Inclusive and integrated risk assessment, risk management and SIF definition under the IEC61508 group of standards.

Clive de Salis, Milford Haven Refinery, P.O.Box 10, Milford Haven, SA73 3JD

Keywords: LOPA, SIL, Functional Safety Management, SIF, HAZOP, IEC61508, IEC61511

The lifecycle requirements under the IEC61508 group of standards includes the stages of identifying the risks to be covered and implementing a Safety Instrumented Function (SIF). These stages include links to risk assessment techniques such as HAZOP study, LOPA analysis, and others. However the IEC61508 group of standards do not clearly show how all the stages of risk assessment and the SIF definition become integrated into a safety management system. Nor do the standards show how to provide tools and resources that are inclusive towards all of those who need to participate at each step of the management process. This paper looks at the experience of developing an integrated and inclusive safety management system in a UK refinery using IEC61508 and IEC61511 standards.

The IEC61508 group of standards cover safety instrumented systems with high integrity. The high integrity requirement is classified by a safety integrity level which is identified as a SIL rating. The application of IEC61508, the master standard, in the process industries is covered in the sector standard IEC61511. All stages of the safety instrumented system lifecycle involve team work. The lifecycle does not rely upon individuals, that is not how the standards are written and it is not what is intended. The hazard and risk assessment stage is undertaken by a team. The allocation of protective layers is agreed by a team. The SIL assessment is undertaken by a team. Verification and validation activities are both team work lifecycle stages. Only really the design of the safety instrument loop might, in the extreme, be designed by an individual but, as soon as the checking and assessment stages are included in that scope then, even that is arguably a team.

When HAZOP study started to be used widely in industry it became apparent that the requirement for team study meant that those with expertise should develop approaches that were inclusive to others on the team. Chemical engineers have long known that the chair in a HAZOP study needs to pay attention to, and work on, the inclusion of the operator's experience. It appears that a number of my fellow experts in the IEC61508 and IEC61511 arena may be needing to make the same discovery. Expertise and knowledge are not the same and some of the quieter team members are often holding valuable knowledge even though they would not profess to be experts in any way. It is not enough to simply identify the roles necessary for a team. It is just as important that we support the team with tools that empower every member of the team. The tools used at each stage of the lifecycle need to be inclusive. The identical issues arise in the safety management system under the IEC61508 group of standards at every stage of the lifecycle. The tools adopted need to be inclusive of all of those who need to be consulted and participate.

The need for effective team working isn't just about getting good results, it's also about minimising time wasted in team meetings. Add up all the manhours involved in team meetings and you soon realise that there is a cost to such working practices even when those working practices are to be fully endorsed and encouraged. Therefore there is an additional benefit to an effective safety management system that works well with teams.

For any safety instrumented system under the IEC61508 group of standards we are going to need to know the corporate tolerability criteria that is going to apply. As we develop tools for safety instrumented systems the application of corporate tolerability criteria must be consistent throughout all tools. The same tolerability criteria must be used when screening for those hazards that need further examination as are used later for the SIL assessment calculation. It makes no sense whatsoever to have screening tools that do not align with the same tolerability criterion as will be used later for the SIL assessment.

Risk graphs are a very effective tool for team work. They are noticeably more friendly than calculation techniques but have limitations of accuracy. IEC61511 Part 3 does include some risk graphs which, as it says in the introduction to Part 3, are offered as illustrations of techniques and are not suitable for use as shown.[1] Risk graphs are supposed to be calibrated mathematically and not just calibrated by altering the words of the definitions.

