

Part 3

Change Fundamentals

Chapter 7

Defining Changes

Before starting work on a change, your organisation must define the aims, extent and context of the change.

7.1 Guidance from volume 1

This is often done in a requirements specification.

If you are in doubt about the aim, extent or context of the change, you will also be in doubt about claims for its safety.

When you define a change you should also find out which authorities will have to approve your safety case.

7.2 General

Understanding the aims, extent and context of a change is fundamental to successful ESM. Any change to the railway can be regarded as introducing a new *system*, or changing an existing one. Understanding the boundary between this system and its environment is a pre-requisite to understanding how the system might contribute to an accident (that is understanding what its *hazards* are).

Figure 7-1 illustrates the relationship between the system boundary, hazards and accidents.

The system or equipment may consist of software, hardware, people and procedures.

The environment consists of any anything that could influence, or be influenced by, the system or equipment. This will include anything to which the system connects mechanically, electrically or by radio but may also include other parts of the railway that can interact through electromagnetic interference, or thermal interchange. The environment will also include people and procedures that can affect, or be affected by, the operation of the system or equipment.

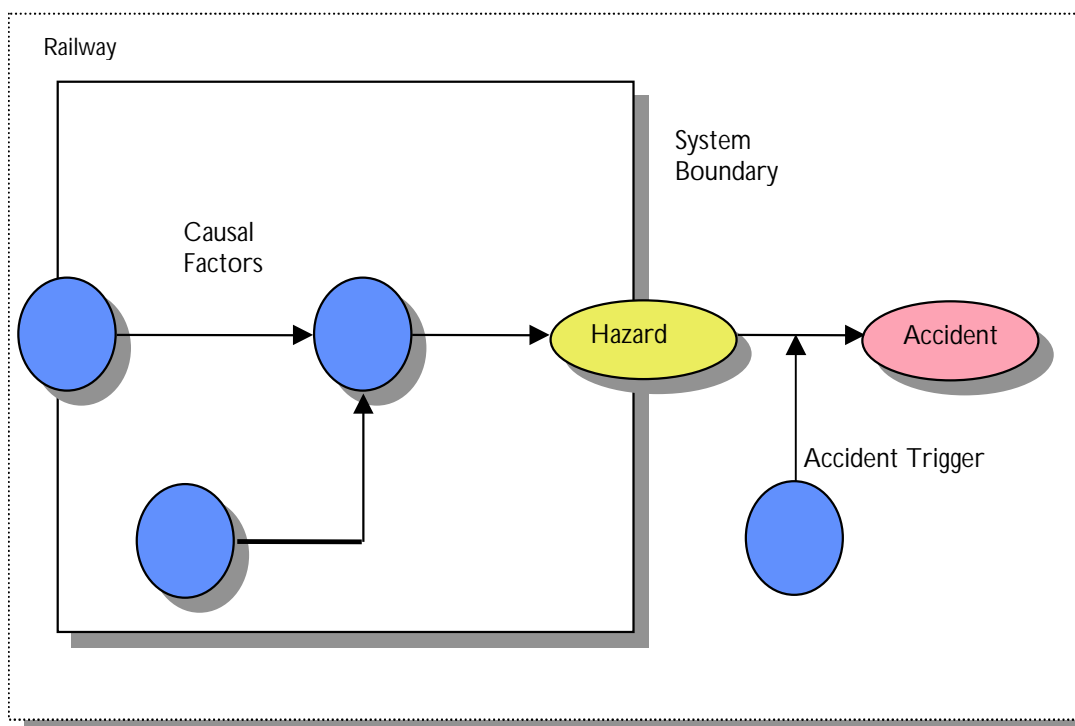


Figure 7-1 - The system boundary in safety analysis

The aims, extent and context of a change should already be defined in a requirements specification but, if they are not, you should clarify them before starting to proceed with safety analysis. If there is not sufficient information available to completely define the change, then explicit assumptions should be made. These assumptions will need to be confirmed at some later stage in the lifecycle of the change and this confirmation should be planned.

The aims, extent and context may change during the life of the system or equipment. You should monitor them for change and, if they do change, you should review all affected ESM activities and rework them as necessary.

This chapter is written for:

- Project Managers, and
- anyone involved in performing or reviewing a risk assessment.

7.3 The supplier chain

Any railway involves a network of stakeholders. The ultimate services to the public are provided by railway operators, including train operators, station operators and infrastructure controllers. However they rely on suppliers in order to do this, their suppliers rely on other suppliers and so on.

It may be the case that the overall safety of the railway depends upon the weakest link in this chain.

Figure 7-2 shows an example of this state of affairs. A train operator relies on a train supplier to provide them with trains. The train supplier, in turn relies on other companies to supply train components, such as the driver's display. Of course this is just a small fragment of a much more complex network of suppliers.




Figure 7-2 - A railway supply chain

There is a hierarchy of systems associated with this network of suppliers, as illustrated in Figure 7-3. This shows that System B (a train, perhaps) is part of the railway as a whole and System A (the driver’s display, perhaps) is part of System B.

