iESM
BROUGHT TO YOU BY ARC

international Engineering Safety Management

GOOD PRACTICE GUIDANCE

APPLICATION NOTE 9
SAFETY INTERGRITY WITHIN ENGINEERING SAFETY MANAGEMENT

Published on behalf of the international railway industry
by Abbott Risk Consulting Ltd
Issue 1.2 May 2022

arc

CONTENTS

Issue 1.2

DISCLAIMER

Abbott Risk Consulting Limited (ARC) and the other organizations and individuals involved in preparing this handbook have taken trouble to make sure that the handbook is accurate and useful, but it is only a guide. We do not give any form of guarantee that following the guidance in this handbook will be enough to ensure safety. We will not be liable to pay compensation to anyone who uses this handbook.

ACKNOWLEDGEMENTS

This handbook has been written with help from the people listed below.

- S Bickley
- M Castles
- Chan WK Yai Kit
- P Cheeseman
- J-M Cloarec
- Dr R Davis
- B Elliott
- S Hughes
- Ng Nelson Wai Hung
- G Newman
- G Parris
- M Roome
- A Russo
- G Topham
- Yu Bernard Tin Fu

These people worked for the organizations listed below.

- Abbott Risk Consulting Ltd.
- Arbutus Technical Consulting
- Bombardier Transportation
- Certifer
- MTR Corporation Limited, Hong Kong
- Ricardo Rail
- Rio Tinto
- SNC Lavelin
- Systra
- Technical Programme Delivery Group
- WSP | Parsons Brinkerhoff

This handbook does not necessarily represent the opinion of any of these people or organizations.

1 INTRODUCTION

This Application Note (AN) is a component of the international Engineering Safety Management Good Practice Handbook, or 'iESM', for short. The handbook as a whole describes good practice in railway Engineering Safety Management (ESM) on projects. It covers both projects that build new railways and projects that change existing railways.

The iESM handbook is structured in three layers:

- Layer 1: Principles and process
- Layer 2: Methods, tools and techniques
- Layer 3: Specialized guidance

The first layer comprises one volume, Volume 1. Volume 1 describes some of the safety obligations on people involved in changing the railway or developing new railway products. It also describes a generic ESM process designed to help discharge these obligations.

Volume 2 provides guidance on implementing the generic ESM process presented in Volume 1 on projects. Volume 2 belongs in the second layer. At the time of writing, Volume 2 was the only document in the second layer but further volumes may be added to this layer later.

The third layer comprises a number of Application Notes providing guidance in specialized areas, guidance specific to geographical regions and case studies illustrating the practical application of the guidance in this handbook.

The structure of the handbook is illustrated in the figure on the right.

This document is Application Note 9. It supports the main body of the iESM handbook, in particular Section 9, by providing guidance on safety integrity as part of ESM. Safety integrity is not always well understood, and there is a tendency to use it incorrectly and inappropriately. This AN attempts to help avoid its misuse and recognise its misleading use by others.

The role of iESM Application Notes is to develop more detail where required under the existing principles and guidance in iESM Volumes (layers) 1 and 2.

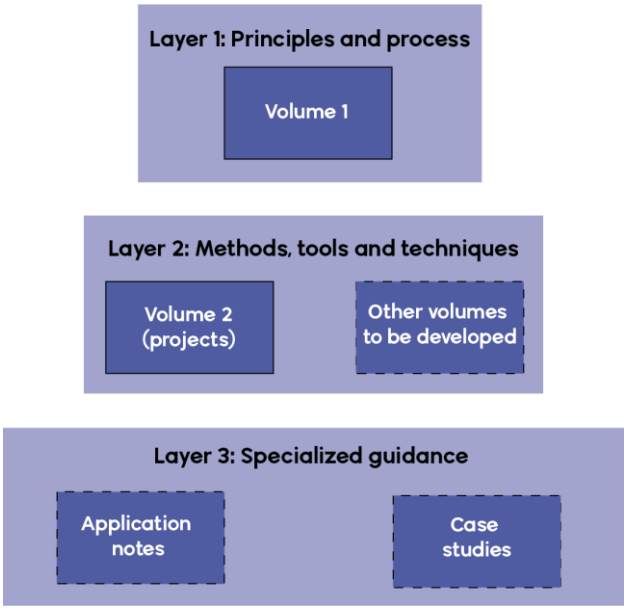

Figure 1 The Structure of iESM Guidance

2 BACKGROUND

2.1 Purpose

This AN expands on the guidance provided on Safety Integrity in Section 9.2.4 of Volume 2 of the iESM guidance. Historically, safety thinking categorized things as “safe” or “unsafe”, “vital” or “non-vital”. Experience shows that the goal of eliminating risk and bringing about a state of absolute safety is not attainable, even more so with complex systems. Safety is therefore not binary but a spectrum and the choice of how much safety is needed is an economic one as well as a technical one. The more significant the impact of a failure, the greater the need for dependability of function and the lower the hazardous failure rate[1] that is acceptable.

