Where are your SIL assessments now?

Jo Fearnley, Jacobs Consultancy, Phoenix House, 3 Surtees Way, Stockton-on-Tees TS18 3HR

Since the introduction of the concept of safety integrity level (SIL) assessments companies have gradually aligned themselves with the process and have completed their assessments such that they consider themselves compliant. However the question people should be asking is are those risk assessments still valid, and how well would the company score if an audit were to be completed today on both the systems supporting the SIL assessments, and the assessments themselves?

Many companies who decided to complete SIL assessments early in the lifecycle of the guidance used the risk matrix approach, which is now not considered best practice. Others who used the layer of protection analysis (LOPA) concept did so with the best intentions, but in many cases the layers considered, and the credit assumed would not now be considered in line with current expectations.

However, more fundamentally, the question is, of the thousands of assessments which have been completed over the last 10 years how many actually reflect the reality of how the plant is installed, operated and maintained?

How many assessment actions have been closed out completely, and the assessment updated to include the action resolution? How many assessment actions have been completed with the SIL rated safety instrumented functions actually installed, validated and maintained to the stated requirements?

How many assessments have been reviewed since originally completed to ensure they are still valid? How many assessments have been checked when considering a plant modification to ensure the risk assessment is not compromised or changed by the plans?

There is no deliberate intention by companies to not follow through every aspect of the SIL assessment process; it is simply that time goes by, priorities shift, personnel change, and the file on the shelf remains gathering dust. Actions that initially seemed critical to complete become less urgent the longer they remain outstanding, and may even be closed out on the basis that if it has not been done yet then it can't be important!

So, be honest with yourselves, are all your SIL assessment records in place, reviewed and up to date? Do they accurately reflect the current status of your plant? Are all actions closed out correctly? Are all testing and maintenance regimes robustly in place? Are routine reviews of the SIL assessments within an action management reminder system? Do your process safety performance indicators include appropriate leading and lagging indicators around the on-going validity of the SIL assessments?

The Health and Safety Laboratory (HSL) run a course on LOPA which includes a section on the pitfalls of LOPA assessments, but this paper goes deeper into the routine failings which occur in the management of these risk assessments, and also on how changes in expectations over the years since they were introduced have been instrumental in making some earlier assessments out of date.

Keywords: SIL, LOPA, risk graph, risk assessment, validity, re-assessment

Being a consultant is a strange life, and not one that suits everyone. I am a principal consultant in process safety and environmental risk assessment, and as such I am involved with a large number of companies, across a range of industries and hazardous processes. I have a valued role, and yet that role is transitory. I come, and I go, and all I leave behind is a report with details of the risk assessments completed; a record of the meetings held and the discussions facilitated; a list of recommendations and actions; data needed to make decisions; options to be considered. I have no responsibility or ownership for deciding whether to implement those recommendations, for completing the actions, for compiling the additional data needed, or for converting the options into a way forward. Therefore, I do not always know the final outcome, and I have no 'right' to do so!

After years of working in a production role, where completing those actions and making those decisions was my responsibility, I am well aware that I had too many outstanding actions, too many to-do lists, and was never fully compliant with the ideal world. So there is no criticism intended, now that I am an outsider who provides advice and facilitation and am seldom involved in the completion aspects. It is purely an observation that I do not believe that all those actions or recommendations relating to risk assessments are fully closed out.

This is valid for all types of risk assessment, but at least if the assessment is a project related hazard study it is typical that the next step in the project is to ensure completion of all the actions outstanding from the previous stage, so hazard study 1 and 2 actions would be completed before hazard study 3, and the latter actions are completed before commissioning at the latest. For safety integrity level (SIL) assessments the project management process is likely to be in place for a new build plant, but when completing retrospective assessments, then the action management process may not be as robust.

My expertise is in the determination steps for SIL assessments, from hazard identification through to the associated risk assessment, with allocation of safety function to the protection layers; or clauses 8 and 9 of the BS EN 61511 safety instrumented system (SIS) lifecycle. I am cognisant of the other clauses or stages within the SIS lifecycle, but I do not offer my services to address those other areas; so this further removes me from detailed knowledge of the final installation.

