SAFETY SYSTEMS

The SCSC Newsletter

Spring 2017 Volume 26 Number 1

Guest Editorial: Twenty-Five Years of Safety - SSS'17: A Personal Perspective

by Dick Selwood

Much has happened in the safety-critical field since the first Safety-Critical SystemsSymposium (SSS) in 1992. At SSS 17 in Bristol these advances were celebrated by several of the keynote speakers, who had themselves been instrumental in the last quarter century's developments (and indeed before).

As it is impossible to compress three days of thought provoking presentations and debate into the space of a newsletter, you may want to buy the 500 plus pages of proceedings, *Developments in System Safety Engineering: Proceedings of the Twenty-fifth SSS,* on Amazon (£10.25) or download individual papers for a small charge (free to members) from the Club website, www.scsc.uk. The speakers' slides are also available from the website, and we will soon post video highlights from the Symposium there too.

Meanwhile, what one member of the audience saw as highlights (not that there were any low-lights) follows.

The historian keynoters were:
Audrey Canning, *Functional Safety: Where have we come from? Where are we going?*

Robin E Bloomfield, Kate Netkachova, Peter Bishop, *Confidence in a Connected World: Safe, Secure, Resilient and Autonomous*

Nancy G. Leveson, *My 36 Years in System Safety Engineering: Looking Backward, Looking Forward*

Dewi Daniels, *From the IBM 29 Card Punch to the Boeing 787 Dreamliner (and Beyond)*

Ron Bell, *Safety-critical Systems - A Brief History of the Development of Guidelines and Standards*.

In this Issue

p.1 Dick Selwood looks back on SSS'17
p.4 Working Group Report: Safety of Autonomous Systems
p.5 Steve Thomas applies structured design principles to Safety Management Systems
p.12 Vicky Gliddon discusses safety culture assessment
p.18 John Ridgway on the implications of vagueness
p.22 Steve Gandy reflects on IEC 61511
p.35 We remember **Prof. John Knight**
p.35 Calls for papers
p.36 Events Diary

Please send articles or news items for future editions to the editor, Katrina Attwood, at newsletter@scsc.org.uk

Together, these talks established the setting for the establishment of the SCSC, and followed not just with a chronicle of the activities that have taken place in the world of system safety engineering, but also an insight into the complexities of developing safety standards. One such is *IEC61508: Functional Safety of Electrical/ Electronic/Programmable Electronic Safety-Related Systems* (published 1998), the foundation stone for many more industry-specific standards, such as ISO 26262 for automotive. In some cases, herding cats may be easier than getting agreement between the differing interests and different countries involved in producing a standard. The amount of sheer hard work that has gone into these documents is astonishing. It wasn't entirely positive: several speakers held that the existing standards, while better than nothing, are a long way from perfect. One speaker said that, "most standards are unverified hypothesis", and another stated that while standards are often the consensus of a large group of experts, they are not evidence-based.

Tom Anderson, the long term Chair of SCSC, gave a wonderful talk entitled, *What can I say?* (One might argue that finding something to say is never a problem for Tom!) The talk brought together several of the most important elements of Tom's life: safety, mountains, trains and single malt whisky, to provide a very personal perspective on history.

The first, non-historical, keynote was particularly thought-provoking and provided a very high benchmark for the following three days. In *Playing Catch–Up: The Fate of Safety Engineering?*, John McDermid compared safety engineering to other engineering and scientific disciplines. Dipping into the philosophy of science, including Popper and Kuhn, he argued that other disciplines, including engineering disciplines, use experience either from experiment or in the field to challenge and modify the prevailing orthodoxy. By contrast safety engineering is always playing catch–up. While this is inevitable, as the design engineer is advancing the state of the art, it should be possible, by establishing a set of principles, and building better links with the design community, to reduce the catch–up distance drastically and so improve safety.

John McDermid mentioned Systems-of-Systems (SoS), one of the latest buzzwords. A paper by Mike Brownsword, Andy German and Ian Mitchell looked at the problems arising from SoS, which they define as the result of combining two or more independently managed elements subject to different management, investment and control policies. Driverless trains, planes/drones and autonomous and connected vehicles are already SoS, and the technology is extending into industrial, medical and other safety-critical applications. The paper discussed various approaches to understanding and identifying the safety problems that are created by SoS.

An approach to the safety of autonomous vehicles was the subject of a paper by John Birch, Mark Cousen and David Ward. The product integrity assurance argument framework for vehicle autonomy builds an MISRA's work and the ISO 26262 standard and is put forward as a starting point for more detailed work.

Related to this are the issues of Reduced Crew Operations (RCO) in the civil aviation industry. Airlines are investigating having one person (instead of two) in the cockpit, and supporting them from the ground through high speed data communications. Kevin Driscoll explored the safety issues, ranging from possible cyber-security problems with communications (there are already documented cases of individuals pretending to be air traffic controllers), through the

possible impact of latency caused by deeply secure encryption and decryption, to the problem of a pilot with a full bladder. He was uncertain that the economic benefits of RCO, once these problems were addressed, were ever actually going to be sufficient to justify the approach.

Cyber security problems in a hospital were the subject of Harold Thimbelby's paper. A hospital has seen nurses successfully prosecuted for falsifying data. A key plank of the prosecution was the evidence of records in a database. In a further trial, Thimbelby successfully discredited the evidence derived from the database, because of basic flaws in the recording process, poor management of the database and even editing between the records being created and the police downloading the data. It was disconcerting to the entire audience that, under English Law, a computer producing an evidential record is presumed to be working correctly.

Accuracy of medical data was the subject of a paper by Tom Adams, Paul Hampton and Mike Parsons. Data has previously been virtually ignored as a factor in safety. By drawing on case studies and personal perspectives, the paper identified data problems and suggested mitigating techniques, leaving the impression that there is still a great deal of work to do. The SCSC has a working group on data safety and its latest report *Data Safety Guidance (Version 2.0)* is available for free download from the Club website.

As well as poor data, another area that has not been well studied is the threat posed by the insider. As Ryan Meeks and Robert Dickie discussed, there are a number of potential threats, ranging from mistakes through espionage, fraud, theft (particularly of IP), to sabotage. They used cases such as Chernobyl, Edward Snowden and the German Wings flight 9525 to make a strong case for always considering the insider as a safety issue.

A new problem that faces safety engineers is the increase in what Rob Ashmore and Elizabeth Lennon call "non-traditional software". Current standards assume that software is created from a set of requirements, which include safety requirements, against which the software can be verified. However machine learning, using techniques such as neural networks to change code, or even create entirely new code, is moving out of the lab and into real applications. Here the link between the initial high level requirements and the executing code is not so easily determined and such systems are not amenable to many existing safety techniques. They see hope in the Four Plus One principles (set out in 2013 by Hawkins et al) and, in analysing how these can be applied to machine learning, came up with an additional principle. So they put forward Four Plus Two as the basis for developing safety engineering for machine learning.

So far, we have been looking at emerging problems and possible ways of solving them, but as is usual at the SSS there were a number of papers reporting work that has analysed existing problems and have either solved them or are actively implementing solutions. These included: human factors in helicopter operations; certifying a multi-core processor architecture; software handling of hardware errors; electromagnetic disturbances; formally verified optimising compiler, time-triggered software architectures; formal proof in DO-833; and sneak path analysis.

The final keynote was by Les Hatton. He didn't look at history, but instead addressed the issue of raging feature-itis - adding more features to a product - one contributor to the 20% per annum growth in system source code size. This is, in part, because there is no data for the system engineer that matches the structural

engineer's ability to say, "A beam that long will break." Instead, all that can be said is "My experience leads me to think that adding that new feature will cause severe unintended consequences." Les has been working on trying to bring measurement to software quality evaluation, but has concluded (not without a great deal of time and effort) that we are more likely to improve the safety and reliability of systems containing software by pursuing the traditional engineering virtues of redundancy and designing for failure than by pursuing studies of the fabric of software.

This is a very brief summary and doesn't include the stimulating after-dinner talk by Air Marshal Julian Young, the panel session nor the to-and-fro around the exhibitors' displays, across the coffee cups and the lunch plates, and during the beer tasting.

Dick Selwood is a freelance writer with a strong interest in safety and related matters.

SCSC would like to thank all speakers, session chairs and organisers for a wonderful 25th Anniversary Symposium. We look forward to seeing you all again next year at the Principal Hotel, York, 6th - 8th February 2018.

Working Group Report: Safety of Autonomous Systems

by Rob Alexander and Philippa Ryan

Building on the success of the long-established Data Safety Initiative Working Group, SCSC has recently established new Working Groups, to address various pressing issues in system safety engineering. To date, three further groups have been established: the Safety of Autonomous Systems Working Group, the Assurance Cases Working Group and the Service Assurance Working Group. *Safety Systems* will be reporting on developments within the Working Groups. In this issue, the focus is on the Safety of Autonomous Systems Working Group (SASWG).

The SASWG has been formed to produce clear guidance on how autonomous systems and autonomy technologies should be managed in a safety-related context, in a way that reflects emerging best practice. The first meeting was held at Conekt in Solihull on 17th January. The discussion was interesting and wide-ranging, and considered, among other things, domains and applications of interest and terms of reference. The group agreed to produce an outline – content list and candidate structure – of a cross-sector guidance document by the end of the year. We plan to hold meetings roughly once a month for the foreseeable future.

The second meeting was held on 9th March at the Transport Systems Catapult in Milton Keynes. A number of domain-specific examples were presented, from automotive, medical, civil aerospace, defence, rail, maritime and space. These were used to establish the zone of interest for the guidance, in terms on internal and contextual complexity.

The third meeting will be held in early May. Location to be confirmed. For further details, please contact Rob Alexander: rob.alexander@york.ac.uk.

See the *Events Guide* on page 36 for details of other forthcoming Working Group meetings and other Club events.

Structured Design of Safety Management Systems

by Steve Thomas

Introduction

The adoption of an explicit and systematic approach to safety by Air Traffic Management (ATM) service providers in Europe was motivated to a large extent by the emergence in the late 1990s of a series of safety regulatory requirements from EUROCONTROL, the European Organization for the Safety of Air Navigation. As a result, air navigation service providers were faced with the need to introduce unfamiliar new activities within their organizations via the development and deployment of a formal Safety Management System (SMS). The SMS would have to be fully compliant with regulatory requirements, but also be practicable in terms of the resources it demands during operation, and be compatible with other business processes. This challenge is common to providers of products or services in any regulated, safety-related industry but leaves the question of how an organisation should go about solving it.

The approach in this paper recognises that, although it is non-physical in form, an SMS is nevertheless a system. Its development therefore can, and should, exploit engineering design principles within a structured approach. As a consequence, there is a need to represent SMS requirements and design at different levels of decomposition using suitable notations. In view of similarities between an SMS and real-time software, the paper proposes an established software design notation for this purpose.

Design Hierarchy

The regulatory compliance aspect can be viewed as an information transformation problem. The regulatory requirements must be analysed and interpreted, and then be progressively decomposed and elaborated until a much larger set of interconnected low-level actions has been defined which is suitable for direct execution by the safety 'actors' within the service provider's organization. The safety actors comprise those operational, technical and management personnel who perform the activities defined within the scope of the SMS.

The intermediate representations within this transformation hierarchy may be referred to as 'Safety Policy' and 'Safety Principles' which, like the regulatory requirements, are usually not directly actionable by safety actors. The derived low-level actions, which can be directly traced back to regulatory requirements via the Safety Policy and Safety Principles, are collected together in a set of procedures, with the actors being defined via a safety organization whose members have specific accountabilities and responsibilities. A typical logical model for the Safety Management Manual (SMM), which can be thought of as the operational requirements part of the SMS, is shown in Figure 1 below.

Some of the design choices to be considered when deciding on the structure of this logical model and the contents of each block are:

- The number of transformation levels that are needed, the amount of detail at each level, and the ownership of each block;
- The optimum grouping of the low-level actions within and among the set of procedures to achieve the best 'size' for the corresponding safety processes, and to minimise their interdependencies;

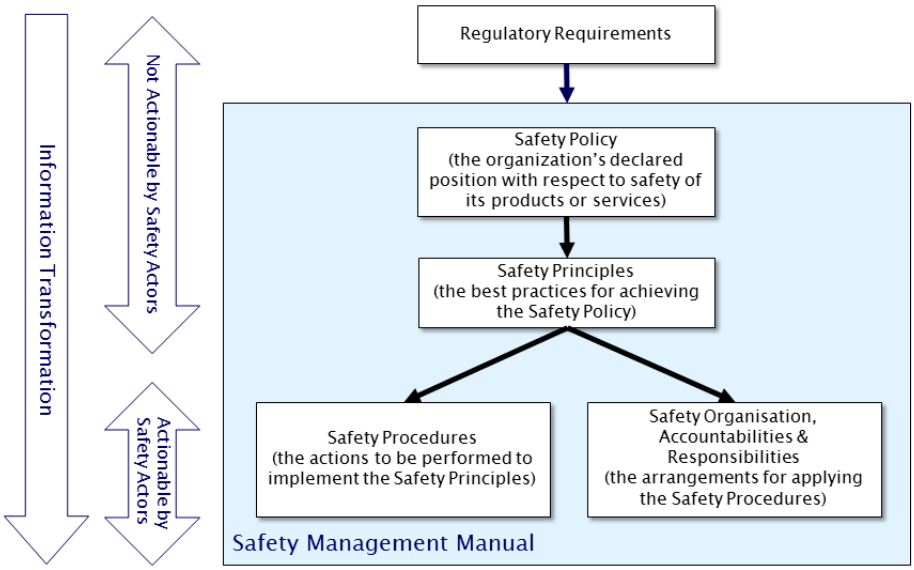


Figure 1: SMM Logical Model

- Balancing prescription (which makes procedures actionable and auditable) against objectives (which permit innovation and flexibility) when defining the safety activities;
- Defining who are the intended safety actors;
- Defining the required actions in a way that is consistent with the capabilities of the intended actors;
- The distribution of information about how to conduct safety activities in and among procedures (mandatory action), guidance (recommended action) and training;
- The integration of pre-existing business processes with the SMS.

All of these are variables influencing the SMS design. It follows that arriving at the 'best' SMS design is an optimization problem as for any technical system design exercise. However, since the scope of an SMS can encompass existing (or modified) business processes in addition to many new safety processes, in practice it will inevitably comprise a large set of processes and interactions even when optimised. Therefore, a way of representing SMS processes, their boundaries, interactions and their relationship to actors is needed in order to facilitate investigation of the different SMS design options around process granularity, interfaces and integration.

It should be mentioned at this point that, in addition to the need for regulatory compliance, there will usually exist other external influences on the initial and ongoing development of an SMS. These can include the desire to satisfy safety standards, and incremental developments as a result of feedback from safety actors, audits, and *ad hoc* suggestions for safety improvement. The SMS should contain specific procedures for capturing this information in order to trigger and inform any necessary SMS design changes. Equally, it should contain a procedure for assessing the impact of any changes to the regulatory requirements. Consequently, a complete SMS will contain procedures which can effect changes to the SMS design itself.

SMS Characteristics

To implement these design concepts there must be, as a minimum, a means to trace and justify the relationship between the low-level safety actions and the regulatory requirements in order to verify compliance. This implies the use of an information hierarchy in which the different levels of definition, as described previously, are mutually consistent. Furthermore, the low-level actions in the procedures need to be articulated as requirements to be enacted by the (human) safety actors.

It is also apparent that an SMS is a real-time system, albeit consisting of human activity (even when the SMS is supported by tools). Such activity has associated inputs, outputs, triggers and time constraints. Moreover, since the SMS encompasses all the diverse aspects of safety management, it should comprise processes that are sufficiently decoupled to allow them to run concurrently and asynchronously.

Therefore, it can be seen that there are some parallels between an SMS and a software-based system, and between the means of their development. For example, there is a need for progressively elaborated representations in the requirements/design part of the development lifecycle (from user requirements down to source code in the case of software-based systems), and a need to model their behaviour (including real-time aspects). Once these similarities are recognised, one can go on to explore whether established methods used in software engineering could be beneficially exploited for the development of an SMS.