2.2 Safety-related failures

There are well-established techniques for assessing and controlling the risk arising from random failures, typically hardware components. These are dealt with in some standards in a somewhat arbitrary way by the use of Safe Failure Fractions (SFFs) which prescribe a set of architecture rules to address the “unknown unknowns”. The SFF is the percentage of non-hazardous failures (i.e. those that fail-safe or are detected) of the total failures, the remainder being the undetected unsafe failures. In practice, this is often dealt with by means of diagnostics like watchdogs to monitor the health of the protection system. No equivalent to SFF yet exists in the railway industry.

Systematic failures are those where the intended outcome is not achieved, usually due to faulty specification or development processes - nothing actually breaks or wears out and the failure is repeatable for the same set of conditions (inputs and context information as appropriate). Risks from systematic failures are controlled in many engineering activities through rigorous checking and the application of standards, codes and accepted good practice. The CENELEC family of Railway Application standards provide processes and techniques to help make a claim of achieving an acceptable level of integrity and hence risk. Note also that some of these processes and techniques also help to reduce random hardware failure. Figure 2 below, taken from EN 50126 [50126] illustrates the place of systematic failure within the requirements hierarchy for a typical railway product or system.

It is important to note that focusing on systematic failure integrity alone does not achieve a safe system. Demonstration of safety integrity relies fundamentally on the correct safety requirements being defined and implemented. That requires a robust ESM process to be applied systematically by competent people with both relevant technical and domain knowledge with suitable support from their organization.

[1] This is not the same as the failure rate that reliability analysts may use. It relates to those failures that lead directly to hazards.

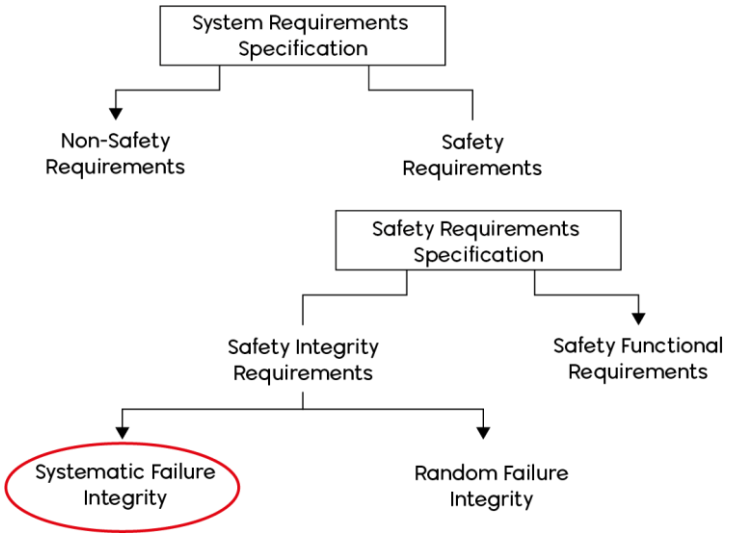

Figure 2 Systematic Failure Within a System

Thus Safety Integrity comprises two elements:

- the quantitative hardware (random) failure rate arising from failed or degraded components for example component ageing failure, and
- the non-quantitative systematic (man-made) failure integrity – hardware and software arising from an inadequate specification or mistake in the development process, for example:
    - Stress failure due to excessive vibration or too high temperature
    - Error due to incorrect software
    - Human interaction failure for example valve left in wrong position repeatedly under similar circumstances

The last of these is what could be termed a maintenance error and emphasizes the importance of clarity of any safety-related application conditions arising from the Safety Case. This is important as it illustrates that systematic failure integrity achievement is not wholly controlled within the design part of the lifecycle. Yet nor can it be dismissed as a reflection on (lack of) maintainer competence. Many accident reports[2] illustrate the effect of poor design leading to mistakes by maintainers and operators particularly under stress.

Initially it may seem sensible to make everything as safe as possible ("safety is a given" is often quoted). However this has a number of drawbacks:

- It is expensive, time consuming and absorbs valuable resources
- The resultant "safe" system may be highly unreliable and therefore unsuitable for its intended use
- When any modifications are made or non-conformities corrected the necessary assurance activities must be repeated. This adds cost and also delays system improvements for reliability or system-level performance of functions with no safety content

However, as the complexity of designs increases, systematic failures contribute a larger proportion of the risk and the effectiveness of checking or other quality assurance techniques, decreases. For software, all failures are systematic. In software and some other areas where designs may be particularly complex, such as electronic design, current best practice is to make use of Safety Integrity Levels (SILs) to manage systematic failures.

[2] See "Three Mile Island" for a dramatic illustration of this https://en.wikipedia.org/wiki/Three_Mile_Island_accident.