Over the years I have become concerned that the final installation intended to address the hazard originally identified is not always exactly as envisaged, and therefore the risk assessment documented may not be a reflection of reality. There are a variety of reasons why this is the case, and it is never a deliberate intent, but the consequence of myriad evolutions and iterations, which are not fully captured.

That the SIL assessments are not perfect is reflected in the Health and Safety Laboratory (HSL) course on 'Layers of Protection Analysis: Practical Application and Pitfalls', which reinforced the concerns I have about current validity of many SIL assessments.

The aim of this paper is to discuss some of the SIL assessments I have been involved with over the years, whether using risk graphs or layer of protection analysis (LOPA), to indicate where the gaps are which have prevented the 'determinations' completed in good faith by the team involved, from being converted into verified and robust reality. Further it highlights some of the common issues arising during risk assessment. Some examples relate to those assessments I have facilitated, some are 'inherited' assessments from third parties, and all are from a variety of clients – with no names attributed!

Probability of Harm

Early adopters of the SIL philosophy developed a calibrated risk graph for their site, as this was seen as a straightforward means of providing a common basis for risk assessments across a range of facilities on site. Checks for clients have shown that provided the original calibration is consistent with that which would be used within a LOPA methodology the resultant integrity level decision will be essentially the same. However there are aspects of the risk graph methodology that has led to higher resultant risk reduction factors than would arise from a standard LOPA.

The first of these is the limited choice of a 0.1(F1) or 1(F2) probability for the presence of personnel close enough to the hazard such that they are harmed. Whilst this is sensible choice of probability for a majority of risk assessments, there are occasions when a lower value is valid, for example if the risk is from the failure of a long pipeline, which could be at height or in an unoccupied area of the site. In cases like this, the probability that someone is close enough to the failure point (which could be anywhere along the length) may be 0.01, so the use of a risk graph over estimates the risk. Whilst it is possible to just include an adjustment factor in the calculation to arrive at the event frequency (W), this was seldom done in the early days as people tried to follow the stated methodology.

Conversely I have seen companies using the LOPA methodology try and use a probability of harm (vulnerability) factor which is so low that I have had to veto it, as justifying that someone will be in the vicinity for a maximum of two minutes per shift is not viable to me. When justifying the vulnerability factor a common sense approach should be used, which considers what other equipment is close to the location under review, not just the actual system, such that operators, maintenance personnel or others, including passers-by, could be within the harm zone. Most of the time using the range of probabilities 1, 0.5, 0.1 and very occasionally 0.01 is the limit I suggest as justifiable.

When companies state a very low value, such as 0.01, for equipment other than a long pipeline, then it is normally because they are going to restrict access to the area around the equipment in order to achieve this. However, when I have subsequently been on the plant I have seen these controls in disarray, so the company is not delivering the layer of protection. Examples are flare-stacks, where local access was supposed to be restricted, but fences were never erected or gates not controlled. In several cases the 'restriction' was a line drawn on the ground, or maybe a chain on low supports, but with no notices to state that access beyond that point is restricted. Often that 'unused ground' is then used as storage for spare equipment or packaged chemicals, which blatantly shows that personnel are accessing the area, and there is no evidence to show that the access was controlled / authorised.