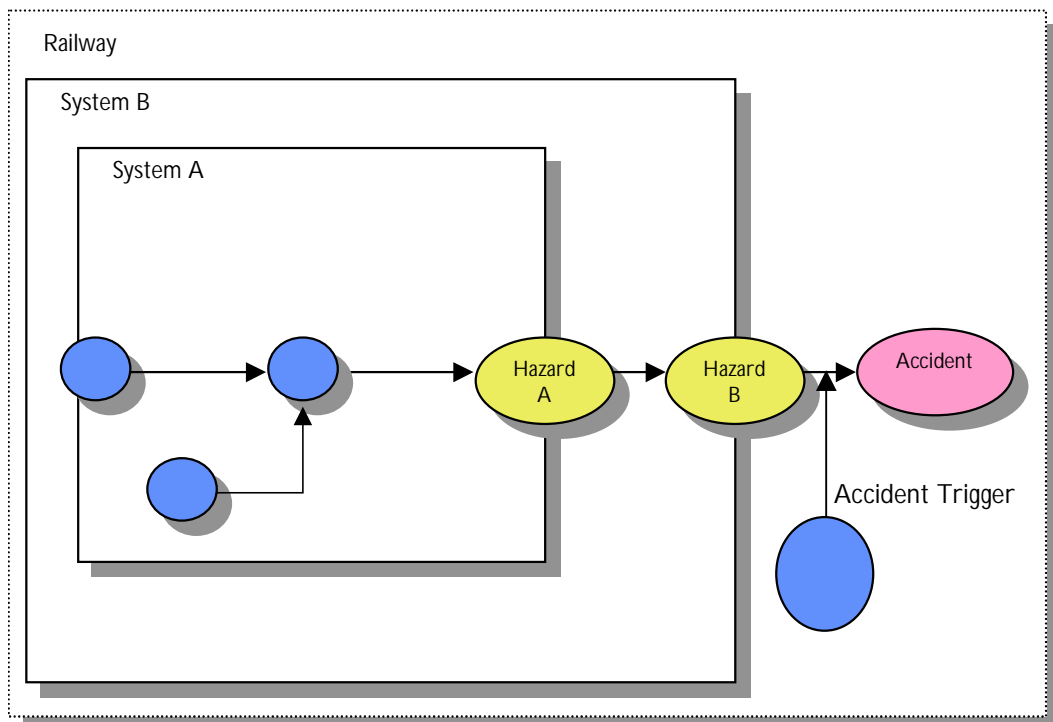


Figure 7-3 - The system hierarchy

The suppliers of both system A and system B need to carry out ESM but they will use different system boundaries and, as a result, concentrate on different hazards.

System B provides the environment for system A. The supplier of system B should, therefore, provide the supplier of system A with information that the latter needs to carry out ESM, including relevant hazards, risks and safety requirements associated with System B. This is discussed further in chapter 6.

7.4 Product development

A product manufacturer may not know all the environments in which their product may be considered for application. In general, they proceed by making informed assumptions (from their own knowledge and by talking to likely customers) about the environment that their product will experience (see Figure 7-4). These assumptions should be made explicit and written down. When it comes to preparing a Safety Case for a specific application, a large part of the work required will be to check that these assumptions hold in the application in question.

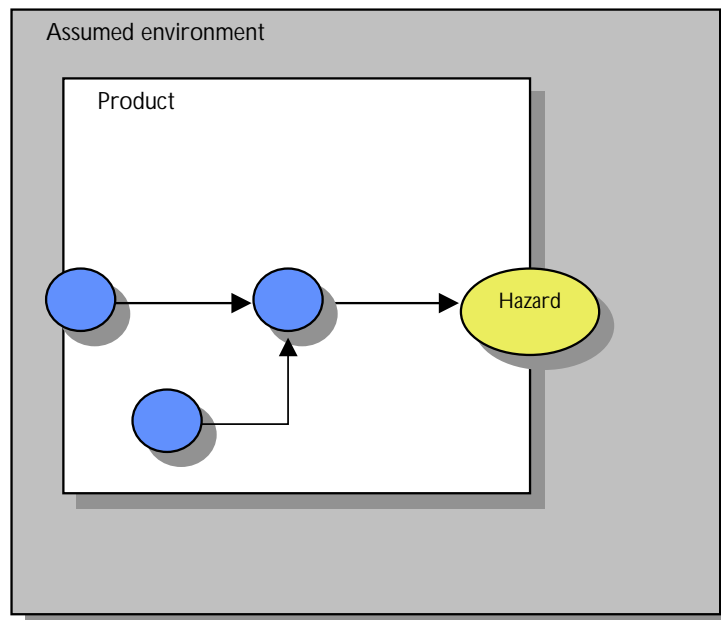


Figure 7-4 - Product development

7.5 Related Guidance

Chapter 6 provides guidance on the information that should be provided to suppliers, to allow them to carry out effective safety analysis.

Chapter 8

Identifying Hazards and Assessing and Reducing Risk

When your organisation considers change, it must make a systematic and vigorous attempt to identify any possible hazards. Your organisation must consider hazards which could contribute to an accident at any time, from introducing the change into the railway to removing it.

Your organisation must assess the effect of any proposed change on overall system risk.

Your organisation must carry out a thorough search for measures which reduce overall system risk, within its area of responsibility. It must decide whether each measure is reasonably practicable and, if so, must take it.

If your organisation finds that risk is still intolerable, it must not accept it.

8.1 Guidance from volume 1

8.1.1 Identifying hazards

Identifying hazards is the foundation of ESM. If you do not identify a hazard, you can take no specific action to get rid of it or reduce the risk relating to it. However, you may be able to take general actions, such as introducing safety margins.

You should not just consider accidents which might happen during normal operation, but others which might happen at other times, such as installation, track-testing, commissioning, maintenance, emergencies, decommissioning and disposal.

You should consider the people who the change will affect, and design it to help them avoid mistakes.

When identifying hazards, you should consider all the effects of the change on the rest of the railway and its neighbours.

You may identify a possible hazard which you believe is so unlikely to happen that you do not need to do anything to control it. You should not ignore this type of hazard; you should record it together with the grounds for your belief that it is so unlikely to happen.

8.1.2 Assessing Risk

There are legal duties to assess risk.

Risk measures the likelihood that an accident will happen and the harm that could arise. You should consider both factors. Your organisation should also consider *who* is affected.

Some changes are made specifically to make the railway safer, that is to *reduce* risk, at least in the long run. You should still assess them in case they introduce other risks.

8.1.3 Reducing Risk

If the risk is in the broadly acceptable region, you need only consider measures which are clearly reasonably practicable.

There are legal duties to do this.