Using a risk graph as an effective screening tool with a team is, in the author's opinion, a good practice, but BEWARE: The risk graph in both IEC61508 2nd edition Part 5 Annex E and in IEC61511 Part 3 annex E (see figure 1 below) is popular but is actually mathematically wrong. The parameter F1 is defined as a less than 10% chance of someone being in the danger area, therefore F1=0.1 and it then follows that F2 is a good chance that someone is in the danger area so F2=1. C2 is defined as death to one person, so C2=1. The graph can be initially analyzed using the conventional approach that, if C2=1 then it follows that, C1=0.1, C3=10 and C4=100. P1 is the probability that the operator can avoid unwanted event and conversely P2 is for when it is difficult to avoid. Taking P1 conventionally as 0.1 and P2 as 1 we now find the following:

a = 0.1 x W3 = non-SIL rated

1 See IEC61511 Part 3 Page 9 clause 1.2

See also IEC61511 Part 3 Page 7 which states:

This standard deals with guidance in the area of determining the required SIL in hazards and risk analysis (H & RA). The information herein is intended to provide a broad overview of the wide range of global methods used to implement H & RA. *The information provided is not of sufficient detail to implement any of these approaches*. Author's emphasis added

b = (C2=1 x F1=0.1 x P1=0.1) = 0.01 x W3 = SIL 1

c = 0.1 x W3 = SIL 1

d = 0.1 x W3 = SIL 2

… and so on.

So whatever value one assigns to W3, mathematically this risk graph says that non-SIL, SIL 1 and SIL 2 are all the same (i.e. they are all equal to 0.1 x W3). This is clearly ridiculous. Indeed if you now alter numeric values to suit other conventional values you will still find that it is mathematically wrong. Amongst members of the committee that writes the standard that particular risk graph has proved to be a mathematical conundrum. The problem lies in the structure of this illustration shown in the standard and so it is worth remembering that the standard itself, in the opening to part 3, says that it is an illustration only, not suitable for use as shown.

The author is not in any way suggesting that a risk graph should not be used as a screening tool but is making clear that it needs to be designed and calibrated to make sense. It needs to be properly aligned with the relevant corporate tolerability criterion. When we develop a risk graph as a team friendly screening tool we need to remind ourselves of the purpose. If you lose sight of the purpose of this stage in the lifecycle you can waste a lot of time and have a flawed system.

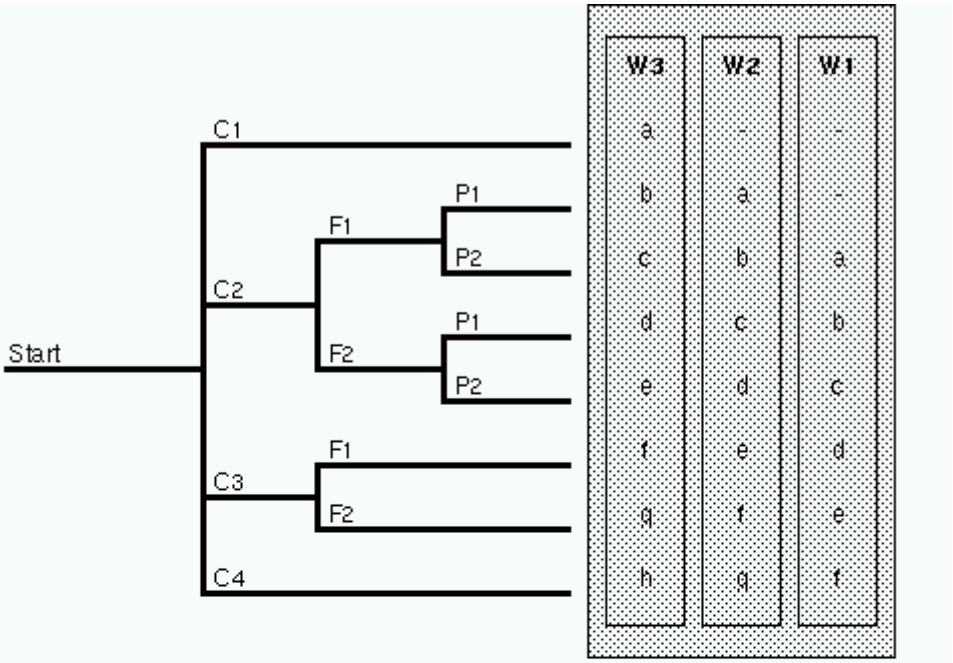

Necessary minimum risk reduction	Safety Integrity Level **(SIL)**
-	No safety requirements
a	No special safety requirements
b , c	**1**

Figure 1: Risk graph from IEC61508 part 5 figure E.2 (the same graph appears in IEC61511 Part 3 figure E1 with AK numbers converted to SIL numbers in figure E.2).