The logic for the value to use for vulnerability has also changed over the years. In the early days the presence of someone in the vicinity was based on the normal operations. So if the event is related to start-up, and people could be in the vicinity for the start-up operations, then the use of probability of 1 would typically have been used. However if the event related to normal operations then a probability of 0.1 that people would be in the vicinity may have been selected. However, as has been highlighted by the HSE over recent years, including within the HSL LOPA training course, there is a potential fault in this logic. The fault arises if one of the other layers of protection considered is operator response to an alarm. If response to the alarm results in an action taken from the control room only, then the vulnerability selected will not be affected. However; in many cases the actual response, especially before taking action to shut down a unit, will be to ask an operator to go and check that the alarm is correct; or maybe the only action which can be taken needs to be done locally, for example to change over a pump. In this case the vulnerability increases to 1, so taking two independent layers of protection for response to an alarm and vulnerability is incorrect. For scenarios not reviewed since this logic became widely accepted there is therefore liable to be some which need revising. Whilst this issue was highlighted within BS EN 61511-2:2004 guidance, section 8.2.1, it was often overlooked:

> *'A further issue when considering consequences will be the number of persons likely to be effected by a particular hazard. In many cases, operational and maintenance staff will only be present in the hazardous zone on an infrequent basis and this should be taken into account when predicting consequences. Care is needed when using this statistical approach since it will not be valid in all cases, such as where the hazard only occurs during start-up and staff are always present. Also considerations should be given to the potential increased number of people being in the vicinity of the hazardous event as a result of investigating the symptoms during the build-up to the event.'*

Routine Re-assessments

The SIL guidance BS EN 61511-2:2004 states under section 5.2.5.3:

> *'A review of the SIS performance should be periodically undertaken to ensure the original assumptions made during the development of the safety requirements specification (SRS) are still adhered to. For example, a periodic review of the assumed failure rate of different components in a SIS should be carried out to ensure that it remains as originally defined. If the failure rates are worse than originally anticipated, a design modification may be necessary. Likewise, the demand rate on the SIS should be reviewed. If the rate is more than that which was originally assumed, then an adjustment in the SIL may be needed.'*

Further the Buncefield Process Safety Leadership Group report (para 17) says:

> *'The format and detail of the LOPA report should facilitate future internal review by the operating company and should also reflect the likelihood that it may be scrutinised by an external regulator and other third parties.'*
>
> http://www.hse.gov.uk/comah/buncefield/fuel-storage-sites.pdf

How many companies have included in their action management reminder systems a routine review (even if only every five years) of all existing SIL assessments within the company, possibly linked to their hazard identification reviews which should pick up new scenarios for assessment, and provide input on incidents, failures and site changes which should also be used within SIL assessment reviews? There are very few companies I know who routinely review their SIL assessments to verify that the layers of protection stated are still valid (whether listed within the risk graph frequency assessment, or in a standard LOPA format). Whilst the BS EN 61511 guidance quoted indicates that this review is aimed at failure frequencies, the other layers credited also need review. A review of the vulnerability would pick up failures in access controls, for example, but should also pick up changes in the area.

I have a concern that few management of change procedures include a requirement to check whether the modification proposed would impact an existing SIL assessment, for example by locating equipment nearby and hence changing the vulnerability of personnel or creating a revised potential consequence. Further, even if the wording on the procedure does require this to be done, when I have questioned people as to how they have checked this, few, if any, have confirmed that they have pulled out all the existing SIL assessments to check whether they would compromise any of the protection layers stated. Even those people who were involved in the SIL assessment process would not remember the layers of protection stated for each event, and in the majority of cases the personnel involved in the management of change process are not the same, which becomes more of an issue as time goes by and personnel on site change.

This change in personnel has further implications, as the ownership for the SIL assessments is lost. This is compounded especially for some of the earlier assessments where the justification for the value stated, e.g. for vulnerability, may not have been fully documented. Therefore if the justification was a result of restricted access to an area, but the exact details are not stated and cross referenced to a procedural control, then it is not surprising that the control is lost. In many cases the SIL assessment raises an action to implement the access controls, for example, to achieve the level of protection selected, but if the close out of that action is incomplete, or is not cross referenced back correctly, then completing regular reviews becomes more difficult.