You should look for:

- ways to get rid of hazards or to reduce their likelihood;
- ways to contain the effects of hazards, if they happen; and
- contingency measures to reduce harm if there is an accident.

You should look for ways of controlling both hazards introduced by the change itself and hazards that are already present in the railway. Even if a change is designed to make the railway safer then you should still see if there are ways that you could make the railway even safer.

8.2 Background

Most railway changes are associated with risk, that is the potential for harm to people. The risk associated with a change can vary from negligible to totally unacceptable.

Risk can generally be reduced, although usually at a cost.

Risk assessment entails a systematic analysis of the potential losses associated with a change and of the measures for reducing the likelihood or severity of loss. It enables losses to be aggregated and compared against the cost of measures.

Risk assessment is tightly coupled with hazard identification and risk reduction. The hazards of a system have to be identified before an accurate assessment of risk can be made. Risk assessment provides, throughout the lifecycle of a system or equipment, both input to risk reduction and feedback on its success.

This chapter presents a single, systematic framework for:

- identifying hazards,
- assessing risk, and
- reducing risk.

The next section provides some further background.

The following sections describe a seven-stage process for hazard identification, risk assessment and risk reduction.

This chapter is written for:

- anyone involved in performing or reviewing a risk assessment.

8.3 Underlying concepts

8.3.1 Concepts and terminology

Risk assessment requires an understanding of potential **Accident Sequences**, the progression of events that result in accidents.

An **Accident** is an unintended event or series of events that results in harm.

A **Hazard** is a condition that could lead to an accident.

Hazards arise from events or sequences of events such as **Failures**, that is, when a system or component is unable to fulfil its operational requirements. An accident sequence may be represented as follows:




Figure 8-1 – Accident sequences

However not every failure results in a hazard and not every hazard results in an accident. Fault tolerant mechanisms may mean that more than one failure is required before a hazard occurs. Similarly, hazards may not result in accidents due to the action of mitigating features.

Failures may be classified into two types:

- **Random.** Failures resulting from one or more of the possible degradation mechanisms in the hardware. These failures occur at predictable rates but at unpredictable (that is random) times.
- **Systematic.** Failures related in a deterministic way to a certain cause, which can only be eliminated by a modification of the design or of the manufacturing process, operational procedures, documentation or other relevant factors.

The distinction is made between random and systematic in order to establish targets for failure mechanisms in the system. Random failure targets can be decomposed as numerical requirements through mathematical methods. However, systematic failure targets are divided into four bands and, for each band, a level of design processes and requirements is defined to reduce the risk until acceptable. These levels are called **Safety Integrity Levels (SILs)** and are discussed further in chapter 9 on Safety Requirements.

Note that SILs are not the only means of controlling systematic failures; they may be controlled through architectural design features as well.

Risk is defined to be the combination of the likelihood of occurrence of harm and the severity of that harm.

The **individual risk** experienced by a person, is their probability of fatality per unit time, usually per year, as a result of a hazard in a specified system.

8.3.2 UK Law and the ALARP principle

We have seen that the '*Health and Safety at Work Act (1974)*' places duties on employers to ensure health, safety and welfare 'so far as is reasonably practicable'. This section gives more guidance on this test. It is based on the HSE publication '*Reducing Risks, Protecting People*' [F.6].

If you are working on a change to the railway, you should first identify the hazards associated with the change. You should make sure that you have precautions in place against each hazard within your control, unless you can show that the risk arising from the hazard is negligible.

You should make sure that your precautions reflect good practice, as set out in the law, government guidance and standards. If the risk is low and completely covered by authoritative good practice, showing that you have followed it may be enough to show that the risk is acceptable. For instance the electrical safety of ordinary office equipment is normally shown by certifying it against electrical standards. However, before you decide that just referring to standards is enough, make sure that:

- the equipment is being used as intended;
- all of the risk is covered by the standards; and
- the standards cover your situation.

If following good practice is not enough to show that the risk is acceptable, you should also assess the total risk that will be produced by the part of the railway being changed. You then need to compare it with two extreme regions.

- An unacceptable (or intolerable) region where risk can never be accepted.
- A broadly acceptable region where risk can always be accepted.

To decide whether or not to accept a risk:

- 1 check if the risk is in the unacceptable (or intolerable) region – if it is, do not accept it;
- 2 check if the risk is in the broadly acceptable region – if it is, you will not need to reduce it further, unless you can do so at reasonable cost, but you must monitor it to make sure that it stays in that region; and
- 3 if the risk lies between these two regions, accept it only after you have taken all 'reasonably practicable' steps to reduce the risk.

Figure 8-2 illustrates the principle described above. This is often referred to as the **ALARP principle**, because it ensures that risk is reduced 'As Low As Reasonable Practicable'.

You should consider ways of making the change less likely to contribute to an accident. You should also consider ways of making the change more likely to prevent an accident. You do not have to consider steps that are outside your control.

You will generally expect the risk to be lower after the change than it was beforehand; if it is higher, it is unlikely that you have reduced risk as low as reasonably practicable.

If you are uncertain about the risk then you should err on the side of caution – uncertainty does not justify inaction.

The principle should be interpreted intelligently. Sometimes it may be necessary to accept a modest increase in risk in the short term to achieve sustained decrease in risk in the long term.