Structured Design

This paper proposes the use of the structured design technique from software engineering, which can potentially aid SMS development by enabling the visualisation, manipulation, and optimization of its data processing elements. One established technique that can readily be applied is the use of data flow diagrams (DFD) with real-time extensions [1], in which processes and their interactions can be diagrammatically represented at different levels of decomposition using the basic symbols shown in Figure 2 below.

Although the original DFD notation has been in existence since the 1970s, hitherto it has been used for representing human activities very rarely [2][3], and there is no published work on its use for development of a complete SMS in any industry sector. Therefore, while recognising that a number of other process modelling notations could have been used, it was decided to adopt the technique for the creation of an SMS for an Air Traffic Control centre in order to judge its merits. In particular, it seemed to offer the means to better manage the growth of the SMS and retain traceability of requirements as its development progressed. It was applied in a pragmatic way, without tool support, by personnel who were not experts in software engineering. The result was a fully-coherent SMS design, depicted as a DFD process model, in which all process inputs, outputs and triggers were represented and traceable.

It should be noted at this point that due to the non-physical nature of an SMS, there are some important differences between the use of DFDs for design of an SMS, and their use for design of software-based systems. One is that safety actors are internal to the SMS processes being represented rather than being part of the environment external to the system, as in the case of equipment. Secondly, since the documented inputs and outputs from safety processes likewise are considered to be an integral part of the SMS, DFD terminators can appear in and among the DFDs, whereas for software they only appear in the top-level DFD (aka

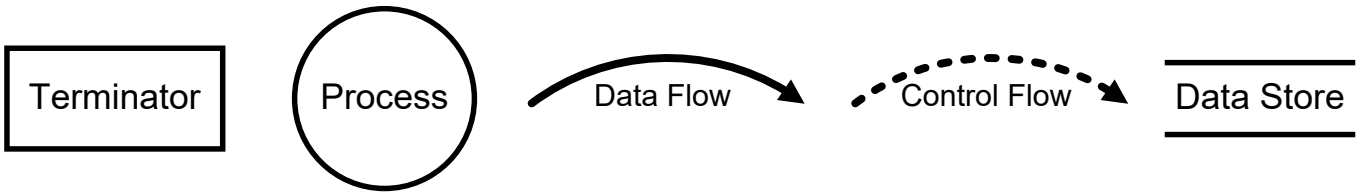


Figure 2: Data Flow Diagram Notation

the context diagram). Therefore, the DFDs herein are an adaptation of the pure notation found necessary to fit the specific application.

Safety Procedures

The generic relationship between the SMS procedures and the DFDs is shown in Figure 3. This shows how elements of the DFD process model are captured in the corresponding procedure. Depending on the SMS design, the procedure might capture the input and output objects (Object_1 to Object_3 and Database_X in the example) instead of the associated data flows (Data_1 to Data_6). For example, it may suffice for a procedure to mandate the document types to be used as inputs and outputs, rather than the individual items of data contained therein.

For the ATM application, the group of SMS processes which govern the reporting and assessment of safety occurrences is depicted in Figure 4. This represents a solution to the regulatory requirement known as ESARR2 [2], which mandates that all accidents and incidents in the airspace are reported and investigated. Investigation and reporting of operational and technical incidents affecting ATM service provision, but which have not led to an accident or incident in the airspace, is also required. These events are termed 'ATM-specific Occurrences'.

Figure 4 depicts the corresponding part of the SMS design in which the airspace accident/ incident investigation process is shown as 'OCC', with separate processes covering technical ATM-specific occurrences, non-technical ATM-specific occurrences, and external reporting. Visualisation has been enhanced by means of an outline colour for a DFD element denoting the source regulatory requirement document, and a fill colour denoting its ownership within the

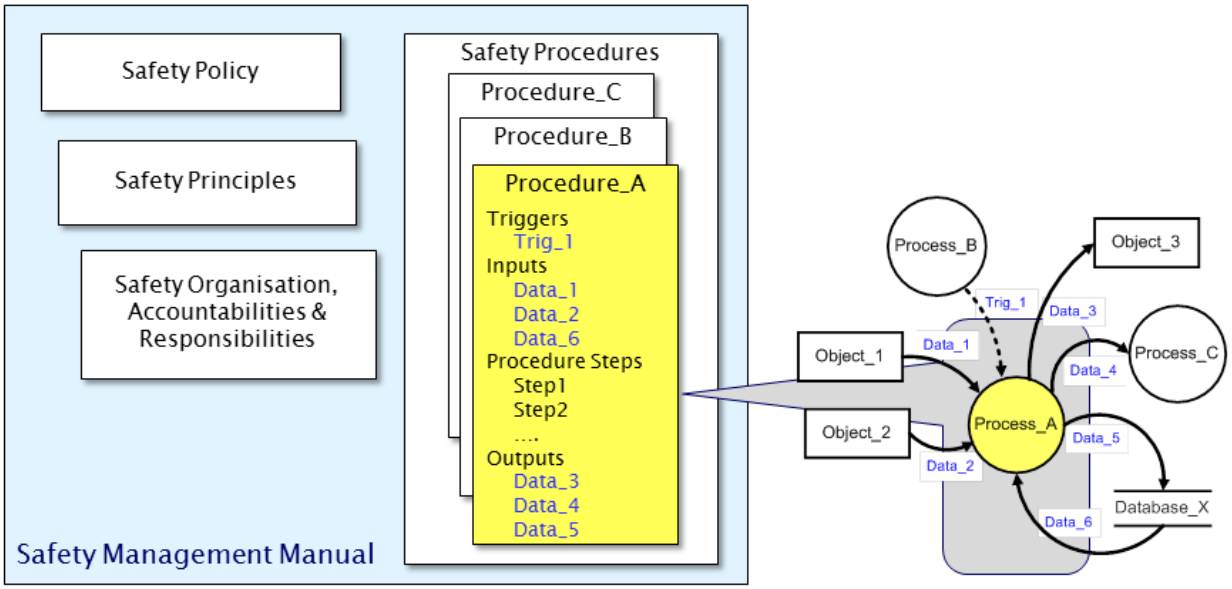


Figure 3: DFD Relationship to Procedures

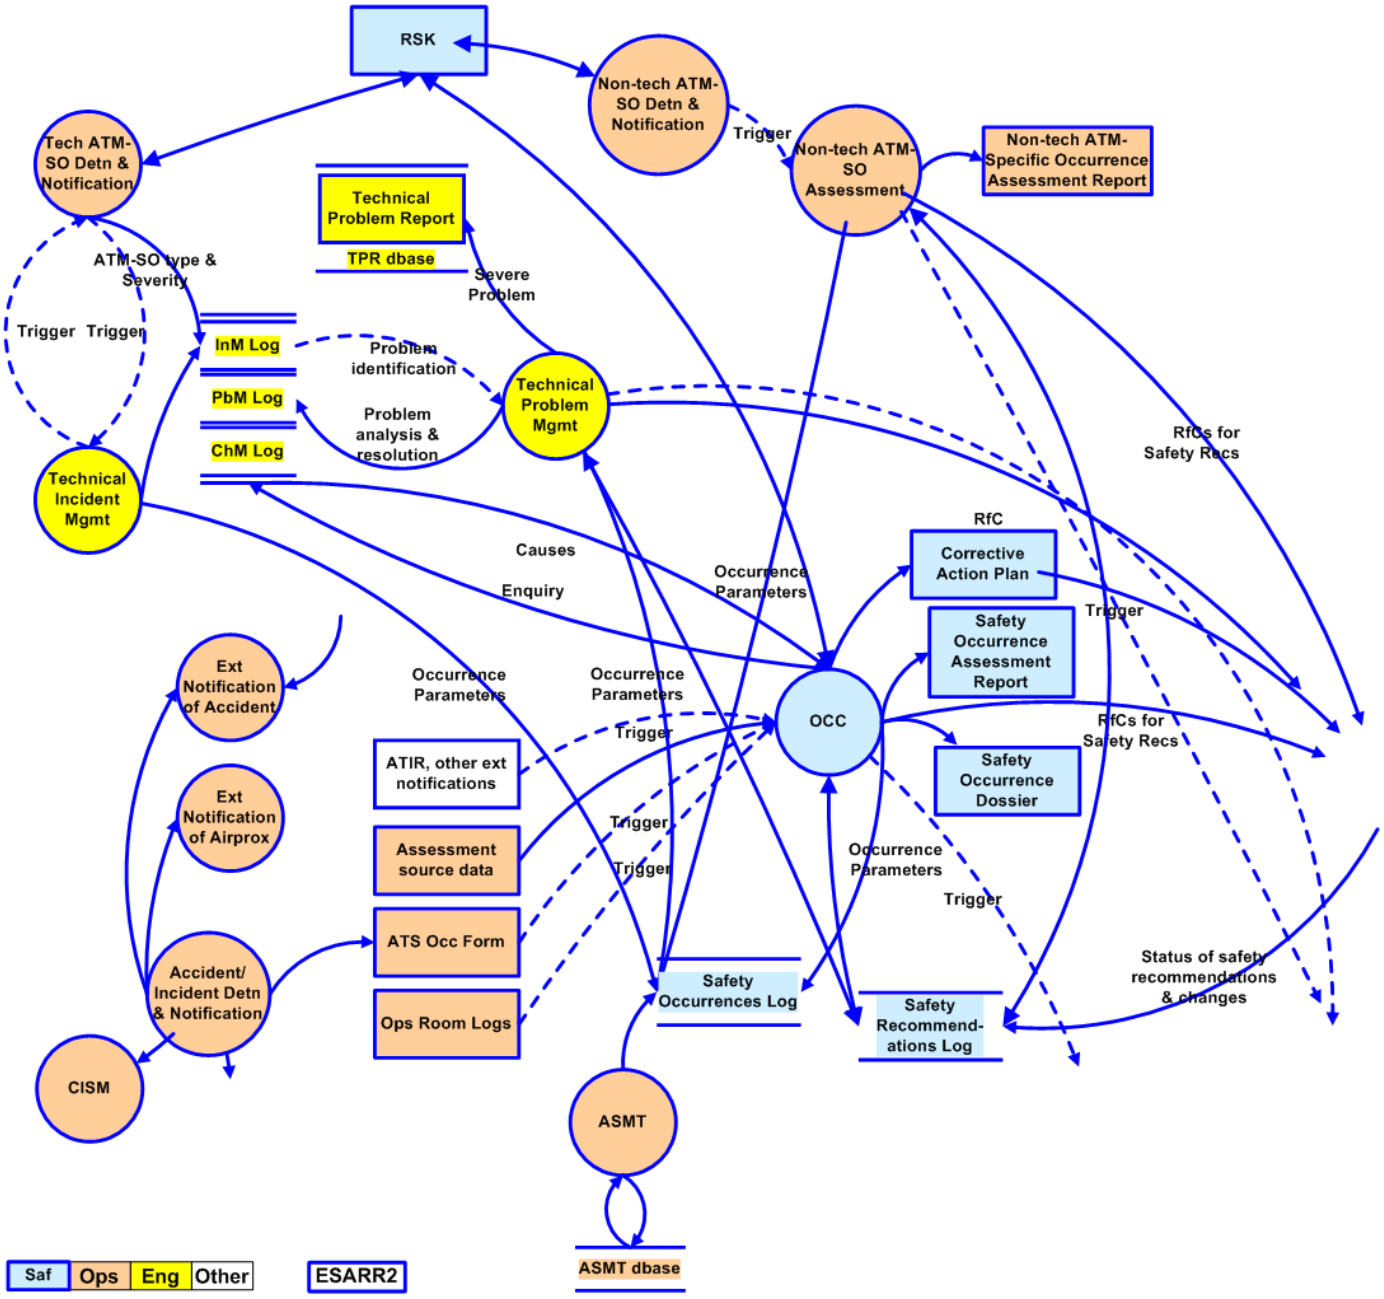


Figure 4: Safety Occurrences Process Group

organization (Safety, Operations, Engineering, Other). Flows which do not appear to terminate actually connect with other groups of processes, which have been masked out from Figure 4 for the sake of clarity.

Diagrams like this are indispensable for communication between SMS developers because they reveal exactly how each process fits into the overall picture and how its information flows are used. Safety data that is produced by a process but not consumed elsewhere can be easily identified, along with other structural inconsistencies. As implied by Figure 3 above, the DFD process model is documented outside of the SMM and is therefore invisible to the safety actors who need only base their actions on conventional, structured free-text procedures and guidance.

As well as safety activities conventionally described within procedures, it should also be recognised that meetings are processes in their own right. Therefore, they can explicitly be included in the SMS design

whether triggered by time (periodic meetings) or by events (ad hoc). Capturing meetings and their inputs/outputs in a DFD forces SMS developers and managers to consider their precise purpose and how they should connect to the rest of the SMS. In turn this can help in formulating their Terms of Reference, memberships and agendas. It will also reveal the precise means used for capturing important decisions from a meeting such as an 'approval to operate', and which other SMS process(es) subsequently use these as an input.

To illustrate these points, the ATM example in Figure 5 shows some of the important interfaces between a periodic Safety Committee meeting, the Safety Monitoring (MON) process, the Safety Surveys (SUR) process, and the regulator. This group of processes is part of the SMS design solution to a regulatory requirement known as ESARR3 [3], which defines the required components of an ATM service provider's SMS (broadly at the same level of abstraction as an SMS Safety Policy and Safety Principles). It can be seen that the Safety Committee makes use of a MON database, part of which comprises the Safety Occurrences Log established to satisfy ESARR2.

Safety Responsibilities

As well as process definitions, the SMS needs to allocate safety actions to recognisable entities within the organization. The actions should be meaningfully defined. Therefore, whereas a safety responsibility such as "the Head of Engineering shall ensure that the Safety Policy is implemented in the Engineering Department" is a worthy statement of intent and can easily be documented, it is not directly actionable (or auditable) because it is defined at too high a level (note 1). In fact, any responsibility statement phrased in the form "xxx shall ensure that yyy" could conceivably be discharged in a number of ways by the responsible person. This is because a number of subsidiary actions arising from the verb 'ensure' (such as 'plan', 'perform', 'verify', 'monitor') are implied. Collectively, these might satisfy the concept of 'ensuring', but have not been prescribed explicitly.

It is also conceivable for SMS developers to include less tangible responsibilities such as "an Air Traffic Controller shall have a pro-active attitude to Safety Management'" Although a desirable consequence of having an effectively operating SMS, again it is not necessarily directly actionable by the

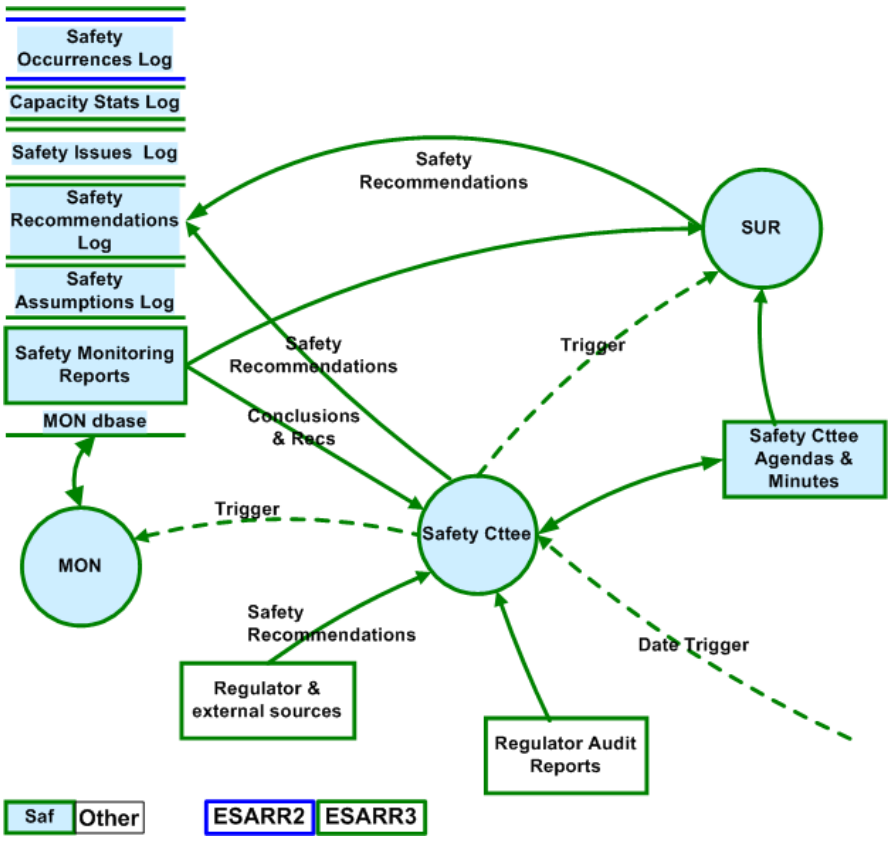


Figure 5: Safety Committee Process Group

responsible person, and may be difficult to verify by objective testing.