Alarms

Another aspect of human response which may be incorrectly considered within the SIL assessment is the actual response to the alarm. There are many assessments I have seen where credit has been taken for this response, but when challenged the reality is that the response would be too late, or insufficient to prevent the hazardous event. An example would be where a tank high level alarm sounds. If the direct response to the alarm does not stop the level rising, or the action required takes so long that the high level trip or overflow activates first, then no credit can be taken. The HSE guidance is a minimum of 20 minutes should be available for the response before a hazard occurs or the next layer of protection activates for it to be considered an independent layer. If manual intervention is required then the actual time taken to respond should be assessed. This is explained in the previously referenced Buncefield PLSG report in paragraphs 222 – 225.

A further example is where the risk is from a polymerisation reaction. If the alarm and required action are to prevent the initiation of polymerisation the credit is valid. However; if the alarm indicates the onset of polymerisation then there may not be a suitable response which can inhibit the polymerisation reaction and stop it safely and reliably: if which case the only response may be to evacuate.

Another example I've encountered which reflects concerns raised by the HSL LOPA course:

> *'erroneous assumptions can be made because many risk assessments are desk based and*
>
> *- lack the necessary information on how operations are performed in reality*
>
> *- the assessment team is not representative'*

In one case the assessment team wanted to take credit for the operations team activating the emergency stop button for a compressor following alarm activation. Having experienced first-hand reluctant to use 'the big red button' I requested that the team rang the control room to confirm the action that would be taken. Sure enough the response was that they would send someone out to look at the machine and would monitor it, but in reality would wait for the trips to activate rather than shut it down themselves. The reasoning that it would then take 24 hours to restart the plant, and involve a lot of work for themselves was the justification for relying on the trips ' because that is what they are there for!'

Another mistake in some older assessments is taking too much credit for alarm, trip and human response inter-dependence. Examples of early assessments I have seen have taken credit for the alarm, the human response to the alarm and an automated action from the same instrument all as independent layers of protection, i.e. three layers for what actually is one (or none if the automated action is being assessed for the required SIF trip integrity). Whilst most teams now recognise that a common instrument means a single layer I still occasionally have to explain to people that there may be two functions from an instrument but both are still subject to a single failure mode. The difference between an automated response from within the basic process control system (BPCS) and an independent trip function still causes confusion, so I try and differentiate for people by using the term 'interlock' for the BPSC action being considered as a layer of protection, and trip for the SIF being considered as the ultimate protection within the SIL assessment. That a SIL 1 SIF can be managed within the BPCS, if correctly designed, is another cause for confusion. It is much easier for people to understand the differentiation if the company policy is that all SIL classified SIF must be independent of the BPCS.

When designing a new plant, ensuring that the BPCS input/output cards are managed to give independence between initiating event failures, layers of protection, interlocks and trip functions (if applicable) is relatively straightforward. However retrospectively confirming whether this is the case for an existing plant, especially an old one, is likely to require a physical inspection. I do have a concern that this is not completed for existing systems, and that some companies trust to luck that this is the case.

Another issue which is evident from the use of risk graphs is that the output is only given as a SIL classification, e.g. a requirement for a SIL 1 or SIL 2 safety instrumented function (SIF), compared to a LOPA output, which provides a numeric value for the risk reduction factor (RRF) required. To be on the safe side for the risk graph it is therefore necessary to aim for the top RRF associated with the SIL classification, which puts an additional onus on the verification process. Whilst the difference is less than an order of magnitude it may result in additional cost to achieve a suitable SIF.

Action Completion

The completion of actions arising is the main area where I have concern that company SIL assessments are not robust, and may not reflect what is actually installed. As layers of protection are considered during an assessment the suitability of an identified protection may be queried. For example where response to an alarm is considered, the routine testing of that alarm, and the verification of the set point versus what is required may be queried. Therefore the team know that a layer is in place, but action(s) are raised to improve the robustness of the layer. These actions may be seen as relatively low priority compared to others raised during a series of assessments, and so they drop down the priority list. As plant life moves on, more and more actions from various sources are piled on to the individual concerned, and eventually the action is passed to someone else. Too many times I've checked the action close out statement within a company management system, and the action is marked as 'closed', when in reality the wording states that they have asked someone else to complete the action. As this transfer of responsibility is not formalised, the action is lost, and never completed. However, the SIL assessment continue to takes credit for the supposedly robust layer of protection, and a superficial glance at the actions implies that the action was completed. In reality, the day the alarm is needed it may not work, as it has never been tested, or the set point is incorrect, and so the potential for a hazardous situation increases.