2.3 Safety Integrity Levels

SILs are described in a number of widely-used standards, including EN 50126 [50126], EN 50128 [50128], EN50129 [50129] and IEC 61508 [61508]. Confusion can arise because standards written for use in other industries do not easily apply to railway system architecture (see Section 2.4). SILs represent different levels of rigor in the design process and are related to approximate probability targets. Five levels are defined in EN 50128 [50128], ranging from SIL 4, the most stringent, to SIL 0, the least stringent. EN50126-2 [50126] and EN50129 [50129], have four levels defined3. SIL 4 has the highest level of safety integrity; Basic Integrity the lowest. Functions which are not relied upon at all to control risk have no SIL and therefore have no part in ESM.

Each level is linked with increasingly stringent processes and techniques. However whilst there are recommendations for specific processes there is no guidance as to which subset of the processes and techniques should be used, nor the level of rigor to be applied. The nature of the product or system will help decide, but it is difficult to offer specific guidance other than consulting suitably experienced and competent people.

SIL is a concept, not a tangible property of a product or system. It relates to the engineering effort applied to the design, validation and maintenance of a system within an application ("doing things right"), not to an actual safety performance achieved ("doing the right thing"). A SIL does not guarantee that the item is bestowed with the attributes necessary for its safe application. It is therefore possible to have a SIL 2 system, for instance, which delivers a safety performance which is consistent with that expected of a SIL 4 system. Unfortunately and dangerously, the opposite is also true.

Note that SIL has nothing to say to occupational health and safety hazards.

The CENELEC guidance document of systematic allocation of safety integrity levels is suggested, although not easy reading and does not cover software SILs [50451].

2.4 Safety & Reliability

The Tolerable Hazard Rate (THR) is the maximum rate of occurrence of hazards (situations with the potential for harm) that would be acceptable due to a product or system. Whilst it may be driven significantly by the underlying failure rate of the system, hazards exist even when everything is working properly. The THR is application-specific (although common requirements will be found in many sets of applications) and is chosen such that the risk arising from the overall system is acceptable.

When a lower THR is required a higher integrity is logically necessary; thus the risk of a shut-down will also increase as any defect, detected will normally result in a move to a safe state (usually on a railway a fail-safe state) as that is likely to be the safest and most practical response. The additional software or circuitry to do the monitoring has its own inherent reliability and may shut the system down when it or the primary system is actually functioning correctly. Care is needed at the specification stage to decide if any detection function is essential to the safe delivery, or whether for example the railway may operate safely without such function for a period of time.

[3] In EN 50129:2003 "SIL 0" was introduced to indicate non-safety-related functions. This term is no longer used and non-safety-related functions are just referred to as such.

SIL is sometimes mis-used as a proxy for system reliability for example by specifying SIL 4 in an attempt to achieve a highly reliable system. This is bad practice and not supported by the CENELEC Railway Application standards or the iESM guidance. This misunderstanding arises because tables in the standards seem to imply that a specified level of reliability may be claimed at each SIL. However by imposing a SIL on functions that have no safety content, reliability may be decreased without any benefit - just the opposite of what was intended. For example, for an emergency communications link it may be availability rather than reliability that is the critical characteristic because satisfactory communication may be possible in spite of intermittent failures.

From a safety perspective it is only the hazardous (unsafe) failures that we are interested in rather than the absolute number (or rate) of failures. Whether a failure is unsafe is application dependent. Most of the techniques and measures offered by IEC 61508 and EN 50128/9 are aimed at reducing the total number of defects regardless of their effect on safety.

Safety and availability requirements must be specified in a compatible way to produce a meaningful set of requirements. If that is not the case the requirements should be reviewed.

3 WARNINGS

The following warnings are given because of the widespread misunderstanding of SIL in the railway industry. They may not apply in all situations but have been highlighted here to reduce the risk of oversight:

- Safety is application-dependent. To attach SILs to products outside the context of the systems in which they will function without further analysis can be misleading – or even dangerous.
- As published at the time of writing this AN, the CENELEC Railway Application standards are inconsistent in their definition of SIL:
    - They are scoped to only apply to railway signalling systems. Although we believe there is no fundamental problem with that approach with benefits arising from applying a consistent approach across the wider railway system care is needed to avoid imposing additional cost and effort without any tangible safety benefit.
    - EN 50128 has introduced the concept of a defined lower safety level of "SIL 0" whereas EN 50126 and EN 50129 used SIL 0 to mean something that is not safety–related. The latest updates of EN 50126 and EN 50129 use the term "basic integrity" for any safety-related system where the TFFR is greater than 10^{-5}.

