F4 are shown in the table of Figure 7and these represent the total financial loss that could result if an alarm fails or is missed by an operator.

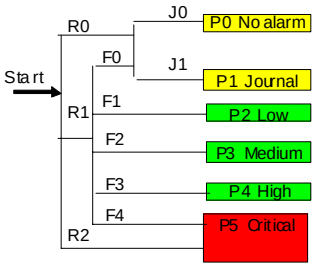

Factor	Financial Qualitative definition
F0	No plant damage or loss of production
F1	No immediate likelihood of plant damage but the possibility of this has increased. Minor loss in productivity or efficiency
F2	Some chance of minor plant damage. Significant reduction in plant output, for example, 10% reduction for 1 hour.
F3	High chance of minor plant damage, low chance of serious plant damage. Significant loss of production, loss of an hour of total plant
F4	High chance of serious plant damage. Serious and prolonged output loss, for example, loss of one day of complete plant output

Figure 7: Financial Consequences Graph

Financial losses are the combination of consequences on production loss, equipment

damage and repairs that could result if an alarm fails or is missed by an operator. They can be described in qualitative degrees of severity or calibrated in quantitative cash terms.

Environmental Graph

The Environmental Graph is shown in Figure 8 and operator response capability parameters R0, R1 and R2 and journal requirements J0/J1are exactly the same as in the Personnel Safety graph.

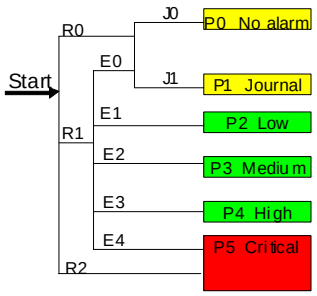

Factor	Environmental Qualitative Definition
E0	No environmental impact
E1	Negligible risk that failure to respond to alarm results in any breach of environmental limits
E2	Remote possibility of breach of environmental limits
E3	Situation with some possibility of breach of environmental limits
E4	Situation with a real potential for serious breach of environmental limits.

Figure 8: Environmental Consequences Graph

Environmental consequences parameters E0 through to E4 are the degree of damage that could be caused to the environment if an alarm fails or is missed by an operator. They can also be described in qualitative degrees of severity or calibrated in limits set by local environmental regulations. Examples of environmental consequence factors for making the environmental impact priority assessment are shown in the table of Figure 8. It is also worth remembering that damaging the environment can have a considerable impact upon company reputation, as many large organisations have discovered after causing an environmental incident. Public reaction can impact on a loss in product sales for a considerable time after an event and this should also be factored into the financial losses.

The Final Priority

The final alarm priority selected is the highest priority from the Safety, Financial and Environmental:

Where priority $P = \max(P_S, P_F, P_E)$ ……(1)

Decisions need to be made on how the actual alarm annunciation in terms of the audio and visual alert that will be activated to alert the operator. Too many variations of both only lead to confusion and, although it is not the intention of this paper to go into the human factors issues, it is important to have a clear set of requirements for alarm presentation. A useful reference for human factor related issues is – Woodson, W.E. Tillman, B. & Tillman, P. 'Human factors design handbook' McGraw Hill, 1992.

Alarm Priority	**Alarm Annunciation and Display Attributes**
P0	No alarm is required.
P1	No alarm annunciation, but the change of state is recorded in alarm and event journals.
P2	A "Low Priority" alarm with an audible tone is generated and the change of state is recorded in alarm and event journals. The alarm is displayed on the Alarm Summary displays if it is within a PCS.
P3	A "Medium Priority" alarm with an audible tone is generated (a different tone to P2 "Low Priority" alarms) and the change of state is recorded in alarm and event journals. The alarm is displayed on the Alarm Summary displays, if it is within a PCS, and the Alarm Annunciation display if configured.
P4	A "High Priority" alarm with an audible tone is generated (a different tone to P2 and P3 priority alarms) and the change of state is recorded in alarm and event journals. The alarm is displayed on the Alarm Summary displays, if it is within a PCS, and the Alarm Annunciation display if configured.
P5	A "Critical" alarm with an audible tone is generated (a different tone to that used for P2, P3 and P4 alarms) and the change of state is recorded in alarm and event journals. **Alarms of P5 priority should not be configured within a PCS environment as the equivalent SIL 1 or greater**. The alarm facility used should have the appropriate PFD. Full consideration should be made to provide automatic protection.