In order to have unambiguous, actionable, repeatable implementation of safety responsibilities, they should instead be stated ar a suitably low level of decomposition. The responsibilities, like the safety actions, therefore need to be elaborated (and usually delegated) in a way which satisfies these attributes.

Safety responsibilitles are in fact formal requirements for action. As such, they should as far as possible employ a common syntax, such as <actor><action><target> where all three components are fully traceable. This would provide the responsible person with clear pointers to the relevant SMS process or process input/output with which the responsibilities are implicated. A number of standard syntaxes can be employed to cover the range of human activity implied by the process model. As with the procedural elements in the SMM, it then becomes possible to derive (or at least to trace) the safety responsibilities from the process model, as illustrated in Figure 6.

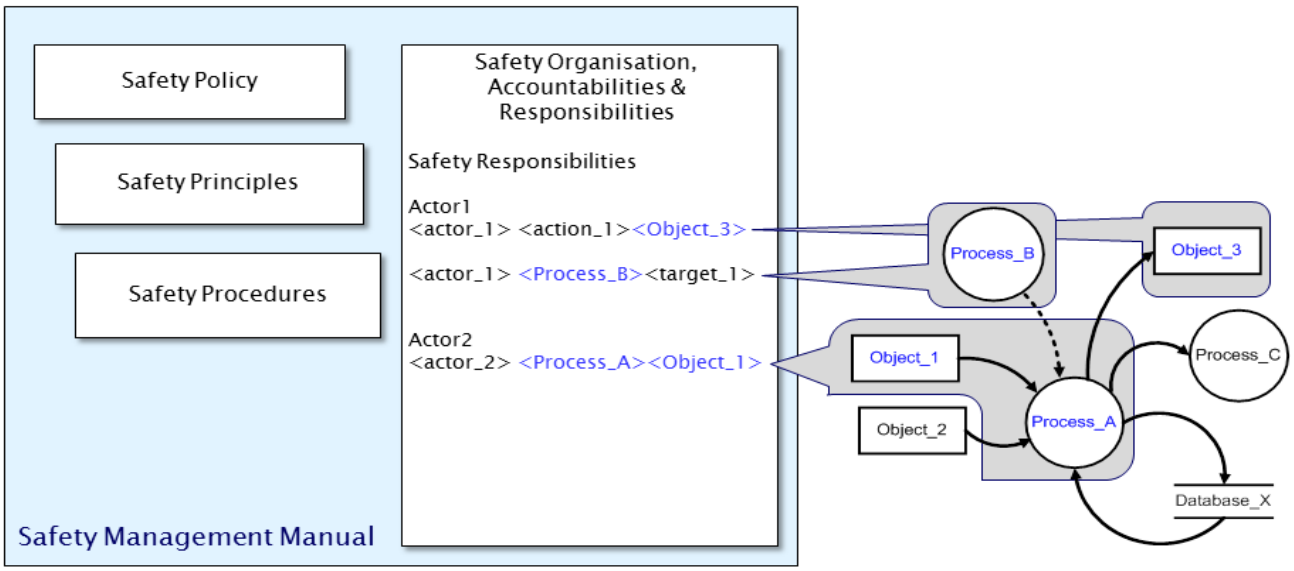
Examples of safety responsibilities using this approach are forms like: *<The Safety Engineer> <shall perform Functional Hazard Assessment of> <operational changes>* and *<The Head of Operations> <shall approve> <Safety Objectives for ATM System changes>* (see note 2 below).

Being able to identify explicit relationships between responsibilities and processes in this way allows the SMS developers to verify that there are no gaps or redundancies in either the process model or the responsibilities. or example, it might expose the fact that there is no process (or measurable outputs) defined in the SMS for the Air Traffic Controller's responsibilities to "have a proactive attitude to Safety Management".

Conclusion

This paper promotes the use of structured design

Figure 6: DFD Relationship to Responsibilities



concepts taken from software engineering for aiding the development of an SMS. Specifically, the use of DFDs ensures that the SMS comprises a properly connected set of processes, and the use of standard syntax for the formulation of safety responsibilities ensures proper connectivity between the safety actors and their required actions. This is illustrated using practical examples from ATM. The approach has the potential to yield the following benefits at various points in the lifecycle of an SMS within any industrial sector:

- Initial design of the SMS is represented in a visual, detailed way that allows interested parties to understand and comment on its structure.
- All levels of design can be visualized and their consistency verified because the complete hierarchy of 'process levels' from regulatory requirements down to the internal structure of individual SMS processes can be represented.
- Development of the SMS can be more easily managed. By representing and managing process interfaces, new safety processes and existing quality management processes can be correctly integrated.
- Validation that a developed SMS is complete and coherent can be performed retrospectively by capturing the design from the suite of SMS documentation, thus providing a means of independently assessing the SMS structure.
- Maintainability is enhanced because the implications on the rest of the SMS of modifying a given procedure can clearly be seen.

Notes

1. In effect, this requires the responsible person to develop their own Safety Principles and procedures.

2. The <> symbols are used here to delimit the three syntactic components and would not be documented in the SMM.

References

[1] Hatley, D. J. and Pirbai, I. A., *Strategies for Real-Time System Specification.* Dorset House Publishing, New York, 1987.

[2] Smith, I. C., 'The Maintenance of Computer Based Safety Systems' *IFAC Synposium on the Safety of Computer Control Systems (SAFECOMP'88): Safety Related Computers in an Expanding Market, 9-11 November 1988,* pp. 27-33.

[3] Riaz, Z., Edwards, D. J., Holt, G. D. and Thorpe, T., 'Data Flow Analysis of Plant and Equipment Health and Safety Management', *Journal of Engineering, Design and Technology* vol 9 no 2, July 2011, pp. 178-203.

[4] EUROCONTROL, ESARR 2, *Reporting and Assessment of Safety Occurrences in ATM,* Edition 3.0, 2 December 2009.

[5] EUROCONTROL, ESARR 3, *Use of Safety Management Systems by ATM Service Providers*, Edition 1.0, 17 July 2000.

Dr Steve Thomas is an Independent Consultant at Entity Systems Ltd. He can be contacted by email at steve@entitysystems.eu

Assessing Safety Culture - Where do I start?

by Vicky Gliddon

Safety culture is an important topic in safety-critical industries, where it has received a great deal of coverage. However, for many it can seem daunting to know what to do and where to start to understand the safety culture of their own organisation. Here, we provide advice from Greenstreet Berman Ltd's practical experience for an effective approach to conducting safety culture assessments, as well as tips for implementing targeted improvement following the assessment.

What is Safety Culture and Why is it Important?

Safety culture can be described as attitudes and perceptions that influence safety behaviour within an organisation.

A positive or negative safety culture can have a huge impact on the safety performance of any organisation. It underwrites the attitudes towards safety behaviour and influences the workforce's opinion on the importance of safety, and thus, how they act with regards to safety. Poor safety culture has been identified as a contributory factor in a number of high profile disasters including Piper Alpha, the Nimrod aircraft crash and the Ladbroke Grove rail crash.

The Benefits of Safety Culture Assessment

There are many benefits to assessing safety culture at whatever stage of safety culture maturity your organisation is at, such as:

• Identifying improvements that can make your organisation as safe as it aspires to be;
• Mitigating the risk of harm to employees through safety management;
• Remaining alert to potential emergent issues;
• Learning lessons from your organisation, both good and bad;
• Gaining insight into how effective your safety management systems are, and how effective these are perceived to be;
• Periodic reviewing the effectiveness of improvements.

Plan Plan Plan ...

It is important to make a plan for assessing safety culture before delving in to the assessment itself. This helps to ensure that the outcome of any assessment, does provide answers to the questions you had in the first place.

A well-planned and well-executed assessment can, for example, achieve response rates to questionnaires of 70% or more. A poorly planned and under-supported assessment may get very low response rates (for example, around 15%), and greatly reduce the value of the assessment.

In addition, creating a plan will help you to outline key assessment questions are, by summarising what you hope to achieve from performing an assessment of safety culture.

The undertaking of an assessment can (by itself) affect staff opinions for better or worse, whilst also providing information on which to make improvements to safety. Consequently, an assessment of safety culture should only be contemplated when senior management effectively support it and follow up on the results.

So, what should you be considering at this stage of planning?

First of all, define the context and purpose. Clarify why you are motivated to assess your organisations's safety culture. This will help you to scope out what you intend to achieve through conducting an assessment of safety culture, and will help you to formulate a strategy and framework to do this effectively. The following are points to get you thinking about the context of your assessment:

• Has something happened that has identified issues in organisational safety culture which you wish to address? E.g. Has a procedure been ignored which played a part in an incident?
• Does a particular department have a poor safety performance record?
• Are you interested in conducting an assessment on the entire organisation or only selected departments or functions?
• Do you want to follow a staged approach, conducting assessments on various business areas one at a time or, conduct a company-wide assessment all at once?
• Do you wish simply to measure what the level of safety culture maturity is

based on peoples' attitudes and behaviours, or to go further to understand why this is by exploring the drivers and reasons behind this?

- Do you want objective evidence of behaviours in addition to subjective measures of attitudes?

Consider these types of questions to set out a clear scope for an assessment. The scope will then guide you to select appropriate methods of assessment.

What Safety Domains are you Interested in?

Past accidents have shown that safety attitudes and behaviours can vary from one safety domain to another. The Chemical Safety Board's (CSB) incident report (the *Baker Report* [1]) stated: "BP has emphasized personal safety but not process safety" and, "…a very low personal injury rate at Texas City gave BP a misleading indicator of process safety performance".

Attitudes and behaviours may be divergent across different areas of performance. In our experience, differences have also been found between occupational health and occupational safety, such as high levels of fatigue being 'tolerated' despite robust occupational safety management. This does not mean that all organisations have domain-specific safety attitudes. An organisation may adopt a holistic view of safety, with equal expectations for occupational and major hazard safety.

From a practical perspective, preliminary work such as a small number of workshops, can verify how the organisation conceptualises safety. Is a singular concept of safety pursued, or are there multiple safety domains? If there is clear evidence that a singular concept is espoused, then a single measure of safety culture may be valid. If there are multiple safety domains, then you may wish to either assess both in parallel or focus on one.

Either way, a definition of the area of safety being assessed should be agreed and stated at the outset of any assessment tool, ideally with examples of safety policy and practices. This ensures all staff are thinking of the same safety area(s) when responding to questions.

Allocate Resource

Once you have a plan in place to define the scope and depth of assessment, you can allocate the appropriate resources.

Do you have the necessary capability in-house to complete an assessment from the planning stage through to targeted improvement? If not, you may want to consider outsourcing all or part of the assessment to an independent organisation with the relevant experience and qualifications to support you in assessing safety culture.

In the case of assessment methods such as, workshops, observations and interviews it can be particularly important to have independent facilitators. Do you think your employees are likely to speak honesty about safety if they are being asked questions by the health and safety manager? Whilst this may not deter some from speaking openly and honestly, you are unlikely to get a true view of your employee's perceptions and your data will not be reliable.

An independent organisation with safety culture expertise can also review your plan and provide advice on your strategy for assessment at the early stages. It is also important to utilise the resource you have in-house. Outsourcing to an independent organisation, doesn't mean you won't need to use your own resources. If you have an internal team who focus on communications or employee engagement, get them on board to help

plan and communicate the assessment.

Communication is Key

Poor communication of an assessment prior to roll out, can be a huge drawback. The engagement of employees and responses for participation are affected by how well you communicate the purpose of the assessment, what it involves, the role of employees, and the aims and objectives of completing an assessment of safety culture.

This communication should be targeted to all levels of the organisation. An initial communication activity should be conducted to inform those at a more senior level of the impending assessment, giving them a chance to ask questions and raise any concerns. Involving them in the early stages can gain their approval, giving the assessment their endorsement and cascading this through the organisation. This senior management support will help minimise resistance and gain buy-in to the assessment across other levels of the organisation.

Select Methods for Assessment

A host of methods can be applied to assess safety culture, and which of these you select will depend on the intent of your assessment, and the nature and demographics of your organisation. Each method has different advantages and limitations. For example:

- Staff questionnaires are a cost effective, repeatable option for measuring and comparing attitudes across departments in large organisations, over time;
- Workshops help to explore staff perceptions and attitudes, elicit examples of behaviours and the drivers behind these behaviours;
- Observation of behaviours and communication, e.g. during meetings, can provide objective evidence of behaviour but do not reveal attitudes or reasons for behaviours and can be intrusive (and limited to observable behaviour);
- Reviews of safety documentation, such as policy and procedures, can indicate management intent, but do not prove implementation;
- Reviews of leading and lagging safety culture indicators, such as rates of procedural non-compliance, can be objective but do not reveal reasons for behaviour and may not measure all aspects of culture;
- Critical incident and root-cause analysis can assess actual behaviours; however, they can be skewed to poor performance and incidents may not have occurred in recent years. This can, to some extent be balanced by seeking examples of positive behaviour.
- Focus groups can explore the reasons for observed and perceived attitudes and behaviours.

Greenstreet Berman Ltd advocates the use of a mixed-methods approach to the measurement, diagnosis and validation of assessment results. Methods such as surveys allow you to measure employee perceptions and attitudes, but do not allow you to fully understand the reasons for those opinions. Diagnostic methods such as workshops allow you to probe into and clarify feedback from other methods, and will give you a richer assessment of safety culture within your organisation. We also recommend having a common set of safety culture elements such as leadership and communications, assessed by each method, to allow results to be collated and compared.

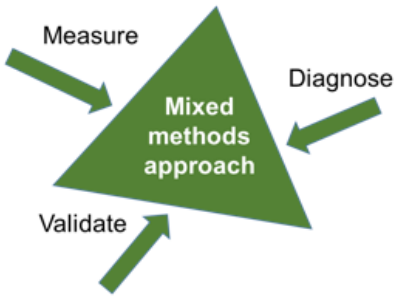


Figure 1: Mixed Methods Approach to Assessment

The validation of results received from one assessment method should be validated by other methods of the mixed-methods approach, allowing you to diagnose as well as measure safety culture. For instance, a workshop can gather people's opinion on safety and how it is managed whereas, observations can validate whether this perception is accurate.

The methods should be selected to be accessible to all participants you wish to be included. For example, if you are doing an assessment of those on the frontline in operations who do not readily have access to a computer, it would be imprudent to issue an electronic survey. More relevant and useful would be a paper survey which can be returned in a sealed box to ensure confidentiality and anonymity of responses, and therefore more likely to yield a higher response rate.

Prepare Methods

Whilst the principles of assessment remain the same, there is no "one size fits all" for assessing safety culture. All assessments should be tailored to the organisation. For instance, a safety culture survey should have content which is relevant and meaningful to employees, such as ensuring that the types of questions it asks are appropriate, and the language and terminology used is understood.

Care should be given to tailor the assessment materials to match the scope of assessment. If you intend to roll out the assessment across the entire organisation but then compare departments to identify similarities and differences, you need to know what department each respondent is within, otherwise this comparison will not be possible.