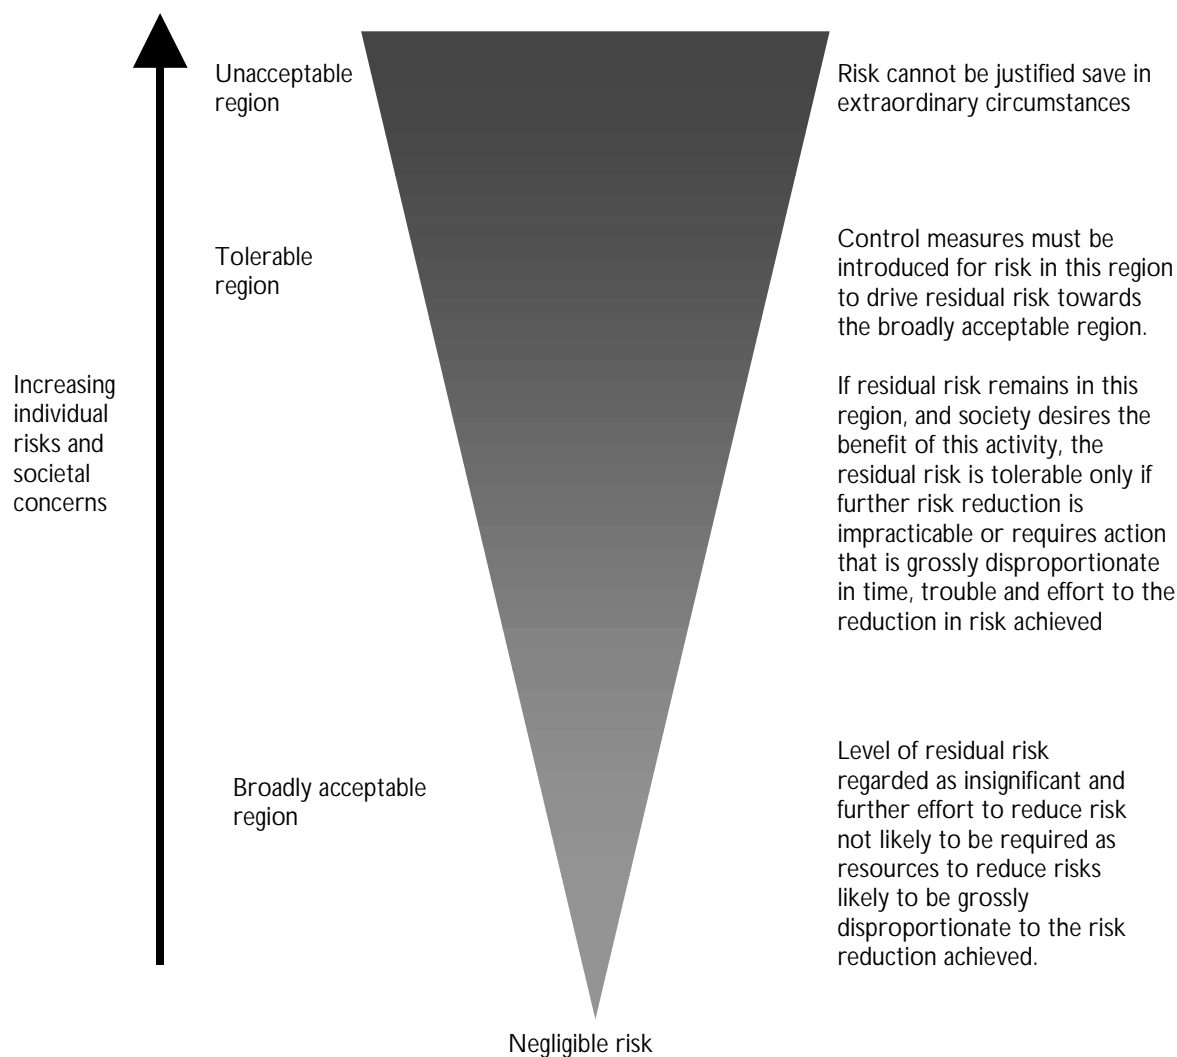


Figure 8-2 – The ALARP Principle

There are requirements to assess risk as well as to reduce it. The *'Management of Health and Safety at Work Regulations (1992)'* are made under the *'Health and Safety at Work etc Act (1974)'* and have the force of law. They require employers to perform *'suitable and sufficient'* assessment of safety risks to all people exposed to the hazards of an undertaking.

To be suitable and sufficient, the sophistication and depth of risk assessment should be proportionate to the level of the risk.

8.4 The seven-stage process – general remarks

8.4.1 Overview of process

The seven-stage process, depicted in Figure 8-3 will form the basis of the guidance in this section.

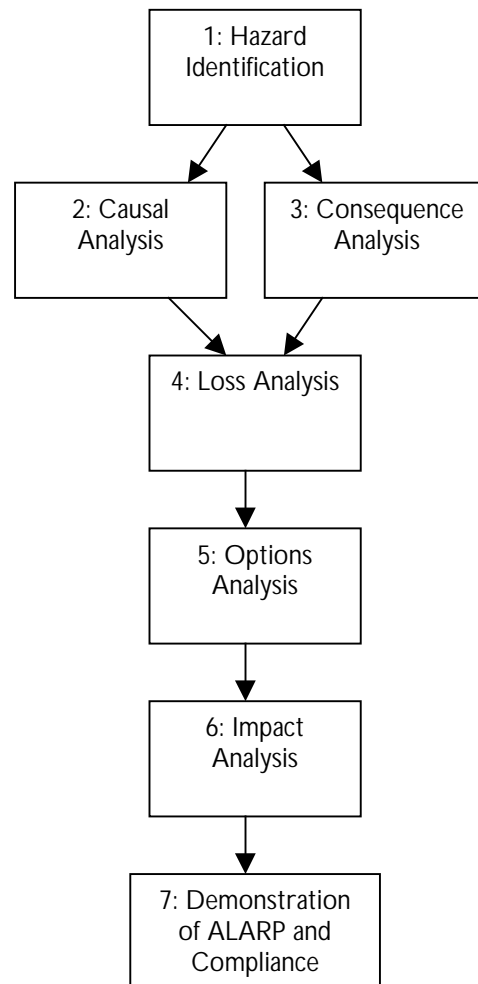


Figure 8-3 – Risk assessment stages

This seven-stage process is the approach recommended by this book. There are alternative, effective techniques.

Hazard Identification involves identification and ranking of hazards.