Other types of actions are also vulnerable to incomplete closure. Often an action is raised to verify information assumed within the assessment, and these are also often considered low priority to close out, whereas they could have a significant impact on the outcome. Team members will frequently firmly believe something about the plant, e.g. how it is designed and/or operated, and as such the assumptions are considered to be true. However; especially for older plants, these assumptions may be folk lore, and the reality would change the risk assessment. Therefore the non-completion of the action is not due to mal-intent, but purely relative priority, and a belief that it is not important.

For many people 'completing the SIL assessments' is seen to be the end of the journey. However, these assessments are only the determination of what is required, and without the completion of the actions even these assessments are not complete. Following on from the determination are all the other steps within the life-cycle of the process, involving significant amount of further work. But when asked 'are your SIL assessments up to date / complete?' some members of the plant management team may reply in the affirmative, because they sat through countless meetings for the 'SIL assessment'. But if those assessments have not been converted into an adequately managed hardware and software system then the facility is still at risk.

Human Error Potential

Another factor considered within the layers of protection where the expectation from the HSE has changed over time is the potential for a human error. BS EN 61511-3:2004 (annex F) guidance provides informative values. Section F6, Table F3, indicated some probability of failure on demand (PFD) values to use for human performance:

'Human performance (trained, no stress)	*PFD 1.0 x 10^{-2} to 1.0 x 10^{-4}*
Human performance (under stress)	*PDF 0.5 to 1.0'*

In many earlier SIL assessments, whether using LOPA or calibrated risk graph, values for human error potential (HEP), for example when following a procedure, were within these stated ranges. However; the latest guidance from HSE, states:

> *'197. In most cases, a human error potential of 0.1 can be considered a conservative or cautious estimate of the risk of human failure. This value can generally be accepted as appropriate for use in order of magnitude tools, like Layers of Protection Analysis (LOPA). However, human factors specialists would still expect to see the duty holder demonstrate a thorough understanding of the tasks being undertaken and the conditions in which they are performed.*
>
> *198. Claims of reliability beyond 0.1 will require significantly more demonstration and justification; typically this is when a site might use a human reliability assessment tool but quantification is not always necessary. The HID Human Factors Specialist Inspectors team advocate using a qualitative approach to ensure the duty holder has a thorough understanding of the issues. Where quantified methods are used, HSE has found that values are often taken from publicly available data sources and HRA methods without any justification or consideration of the site specific conditions that might influence their applicability. For example, documents such as 'BS EN 61511-3:2004 (annex F) and the Center for Chemical Process Safety (CCPS) book on LOPA have tables that provide examples of HEPs. While these values are probably appropriate in many situations, the associated text to describe the context is extremely limited; duty holders need to consider how applicable the data are to the situation being assessed and to justify their use. If a duty holder has adequate site-specific performance data regarding human reliability, this data could be used to support HEPs obtained from HRA methods and other sources. This historical data can be considered adequate if it has been collected over a sufficient timescale to be statistically significant. However, in many cases such data are not readily available and duty holders, having decided on a quantitative analysis, must draw upon their knowledge of the task to work through a HRA method.'*

http://www.hse.gov.uk/landuseplanning/failure-rates.pdf

As many companies will not have reviewed their SIL assessments since this updated HSE guidance was introduced (28/06/2012) there is likely to be credit taken which is not suitably justified. Whilst the value used may remain valid, the documented justification needs to be provided. Providing this may require significant work if the related procedure, for example, has not been reviewed in line with human factors best practice for a safety critical procedure. Conversely, changing the value used to 0.1 for the layer of protection may result in an increased RRF required.