- Even in complex systems, SILs are not the only means of controlling safety integrity. It may be controlled through architectural design features, simulation modelling, non-destructive testing as appropriate to the nature of the work.
- In the CENELEC standards the SIL is with reference to a function, not a system. It follows that it is misleading to call a particular product "SIL 1" or "SIL 2." This is because the entire system and its application must be taken into account. It is accurate to say that a product can be used to perform a SIL 2 safety-related function in a specific application.
- Despite what many people say, safety is different from reliability. A reliable system is not necessarily safe, if its requirements are wrongly defined, wrongly implemented or it is wrongly used. Nor is a safe system inherently reliable as fail-safe features prioritise safety over functionality. Safety without availability is pointless and also usually less safe when viewed from a practical perspective. A railway where all the trains are stopped is safe (at least for a while) but also useless. Moreover, in real systems if a system fails safe, alternative, usually less safe, means may be employed to continue operations. An example is where operational rules are used to move trains after a signalling system has failed to a safe state, even though operational rules do not provide the same level of safety.

4 GUIDANCE

4.1 Safety functions

In railway applications the SIL applies specifically to a safety function. A safety function is a property of the product or system that must be delivered with a specified level of integrity (dependability). For example, an over-temperature protection device, using a thermal sensor in the windings of an electric motor that could de-energise the motor before the windings overheat, is an instance of functional safety. Conversely, providing specialised insulation coating to withstand the high temperature on the motor windings if they do overheat is not an instance of functional safety (although it is still an instance of a safety requirement and could protect against exactly the same hazard).

Note that it matters whether the safety function is continuous or 'on-demand', and if 'on-demand' the conditions of use and the required deployment time should be stated e.g. "when a specific event happens…" or "…within so many seconds". Low demand is commonly accepted as no greater than once per year[4]. A SIL requirement or offer without these qualifiers is meaningless, or potentially dangerous, or both.

It is bad practice to use incorrect shorthand when describing products or systems as "having a SIL of x". To do so can create unintended commitments or expectations that may be difficult or impossible to deliver. Areas to be specified should include:

- Indications of the rigour that has been applied throughout design and development
- The nature and the independence of the assessment to which the design and development processes were subjected
- The testing and test results which provided confidence in the developed product
- The history of use, if any, which confirmed that confidence

In other words, we need a safety argument to go into the overall Safety Case wherever a SIL is claimed.

Avoid things like 'Product A' 'is SIL4' but use wording such as: 'is designed to provide specific functions at an integrity that meets or exceeds the safety integrity requirements needed to satisfy EN50128/129 SIL4 when correctly installed, integrated, operated and maintained'. Do not say 'the system will be SIL2', rather use wording such as 'the Control Centre software shall be written and assured using appropriate techniques from EN 50128 for functions requiring an assured safety integrity of SIL2'. These are subtle but important differences.

Assigning the desired probability, or degree of confidence, that the system should achieve in performing the required safety functions, requires an assessment of the risks of the specific application. It is subject to all the uncertainties inherent in reliance on human judgement and fuzzy and incomplete knowledge of risk. It is discussed further in Section 4.2.

Demonstrating that a procured system does or may be able to meet the expected performance is typically based on showing that the system and any software has been developed in conformity with one of the sets of processes and techniques recommended for the required SIL by the CENELEC Railway Application standards in each step of the lifecycle as discussed in Section 4.3 below.

[4] A year being roughly 10^4 hours which accounts for the difference in the figures between the two tables in IEC 61508

4.2 SIL Requirement

4.2.1 System Considerations

Consider a train that is equipped with automatic driving. If the train is designed to be periodically driven manually, then a human driver cannot be assured to meet any SIL. Human factor errors should be used in such analysis and whilst they may look like they fit into a SIL range are in fact a mixture of both systematic and random error. If the Automatic Train Operation (ATO) system does contain safety functions such as door control or roll-forward protection then those functions can be allocated a SIL if necessary; the non-safety functions cannot be. Similarly most control centre functions have their safety assured by the signalling interlocking that sits behind them. Anything associated with the execution of a manual operational command cannot have an overall integrity better than the manual input that initiates it.

Note that some railways with their long history of largely excellent safety performance and strong safety culture do not consider control centres to be safety systems. That said, certain commands like 'emergency signal replacement' or 'train operated route release' can have a safety function and a SIL may appropriately be assigned to ensure that such a function is correctly executed to a defined level of certainty once a command is given. Wherever possible, these functions should be segregated in the overall architecture, thus minimising the impact that the subsystem will play in the overall safety assessment and ongoing safety management. Complex non-safety-related functions such as data analysis and decision support, often required as part of modern control centers, need the functionality of High Level Operating Systems such as Windows™ or Linux and the imposition of unnecessary SILs in these areas will be in conflict with delivering such functionality cost-effectively, or indeed at all.

4.2.2 Assignment of SIL Requirement

A quantitative approach to SIL assignment is the most rigorous technique as described in EN50129 [50129].

The hazard likelihood is calculated by modelling the combined influence of the potential causes, control measures and any external events that are required for the consequence to be realized. This information can be used to construct and quantify a Fault Tree or Event Tree as described in the iESM Guidance and Application Note 3.

The Event Tree can include probabilities of events such as immediate or delayed ignition or that people are present and subject to the hazardous situation. This requires a thorough understanding of the event sequences and data for each basic event or node.