Causal Analysis involves establishing the primary causal factors which may give rise to a hazard and estimating the likelihood of occurrence of each hazard.

Consequence Analysis involves establishing the intermediate conditions and final consequences, which may arise from a hazard, and estimating the likelihood of accidents arising from each hazard.

Causal and Consequence Analysis may be undertaken in parallel.

The consequences of each hazard may be associated with a range of losses (that is harm to people, damage to the environment or commercial detriment). **Loss Analysis** requires estimation of the magnitude of the safety losses (that is harm to people), before considering options to reduce risk.

Risk reduction and control requires identification of a range of potential risk reduction measures for each hazard. **Options Analysis** comprises determination of such measures and assessment of their implementation costs.

Impact Analysis involves assessing the net benefits associated with implementation of each risk reduction measure, in terms of the reduction in risk. This is achieved by revising the previous stages to allow for the effects of the measure.

Demonstration of ALARP and Compliance involves determining which risk reduction measures should be implemented and justifying the acceptance of any remaining risk. This is done by selecting those that are required by the ALARP principle or by safety targets imposed by the railway operator.

8.4.2 Scope of application

If you are faced with a decision that involves risk, you will generally have to do two things:

- 1 Establish the facts on which you have to take a decision – what the hazards and risks are. This is generally a technical and objective process.
- 2 Establish and apply decision criteria to the facts. These are always based upon values and hence have a subjective element.

The seven-stage process provides a generally-applicable framework for the first stage and a framework for applying certain published decision criteria to justify a claim that risk has been reduced ALARP. However you should be prepared to tailor it to your specific situation.

To understand the sort of tailoring that may be required, it is convenient to refer to some definitions from the UK Offshore Operators Association's *Industry Guidelines on a Framework for Risk Related Decision Support* [F.7]. This document explains how risk related decisions can be placed in a spectrum running from:

- **technology based** decisions for risks that are well understood, uncontroversial and with low severity consequences; to
- **values based** decisions where there is significant novelty, public concern or potential for catastrophic consequences.

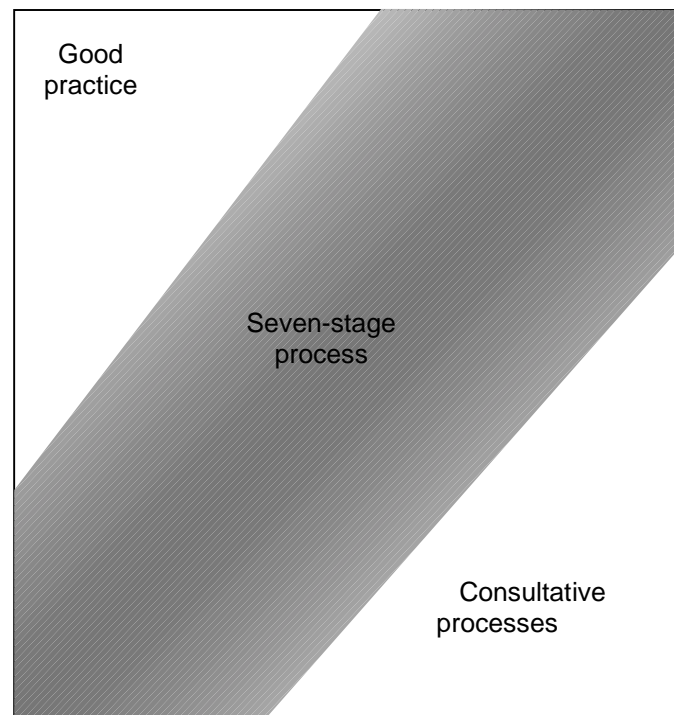
If you are faced with decisions towards the technology based end of the spectrum, you can replace some of the stages in the seven-stage process with reference to authoritative good practice (see section 8.3.2). Essentially the good practice embodies the results of analysis that has already been done which you do not need to repeat.

Even if you use the full seven-stage process, you will still want to show that you have used good practice, unless you have moved so far from the technology based end of the spectrum that there is no established good practice for what you are doing.

As you move towards the values based end of the spectrum, you are likely to find that the process of establishing the facts becomes an increasingly smaller part of the problem, and that establishing decision-criteria becomes the larger part. For these decisions, you will need to supplement the seven-stage process with significant additional activities to consult stakeholders in order to arrive at justifiable decisions.

Figure 8-4 illustrates the parts that good practice, the seven-stage process and stakeholder consultative processes might play in different sorts of decisions. The width of each band gives a rough indication of the relative significance of each type of activity.

Technology Based



Values Based

Figure 8-4 – Approaches to different risk decisions

8.4.3 Quantitative and qualitative analysis

The seven-stage process presents a uniform framework for assessment of the full range of risks associated with any given undertaking. Within this framework, the analysis may be performed to different depths. Qualitative risk assessment is appropriate for the smaller risks and quantitative risk assessment for the larger risks. It is also possible to adopt hybrid approaches.

It is acceptable, in both approaches, to adopt approximations provided that they are *conservative*, that is that they do not under-estimate risk.

Qualitative risk assessment relies mainly upon domain expert judgement and past experience. It addresses the risks of an undertaking in a subjective and coarse manner. There is not a complete lack of quantification but order of magnitude estimates are generally used. Its advantages are that:

- it does not require detailed quantification, data collection or analytical work,
- it is relatively simple, and
- it is less expensive than quantitative risk assessment.

Its disadvantages are that:

- the assumptions require thorough documentation, and

- it is inadequate as the sole basis for assessment of major risks, including those arising from low loss incidents of high frequency, as well as from low frequency incidents associated with high losses.

Quantitative risk assessment employs rigorous analytical processes. Whilst based upon the same fundamental principles as qualitative risk assessment, quantitative risk assessment will typically employ modelling, using objective and validated data; explicit treatment of the uncertainty associated with input data; and explicit treatment of the dependencies between significant factors contributing to risk.