Ignition Probability

The use of ignition probabilities is an area where I have seen increased conservatism over the years. Many early assessments used flammable ignition probabilities based of historic data for small, medium and large liquid or vapour releases, which related to offshore releases into zoned areas. Following the Buncefield incident, it is not acceptable to use relatively low ignition probabilities for larger releases, especially if the cloud could drift into an un-zoned or off-site area. Re-assessing ignition probabilities used for early risk assessments may require additional layers of protection.

Demand Rates

Once the systems are in place and operating for the safety instrumented functions across a site, I still have concerns that some of the layers of protection, and the safety functions themselves are actually being used for process control. If a relief valve is lifting routinely and essentially acting as pressure control valve rather than a relief valve, then taking a relatively low PFD for it is not good practice and the PFD should ideally reflect that used for a control device. The ultimate action in this case is to address the cause of pressure fluctuation leading to the relief, and if the root cause cannot be easily addressed then a properly designed control function should be added to the system. Further, feedback from the assessment teams too frequently indicate that the trip function which is being SIL assessed activates routinely, in one case almost weekly. In this situation it is acting as a control not a safety function and should not be considered as a suitable SIL rated SIF. Again root cause analysis and redesign are liable to be necessary, rather than accepting the status quo.

The use of well-defined leading and lagging process safety performance indicators (PSPI) would highlight these issues, but although the sites may be aware of the frequency of demand they do not seem to always recognise the significance.

BS EN 61511-2:2004 section 9.2.3 states:

> *'There are some applications where demands are frequent (for example, greater than one per year) and it is more appropriate to consider the application as continuous mode because the probability of dangerous failure will be primarily determined by the failure rate of the SIS.'*

However, rather than adapt the system to the demands required for a continuous mode SIF most companies just carry on, whilst trying to address the issue, not realising the risk they are therefore accepting.

'SIL Rated' Equipment

Purchasing suitable equipment is another area where there is a risk of compromising the good intentions of the related SIL assessment. As the expectations to comply with SIL guidance becomes more commonplace, equipment suppliers are 'helpfully' supplying SIL rated equipment.

The wording from BS EN 61511-2:2004 section 9.2.3 states:

> *'The targets for average probability of failure on demand or frequency of dangerous failures per hour apply to the safety instrumented function, not to individual components or subsystems. A component or subsystem (for example, sensor, logic solver, final element) cannot have a SIL assigned to it outside its use in a specific SIF. However, it can have an independent maximum SIL capability claim.'*

Despite this there is a fundamental failure to understand the implications of purchasing such items, for example valves or instruments which are sold as having a 'SIL rating'. I am concerned by the number of project and site based people I have talked to who think that if they purchase such items then the installation will be suitable to meet the indicated SIL rating. They seem to miss the fact that the whole of the instrumented function loop – from initial input to final element needs to be verified as a single entity to meet the RRF; instead they are latching onto the purchase of suitable items of equipment within that function as an easy way out.

Equipment Suppliers

When dealing with equipment suppliers who have been asked to provide SIL rated equipment as part of the contract, it is fortunately becoming less common that this request is met with a blank look, or a response that 'no-one else asks for this!'. The former is especially the case if the vendor is non-European; the latter when equipment is suitable for use in a range of industries, for example refrigeration also supplied to the food or agricultural industry. When pushed, a SIL assessment for their equipment is sometimes produced, which is entirely generic and has no link to the end-use of the equipment, the location it will be used in, population levels on site, or any hazardous events initiated external to their equipment. Typically the assessment has also used a standard, non-calibrated risk graph, so the assessment conclusion that 'one SIL 1 rated high temperature trip is needed which stops the machine',

bears no relation to the risk of using the machine in the proposed plant setting, and is not consistent with the tolerability targets of the purchasing company.

Over the years I have facilitated several SIL assessments between the equipment supplier and the purchaser, and they can be tortuous as they are often starting with a reluctant and frustrated supplier who is worried that they will lose the job if they can't deliver what is needed, and are entering the unknown, as they have not gone through the process before. I still find that an order has been placed without having specified the SIL requirements, so dependent on the contract wording either the supplier or the purchaser is going to suffer unexpected cost implications.