It is very difficult to prove functional independence within a sub-system and so it is important to take care in assigning functions to sub-systems.

For a meaningful safety assessment to take place an overall system safety target needs to be set and then THRs derived for each of the subsystems contributing to it. A system of SIL4 subsystems will not necessarily achieve SIL4 functionality unless the assessment, the architecture and system design that underpins it are done correctly.

If possible, functions with differing SIL requirements it may be more cost-effective to separate them, either physically or logically, so you don't have to apply the requirements for a higher SIL to lower SIL functions.

Practitioners have successfully justified designs with software functions of different SIL on the same processor, although EN 50126 does not provide any support for this practice. To be able to use software functions of varying SIL on the same processor, it must be possible to produce a safety argument that demonstrates that the lower SIL functions cannot influence the behavior of those with higher SILs. This may be through mechanisms that prevent interference such as memory protection or by analysis of the code, for example by demonstrating that no part of the code will write to memory outside its designated area. However, this can be difficult to do, and the effort required may be excessive compared with other solutions to the same problem.

Numerical safety targets are normally derived from a Fault Tree to:

A. derive numerical accident limits which conform to the legal or other criteria for acceptable risk;
B. derive hazard occurrence rate and/or unavailability targets which are consistent with (A);
C. if applicable, derive SILs for the system functions that are consistent with (B).

The requirements may be apportioned further to sub-systems of the hierarchy and aligned with the system design. In general, targets for systematic failure should not be set below sub-system function level.

For a single safety function with a THR within a specified range, one widely accepted association is shown in Table 1, which is derived from EN 50129 [50129]. Note this table should only be used in the following way - for a rate of hazardous failures, the coinciding SIL, is searched in the table, i.e. the table is only used from left to right (rather than used from right to left to derive a metric for the specified SIL). Then, specified processes and techniques have to be applied during the design process according to the relevant standard.

THR per hour	SIL Assignment
$10^{-9} \leq THR < 10^{-8}$	4
$10^{-8} \leq THR < 10^{-7}$	3
$10^{-7} \leq THR < 10^{-6}$	2
$10^{-6} \leq THR < 10^{-5}$	1

Table 1 THR and SIL

Note in rail system architecture, there are often multiple unsafe failures of safety functions that can lead to a hazard (unlike some other industries). In such cases it may be necessary to apportion the THR across the set of unsafe failures (called in the latest CENELEC Railway Application standards, the Tolerable Functional unsafe Failure Rate (TFFR) and in earlier documents THR_i). Each TFFR will be a proportion of the total allowable THR.

Note also that unlike IEC 61508 [61508], the SIL table in the CENELEC Railway Application standards only has one column for frequencies (in an earlier version of the standards they were identified as high demand or continuous mode) and does not have a column for failure probabilities on demand (formerly called demand mode). The reasons to restrict to one mode are said to be:

- Less ambiguity in determination of SIL
- All demand mode systems can be modelled as continuous mode systems
- Continuous control and command signalling systems are clearly the majority in modern railway signalling applications.

Sometimes a quantified analysis derived from a top level target is not possible or not available and a SIL is assigned in a qualitative way. In that case the most restrictive THR of the relevant range shall be specified and demonstrated. For example, for a function with a qualitative requirement of “SIL2” a THR of 10^{-7} per hour should be used for any quantitative demonstration.

A better qualitative approach would be using a "reference system" but which matches an equivalent THR or unsafe functional failure rate rather than the SIL itself, which can be re-derived for the new application.

Nor is it actually necessary to specify a SIL at all. 10^{-6} hazardous failures per hour is an adequate statement or a numeric Probability of Failure on Demand can be just as clear. Some railways find it simpler to identify the relevant hazards, and specify the main activities that a supplier should carry out, instead of stating a certain SIL. This can be useful if the system under design has not historically had a SIL allocation to its functions.

4.2.3 Apportionment of SIL Requirement

Having set a SIL for a function to achieve the necessary probability target, it may be necessary to apportion this between lower-level functions. By default the lower-level functions will inherit the highest SIL of the top-level functions that they support.

SILs for the functions of a sub-system may generally be set according to the target rates of hazardous failures of the system. This may allow some sub-systems to be provided at a lower SIL than that of the overall system by building in back-up or protection systems. However EN 50126 [50126] recommends restrictions on assigning a lower SIL to sub-systems in the cases where the functions of two sub-systems control the same hazard and there is some interaction between the two functions – for further information on this, refer to EN 50126 [50126].

If the architecture ensures that a top-level function can only fail if both a main and back-up function fail and the two functions are fully independent (including software), then the SIL of the top function may sometimes be higher than that of the main or back-up function. In some cases, there may also be a combinator function (for instance, a voting arrangement), which combines the results of the main and back-up functions. Table 2 shows some combinations which are generally regarded as valid, provided that:

- the lower level functions are physically separated and built using different design principles, and
- the combinator suppresses any hazard for any failure of one lower level function.