For an inclusive approach we must remind ourselves that at the screening stage we are NOT trying to decide the SIL rating, we are trying to collect the team understanding of the risk as well as deciding if the hazard needs further examination.

A risk graph for a screening tool should allow all members to provide the team view of:

a) The available information on the initiating event frequency (NOT the demand rate)

b) The number of layers of protection and what they are.

c) The available alarms and warnings to assist the operator to avoid the event

d) The ultimate unwanted consequence

e) The probability of people being in the danger area.

The risk graph is the data gathering tool for the above. The information listed above is the information needed from the team for later examination in more detail by LOPA or other calculation techniques. The risk graph, in this case, doesn't decide the SIL rating but it is used to check and to help to finalise the safety requirements allocation: That is the purpose of the risk graph. Figure 2 shows a segment of the type of risk graph used as a screening tool at the Milford Haven Refinery. The risk graph asks for the team to record their knowledge of the initiating event frequency. The time intervals selected of 1 (i.e. every year), 3, 10, 30 and U (i.e. unknown or never experienced) also match available data sources on site so that, after the team meeting, data can be checked against such resources as 3 year trends or inspection records. This approach helps gather evidence to support the data. The team is also asked to record not just the number of layers of protection but also their details. The number of layers of protection is then applied to the alpha character, N (none), O (one) and so on.[2]