Tailoring demographic questions to your organisation such that the data collected can be segmented for analysis purposes, is critical at the early stages. Work with your Human Resource team to define the demographics present across the organisation and which will be included and captured within the assessment. Ensure that assessment methods are planned in advance to allow the time to for logistics to be arranged. This includes, ensuring there is sufficient printing time for surveys, to ensuring participants are available and locations booked in time to run workshops. We tend to advocate a few weeks for planning, a few weeks for communications, about six weeks for conducting the assessment and one month for analysis and reporting.

Apply Methods for Assessment and Diagnosis

In addition to a well-planned and properly communicated assessment, it is important to track response rates during the course of the assessment. There are many practical reasons for participation levels varying, such as unforeseen operational demands or unforeseen problems in responding. Accordingly, response rates should be tracked by department and grade during the course of the assessment. In the event of under-representation of any one grade or department, additional actions may be needed, such as extending the time period for questionnaire responses, re-issuing requests to respond to questionnaires, or running additional workshops.
In addition, it is important to check the representativeness and validity of responses. Have all departments achieved a similar level of response? Do examples of behaviours cited in workshops match the assessed level of safety culture maturity? You need to understand what has been assessed, how representative the results are and the reasons for any skews in results. With this understanding, it is possible to weight the results from each form of assessment according to your confidence in each method.

Post-Assessment Analysis

Once the assessment is complete it can be like walking through a minefield to understand what the results mean, and to know where to start with a deluge of data, information and findings.

Go back to the planning of your assessment and use this to guide you through your analysis. Was the purpose of your safety culture assessment to assess the entire organisation? If so, it would be sensible to compare the findings from various departments and business areas to identify any similarities and differences. Did the plan outline one safety domain or more? If it included more than one type of safety (occupational safety and process safety) assess whether there was a difference in findings. Was an issue with safety culture identified prior to the assessment, and was this backed up by the assessment findings?

Use this analysis to identify trends in safety culture and evaluate the findings to show both positive findings and areas for improvement. Focus on key themes and major differences in results between departments and grades.

Assessing the Results

Greenstreet Berman Ltd uses a pre-defined evaluation scale for evaluating the results of the assessment. The use of a scale is good practice and also advocated by other safety culture tools. Reviews of safety culture models and feedback from industry has led us to adopt a variation of the Pathological to Generative safety culture maturity scale, as per Figure 1. This is a widely used and intuitive scale underpinned by research.

The use of the evaluation scale is beneficial to identify where the organisation sits at in terms of the level of safety culture maturity, and to define where the organisation should practically aim to be. The results are typically collated for the organisation as a whole and split by the demographics (e.g. department) as outlined at the scoping stage to make comparisons.

Feedback Results

It is imperative that the findings are communicated back to the organisation as was identified at the communication stage. This will encourage continued engagement in the assessment and for the improvement of safety culture. The level of information shared with employees at various levels of the organisation is at the discretion of the organisation, and any promises made about the dissemination of results should not be broken.

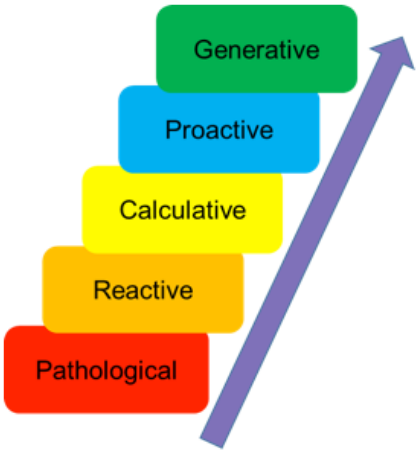


Figure 2: Generative to pathological culture maturity scale

Create a Strategy for Improvement

Once the organisation has determined its level of safety culture maturity, it should articulate what they wish to achieve. Whilst there is no definitive guidance on the required levels of safety culture, our experience suggests that safety critical organisations seek a calculate (managing) culture as a minimum, and more commonly aspire to a proactive culture.

Upon articulating your objective, you can begin to create a strategy for improvement to define how you will reach this level.
The results will guide you to decide where to focus your efforts. For instance, if a particular department was rated to have a lower level of safety culture maturity than others, you may choose to focus on making improvements within that department first of all.

Elements of safety culture, such as leadership versus communications, could be acted on in order of priority by the level of safety culture maturity they were found to be at. Qualitative understanding of the drivers for behaviour can help identify specific improvement actions.

It is also important to identify the stronger areas identified through assessment as this provides lessons learned that can be applied to other departments or other elements. Creating a strategy for improvement will provide direction to address problems and build on the strengths of safety culture in your organisation.

Reassessment

Safety culture should be periodically reviewed to check performance and the effectiveness of any actions coming out of the safety culture assessment. Some safety critical organisations assess safety culture every year, comparing trends from one year to the next. In our experience, the extent and frequency of assessment depends on the level of safety culture maturity and the need for improvement. A more mature and possibly stable, culture may not require annual review or in-depth assessment. A less mature culture, or one with a major improvement programme, might benefit from more frequent and more in-depth reassessment.

References

[1] J. Baker et al, *The BP U.S. Refineries Independent Safety Review Panel*, January 2007. http://www.csb.gov/assets/1/19/Baker_panel_report1.pdf. Accessed December 2016.

Vicky Gliddon is a Senior Consultant at Greenstreet Berman Ltd. Vicky is qualified in psychology and ergonomics. In addition to her work conducting human factors and ergonomics assessments for safety-critical organisations, she has extensive experience of safety culture assessments and supporting targeted improvement within high-hazard sectors.



Let Me Make Myself Unclear

by John Ridgway

One of the very first books I read as a child (apart from the 'Janet and John' books) was a volume proudly retrieved from the local library. It was called *How Big is Big?* and it sought to demonstrate mankind's scale within the Universe by illustrating a series of objects of ever-increasing size (firstly terrestrial, then extra-terrestrial), finally arriving at the Universe in its totality. It then proceeded to illustrate objects of ever-decreasing size, culminating in the sub-atomic realm. It remains to this day the most influential book I have ever read and, no doubt, set me on a lifetime of scientific curiosity. However, despite my advancing years, I have to admit that I still can't answer the basic question posed by that influential tome. I still can't tell you how big 'big' is.

The problem lies not in my education, but in the inherent vagueness of the term 'big'. It's one of a very large class of words known as degree adjectives. All natural languages are full of them. In fact, most languages would be a pale shadow of themselves if one were to remove all the inherently

vague words from their lexicon. Furthermore, the majority of early words learnt by a child will fall into the category of 'vague' ('Mammy' and 'Daddy' are notable exceptions). Despite all of this, however, the word 'vague' is seen very much as a pejorative term, particularly in the hard-nosed world of high-integrity systems development. So how can it be that words that have such an undeniable ubiquity and (presumably) linguistic merit should be treated as swear words when it comes to developing safety cases?

In this article, I propose to have some fun by challenging the obvious notion that all vagueness is bad. I would also question the assumption that the alternative should be any better. But firstly, to avoid the embarrassment of writing a vague article on the subject of vagueness, I need to define my terms.

What Precisely is Vagueness?

According to the dictionary, 'vague' is an adjective used to describe words and language that are inexact or imprecise in meaning. However, here I wish to use the term in a more narrow and technical sense. There are a number of ways in which one may be imprecise but the one that interests philosophers of vagueness (yes they do exist) is the imprecision that results from using expressions that have borderline cases, i.e. marginal cases for which the expression's applicability is unclear (see note 1). These problems will be inherent within the expression being used. One may try to resolve the difficulties by introducing a defined boundary but this will have two implications. Firstly, the expression is no longer being used in the general sense required by its definition. Secondly, it is still necessary to deal with paradoxes that arise at the threshold you have defined.

For example, the term 'warm' is vague because it is not clear just where 'cool' ends and 'warm' begins. In an attempt to introduce clarity, one might agree that 18°C (or above) constitutes 'warm' but the expression 'warm day' would now be potentially ambiguous: if others choose to define the threshold differently, one would have to be sure which definition was being used. Furthermore, the applicability of the term 'warm' will always be context-specific – what might be a cool summer's day could be considered warm for winter. Besides which, arbitrary thresholds are rarely realistic. If the actual temperature were fluctuating slightly about the threshold, one would be unlikely to refer to the weather as vacillating rapidly between warm and cool.

Worse still, one would now have an expression that is susceptible to the sorites paradox. This paradox gets its name from the word 'soros', which is the Greek for 'heap' or 'pile', and it originally alluded to the difficulty one has in determining exactly when a pile of sand ceases to be a pile as one removes its grains. The problem stems from the vagueness of the word 'pile' and the lack of a convincing cut-off.

To fully appreciate the significance of the paradox, consider the following example: If 10°C is defined as 'cool' then a temperature in excess of 10°C must also be considered cool if it is practically indistinguishable from it. By the same logic, any temperature that exceeds this new temperature, and is practically indistinguishable from it, will also qualify as being cool. And so on – one can repeat the process any number of times to any arbitrarily high temperature, and yet still require that the temperature be defined as 'cool'. This can't be right. It's like the frog which, placed in a pan of cold water, boils to death because it cannot perceive a gradual rise in temperature (see note 2).

It is because of all of the

above that linguists worry about adjectival vagueness. However, I should point out that adjectives do not hold the monopoly, since vagueness can also be encountered when dealing with nouns. For example, when I'm driving on the motorway, I am never quite sure if I have just been dangerously cut up by a lorry or a van. Where does 'van' end and 'lorry' start? Here, the question is the extent to which an object exemplifies a particular quality, such as 'lorryness'. In the real world, things are often not clear-cut and one has to accept degrees of truth.

The Acceptable Face of Vagueness

A number of philosophical theories have been developed in an attempt to address such problems but we needn't go into them here (see note 3). I suspect you won't find them of too much help when writing your safety case or deciding upon your safety policies. I do, however, think that, despite these problems, it would be wise to reflect upon the benefits of vague expressions before you assign them to Room 101. We need to ask whether there are times when you would be better off remaining vague, and whether the cure can be worse than the disease. To address these matters I will need to subject-hop for a while, so please bear with me.

Promoting Reusability

One of the problems with vague expressions is that their applicability changes depending upon context. But actually, that is the main reason why they are popular – they have semantic utility. And, to the extent that re-usability is a positive feature for safety case elements, one may be looking for expressions that have such utility.

Take, for example, the term 'reliable'. This is a vague term because it allows for borderline cases and its precise meaning will change with circumstance. However, this allows it to be used in a variety of contexts whilst still retaining its transcendent meaning, i.e. 'dependable enough for purpose'. An adaptable safety case element can therefore be constructed using this term, as long as one is prepared to validate the anticipated failure rate within the context of each of the safety case element's instantiations. For example, one can construct a generic argument for the safety of a utility component by providing evidence for its predicted failure rate and then claiming that the component is therefore 'reliable'. This argument necessarily remains vague until the case element is placed within the context of a particular system's safety case, i.e. once the application-specific requirements and design context have been taken into account.

Reducing Pointless Detail

Then there is the *maxim of quantity*, which requires that statements should be as informative as possible whilst being no more informative than is necessary. By this maxim, vagueness is allowed as long as it isn't gratuitous, and tackling vagueness by being more precise than is required (or justified) is the worst thing you can do. For example, we engineers love numbers and, as soon as we see them, we breathe a sigh of relief. No more naughty vagueness, we declare, whilst quietly overlooking the need for error bars. But we give numbers far too much credit, often quoting decimal places just because our calculator produced them.

And whilst we are talking about numbers, remember that quantified thresholds can still boil your frog. If 1000 Killed or Seriously Injured (KSI) is tolerable, then why not 1001? And if 1001, why not 1002? You may think you can tame the vagueness with numbers and limits, but this is all part of the illusion.

Handling Subjectivity

This brings me onto the

unquantifiable. Don't go pretending that it doesn't exist. It is all around us and its assessment is always problematic, rendering the use of vague expressions all but inevitable. Thankfully, that's what we have experts for. Most vagueness in safety cases is hiding behind the over-confident proclamations of a revered expert.

Vague expressions are also popular because they can convey an opinion. The terms 'acceptable' and 'safe' come to mind. We can use specific targets, such as '100 KSI', but the targets mean nothing until they are equated with the expression 'acceptably safe'. The problem here is not so much the vagueness; it is the fact that opinions can differ.

Avoiding the Arbitrary

Another major benefit of vagueness is that it avoids the tyranny of the arbitrary cut-off. Take, for example, the definition for KSI. In order to have been killed, one has to have died within 30 days of the accident. So, presumably, if you die just one minute later then you were not killed. The conditions required to qualify as 'seriously injured', on the other hand, are not as clear-cut; which is perhaps a good thing.

Legislating

Vagueness will often be deliberately written into legislation, thereby requiring someone to interpret the application of the law in test cases. Here, again, the vagueness is being used to bestow utility. The classic example within safety legislation is the concept of reasonably practicable. It seems we expect our judges to know how big 'big' is, and we are expected to be able to second-guess them. In fact, the ALARP principle doesn't just tolerate vagueness. It virtually worships it.

Hedging Your Bets

Last, but not least, no one wants to be caught out making bold predictions that Nature then disrespects, so why not make a vague claim that will be easier to defend with hindsight? After all, you are just offering a lower Safety Assurance Level (SAL), aren't you? But what if there is so much riding upon the acceptance of the safety case that you just slip in the odd vague term in the hope that no one notices? Or was that just me?

Utility Versus Precision

There are some who view vagueness in language as a linguistic anachronism - a legacy of our primitive past, when sophistication and precision provided little evolutionary advantage. I disagree. Our love of vague terms is not a failure to divest ourselves of a bad habit. On the contrary, the vagueness lying at the heart of our language provides utility that enables us to operate effectively and efficiently in an uncertain world.

As systems safety engineers, we dislike vagueness because it reflects uncertainty. Nevertheless, as long as we continue using degree adjectives such as 'safe', 'reliable' and 'predictable' in our safety arguments, we will have to accept that linguistic vagueness is playing a central role. And rather than suffering unnecessary embarrassment, we should recognise the importance of this vagueness and work with it. Safety cases are written for an audience that was brought up on natural language and the one thing that linguists can say for sure is that natural language is riddled with vagueness. Yet, despite this, we are not frogs who boil in a bucket. We are perfectly capable of constructing a reasoned argument and acting upon it, enabled rather than challenged by the subtleties and mysteries of linguistic vagueness.

Crisply defined boundaries and quantified limits are all very well, but they can be false props that bring with them their own problems. So I invite you to consider the possibility that a safety case that is self-knowingly

vague in all the necessary places can be a lot better than one that deceives and bemuses the reader with unjustified precision and certitude. If you really don't know how big 'big' is, you are doing no one any favours by pretending that you do.

Notes

1. As such, a vague expression is not to be confused with an ambiguous one. An ambiguous expression is one that has two or more interpretations, and for which a context is required to achieve disambiguation. Such ambiguity can be particularly problematic when it is still not obvious from the context which of the equally valid definitions is meant to apply. It is also important to understand that a loosely qualifying expression is not necessarily vague. For example, 'less than ten', is not vague, since its applicability is straightforward. However, 'nearly ten' is vague, because the applicability of the term 'nearly' is not clear-cut.

2. This is actually an urban myth. Cruel scientists have determined that real frogs aren't that inattentive.

3. These include ideas such as Supervaluationism, Subvaluationism, Contextualism, Epistemicism, Fuzzy Plurivaluationism and Many-valued Logic. If long and horribly contrived words are your thing, see *Further Reading*.

Further Reading

K. van Deemter, *Not Exactly - In Praise of Vagueness*, Oxford University Press, 2012.

R. Keefe, *Theories of Vagueness*, Cambridge University Press, 2006.

'Vagueness', *Stanford Encyclopaedia of Philosophy*, http://Plato.stanford.edu

R. van Rooij, 'Vagueness and Linguistics', *Researchgate.net*

T. W. Grinsell, 'Avoiding Predicate Whiplash: Social Choice Theory and Linguistic Vagueness', *Proceedings of the 22nd Semantics and Linguistics Conference (SALT 22), University of Chicogo 18-20 May 2012*, pp. 424-440.