Table 2 cannot be repeatedly applied to allow a SIL 4 system, say, to be made of many SIL 1 systems.

IEC 61508 [61508] has no process for creating a system of one SIL by means of a number of components of lower SILs.

Top Level SIL	SIL of Lower Level Function		Combinator (if used)
	Main	Back up	
SIL 4	SIL 4	None	None
	SIL 4	SIL 2	SIL 4
	SIL 3	SIL 3	SIL 4
SIL 3	SIL 3	None	None
	SIL 3	SIL 1	SIL 3
	SIL 2	SIL 2	SIL 3
SIL 2	SIL 2	None	None
	SIL 1	SIL 1	
SIL 1	SIL 1	None	None
	SIL 1	SIL 1	SIL 2

Table 2 Combination of SILs

Note that any combinator always inherits the top level SIL, although it is typically a simple device. Redundant sensors can increase the overall SIL achievement e.g. it is often stated as "1 as 1oo1 /2 as 1oo2," meaning: SIL 1 if the device is one-out-of-one device used, SIL 2 if it is one-out-of-two devices with voting is used.

4.3 SIL Achievement

4.3.1 Processes and Techniques

The CENELEC Railway Application standards suggest a large number of processes and techniques in order to gain confidence that a claim of achievement of a specified SIL may be made. SIL 1 demands basic sound engineering practices, such as adherence to a standard quality system, repeatable and systematically documented development processes, thorough verification and validation, documentation of all decisions, activities and results, and independent assessment. Higher SILs, in turn, demand this plus further rigour.

TECHNIQUE/MEASURE	Ref	SIL 0	SIL 1	SIL 2	SIL 3	SIL 4
1. FMEA/FMECA	H.2.30.4 H.2.30.5 See Note 1	-	R	R	M	M
2. Bent Pin Analysis / Cable Failure Matrix Analysis	H.2.14.1	-	R	R	HR	HR
3. Electromagnetic Compatibility Analysis	H.2.14.2	-	R	R	HR	HR
4. Energy Trace and Barrier Analysis	H.2.14.3	-	R	R	HR	HR
5. Materials Compatibility Analysis	H.2.31.1	-	R	R	HR	HR
6. Fault Tree Analysis	H.2.30.7 See Note 1	-	R	R	HR	HR

Figure 3 Typical Extract from CENELEC EN50129 Railway Application Standard (for illustration only)

Figure 3 shows processes and techniques that are:

- R – recommended
- HR – highly recommended (non-compliance would need to be justified)

Other similar tables show:

- M – mandatory
- NR – not recommended techniques.

These are not "magic" SIL processes and techniques but typical well-established design methodologies that are mostly used routinely. When considering SIL they are given special attention because of their ability to mitigate systematic failures. The staff using the technique must be a competent and well-versed practitioner or the benefits may not be obtained in practice.

EN 50128 provides no method for estimating the probability of software failure as there is no random element in software. The practice of using the worst-case probability associated with the SIL of the software is not supported by the standard. We are aware that this practice has been followed on some railway systems, nonetheless. We do not consider it to be an unreasonable approach as the requirements of the standard would be open to challenge if they routinely resulted in software that failed more often than this limit. It should however be supported by actual tests results that support the claim i.e. if the software has met the integrity requirements then very few if any, errors should be apparent at system testing.

Without estimating the probability of software failure it is not possible to estimate the overall probability of failure of a system containing software. It is possible, however, to estimate the probability of system failure from non-software causes and to present this figure, carefully explained, together with the SILs of the system function in the Safety Case. With Fault Trees, the probability of system hazardous failure from non-software causes can be calculated by setting the probabilities of software failure to zero, although it must be understood that this is a device for excluding software failure from the calculation, not an assumption that software does not fail. Be careful however if the software includes functions that protect against other hazard causes. Setting the probability of failure of such functions to zero can result in a zero estimate for the probability of the hazard. In these circumstances it may be necessary to provide probabilities for nodes in the fault tree below the top event, to arrive at a meaningful answer. It may be helpful to split the failure equally (albeit arbitrarily) between software and hardware for the purposes of analysis.

Thus having followed the processes and techniques required (or a justified sub-set of them) for a specified SIL, the THR due to **systematic causes** should lie within the stated range. However knowing what processes and techniques were followed is not assurance enough on its own. Good practice suggests that evaluation criteria should be defined and related back to integrity requirements allocation, to assess how thoroughly the processes were followed and how thoroughly the output was checked. This is particularly important for the SILs where the same set of processes and techniques are specified e.g. SIL 1/2 and SIL 3/4. An alternative approach might be to look at service performance. Note that for highly reliable systems this can only be a claim or assumption, but cannot be successfully proven. In order to demonstrate that a probability of dangerous failure of 10^{-9} per hour had been achieved it would be necessary to have something in the region of 10^{10} unit hours (that is ten thousand million hours, or about one million years) of operation free of dangerous failure. This may be difficult to achieve.