Table 1 - Alarm priority and attributes

The alarm priority types and their respective configuration attributes are described in Table 1. They indicate a range of alarm requirements from no alarm required 'P0', through to the most stringent priority 'P5'.

The alarm priority annunciation and display attributes that are provided are offered only as guidance for implementation, and individual interpretation is possible providing the design ensures that there is appropriate segregation and presentation at the human machine interface.

2. Summating Consequences:

This EEMUA method is a more complex approach that requires far more detailed consideration. The consequences of an alarm failure, with respect to safety (CS), environmental (CE) and financial (CF) consequences, all have to be converted into common units. In order to achieve this common unit of conversion the safety and environmental consequences have to be expressed mathematically in terms of risk (e.g. Safety: injury consequence/year and environmental: environmental consequence/year). An example of the conversion factors could be as follows:

$CS = 10^6$ x (safety consequences in terms of risk of injury)……(2)

$CE = 10^6$ x (environmental consequence in terms of risk of injury)……(3)

CF = 1 x (financial consequences in pounds)……..(4)

These are then summed:

C1 = CS+CE+CF (this is a numerical value and has no units)……..(5)

An assessment is then made as to whether the alarm is 'time critical' based on whether there the required operator response time and this is used to increase the weighting on time critical alarms e.g. if the required operator response was within 3 minutes:

IF (alarm is time critical) THEN

C2 := 3 x C1 (weighting on response time applied)……(6)

ELSE

C2 := C1……(7)

The final priority distribution determines the priority and examples of the weighted total consequence are shown in Table 2. Please note that the numerical weighted values do not have any units:

Weighted total consequence, C2	**Priority**
C2 <900	Low
900 < C2 < 6,000	Medium
6,000 < C2 < 150,000	High
C2 > 150,000	Critical

Table 2 - Weighted total consequence and priority

The 'Summating the Consequences' method is far more time consuming than the consequence graph method and this could preclude it from very large alarm reviews. However, the methodology can be implemented as a software application with facilities for user specific calibration and this significantly reduces the time involved. Figure 9 shows a data input sheet for populating the conversion algorithms which are then implemented in software to simplify the prioritisation process.

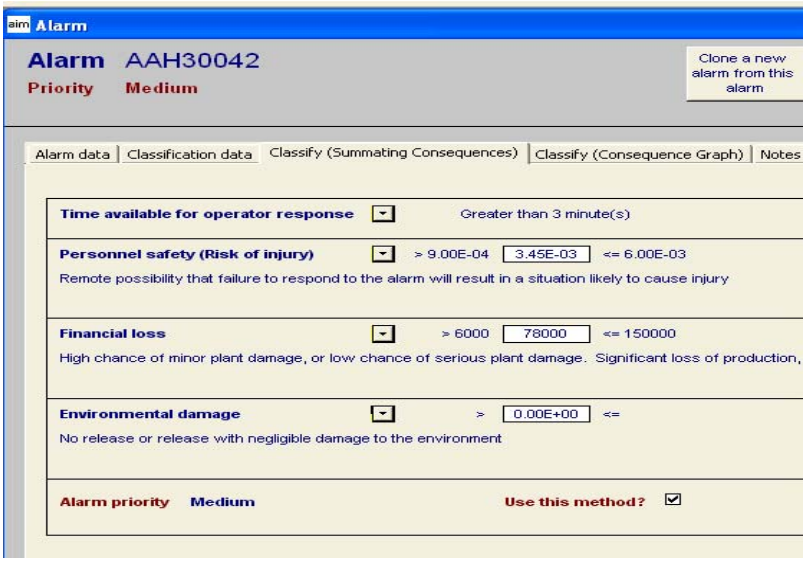

Figure 9: Summating the Consequences Data Sheet

Most process control system (PCS) vendor packages have no tools or structured approach to setting alarm priorities, and they tend to be far more focussed on dealing with suppression techniques and standing alarms than getting down to the root causes. Discussion on these points follows.