How Big is Big? by an author unknown, and long out-of-print.

John Ridgway is a retired analyst who spent most of his career working for Serco Transportation Systems. Having failed to change the World, he now lives in quiet retirement with his wife and two dogs. He can be contacted at j.ridgway4@ntlworld.com.

Conforming to IEC 61511: Operation and Maintenance Requirements

by Steve Gandy

It is hard to believe that the IEC 61511 standard has been in existence since 2003 and most companies operating in the process, chemical and refining industries (or any other process manufacturing) have adopted its practices. It is also significant that any plants with a Safety Instrumented System (SIS) will now be halfway through their useful life. It therefore seems opportune to ask how well companies have been recording the performance of their SISs, in terms of failures, spurious trips, time to repair/restore and proof testing results. Furthermore, the new 2016 edition of IEC 61511 emphasises the need for assessment of SISs more strongly, soeaking in terms of preventing systematic issues through procedures and competency. This paper highlights how testing and documenting the performance of the SIS is an essential part of ensuring that it is able to fulfill designated functional safety requirements. This is especially true as the SIS approaches the end of useful life.

Introduction

Over the past decade or so, automation has been one of the dominant factors

in enabling end users in the process, chemical and petro-chemical industries to be able to streamline their costs and improve efficiency; often at the expense of personnel. Most modern plants today have less manpower than plants of 1980 and even the 1990s. This means that the burden of running and maintaining a modern plant has fallen on fewer and fewer plant personnel. Coupled with the shortage of skilled employees, this places a significant burden on plant personnel to maintain and improve their skill set in order to maintain the technologically more complex instrumentation and automation systems. Aside from the Basic Process Control System (BPCS), there's the Plant Safety System (referred to as the Safety Instrumented System (SIS).

The advent of the IEC 61511 standard [1] for the process industries has provided a path to improve safety by introducing the concept of a Safety Lifecycle (SLC) and moving away from a strictly prescriptive methodology to a more performance based methodology, with the emphasis being placed on reducing risk and mitigating the potential for hazards that could lead to the loss of life, destruction of property and plant assets. The purpose of this paper is not to define the application of the standard but to examine one important aspect of the SLC: the Operations and Maintenance requirements for the plant SIS.

IEC 61511-1 Clause 16: SIS Operations and Maintenance

The term 'SIS' has been used rather than the term 'safety system' as there are many safety systems, not all of which are intended to comply with IEC 61511. Only safety instrumented functions that are part of a Safety Instrumented Function (SIF) are required to comply. Figure 1 has been excerpted from ANSI/ISA 84.91.01-2012 [2] to provide an illustration.

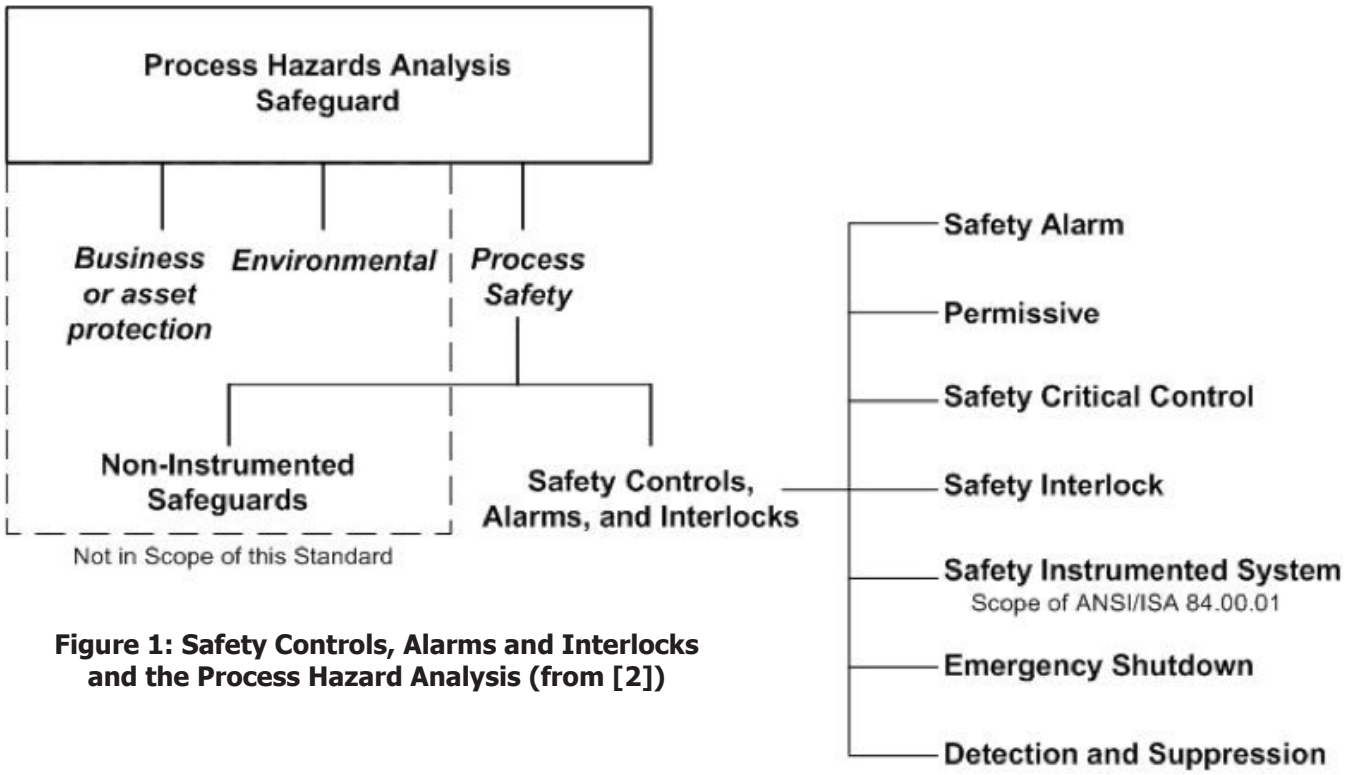


Figure 1: Safety Controls, Alarms and Interlocks and the Process Hazard Analysis (from [2])

The term 'SIS', as defined in IEC 61511-1 Clause 3.2.72, refers to a Safety Instrumented System, i.e. an instrumented system to implement one or more Safety Instrumented Functions (SIFs), which is composed of any combination of sensor(s), logic solver(s) and final element(s), as illustrated

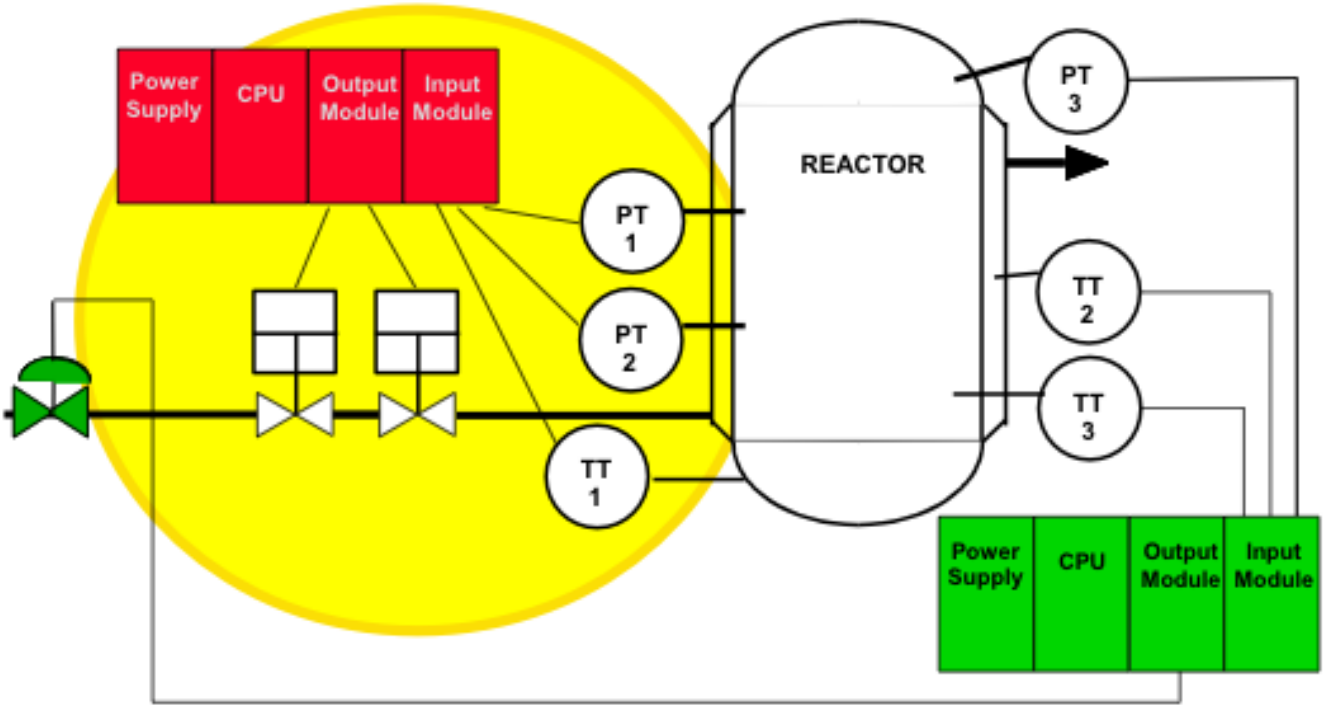


Figure 2: Example of a Safety Instrumented System Block Diagram

in Figure 2. A SIS can include safety instrumented control functions or safety instrumented protection functions, or both. A SIS may or may not include software (i.e. could be solid state or hardwired with electro-mechanical relays).

As mentioned in the introduction, the IEC 61511 standard is based around a Safety Lifecycle (SLC). Figure 3 illustrates a simplified version of the SLC and highlights the Operations and Maintenance Section of the lifecycle.

In order to fulfil the requirements of IEC 61511-1 Clause 16, the end user is required to have a properly and well defined Operation and Maintenance Plan to ensure that the required Safety Integrity Level (SIL) is maintained during operation and maintneance tasks to ensure that the SIS maintains its functional

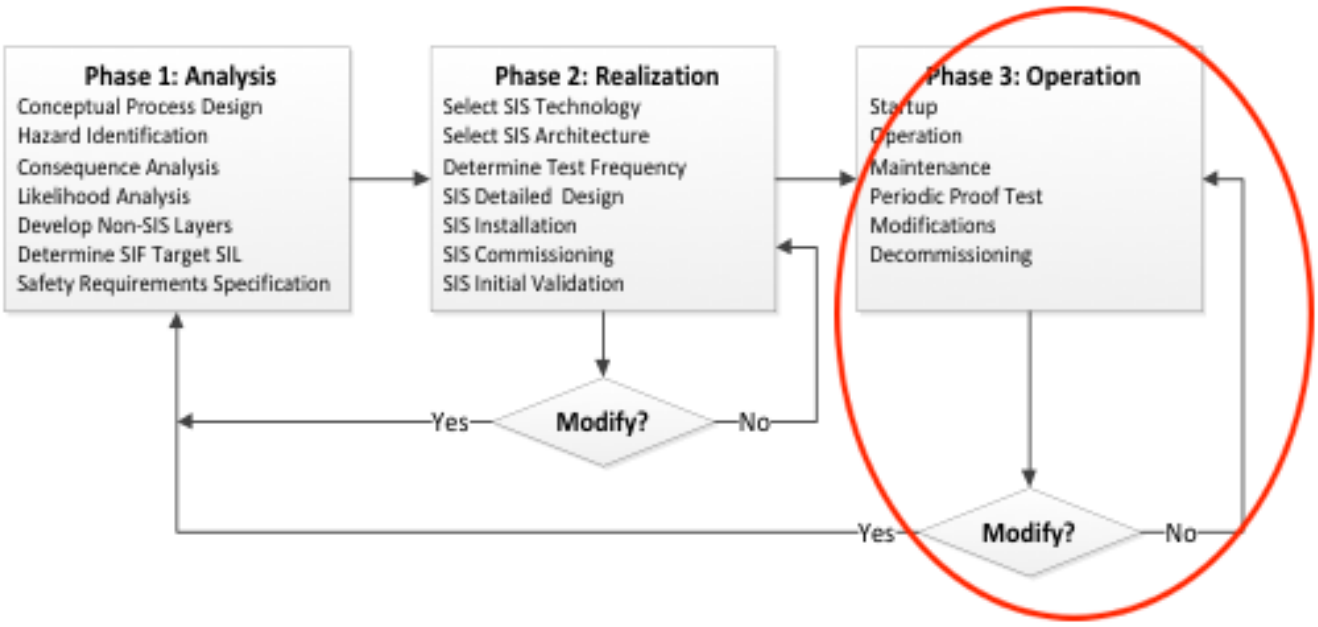


Figure 3: Simplified Functional Safety Lifecycle Diagram (Operations and Maintenance Phase Highlighted)

safety integrity throughout its entire lifetime.

The SIF and associated SIL are determined at the front end of the Lifecycle during the Analysis Phase and are beyond the scope of this paper.

The Importance of Leading and Lagging Indicators

IEC61511 is a "performance-based" standard that requires the owner/operators to undertake "periodic" assessments. This means that recording "lagging" data is essential. Lagging data would include such things as:

- Near misses;
- Trips, real and spurious;
- Faults, both random and systematic;
- Process upsets.

The purpose of lagging data is to be able to help in preventing future problems or events and/or developing training programs, procedure improvements, etc.

The purpose of "leading" indicators is to help predict future events. Examples of leading indicators would be:

- Inspections that were overdue;
- Maintenance that was not carried out in a timely manner (e.g. repair times exceeding the limits defined in the SRS).

By understanding both the leading and lagging indicators, operators can improve overall process safety, thus preventing potential spills, fires and explosions.

Operation and Maintenance Plan

The operation and maintenance plan is a working document that is designed to ensure the SIS is maintained to meet its designed functional safety and will need to cover:

- Routine and abnormal operation activities;
- Proof testing, preventative and breakdown maintenance;
- Procedures, measures and techniques to be used for operation and maintenance;
- Verification of adherence to operations and maintenance procedures;
- When these activities shall take place;
- The persons, departments and organisations responsible for these activities.

Procedures covering each of these aspects need to be developed and maintained, especially following any updates and/or modifications to the SIS. This could be viewed as developing a "safety checklist" that will help eliminate human (systematic) errors from creeping in during maintenance activities.

Operation and Maintenance Procedures

IEC 61511-1 Clause 16.2.2 states that the operation and maintenance procedures shall be developed in accordance with the relevant safety planning and shall provide the following:

- Routine actions which need to be carried out to maintain the "as designed" functional safety of the SIS, for example, adhering to proof test intervals defined by the SIL determination;
- Actions and constraints that are necessary to prevent an unsafe state and/or reduce the consequences of a hazardous event during maintenance or operation (for example, when a system needs bypassing for testing or maintenance, what additional mitigation steps need to be implemented);
- The information which needs to be maintained on system failure and demand rates on the SIS;
- The information which needs to be maintained showing results of audits and tests on the SIS;
- The maintenance procedures to be followed when faults or failures occur in the SIS, including: Procedures for fault diagnostics and repair, procedures for revalidation, maintenance reporting requirements

and procedures for tracking performance;
• Assurance that test equipment used during normal maintenance activities is properly calibrated and maintained.

These requirements place a heavy burden on the Operations and Maintenance (O&M) personnel, who need to have the requisite skill set in order to be able to maintain the SIS. IEC 61511-1 Clause 16.2.4 lists the skills criteria that have to be addressed:

• Understanding how the SIS functions (in terms of trip points and the resulting action taken by the SIS);
• Understanding the hazard the SIS is protecting against;
• Operation of all bypass switches, and under what circumstances these bypasses are to be used;
• The operation of any manual shutdown switches and manual start-up activity and when these manual switches are to be activated (may include system reset and system restart);
• Expectation on activation of any diagnostic alarms (for example, what action shall be taken when any SIS alarm is activated indicating there is a problem with the SIS).