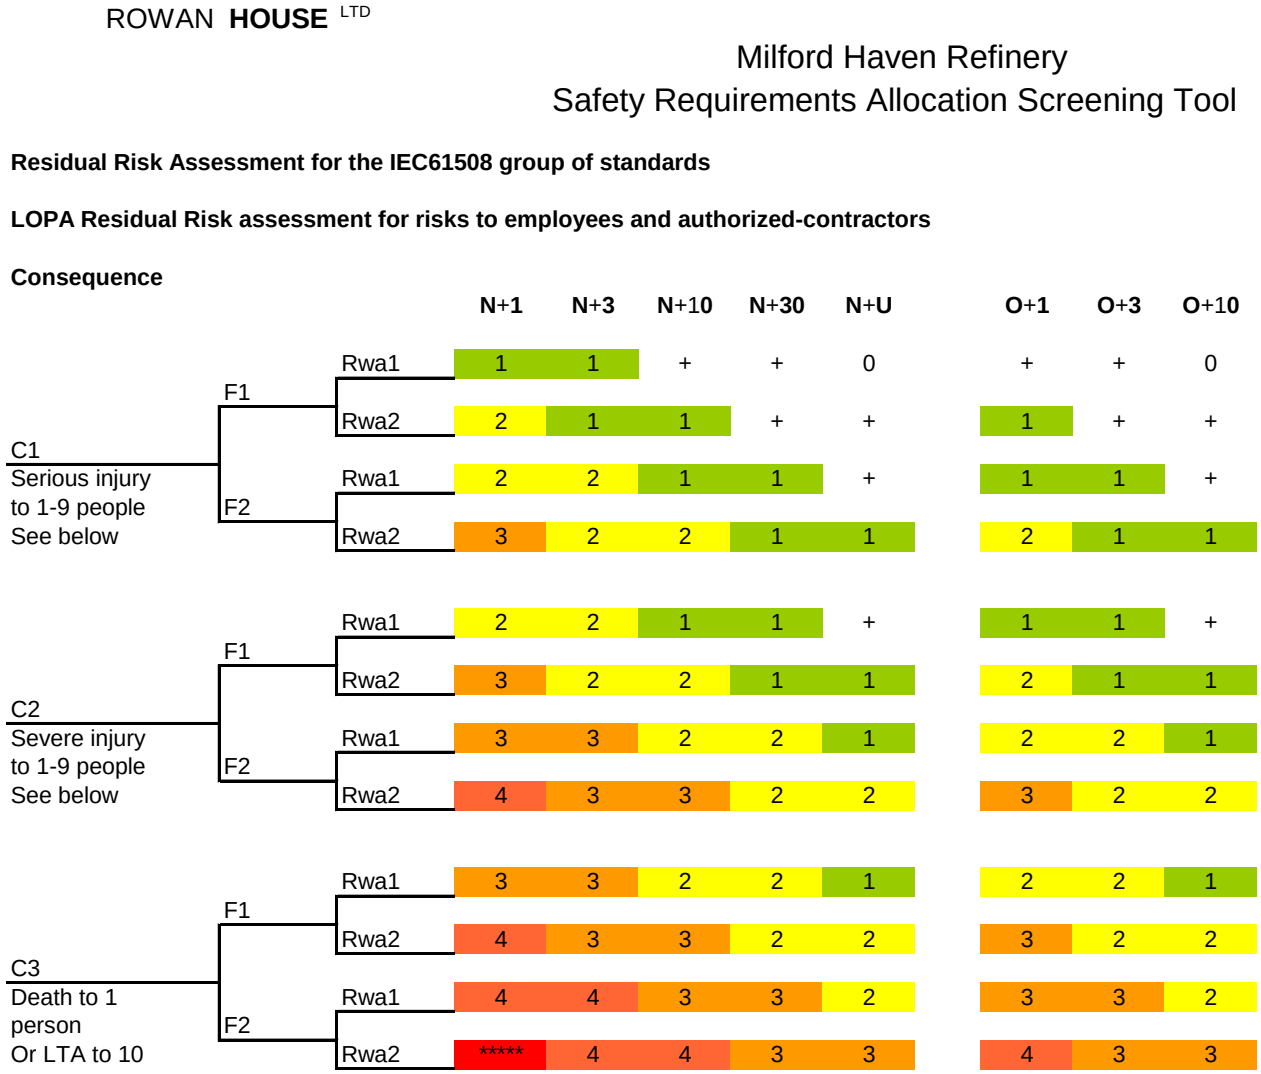

Figure 2: Extract from the risk graph screening tool used at the Milford Haven Refinery

[2] For details of how to design and calibrate such risk graphs see "Using risk graphs for Safety Integrity Level (SIL) assessment - a user-guide for chemical engineers", Clive de Salis, Publisher: IChemE, ISBN: 9780852955543

The output of the risk graph should indicate to the team that the proposed design has the *potential* to give rise to the need for a SIL rated safety instrumented system if left unaltered. It is important to note that the risk graph approach now empowers and enables the team to discuss what else can be done to reduce the risk and see the potential effect. The risk graph is then an inclusive tool for the team to examine options for risk reduction and help seek out ways towards deciding the optimum allocation of protective layers.

Training courses on SIL assessment have propagated a false dichotomy for SIL assessment of either use a risk graph or use a calculation technique. Here the risk graph is used to assist the calculation. These are tools working together as an integrated approach. Annex B of IEC61511 Part 3 describes a process in which the tools work together. It has always been in the standard and so the spreading of the idea that SIL assessment involves "either" a risk graph "or" a calculation is not supported by the standard, and never has been.

As the reader can see above, a risk graph is not about deciding a SIL rating, it is about allocation of protective layers. In IEC61508 this stage is called "Safety Requirements Allocation". It is the stage in which the available layers of protection are optimised, decided and agreed by the team. Thus the risk graph is a team tool for safety requirements allocation (in IEC61511 the matching step is called "Allocation of protective layers").