Its advantages are that:

- it is more accurate than qualitative risk assessment,
- it helps identify hidden assumptions, and
- it provides a better understanding of the potential causes and consequences of a hazard.

Its disadvantages are that:

- it is complex,
- it requires expertise,
- it requires a lot of objective data,
- it is difficult to quantify the probability of systematic failures,
- it is more expensive than qualitative risk assessment, and
- it can require significant computing resource.

Qualitative risk assessment is likely to suffice for most hazards. However, hazards, with the potential to lead to major or catastrophic consequences, may require quantitative risk assessment. A quantitative approach may also be justified for novel systems where there is insufficient experience to support an empirical, qualitative approach.

Quantitative risk assessment is more expensive than its qualitative counterpart and should only be applied if it is justified by the increased confidence achieved.

8.4.4 Iteration and preliminary hazard analysis

Safety analysis is iterative: as the design progresses, the analysis should be repeated to take account of change and extended to cover the extra detail. The design can then be modified to avoid hazards or reduce risks as soon as they are identified. The process should start as soon as a high-level description of the system is available.

A **preliminary hazard analysis** should be carried out at the start of a project to determine a measure of the scope and extent of the risk presented by the change.

Preliminary hazard analysis is a first-pass hazard identification and risk assessment intended to determine:

- a) the scope and extent of risk presented by a change, so that ESM may be applied to an appropriate depth;
- b) a list of potential hazards that may be eliminated or controlled during initial design activity.

At the start of a project, design detail will almost always be limited, so the results of preliminary hazard analysis (in particular the depth of application of ESM) should be backed up and re-assessed by carrying out a full analysis and risk assessment as soon as detail is available.

Preliminary hazard analysis should be carried out before any significant design activity begins. It requires a full high-level description of the system's function and construction and its interfaces to people and other systems.

The risk assessment activity carried out during preliminary hazard analysis should consist of annotating identified hazards with an initial appraisal of their severity and likelihood. Ideally, the preliminary hazard analysis should support the process of initial safety requirements setting and, therefore, should provide targets for the likelihood of each of the identified hazards.

The results of the preliminary hazard analysis should be used to decide where further quantified analysis is required.

The findings of preliminary hazard analysis and the decisions that result should be documented in a report.

8.4.5 Use of historical data

Risk assessment always relies on some form of extrapolation from the past to the future. Historical data is used at many stages but it should be used with care. The reasons for this include the following:

- Insufficient information may be available to determine whether historical figures are relevant to the circumstances of concern, particularly regarding rare major or catastrophic accidents and the circumstances surrounding previous incidents.
- Secondary effects arising from an incident are likely to be difficult to reliably determine (for example fires, derailment or exposure to harmful substances).

Inappropriate use of historical data can undermine the analysis, and significantly reduce the accuracy of risk assessment.

Where historical data is employed in an assessment, a clear argument should be presented that its use provides an accurate forecast of the losses associated with the particular circumstances under study.

8.4.6 Documenting the process

Typically, the results of a risk assessment study will be compiled into a risk assessment report so that they can be subject to review and endorsement.

Once risk assessment results have been reviewed and endorsed they should be immediately incorporated into the Hazard Log which is described in chapter 13.

8.4.7 Division of work

The seven-stage process provides an overall framework for controlling risk and demonstrating compliance with legal obligations. In practical application it is often the case that different parts of the process are performed by different organisations.

Any change to the railway can be regarded as introducing a new system or changing an existing one.

Performing the entire process requires expertise on both:

- the system, its function and design, and
- the railway environment in which the system will run.

Typically the former expertise is provided by the **system supplier** and the latter expertise is provided by the **railway operator**, that is the infrastructure controller, train operator or station operator. Table 8-1 shows the typical division of responsibilities, across the steps.

As a result of the analysis performed, the railway operator will typically define **tolerable hazard rates** for common applications of common systems, that is maximum acceptable rates for the occurrence of these hazards which are consistent with their legal and regulatory constraints and corporate safety objectives.

Step	Railway operator activities	System supplier activities
Hazard Identification	Provides initial hazard list	Confirms and extends hazard list
Causal Analysis	Reviews analysis	Performs analysis
Consequence Analysis	Performs analysis	Reviews analysis
Loss Analysis	Provides initial modelling data	Performs analysis
Options Analysis	Reviews analysis	Performs analysis
Impact Analysis	Provides initial modelling data	Performs analysis
Demonstration of ALARP and Compliance	Derives acceptable/tolerable hazard rates	Demonstrates achievement of acceptable/tolerable hazard rates Demonstrates ALARP

Table 8-1 - Division of work

All parties work within overall safety targets and criteria set by the **railway authority**, the body accountable to the safety regulator for the safety of the railway.

8.4.8 Using likelihood-severity matrices to simplify repeated assessments

If you have to carry out a series of risk assessments of applications of a system which are similar, then you may find that a **likelihood-severity matrix** can save repeating the same work. The matrix may be produced by the railway operator or by the system supplier from information provided by the railway operator or railway authority.

A likelihood-severity matrix has the following general format:

Likelihood	Severity			
	Insignificant	Marginal	Critical	Catastrophic
Frequent				
Probable				
Occasional				
Remote				
Improbable				
Incredible				

Table 8-2 - Example format of likelihood-severity matrix

Table 8-2 is only an illustrative example. It shows the column and row headings suggested in EN 50126 [F.8]. Other headings may be used. See for instance the guidance provided in Railway Group Standard GK/RT0206 [F.9] and Railway Group Code of Practice GK/RC0701[F.10].

The two components of risk – frequency (or likelihood) and consequence (or severity) – are partitioned into broad order or magnitude categories which are then used to index the rows and columns of a matrix. Each cell within the matrix then represents a broad region of risk. The example above is empty but, in a real matrix, a risk acceptability category is written into the cell.