The standard states that maintenance personnel should be trained as required to sustain full functional performance of the SIS (hardware and software) to its targeted integrity.

In addition, the O&M personnel will be required to follow a written proof test procedure as defined in IEC 61511-1 Clause 16.2.8, whereby a proof test procedure has to be developed for every SIF to reveal dangerous failures that are not detected by the SIS diagnostics. These written test procedures will need to describe the following steps:

• The correct operation of each sensor and final element;
• Correct logic action;
• Correct alarm and indications.

The proof tests will need to cover the entire SIS including the sensor(s), the logic solver and the final element(s) (e.g. shutdown valves and motors). In addition, the proof tests will need to be carried out at intervals that were specified and used to calculate the PFDavg for the SIF.

This implies that O&M personnel will be required to undergo regular training, competency audits and competency assessments, especially when new and/or updated components of the SIS are being incorporated and/or old or worn out components are being replaced. Personnel training is a key element in ensuring that the SIS can be maintained and operated correctly.

What Happens in Practice?

In order to be able to maintain and follow the requirements set forth in IEC 61511-1 Clause 16, the end users have to ensure that they have adequate procedures, as well as an adequate documentation and tracking system. Recording spurious trips, Process demands, failure data, audit results, test results, etc., requires a well-organized and maintained documentation system. It also requires the O&M personnel to be diligent in recording this information.

Of course it remains to be seen how diligent the personnel are at recording this data, since it is highly dependent upon the safety culture of the plant. Sadly, the recent Tesoro incident in 2010, which resulted in the deaths of 7 workers, as reported in the draft US Chemical Safety Board (CSB) findings [3], points to:

"..a deficient refinery safety culture, weak industry standards for safeguarding equipment, and a regulatory system that too often emphasized activities rather than outcomes. The conclusion of which suggests the need for refinery safety reforms."

The *Tesoro Report* by the CSB also states that in 2012 alone, the CSB tracked 125 significant incidents at U.S petroleum refineries. The draft report examines the effectiveness of refinery and chemical facility regulatory oversight, noting that Washington State's Department of Labor and Industries (L&I) does not have sufficient personnel resources to verify that process safety management requirements are being implemented adequately.

The inference is that end users need to be more vigilant in how they are maintaining and operating their plants. If end users follow the requirements of IEC 61511-1 Clause 16.3.1 and 16.3.2, regarding proof testing and inspection of the SIS (which would include visual inspection of piping) to ensure no observable deterioration and/or any unauthorized modifications have occurred, then incidents such as happened at the Tesoro plant, might have been preventable.

Although IEC 61511-1 Clause 16.2.5 dictates that maintenance personnel need to be trained, it doesn't define how frequently this training should be carried out and how competency is measured.

How Data is Recorded

The problem faced by many O&M personnel is how to record and archive the data. Most BPCS systems will have an historian for archiving plant data, which includes trips, alarms, diagnostic faults, etc. Normally, this type of data associated with the SIF is also recorded by the same and/or a separate historian. Proof testing and inspection are critical tasks that have to be performed as per IEC 61511-1 Clause 16.3. The purpose of proof testing is to reveal undetected faults and the proof tests shall be conducted in accordance with a written procedure. The sole purpose of this is to detect defects and/or faulty equipment prior to a demand being placed on the SIS. Proof test coverage is another important aspect as it's nearly impossible to achieve 100% proof test coverage, therefore, the frequency and thoroughness of manual proof testing is essential to maintaining the SIS.

IEC 61511-1 Clause 16.3.3 defines what needs to be maintained for record purposes. The clause defines that the user shall maintain records that certify that proof tests and inspections were completed as required. These records shall include the following information as a minimum:

- Description of the tests and inspections performed;
- Dates of the tests and inspections;
- Name of the person(s) who performed the tests and inspections;
- Serial number or other unique identifier of the system tested (for example, loop number, tag number, equipment number, and SIF number);
- Results of the tests and inspection (for example, "as-found" and "as-left" conditions).

The standard does not specify how and by what means these results are recorded and most end users will have a plant enterprise system or asset management system, which may require O&M personnel to upload this information. In most cases the O&M personnel will be using paper-based systems for recording, in which case it will require additional effort to scan these documents into a format that can be transferred to the plant's main system.

This would clearly place a further burden upon the O&M personnel, which could lead to short cuts being made and/or missing data because the O&M personnel didn't have time to do this. Statistically, most plant incidents occur during start-up and/or plant shut-downs, when the possibility of spurious trips, alarms and/or faults would be the highest, especially if a start-up was being implemented as a result of maintenance work and/or plant modifications. If there is a spurious trip then the O&M personnel will be

under pressure to get the plant and/or process line back up and running as quickly as possible. Does this mean they'll have the time to properly record all the data required (as listed above)? This is a very valid and pertinent question.

Technology Can Help

Advances in technology can now provide the means for O&M personnel to record data via handheld tablets in electronic format. However, having a dedicated tool that has been specifically designed for this purpose is the issue. Most O&M personnel will be recording their data in an excel spreadsheet or some form of database, if not using a paper-based system. There are some tools on the market that address part of the requirements but having one that addresses all the requirements is rare.

The O&M personnel would need a tool that can record functional safety related statistics/performance metrics, as well as record life events such as:

- Demands – both real and spurious.
- Inspection and Proof tests results.
- Maintenance activities (e.g. calibration).
- Failure reporting

Recording demands, such that the user is able to determine which layer of protection failed is another important aspect because the information provided would enable the user to determine the demand frequency of the hazard scenario (i.e. a specific hazard with all applicable protection layers - in sequence of operation). This would enable any discrepancies between the expected behavior and actual behavior of the SIS to be analyzed and any necessary modifications to be implemented to restore the SIS to its designed safety level. This is required per IEC 61511-1 Clause 16.2.6.

Having a tool that enables the O&M personnel to be able to record demands, such that they could identify which protection layer was successful protecting against a demand for a given hazard would be highly desirable. Figures 4, 5 & 6 below illustrate example templates for

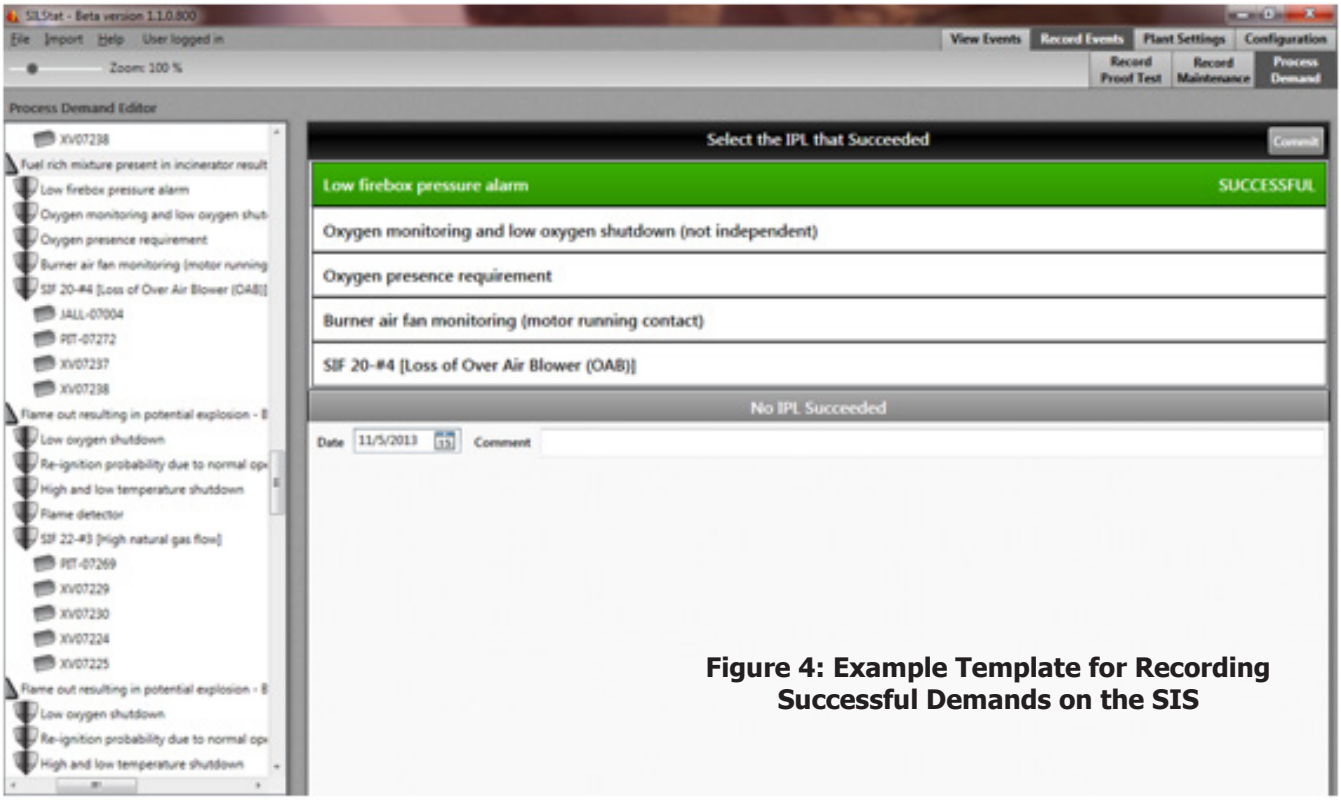


Figure 4: Example Template for Recording Successful Demands on the SIS

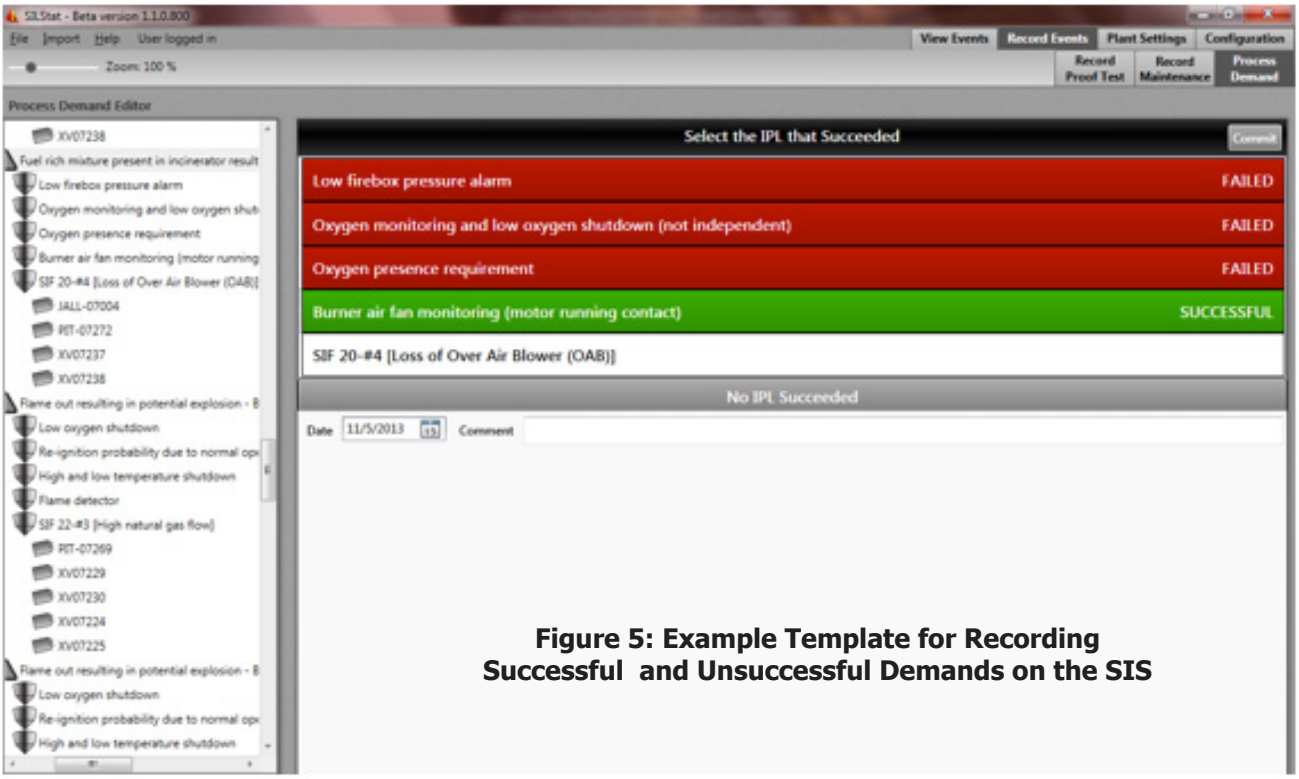


Figure 5: Example Template for Recording Successful and Unsuccessful Demands on the SIS

recording such events. The information could also be used to determine the demand frequency of the hazardous event. Having a tool that enables the physical devices of the SIS to be stored in a database and identified by their associated tags and/or descriptions will enable O&M personnel to be able to carry out effective maintenance and/or replacement procedures. Figure 7 illustrates an example template for recording and entering device information.