Beware of thinking that in demonstrating a system of a given SIL for one application it will be effective in another application that calls for a system of the same SIL. The safety of an item is application-specific and a single, seemingly trivial or even unrecognised, variation in the design or use between the new and old applications can have considerable safety implications. Reuse can be dangerous so there is an increasing need for care.

4.3.2 Software

If software is already in use, it may be possible to collect some evidence for its Safety Integrity from its service record. It may be possible to make a direct claim for the frequency of hazardous software[5] failures without recourse to SILs from records of its operation in service, provided that there is evidence of the following:

- The records of failures are thorough and accurate
- The software is under change control and the version for which the claim is made is substantially the same as the versions for which records were kept
- The software will be subject to a similar pattern of use (duty cycle) to that for which records were kept
- The total time in operation of the software is known

[5] And in fact hazardous hardware failures.

The data used needs to be either complete or a statistically valid subset. Any bias in the collection of data will invalidate conclusions drawn from it. The data needs to include information about the environment that the system was operating in, and the manner in which it was being used. If basing part of a safety argument upon such data, it is necessary to demonstrate that the data used is of a high enough quality. This may require that the party providing the data also provides details of the collection method, sampling techniques and storage regime.

For relatively simple systems, it may also be possible to make a direct claim for the frequency of hazardous software failures, without recourse to SILs, from records of testing, provided that:

- The test inputs were random; and
- In use the software will be subject to a similar pattern of use to that for which it was tested.

However, it is not generally statistically valid to claim that the mean time between a hazardous failure is more than one third of the total duration of use or testing for which records were kept, and then only if no hazardous failures occurred. It is also impossible to prove the absence of faults - it would require an impracticably long time to derive high confidence in reliability from testing alone. In practice it is difficult to make claims for a Safety Integrity better than SIL 2 using a service record or testing data and even that can be challenging.

4.3.3 Example

The following example is taken from iESM AN5 for a train door interlock which is specified to meet the following safety requirement:

"The door release and traction enable functions of the door controller shall be designed and constructed to achieve safety integrity of SIL 4."

The demonstration of achievement is done in two parts:

- It is shown, using Fault Tree analysis that the rate of random failures, which in the case of the door controller means hardware failures, is less than or equal to 10-9 per operating hour. That meets the safety obligation in the AN5 example but could also have shown to be lower than a specified THR.
- Recognizing all systematic failures of the door controller are failures of the embedded software, it is shown that the software has been developed to deliver a safety integrity of SIL 4 for the traction interlocking function. This is done by selecting appropriate processes and techniques from EN50128 [50128] and providing evidence including the competence and experience of the people.

4.4 Making a Safety Argument

Applying the designated SIL processes and techniques is necessary but not sufficient to declare that a system containing software is 'safe' for a particular application.

Many other things may be needed for any safety functions. Where a few particular functions of a system such as a control centre do have safety integrity requirements, these are best segregated or delivered using such techniques as 'click and confirm', which use operator actions and feedback from the high integrity interlocking to confirm correct operation.

It is worth remembering that humans are prone to error and cannot be relied upon to perform accurately at an equivalent level to any SIL. Thus to require or try to demonstrate a SIL for anything that simply passes on a human command is usually pointless. Techniques like 'click and confirm' are good in that they potentially reduce human error rates by forcing confirmation (within a limited time frame) of a required action, giving a chance to correct any error.

Further, what is thought to be a tolerable level of risk at the specification stage of a project is not necessarily what is (or would be) deemed to be acceptable later. Where a risk acceptance principle such as As Low As Reasonably Practicable (ALARP) is used, the risk must be reduced not merely to a pre-defined level but to a level which is demonstrably as low as reasonably practicable. This may not be obvious until the design stage is undertaken.

Here is an example of a robust and carefully worded claim:

“The requirement was for a SIL X function, and good practice required that we adhered to the standard's processes for SIL X. In doing so, we have generated the evidence appropriate to a SIL X function and independent assessment of the evidence has found that we have adhered to the defined processes.”

5 OTHER INDUSTRIES

5.1 Generic Programmable Electronic Systems

IEC 61508 [61508] is the generic industrial standard for programmable electronic systems and IEC 61511 [61511] is the process control industry application of it. It, and its related standards, are widely and successfully used in high hazard industries. The EN set of CENELEC Railway Application standards derives from it but there are some important differences, particularly in the area of SIL.

In outline the IEC 61508 approach follows the iESM process closely through the definition of a system, the identification of hazards, risk assessment and control. The difference arises if it is established that the resultant risks are too high. Then a Safety Instrumented System (SIS) is required by IEC61508 [61508]. The SIS monitors the main process activity (the Equipment Under Control (EUC) is the terminology in the standard) and intervenes if it detects an unsafe output. Typically it consists of a sensor, a processor and an actuator to detect, decide and act if necessary. This is not a typical architecture of a railway system which tend to rely on single channel highly dependable products and systems or identical parallel processes (e.g. 2oo2) with comparison. An Automatic Train Protection (ATP) system is an example of a similar protection concept in railways but it is not specified in the same way as an IEC 61508 [61508] SIS.