It is not possible to create one general-purpose matrix that will suit all railway applications. A matrix should be designed with likelihood, severity and risk acceptability categories that are appropriate to the situation in hand. The matrix should be associated with:

- definitions of the likelihood, severity and risk acceptability categories used;
- an explanation of how the risk acceptability categories relate to the Intolerable, Tolerable or Broadly Acceptable categories of the ALARP triangle and to any overall safety targets set by the Railway Authority;
- assumptions on which the matrix is based; and about the system, its hazards, its environment, its mode of use and the number of systems in service;
- guidelines for the use of the matrix.

It is common practice to employ three categories (Intolerable, Tolerable and Broadly Acceptable). An additional categorisation may also be found useful, in which the Tolerable category is split into two, one towards the Intolerable end of the range and one towards the Broadly Acceptable end.

Before using the matrix, you should show that it meets all the following criteria:

- If all hazards of the system are assessed as Tolerable then it follows, using the explicit assumptions, that the total risk presented by the system to any affected group of people falls in the tolerability region and is consistent with overall risk targets set by the Railway Authority.
- If all hazards of the system are assessed as Broadly Acceptable then it follows, using the explicit assumptions, that the total risk presented by the system to any affected group of people falls in the broadly acceptable region.

- The matrices can be used to support a justification that risk has been reduced ALARP. The guidelines should emphasise that the final judgement on ALARP relates to the total risk arising from the system as a whole, and, in particular should advise that:
 - Partitioning the risk across hazards and evaluating each hazard against a chosen matrix alone may lead to each hazard being considered as Broadly Acceptable or Tolerable, whereas the total system risk may be in a higher category.
 - The total risk should be reduced so far as is reasonably practicable. So, if the total risk is in the Tolerable region but the classification from one particular hazard is Broadly Acceptable, the risk from this hazard should still be reduced further if it is reasonably practicable to do so.

When using the matrix, you should provide justification of the likelihood and severity categories assigned to each hazard.

To avoid possible later problems with use of the matrices, you should submit the matrix with your justification that it meets these criteria for endorsement by any Safety Authority whom you may later ask to endorse a safety argument using the matrix.

8.4.9 Risk assessment and broader decision making

Risk assessment is focussed on demonstrating compliance with legal safety obligations and these are phrased in terms of harm to people. These obligations place constraints on the alternatives that may be followed. The seven-stage process will assist you in eliminating alternatives which do not comply with your obligations. The seven-stage process can be extended to help control non-safety losses (such as environmental and commercial losses) but that is beyond the scope of this book.

In broader decision making, it is appropriate to consider non-safety losses, such as environmental and commercial harm as well as the opportunities for reaping benefits of many different sorts. Techniques such as Weighted Factor Analysis [F.11] provide a basis for balancing the factors in such decision making.

8.5 The seven-stage process – stage by stage

8.5.1 Stage 1: Hazard Identification

8.5.1.1 Introduction

Before conducting hazard identification, you need to understand the boundary of the system concerned and its interactions with its environment. This is discussed in chapter 7. When performing hazard identification, you should always look out for interactions that have not been identified and which have the potential to be implicated in hazards.

Hazard Identification is fundamental to the risk assessment process. Absence of a systematic and comprehensive Hazard Identification phase can severely undermine the risk assessment process. In the worst case this can create an illusion of safety and a false sense of confidence.

When identifying hazards, you should not restrict yourself to the steady-state operation phase but consider all aspects of the systems lifecycle from the point at which it is installed on the railway to its final decommissioning, including maintenance and upgrade.

Systematic identification of hazards may be performed empirically or creatively.

8.5.1.2 Empirical hazard identification

Empirical hazard identification relies largely upon knowledge and experience of the past to identify potential hazards. Whilst it is sometimes sufficient for routine undertakings, novel or modified undertakings will generally also require a more creative form of hazard identification.

Empirical hazard identification methods include:

- checklists (see appendix C), and
- structured walkthroughs.

The following more rigorous empirical methods may also be used:

- Failure Mode and Effects Analysis (FMEA) for equipment and systems (see appendix E), and
- Task Analysis for man-machine interfaces (see [F.12]).

These latter techniques identify particular component failures or human errors, which may lead to hazardous circumstances. They do, however, require a detailed knowledge of the failure modes of components and sub-systems, including human actions and likely errors.

8.5.1.3 Creative hazard identification

Creative hazard identification methods provide systematic techniques to encourage lateral and imaginative creative thought. Ideally they should employ a team-based approach to exploit the diverse and complementary backgrounds of a range of individuals. They include:

- brainstorming,
- Hazard and Operability Studies (HAZOP) (see appendix E).

Empirical and creative hazard identification complement one another, increasing confidence that all significant hazards have been identified.

8.5.1.4 General remarks

Once identified the hazards should be listed. The record of hazards is usually maintained in a Hazard Log (see chapter 13).

Each hazard is usually associated with several causes. If you have identified a large number of hazards, you should check to see that you have not separately identified multiple causes of a single hazard.

To focus risk assessment effort upon the most significant hazards, the hazards should be ranked. The subsequent stages of risk assessment, as detailed in this document, should be applied on a prioritised basis, beginning with the highest ranking hazards. The relative rank of each hazard should be used to guide the breadth and depth of its further analysis. A simple matrix should be employed. A sample ranking matrix is presented in Appendix D.

8.5.2 Stage 2: Causal Analysis

8.5.2.1 Introduction

Once you have identified and ranked the hazards you should determine those factors contributing to the occurrence of each hazard, in order to:

- enable accurate assessment of the likelihood of occurrence of each hazard; and
- help identify measures to reduce the likelihood of its occurrence.

Causal Analysis requires domain knowledge of the system or equipment. Causal Analysis generally assumes that the design material is organised as a **functional hierarchy** which shows how the overall system is broken down into ever smaller components.