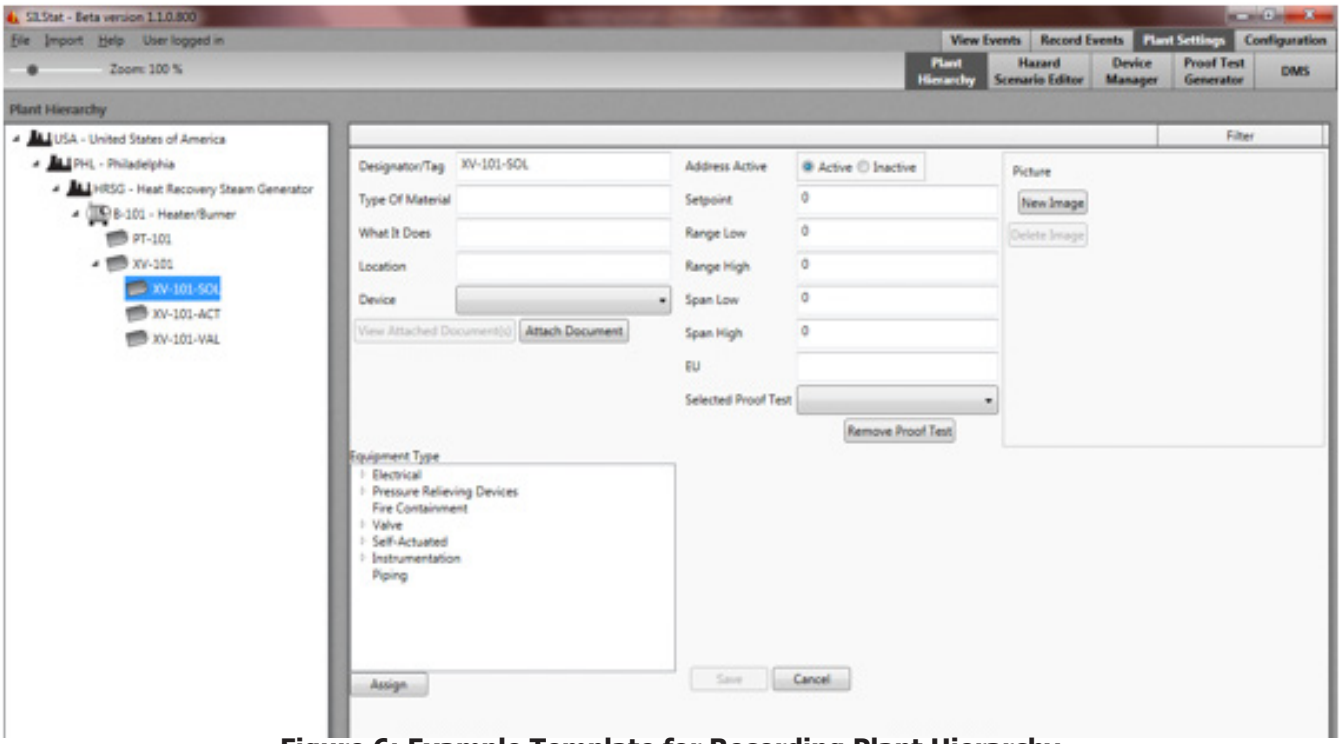


Figure 6: Example Template for Recording Plant Hierarchy

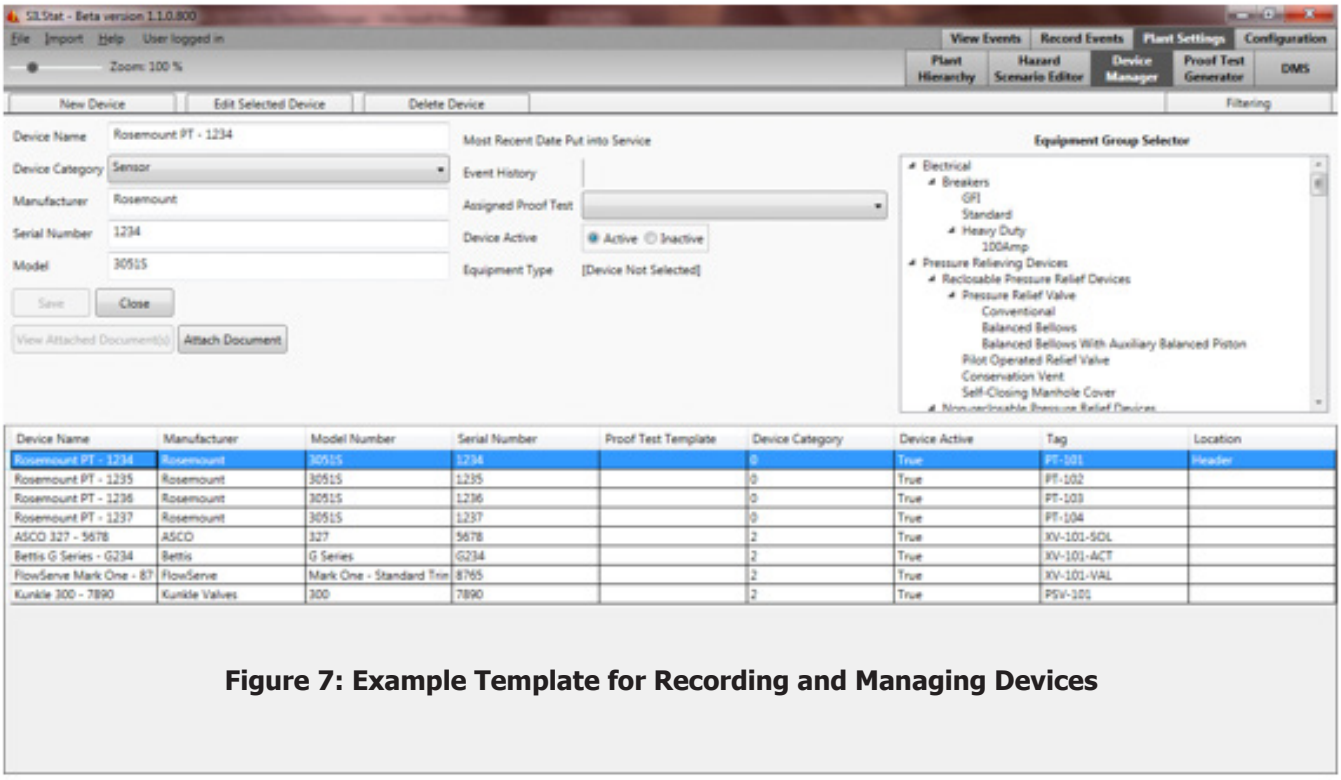


Figure 7: Example Template for Recording and Managing Devices

As mentioned earlier, proof testing is a very important step that has to be carried out in accordance with the PFDavg calculation that is included for each SIF within the Safety Requirements Specification (SRS), although different parts of the SIS may require different test intervals (e.g. the logic solver may require a different test interval than the sensors and/or final elements). Enabling the O&M personnel to have an automatic proof test generator that allows them to specify individual proof test steps, with pass/fail criteria, would be a significant benefit. This would allow the O&M personnel to record only

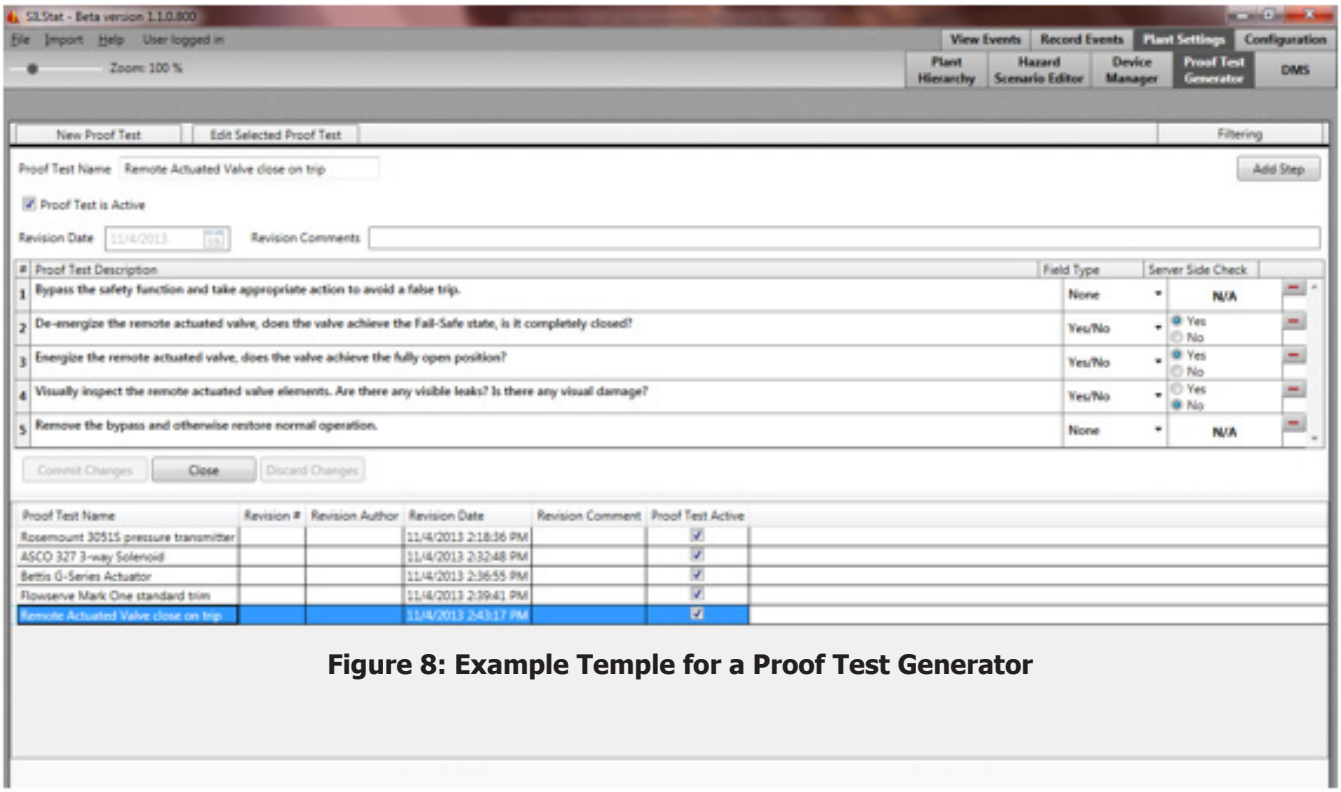


Figure 8: Example Temple for a Proof Test Generator

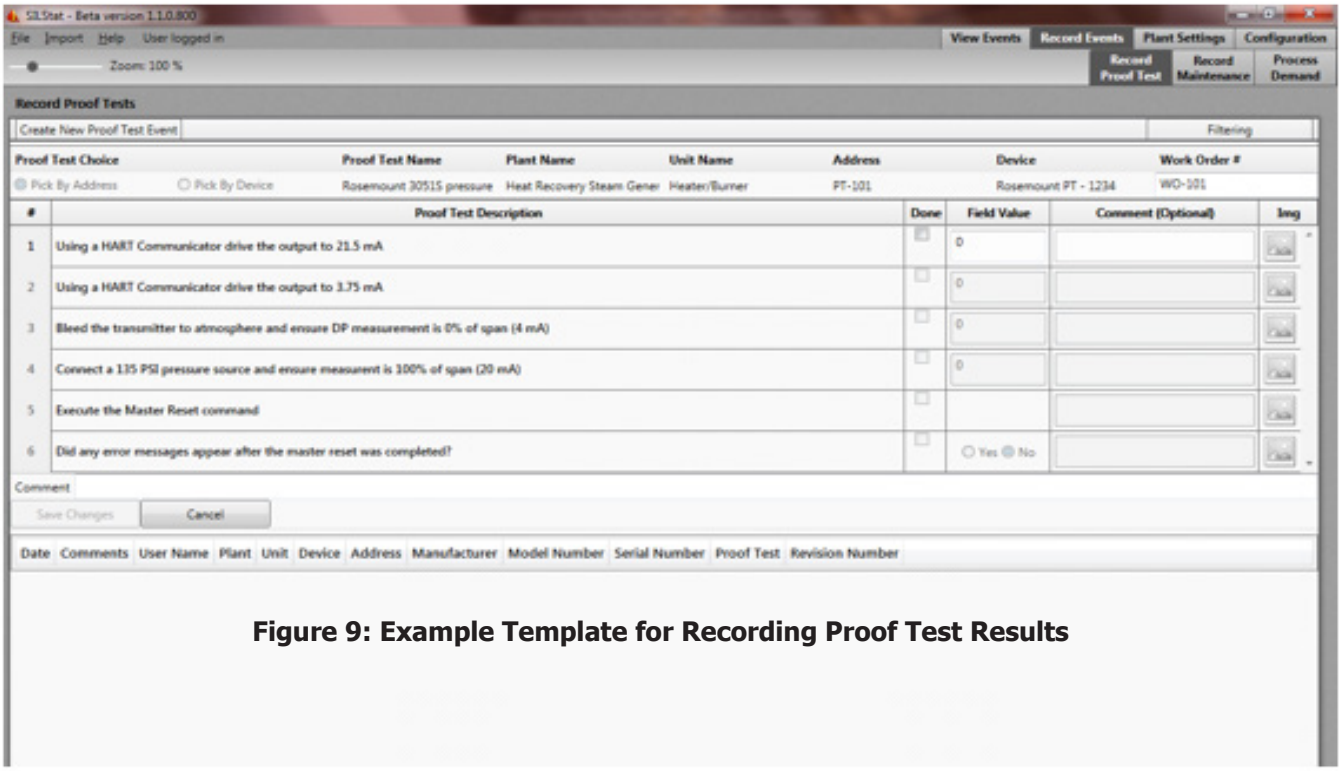


Figure 9: Example Template for Recording Proof Test Results

factual data during a proof test. Figure 8 illustrates an example template for a proof test generator.

Any problems found during proof testing will need to be repaired in a safe and timely manner, as defined in IEC 61511-1 Clause 16.3.1.4. Although the standard doesn't specify any particular time period, the Mean Time To Restore (MTTR), as used during the SIL determination of the SIF(s) and for the PFDavg calculation, must be adhered to in order to return the SIS to its safe state as soon as possible. Having the ability to identify and rectify any deficiencies quickly and effectively is the key.

Therefore, having a tool that enables the O&M personnel to identify the tag address of a device that is required to be tested, based upon steps defined for the proof test, which automatically determines a pass/fail condition for the test will save time and improve accuracy. Figure 9 illustrates an example template for recording proof tests.

Furthermore, being able to record these maintenance activities via a hand-held and/or mobile device would simplify the O&M personnel's job and enable a quick upload of all maintenance data to a central server where it can be reviewed and analyzed by the plant's safety or reliability team.

Essentially, being able to select and locate a device from the plant's hierarchy tree, for maintenance and/or replacement, via the tool, will save time especially if the O&M personnel can record the cause and any comments (the "as found" and "as left" conditions). Figures 10 & 11 illustrate example templates for recording maintenance tasks.