The allocation of protective layers is an important stage in the lifecycle. After the SIL rating has been calculated and determined, the SIL rating is only true if the layers of protection are also managed and maintained. So this stage in the lifecycle for safety instrumented systems is just as important as the calculation of the required SIL rating that comes later on. The agreed layers of protection are not something nominal that can be forgotten later.

When developing a holistic approach to safety instrumented systems it is better to start by putting down the standard altogether and begin by getting the basic layers of protection right as the first step. We develop the basis of safety. Once the layers of protection are agreed and included within the design we then evaluate them all and see if the probability of an unwanted is greater than, or less than, the corporate tolerability criteria allows. If there is a gap, then the safety instrumented system is going to fill it.

In the UK the Dangerous Substances and Explosive Atmospheres Regulations, DSEAR, applies to all sites everywhere, irrespective of whether a site is a COMAH site or not. We can develop a basis of safety approach that simultaneously answers the requirements of DSEAR and, funnily enough, Regulation 6 of DSEAR gives the same sequence of considerations for safety that chemical engineers have always followed.

Regulation 6 of DSEAR asks the following questions (slightly paraphrased):

6(1) Can we eliminate the risk?

6(2) Can we substitute one dangerous substance for a safer one?

6(4a) Can we reduce the quantity of dangerous substance?

6(4b) Can we minimise (or avoid) the risk of a release?

6(4c) Can we control any release at source?

6(4d) Can we prevent the formation of a hazardous atmosphere or hazardous conditions?

6(4e) Can we ensure any release is safely contained or otherwise made safe?

6(4f(i)) Can we avoid ignition sources?

6(4f(ii)) Can we avoid the adverse conditions which could give rise to harmful effects?

6(4g) Can we segregate incompatible dangerous substances?

We can unpack some of these regulation requirements to make them easier to answer but so long as we document our decisions in the same order then our documentation trail will also demonstrate that our design is in compliance with this part of DSEAR. It is helpful to call this "The Basis of Safety".

I would suggest the following sequence of questions:

6(1) Can we eliminate the risk?

6(2) Can we substitute one dangerous substance for a safer one?

Can we obtain the substance in a safer form?

Can we use an alternative, safer, process design?

6(4a) Can we reduce the quantity of dangerous substance?

6(4b) Can we minimise (or avoid) the risk of a release?

6(4c) Can we control any release at source?

6(4d) Can we prevent the formation of a hazardous atmosphere or hazardous conditions?

6(4e) Can we ensure any release is safely contained or otherwise made safe?

6(4f(i)) Can we avoid ignition sources?

6(4f(ii)) Can we avoid the adverse conditions?

6(4g) Can we segregate incompatible dangerous substances?

The questions we have added to 6(2) above is part of the chemical engineer's working practice of looking for inherent safety. Substitution isn't just about swapping one chemical for another: At the Milford Haven refinery one of the process units uses a dangerous chemical which spontaneously combusts in the presence of air and equally generates heat in the presence of water. It is nasty stuff, but instead of swapping it for another chemical we found that we could have it delivered pre-mixed with an alkane that made it safe to handle. The alkane takes no part in the reaction, it is simply a carrier, but it stops any risk of either spontaneous combustion or of significant heat release in the presence of water.

In the water industry some digesters can be fed with methanol to improve the digestion and biogas production. However methanol is highly flammable with a low flash point and an invisible flame. Since the digesters are fed with a methanol / water solution containing only 5% methanol it would be inherently safer to receive pre-diluted methanol that was neither flammable nor capable of flashing at ambient temperatures. This does not mean that a vast number of road tankers are required because dilution to levels around 30% to 40% methanol minimise the number of road tankers whilst optimising inherent safety. In this example "substitution" involves considering dilution.

The author was involved in the investigation of the first major sewage sludge dryer explosion in the UK. The laboratory work showed that sewage sludge has a strong exotherm that begins at around 115 degC. In order to dry the sludge it is necessary to heat to 100 degC. The drying process selected used hot oil for heat transfer at over 250 degC. Inherent safety would clearly have been improved by using saturated steam at 105 degC instead of hot oil. This is an example of how the choice of process itself includes safety considerations. In this example "substitution" involves substituting alternative processes.