A positive outcome for all is often achieved, with equipment suppliers proactively suggesting modifications to their basic models to achieve the identified requirements, sometimes simply by changing an instrument type, a fitting or a relay switch to provide the reliability needed. In one case a burner management system (BMS) supplier sat through a two day risk assessment, having provided a totally unsuitable generic SIL assessment, only to pronounce at the end of it that they already had an independently verified SIL 3 rated BMS 'black box' which they could retrofit to replace the non-SIL rated equipment that had been supplied and deemed unsuitable by the Health and Safety Executive (HSE). I think they had just wanted the learning experience of sitting through the assessment to understand the process for future enquiries.

The key learning point is to ensure that the purchasing department, and those who write specification sheets, are familiar with the principles of safety integrity, if not the details, and are aware that accepting the basic vendor package may be suitable for a low hazard use, but will almost certainly need to be upgraded for use in a high hazard facility. Planning safety integrity function discussions in at an early stage with potential suppliers can remove the risk of installing unsuitable standard vendor packages which then need expensive retrofits to be fit for purpose.

Tolerability Targets – human harm

Fundamental, of course, to a good quality SIL assessment process is the use of suitably defined tolerability targets for the facility concerned. Most corporate tolerability targets I have seen range between $1x10^{-5}$ and $1x10^{-6}$ per annum for a single fatality, with the use of the former typically only suitable for a relatively low risk installation, whilst using the latter will not lead to a challenge to justify its use. However as the potential number of fatalities rise there is a much greater discrepancy between values used. I would prefer to see the first differentiation between 'numbers' of fatalities to be 'up to 2' and 'more than 2' as so often operations on site, especially those involving maintenance personnel or associated with start-up or shut-down, are liable to impact more than one person, so there is often a lot of debate regarding the actual numbers impacted. The tolerability targets used for SIL assessments should also be considered as part of the routine reviews, verifying those used are still in-line with the in-house and HSE expectations.

Tolerability Targets – environmental harm

The recent issue of the Chemical and Downstream Oil Industries Forum (CDOIF) 'Guideline Environmental Risk Tolerability for COMAH Establishments' will require companies, especially those subject to the Control of Major Accident Hazards (COMAH) regulations, to review their significant environmental release scenarios.

http://www.hse.gov.uk/aboutus/meetings/committees/cif/environmental-risk-assessment.pdf

This document provides detailed guidance on whether a release could be a potential major accident to the environment (MATTE), or whether it would be classed as sub-MATTE. In conjunction with decision tables for the MATTE type, the severity of the MATTE, and the duration of harm, the target tolerability for the intolerable and broadly acceptable boundaries is provided. Tolerability targets for significant environmental harm were not previously so clearly stated, so this guidance may change SIL assessment targets used, and hence require re-assessment of those scenarios which have MATTE potential. Whilst most hazard scenarios subject to SIL assessment are not deemed to result in a potential MATTE this clarification is a useful tool. It is worth noting that the harm potential assessment is prior to any remediation, such as oil skimming off-site (though not to mitigation layers such as effluent treatment or other on-site facilities), so if remediation was ever claimed as a mitigation layer then this needs to be removed from the assessment.

Summary

Whilst this paper only discusses a few examples of potential concerns for the validity of existing SIL assessments, the intent is to highlight that all existing SIL assessments should routinely be reviewed carefully to ensure they are fit for purpose. Where changes are identified this should not be seen negatively as the fault of the original assessment team, but simply as part of the continuous improvement progression expected for process safety and environmental risk assessments. For all new processes introduced there is a learning curve, and it is to be anticipated that expectations and standards changes over time, especially as experience is gained by those involved. It may raise a few questions if a previously acceptable risk is re-classified as requiring an additional layer of protection, but better that than a future major accident hazard occurring.