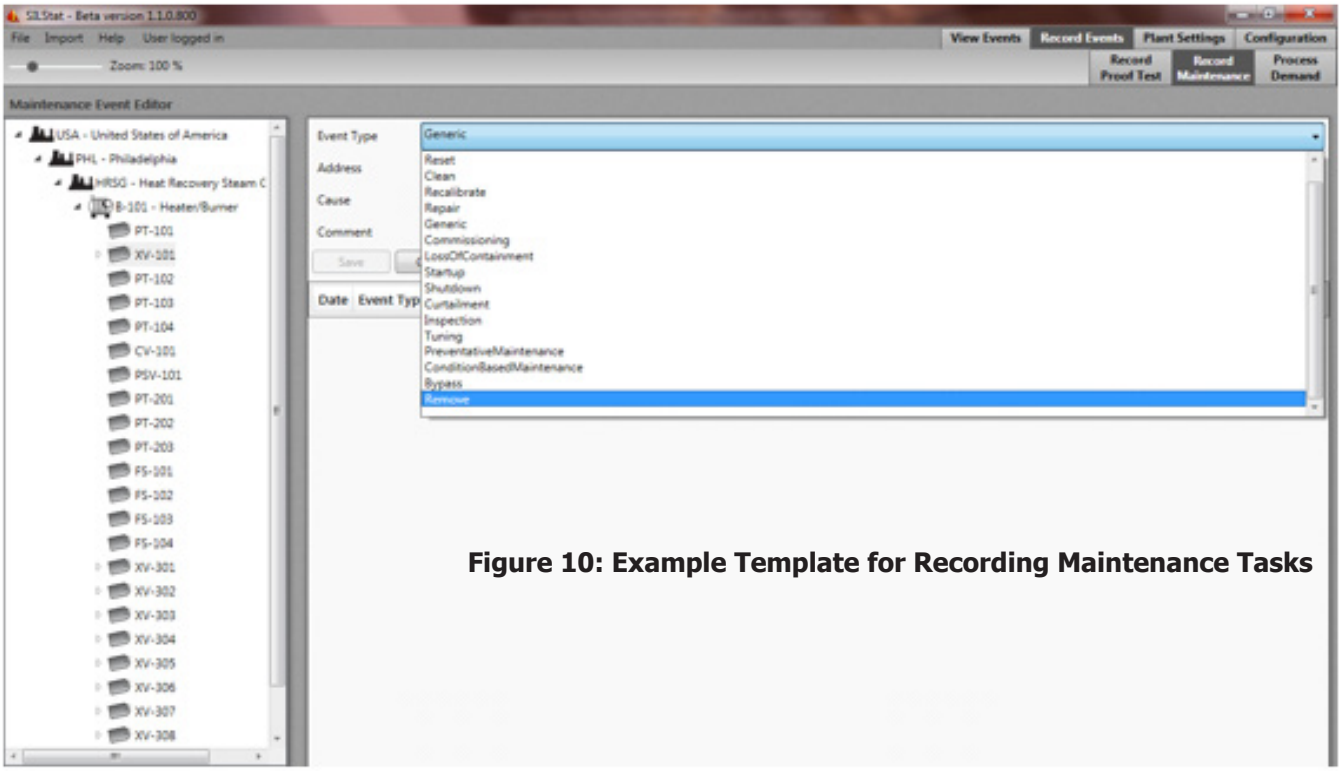


Figure 10: Example Template for Recording Maintenance Tasks

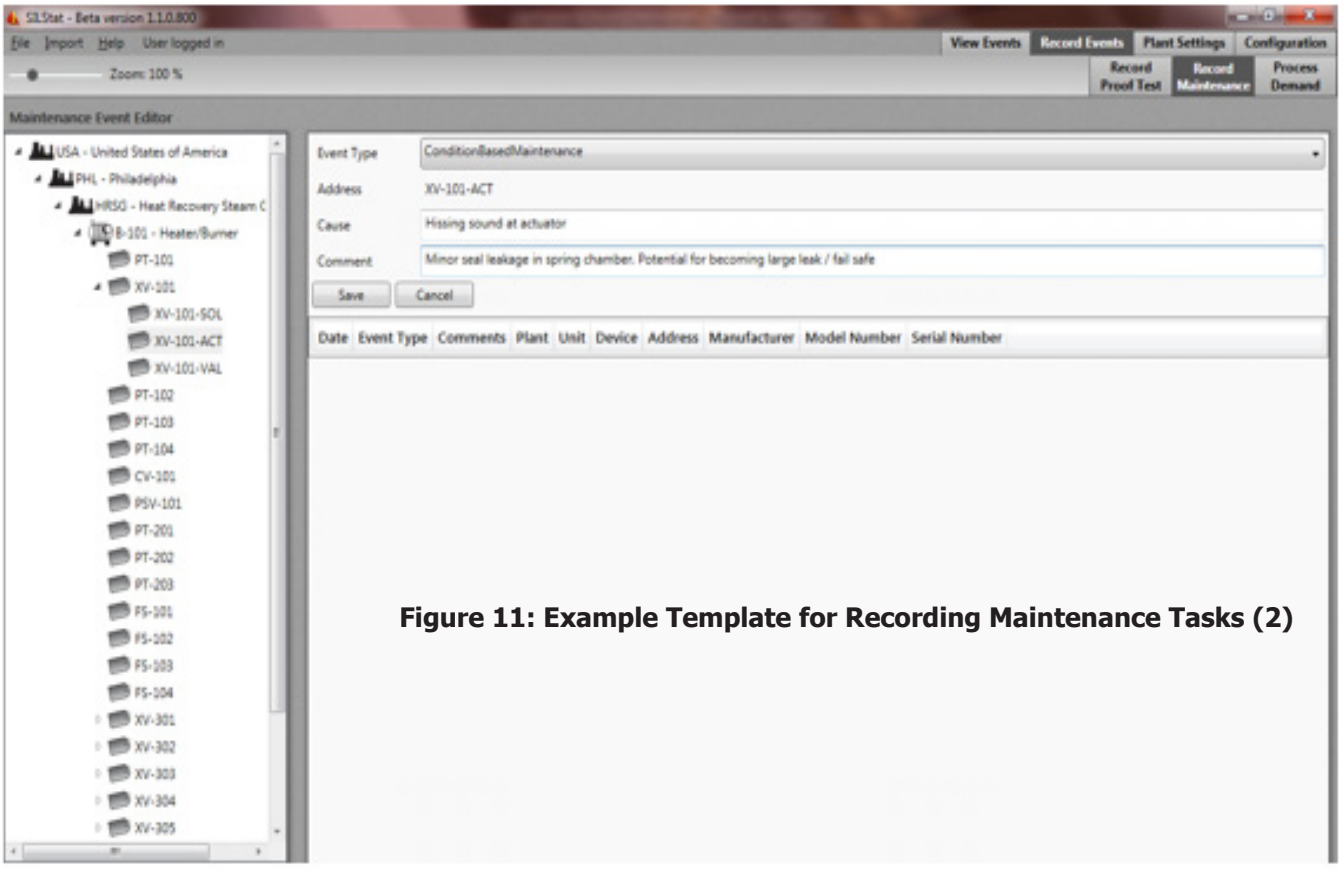


Figure 11: Example Template for Recording Maintenance Tasks (2)

Another benefit would be for the O&M personnel and the plant's safety manager or team to be able to view the events that had taken place, including the time and outcomes. Figure 12 illustrates an example template for displaying events that had occurred. The benefits to be gained from having a well-structured, defined and automated recording system can be characterised

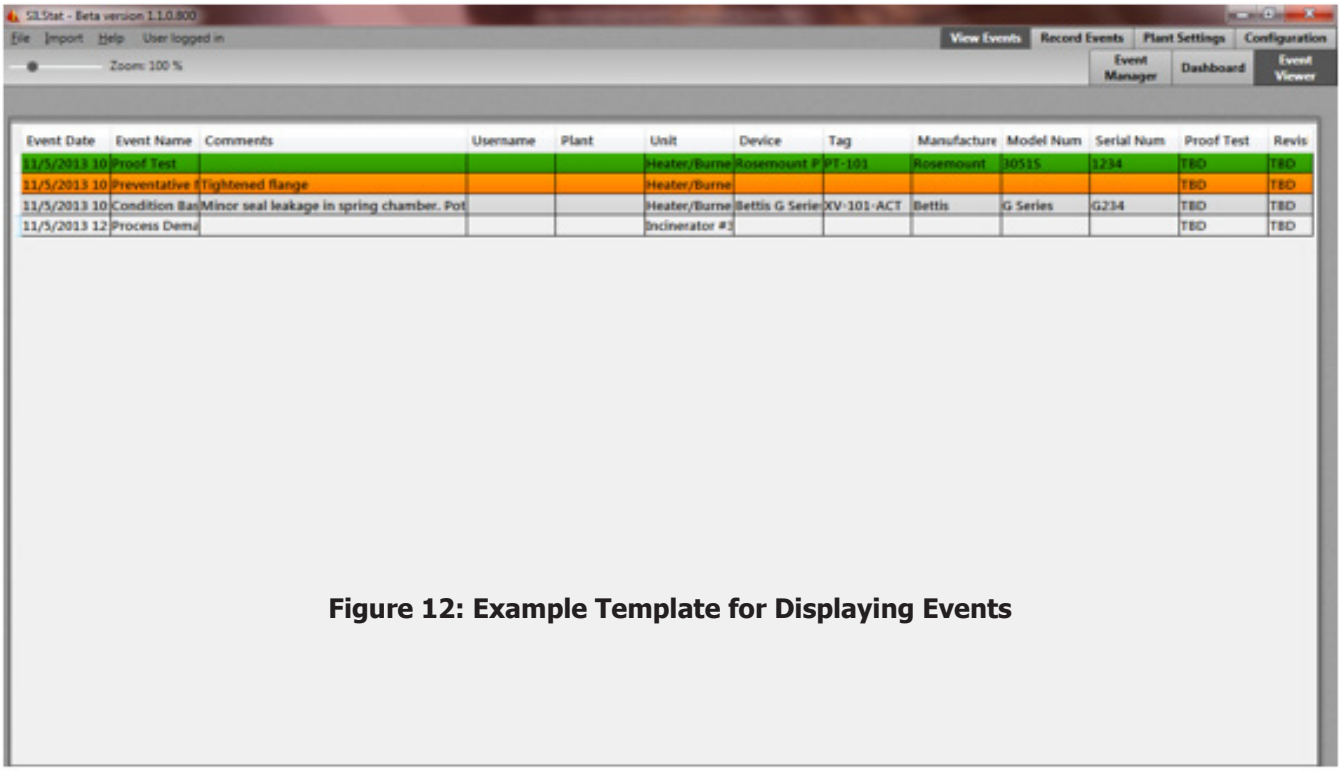


Figure 12: Example Template for Displaying Events

as the provision of:

- Detailed analysis of failures that can identify problems related to:
  - Specific addresses,
  - Specific devices,
  - Specific device types;
- The frequent benefit of reducing false trips of the plant;
- The ability to compare actual performance with assumed performance;
- The ability to ensure that the risk reduction is adequate;
  - Continue data collection;
- The ability to identify if the risk reduction is inadequate, i.e. if hazard risk is higher than expected and there is a need to return to Analysis phase of the SLC for redesign and implementation;
- The ability to identify if the risk reduction is more than adequate, i.e. if the system is overdesigned to prevent the hazard - this allows modification of design practices for future applications and potential for cost savings;
- Data for future Safety Lifecycle tasks, such as:
  - Risk Assessment
  - Layer of Protection Analysis
  - SIL Target Selection
  - SIL Verification.

The information gathered will enable the plant safety personnel to re-evaluate the frequency of proof testing, based upon the historical test data gathered, plant experience, hardware degradation and software reliability. The period for reviewing this will need to be determined by the plant safety manager.

In addition, a software tool can help solve any communication problems within the Plant that exist between the various "Managers" and their different departments. The following organizations are involved in O&M and most likely have different operational objectives:

- Operations;
- Maintenance;
- Process Safety (EHS);
- Reliability;
- Engineering.

Although these different groups will require their own particular data from this system, having all the information in a database that is easy to access and contains relevant, accurate data will ensure

consistency.

Summary

The paper has outlined some of the key issues involved in following the requirements of IEC 61511 Clause 16 for Operation and Maintenance of the SIS. As mentioned at the outset, the paper highlights how testing and documenting the performance of a SIS is an essential part of ensuring that the SIS is able to fulfil its designed functional safety requirements, as defined in the SRS. In addition, the paper outlines how taking advantage of technology and/or software tools to help with documenting and automating maintenance activities can help improve efficiency and reduce errors.

In summary, the key points are as follows:

- There is a need to manage risk - not ignore it.
- There is a need to adopt an appropriate safety-first culture to ensure O&M personnel are trained and competent to maintain the plant SIS.
- Recording lagging and leading indicators is an important part of maintaining and improving process safety.
- Having the proper Operation and Maintenance Procedures in place is vital to ensuring a safe and well-maintained SIS.
- Developing a "safety checklist" will ensure consistency in approach and methodology that can be adopted over multiple sites.
- Undertaking regular employee competency assessments is crucial in preventing mistakes that could ultimately lead to accidents and/or spurious plant trips.
- Ensure that proof testing is conducted in accordance with the Safety Requirements Specification of the SIS (i.e. using the same test interval as used in the PFDavg calculations).
- Recording all maintenance activities accurately and faithfully in accordance with IEC 61511 Clause 16.3.
- Use of software tools/technology to assist in recording and auditing maintenance activities, spurious trips, SIS demands, calibration and faults will save time and improve effectiveness.
- Use of software tools/ technology to help analyse failures, false trips and actual performance of the SIS compared to assumed performance will help in meeting IEC 61511 Clause 16.2.6. Any discrepancies will need to be assessed.
- Well recorded and accurate SIS performance data will enable plant safety personnel to be able to re-evaluate the frequency of proof testing, based upon the historical data gathered, plant experience, hardware degradation and software reliability.

References

[1] [IEC 61511-1] International Electrotechnical Commission (IEC) 61511-1 Functional Safety – Safety instrumented systems for the process industry sector – Part 1: Framework, definitions, system, hardware and software requirements:2016 edition

[2] [ANSI/ISA 84.91.01-2012] American National Standards Institute ANSI/ISA 84.91.01-2012 – Identification and Mechanical Integrity of Safety Controls, Alarms and Interlocks in the Process Industry

[3] [CSB] US Chemical Safety Board draft report 2010-08-I-WA JANUARY 2014 on April 2010 fatal explosion and fire at the Tesoro Anecortes Refinery Washington.

Steve Gandy is Vice President for Global Business Development and Director of the End User Service Business for Functional Safety and Cyber Security at exida Consulting LLC. He has nearly 40 years' industrial experience in training, senior, corporate and R&D management, having started his career as a hardware and software developer for fire protection systems. Steve is a former Board member of the IET, and a Certified Functional Safety Professional, and is in high demand as a speaker on safety and operational issues.

Remembering John Knight

SCSC members will be saddened to hear of the passing of Professor John Knight, who died on February 23rd 2017.

John was a great friend to SCSC, and had a long, distinguished and highly influential career in safety-critical systems and software engineering, particularly in the aerospace and automotive domains. After a spell at NASA's Langley Research Center, John moved to the University of Virginia in 1981. He became Professor of Software Engineering at Virginia, and headed the influential Dependability Research Group there. Members will remember the blend of erudition, practical common sense and wit which characterised both his writings and his presentations at Club events over many years, as well as his joviality and generosity.

A fuller obituary can be found at http://www.dailyprogress.com/obituaries/knight-dr-john-c/article_0509831e-6 1 3 f - 5 1 4 b - 8 2 9 f - 2252ab36730b.html

Calls for Papers

9th International Workshop on Software Engineering for Resilient Systems (SERENE 2017)
September 4-5, 2017, Geneva, Switzerland
http://serene.disim.univaq.it/2017/home
The SERENE workshop has a long tradition of bringing together researchers and practitioners to discuss advances in engineering resilient systems. Since 2015 SERENE has been part of a major European dependability forum - European Dependable Computing Conference - EDCC. This year SERENE will be held before the main programme of EDCC 2017.

SERENE 2017 will provide a forum for researchers and practitioners
to exchange ideas on advances in all areas relevant to software engineering for resilient systems.
The conference proceedings will be published by Springer in the LNCS series.
The submission site is:
https://easychair.org/conferences/?conf=serene2017

Important dates:
Paper submission:
June 1st 2017
Author notification:
June 24th 2017
Camera-ready papers:
July 1st 2017

International Conference on Reliability, Safety and Security of Railway Systems: Modelling, Analysis, Verification and Certification (RSSRail 2017)
November 14-16, 2017, Pistoia, Italy
https://conferences.ncl.ac.uk/rssrail/

The RSSRail conference aims to bring together researchers and engineers interested in building critical railway applications and systems. This will be a working conference in which research advances will be discussed and evaluated by both researchers and engineers focusing on their potential to be deployed in industrial settings.

The conference proceedings will be published by Springer in the LNCS series.
The conference submission site is:
https://easychair.org/conferences/?conf=rssrail2017

Important dates:
Abstract submission:
June 1st 2017
Paper submission:
June 8th 2017
Author notification:
July 8th 2017
Camera-ready papers:
August 18th 2017

Events Calendar

The following list contains details of all forthcoming SCSC events, together with other events which may be of interest to members. Where possible, we have given web addresses for further information. Every effort has been made to ensure that the information here is correct, but readers are advised to confirm with event organisers.

Club Events

April 6th 2017
SCSC Seminar: New Thinking in Human Factors for Safety
London, UK
This seminar eplores insights from the Human Factors field, with the aim of improving co-working of humans and systems and the design of autonomy-human interfaces. Topics to be covered include: learning lessons from human workarounds to accommodate poor system design, the safe handover of control, ways in which system behaviour influences human behaviour, handling malevolant operators, and designing systems to cope with human failings. Speakers include Chris Parsons (Atkins Global), Neil May (NATS), Chris Lowe (Liv Systems), Carl Sandom (iSys Integrity), Alexander Eriksson (University of Southampton), and Christian Fairburn (Reliable Interaction).

April 25th 2017
Data Safety Initiative Working Group
Bristol, UK
33rd Meeting of the DSIWG. Please contact Mike Parsons (mike.parsons@scsc.org.uk) for details.

April 27th 2017
Service Assurance Working Group Meeting
London, UK
Second meeting of the SAWG. Please contact Mike Parsons (mike.parsons@scsc.org.uk) for further details.

June 8th 2017
SCSC Seminar: Risk, Proportionality and ALARP: How do we do enough for safety?
London, UK
This seminar considers the difficult topic of how to decide when we have done 'enough', i.e. the appropriate amount of effort in safety engineering activities. It will include discussion of safety assurance levels, mitigation approaches, the practice of risk management, legal structures (ALARP, SFAIRP etc), risk acceptance, tolerability and societal expectation, what is considered appropriate for different sectors, and arguing sufficiency. Speakers include Felix Redmill (Redmill Comsultancy), Mark Machin (CGI UK), Tim Clement (Adelard LLP), Ron Bell (ESC), Neil Beveridge (CRARisk), Andy Lidstone (Risktec).

All Club events can be booked directly from our website www.scsc.uk

February 6-8 2018
York, UK
SSS'18: Safety-Critical Systems Symposium

Other Safety and Dependability Events

April 18-21 2017
Pittsburgh, USA
SCAV 2017: Safe Control of Connected and Autonomous Vehicles
www.scav.in.tum.de

4-8 September 2017
Geneva, Switzerland
EDCC 2017: European Dependable Computing Conference
edcc2017.unige.ch

October 17th 2017
Bristol UK
HIS 2017: High-Integrity Software Conference 2017
www.his-2017.co.uk