Answering the questions for Regulation 6 of DSEAR doesn't just help our compliance with regulation, it also organises our thinking for the available layers of protection.

At the Milford Haven Refinery we now have an SRA form (named after the "Safety Requirements Allocation" lifecycle stage in IEC61508 only because too many oil industry acronyms begin with AP which would be the acronym for the same stage in IEC61511). The SRA form documents the final decision for the Allocation of Protective Layers (IEC61511 terminology) and records the source hazard and risk assessment. The latter record of the source risk assessment provides the traceability of documentation required under IEC61511 Part 1 Clause 8 paragraph 8.2.3, whilst simultaneously providing the source data for the SIL assessment calculation by LOPA or other appropriate technique.

If we get the layers of protection right then we avoid unnecessary reliance on high integrity complex safety instrumented systems and we avoid SIL ratings becoming a sticking plaster over poor process design.

Now the team can undertake the hazard and risk assessment in whatever form is appropriate for the risk being considered. In IEC61511 Part 3 Annex F Layer of Protection Analysis ("LOPA") is shown to be based upon HAZOP study. In reality LOPA doesn't always need to be based upon HAZOP, it can be based on other forms of hazard and risk assessment. Nonetheless, the annex F in the standard for LOPA shows the team required for the assessment in clause F.2. LOPA is a team assessment, not a calculation by a clever individual working on their own. Herein lies a problem to which we alluded earlier. We want the whole team to offer their knowledge of the hazard and risk but for some in the team a calculation is often not the right vehicle with which to obtain that information. When HAZOP study started to be used widely in industry it became apparent that the requirement for team study meant that those with expertise should develop approaches that were inclusive to others on the team. Chemical engineers have long known that the chair in a HAZOP study needs to pay attention to, and work on, the inclusion of the operator's experience. The same problem occurs when trying to do LOPA as a calculation with a team. Not everyone fully appreciates the numbers or even gets to grips with their meaning. This is where the risk graph approach as a route to gathering information and screening is successful. The decisions made by the team at the Milford Haven Refinery are written into the SRA form and the chair of the team guides them through potential SIL rating in order to identify if anything else can be done that is practical and effective at reducing the risk. The LOPA calculation is then completed from the information on the SRA form.

In IEC61508 Part 1 the lifecycle shows hazard and risk assessment at step 3 but the LOPA is not until step 5. The stage in between is the safety requirements allocation. Thus the new system at the Refinery documents the decisions made with team friendly tools guiding everyone through the process. Layers of protection are checked against DSEAR and simultaneously hazard and risk assessed. The output following resolution of any recommended actions is documented in the SRA form and finalised. The LOPA, or other appropriate calculation, is then based upon the decisions documented in the SRA form. The whole process has a document trail as required by Part 1 of IEC61511.

In the LOPA calculation, the SRA form for each hazard and risk that will be covered by a common SIL rated function, are brought together to calculate the SIL and PFDavg3 required of the safety instrumented function to cover all of the risks.

The standards, IEC61508 and IEC61511, talk about a safety requirement specification. In order to facilitate team work we ask the process engineers to write a process safety requirement specification which is an input to the final safety requirement specification written by the instrument engineers. The process engineers are not just asked to describe and identify the process risks to be covered by the safety instrumented system and the relevant relevant layers of protection, they are also asked to consider and identify ways in which the safety loop can be fully tested end-to-end safely. There are often opportunities during a start-up or shutdown of a process unit in which a safety loop can be tested end-to-end without upsetting the process plant safety. It is usually the chemical engineer who is best placed to answer such questions, not the instrument engineer. The process engineer can help the instrument engineer find effective ways to fully test the safety instrumented system. It is another example of team work in the design. It is usually also the process engineer who can identify what is required for safety: Is a valve required to close to be safe or is leak-tight shut-off required? These two are not the same. For an application such as Fuel Storage Overfill which occurred at Buncefield safety can be achieved by

3 PFDavg = Probability of Failure on Demand average.

closing the valve. Leak-tight shut off is not actually necessary to make the overfill safe. To an instrument engineer designing the safety loop the difference is significant.