The SIS plays a vital role in providing a protective layer around the Equipment Under Control, rather like a watchdog does in some railway systems. It may take the form of an emergency or safety shutdown system or a safety interlock known as the Safety Instrumented Function (SIF). Its purpose is to take the EUC to a "safe state" when pre-determined events happen or when safe operating conditions have been exceeded. It has a very clearly defined safety function to perform (for example "turn off"). The level of integrity required of the SIF will depend on the level of risk that arises from the EUC (both the EUC itself AND the SIS must fail dangerously before a hazard exists). For example, if the tolerable risk is deemed to be 10^{-9} dangerous failures per hour, and the EUC is calculated to have a probability of 10^{-2} dangerous failures per hour, the "missing" safety – the difference between the tolerable risk and that calculated, must be achieved by one or more safety functions. There also needs to be a strong argument that the dangerous failures of the EUC and those of the SIS are independent. The SIS may operate continuously like a car's steering or 'on demand' like a car's air bag. IEC 61508 [61508] defines the levels of risk reduction that may be achieved by applying each set of SIL processes and techniques.

It is noteworthy that a "SIL 4" SIS would be unacceptable in most non-railway industries. If the Equipment Under Control creates a level of risk needing that much risk reduction then the main process would need to be redesigned to reduce its inherent risk (or be abandoned).

IEC 61508 is therefore not easily applied to systems in which the control and protection functionality are combined, as is often the case in railways. Operational railway safety performance is usually assured by the design of the system as a whole, not by an identifiably separate "protection" element or SIS.

Other systems found on railways, typically using more standard industrial equipment, for example fire alarm or ventilation systems may be more amenable to being treated as described in IEC 61508.

Note that the figures associated with each SIL in IEC61508 differ from those in the CENELEC Railway Application standards. The figures relate to a single PES implementing a single SIF. On railways it is more common for a single system to implement several safety functions.

5.2 Motor Industry

In the motor industry ISO26262 defines an Automotive SIL (ASIL) from A to D (the highest). This is based on the consequence of failure of the system in question. For motor vehicles, the ultimate consequence of a system failure (in terms of accidents and their possible outcomes) is speculative, so in the guideline the consequence of failure is defined in terms of something more predictable — the controllability of the vehicle by its driver. So the guideline advocates that designers carry out a hazard identification and analysis and determine the worst possible result of the failure of their system in terms of the controllability of the vehicle. Five levels of uncontrollability are defined and ASIL values and qualitative acceptable failure rates are defined for them.

It is usually much simpler to define the consequences of a hazard on a railway system using some form of Event Tree Analysis. ISO26262 is therefore not immediately useful for railways.

6. REFERENCED DOCUMENTS

This section provides full references to the documents referred to in the body of this document.

[50126]	EN 50126, Railway Applications – The Specification and Demonstration of Dependability, Reliability, Availability, Maintainability and Safety (RAMS). Also issued as IEC62278. *At the time of writing the current issue of EN 50126 was dated 1999 but the standard was being revised to cover the scope of EN 50128 and EN 50129 and other railway systems. As far as we can, we have aligned this guidance with the emerging issue. If an issue of EN 50126 dated later than 1999 is available at the time of reading then this issue should be consulted. If no issue of EN 50126 dated later than 1999 is available then the reader may find it useful to consult the current issues of EN 50126 and EN 50129 but may not find the information referred to in any particular citation of the standard.*
[50128]	EN 50128:2011 Railway Applications - Software for Railway Control and Protection Systems
[50129]	EN 50129: 2003, Railway Applications – Communication, Signalling and Processing Systems — Safety Related Electronic Systems for Signalling
[50159]	EN50159:2010 Railway Applications, Communication, signalling and processing systems, Safety-related communication in transmission systems
[50451]	PD CLC/TR 50451:2007 Railway Applications. Systematic Allocation of Safety Integrity Requirements
[61508]	ISO / IEC 61508:2010, Functional Safety of Electrical/Electronic/Programmable Electronic Safety-related Systems
[61511]	ISO / IEC 61511:2016 Functional safety - Safety Instrumented Systems for the Process Industry Sector.
[62443]	ISA/IEC 62443:2010 Security Technologies for and Control Systems (formerly issued as ISA99:2007)

Note: This revision (Issue 1.2) of the Application Note has not modified any of the technical content present in the previous revision. Some of the standards referenced may have been revised. A full technical review is planned to be undertaken of this Application Note prior to its next revision.

iESM

BROUGHT TO YOU BY ARC

international Engineering Safety Management

Published on behalf of the international railway industry
by Abbott Risk Consulting Ltd.
Issue 1.2 May 2022

arc