Dealing with standing alarms

From the P&IDs, identify items of equipment that are started and/or stopped automatically. Review whether the tags that perform these automatic functions will generate nuisance alarms. If this is the case, identify a way of avoiding this so that normal operation does not generate any alarms, but deviation from expected operation is detected and alarmed.

Review each alarm to establish whether it can be reduced to "Journal" priority. If this is not possible, then set the priority in accordance with the consequence of failure method. If any alarms remain with "Emergency" or "High" priority, review whether it is acceptable for these alarms either to be left as standing alarms or inhibited by the operator under procedural control. If neither is acceptable, then identify another PCS tag or combination of tags that indicates that the piece of equipment is out of commission and use this as a criterion for masking the alarm. Confirm that this will not mask the alarm under any other undesirable circumstances. Document the proposed changes and any associated logic.

Assessing the results and implementing the findings

It is important to review the results and finding of an alarm review to check that the right balance has been made between alarm numbers and the respective distribution of priorities. Sort the alarms into their priority groups and compare the results with the distribution in existence prior to the review. Figure 10 shows a distribution of alarm priorities generated from the results of a typical alarm review using a database application software tool.

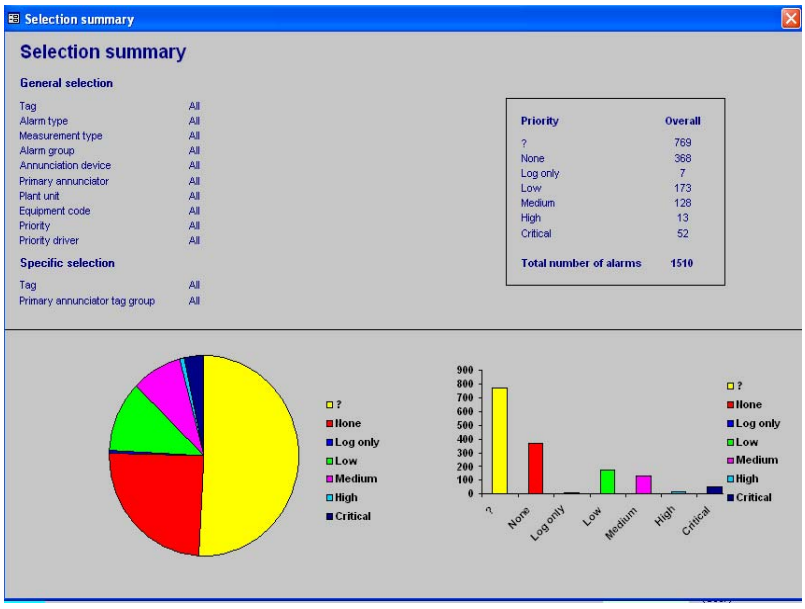

Figure 10: Alarm Priority Distribution

If the alarm priority distribution indicates an imbalance of higher priority alarms this can be normalised by selecting groups of alarms and applying a revised priority assessment to a whole group.

The EEMUA guidance for new alarm system design advises a target distribution as shown in Table 3.
All alarms with a P0 priority can be removed from the alarm configuration. All 'P1' alarms can have their operator annunciation capability removed but they retain their status change as a record in the alarm journal.

Experience shows that the majority of alarms configured into a typical PCS will fall into the 'P0' (no alarm required) and 'P1' (journal only) following a prioritisation review. Thus the major part of the simplification and rationalisation of the alarm configuration is going to be achieved by attending to these two alarm types.