Asking the process engineer to write a short process safety requirement specification is another tool that helps optimise the safety instrumented system design.

The designer of the safety loop must decide the proof testing requirements. However it is not the designer who will do the proof testing. Therefore at Milford Haven Refinery we ask the designer to draft the proof testing requirement and to include in the document the proposed Mean Time to Repair (or Replace). This allows the maintenance department to give their feedback on the proof testing requirement and organise the document into a properly understood format. The maintenance team can also give feedback to the designer on whether or not the MTTR is even achievable or realistic. This is done before training is given on the safety instrumented function.

The use of outside contractors adds a degree of complexity to Functional Safety Management. The HSE guidance for Competency Management for Safety-related Systems (Published July 2007) shows in Part 2 that we are expected to not only assess our contractors but also their sub-contractors. CASS stands for Conformity Assessment of Safety Systems. It is not a certificate but it is a methodology for demonstrating compliance. It is a free system in the UK, set up originally with backing from the Department of Trade & Industry. The author is the new chair of CASS. There is a free downloadable document available from the website www.61508.org that allows the user to write down and document their Functional Safety Management. There are help pages provided to guide the user and the system is free and is not tied to any certification body. Indeed certification is not required under either IEC61508 or IEC61511. If users need further assistance the website www.61508.org also lists consultants and firms that can provide support and advice.

There are three parts to the CASS32 FSM declaration: Part 1 is the business address, the nominated personnel responsible for safety management and the location of the work etc. Part 2 is the scope of work undertaken by the business making the declaration and Part 3 is the details of the FSM (Functional Safety Management) that is used to manage the scope of work that is shown in Part 2.

At Milford Haven Refinery we send out the CASS32 form to all contractors. Any contractor, or their subcontractor, can offer certification by a recognised authority. If proof of Functional Safety Management is available then we are happy to accept it. Any Contractor, or their sub-contractor, does not need to offer certification (as certification is not required under the standard), they can, instead, complete the CASS32 FSM declaration.

The Milford Haven Refinery offers contractors a choice: We make a copy of the CASS32 FSM declaration in which Part 2 has been replaced by the scope of work for the job. A contractor can either complete the declaration for just the one job or they can complete the full, generic CASS32 FSM form. If they complete the form that is specific to the job then they will need to complete a fresh FSM declaration for each project undertaken at the refinery. If they offer the complete standard CASS32 FSM declaration then it can be re-used and applied to other projects.

The CASS32 FSM form is actually better than most currently available certification because the methodology and completed form shows compliance with:

- IEC 61508 2nd edition Part 1 Clause 6
- IEC 61511 1st edition Part 1 Clause 5
- HSE Competency Management for Safety related Systems

… all in a single document. Indeed the contents will soon include IEC 62061 Part 1 FSM as well. So a single document shows compliance with all necessary standards for the refinery. Incidentally, any user that has completed the form can obtain certification that is UKAS accredited if that is what they want to do for other customers. The work undertaken to complete the form is never wasted.

If we have a contractor who can't fill in the FSM declaration, and they still can't fill it in by following the help pages, and they still can't fill it in following advice from a third party … then there really isn't a better clue that the contractor hasn't got Functional Safety Management and doesn't actually comply with standard.

Demonstration of FSM is mandatory under IEC61508 2nd edition Part 1 Clause 6. It's not optional. Anyone can freely fill in the form and the website lists a wide range of notified bodies with whom the declaration can be lodged under the same rules as ATEX declaration. As a refinery that is our safeguard. We suggest you consider this free approach to FSM too. It is an effective way of using good suppliers that are known to you without putting up unnecessary barriers such as certification when no such thing is required under the standard.