Priority band	Alarms configured during design
Critical	About 20 altogether
High	5% of total
Medium	15% of total
Low	80% of total

Table 3 - Priority distribution for system configuration - EMMUA

Alarms with priorities 'P2–P4' should be implemented with the configuration attributes as described in Table 3. Alarms with the 'Critical' priority 'P5' require careful treatment since they indicate a SIL requirement of SIL1 or greater. The provision of automatic protection should be considered to replace the dependence on the alarm and successful operator action for each P5 alarm. If automatic protection is not possible, then the alarm must be engineered to achieve the appropriate probability of failure on demand (PFD) in line with IEC 61508 requirements. The EEMUA accepted PFD_{avg} for an alarm with operator intervention is < 0.01 which is only equivalent to SIL 1.

Other Considerations

Suppression techniques tend to be the main focus of PCS vendors since they invariably have all kinds of smart software techniques to implement suppression, and they are keen to sell this capability to their clients. This paper does not intend to cover the options of static and dynamic alarm suppression in detail because it is a complex subject area that requires careful implementation and, in the author's opinion, should only be considered as a last resort. It is anticipated that the requirement to consider these options will be minimised, if not eliminated, by a thorough alarm review.

Static Alarm suppression

If operators still find difficulty with the qua-
ntity of standing alarms present, following the alarm rationalisation, then consideration can be given to the implementation of static alarm suppression. This will minimise the number of standing alarms that are generated when a process unit or large piece of equipment is shut down. In order to suppress the selected alarms, a defined set of process permissives have to be satisfied in conjunction with a with an 'enable' static suppression status condition.

Dynamic Alarm Suppression

If the number of alarms generated following a trip is still unacceptable, following the alarm rationalisation, then dynamic suppression can be considered so that the first up alarm in a pre defined group audibly alerts the operator, registers in the alarm list and is printed. All other subsequent alarms in the group do not activate any audible alert; they do not register in the alarm list and are not printed. Dynamic suppression on an alarm group should automatically de-activate after a defined period of time following the first up alarm, so that any new alarm that then follows alerts the operator and re-starts the suppression period.

Conclusions

There is no doubt that alarm management requires considerable effort, commitment and tenacity. However, by focussing effort on an alarm review and rationalisation process, then experience shows that this will return a significant, and perhaps even a dramatic, improvement in alarm performance to achieve better than 90% of the available benefits. There will possibly be some alarms that need further analysis but having filtered them out by the review process they should only represent the remaining 10% of effort.

Avoid using suppression techniques for unwanted alarms until a full alarm review and rationalisation study has been completed, since it may not be necessary once the system configuration has been improved.

The review and rationalisation methodology has been applied on facilities with more than 15,000 configured alarms where individual operators were required to accept alarms every 15 seconds under steady production conditions. Typically, over 50% of the configured alarms have been removed, and by also attending to the nuisance alarms the alarm rates have been reduced to 15-20 minute intervals. Perhaps most importantly the spurious shutdowns resulting from missed alarms have seen improvements from weekly plant trips to no trips for over six months. This results in a very, very significant payback.

Making full use of software tools will substantially reduce the effort involved in an alarm review by typically >50%, whilst the electronic version of an alarm database will provide data integrity along with ease of maintenance and change control over the full life cycle.

References

BS IEC 61508, 1998-200, Functional Safety of Electrical / Electronic / Programmable Electronic Safety-Related Systems.

BS IEC 61511, 2003, Functional Safety: Safety Instrumented Systems for the Process Industry Sector.

Timms C.R., 1999, A Methodology for Alarm Classification and Prioritisation – IEE People in Control Conference 1999.

Health & Safety Executive 'The explosion and fires at Texaco Refinery, Milford haven, 24 July 1994, HSE, 1997'.

The Engineering Equipment and Materials Users Association (EEMUA), 1999, Alarm Systems , a Guide to Design Management and Procurement – EEMUA publication No. 191, ISBN 0 85931 076 0

U.K. Health & Safety Executive, 2000, Better Alarm Handling - UK HSE information sheet, Chemicals Sheet No. 6.

Wilkinson J., and Lucas D., 2002, Better Alarm Handling – A Practical Application of Human Factors, Measurement and Control Volume 35.

Woodson, W.E. Tillman, B. & Tillman, P. 'Human factors design handbook' McGraw Hill, 1992.