Before the Causal Analysis can be completed, the analyst should have seen a complete set of design material, normally including but not limited to:

- physical drawings of the system,
- component lists, and
- operating and maintenance instructions.

The key factors to consider in the analysis process are:

- identification and modelling of common cause failures,
- interdependency of some errors and failures, and
- the correct logical relationships.

Most Causal Analysis techniques employ a diagrammatic representation of the errors and failures leading to a hazard. This helps to understand and communicate the relationships between the causes of a hazard and is therefore recommended.

Causal Analysis may be done qualitatively or quantitatively.

8.5.2.2 Qualitative analysis

Qualitative Causal Analysis should be done to a depth sufficient to enable a realistic subjective estimate to be made of the likelihood of the hazard. It may not be necessary to go to the level of detail of failures in basic system elements in order to do this.

8.5.2.3 Quantitative analysis

Quantitative Causal Analysis of a hazard should continue until all the fundamental causal factors have been identified, or until there is insufficient reliable data to go further. Fundamental causal factors include basic component failures and human errors.

Accurate quantification of causal models requires an objective assessment of the frequency or probability of occurrence of fundamental causal factors. These are then combined in accordance with the rules of probability calculus to estimate the probability of occurrence of the hazard.

Key issues are:

- obtaining reliable and accurate data;
- appropriate treatment of uncertainty in the data;
- sensitivity analysis; and

- ensuring that different causal factors are combined appropriately to yield consistent results (for example ensuring that two frequencies are not multiplied to yield units in terms of per time squared).

The depth of treatment of uncertainty in data sources should vary according to the nature of the hazard being assessed. For example, consider a hazard with potentially significant consequences. Suppose that a causal factor is identified whose occurrence leads to a high likelihood of realisation of the hazard. Significant uncertainty in estimates of the frequency of the causal factor are likely to result in significant uncertainty in the frequency determined for the associated hazard (and may, in turn, lead to significant underestimates of potential losses). In such cases, further analysis of the likely frequency of the causal factor is warranted.

Quantitative analysis should aim to minimise the significance of uncertainties. The nature and implications of all uncertainties should be carefully documented.

Where the frequencies of causal factors are specified with confidence intervals, accurate estimation of the likely mean and distribution of the frequency of occurrence of a hazard requires use of statistical simulation techniques.

Quantitative Causal Analysis techniques are generally based upon formal mathematical foundations and are supported by computer based tools. However, they cannot generally handle variation in the frequencies of causal factors over time.

Since the causal models are usually generated with the assistance of individual domain experts, they should be subject to peer review in order to enhance confidence in their integrity and correctness.

If a particular hazard occurs frequently, and reliable statistics are available concerning the probability of its occurrence, detailed quantitative Causal Analysis may not be necessary, but it may still be useful in determining the causes of the hazard and helping to identify potential hazard prevention measures.

8.5.2.4 General remarks

Fault Tree Analysis and FMEA are techniques which may be used to perform Causal Analysis, see section appendix E. ENV 50129:1998 [F.13] provides guidance on identifying the failure modes of hardware items which may support these or other techniques.

8.5.3 Stage 3: Consequence Analysis

8.5.3.1 Introduction

In contrast to Causal Analysis, which is aimed at determining the factors which lead to the occurrence of a hazard, Consequence Analysis involves determining the possible effects of each hazard. The results of Consequence Analysis should provide an estimate of the likelihood of occurrence of each incident following realisation of the hazard in order to:

- support accurate assessment of the likely losses associated with a hazard; and
- help identify control measures for the hazard.

Like Causal Analysis, Consequence Analysis is mainly empirical, requiring domain knowledge of the system's environment. It is generally applied to each hazard in a bottom-up manner until all potential consequences (incidents and accidents) have been determined. This leads to identification of several other intermediate states and consequences.

Key issues are:

- developing a clear understanding of the hazard; and
- determining existing physical, procedural and circumstantial **barriers** to the escalation of the hazard.

Most Consequence Analysis techniques employ a diagrammatic representation of the lines of cause and effect and this is encouraged.

Consequence Analysis may be done qualitatively or quantitatively.

8.5.3.2 Qualitative analysis

Qualitative Consequence Analysis should be conducted to a depth sufficient to enable a realistic subjective estimate to be made of the likelihood of occurrence of an incident or accident. As a general rule, the analysis should be continued until all potential incidents and accidents arising from a hazard have been identified.

Note that identifying all barriers to escalation of a hazard may sometimes be used to provide only an understanding of how each incident can arise. It may not be necessary to quantify the probability of success of each individual barrier in order to estimate the likelihood of occurrence of each incident. Rather, it may be possible to make a simple conservative estimate of the likelihood of each incident based upon the understanding gained by consequence modelling.

8.5.3.3 Quantitative analysis

Consequence Analysis techniques typically present the results of analysis in the form of a logic tree structure. Such trees lend themselves to quantification in order to obtain an assessment of the likely frequency of predicted incidents and accidents. Event Tree Analysis and Cause Consequence Diagramming are such techniques. The latter is described in appendix E.

Quantification of consequence trees requires an objective assessment of the probability of success of each barrier to escalation of a hazard (that is an assessment of the barrier 'strength'). Such assessment may be based upon historical data, the results of specific causal analysis or, where no objective data can be obtained, on the basis of expert opinion.

Key issues are:

- obtaining reliable and objective data sources for the assessment of barrier strengths;
- appropriate treatment of uncertainty in the data sources; and
- sensitivity analysis of barrier strengths.

The depth of treatment of uncertainty in data sources should vary according to the nature of the hazard being assessed. For example, consider a high frequency hazard with potentially significant consequences (major incidents or accidents). Uncertainty in the estimate of the strength of a barrier may lead to uncertainty in the likelihood of occurrence of a major incident. In such cases, further analysis of the barrier strength is warranted.

Sensitivity analysis performed upon the barriers to escalation of a hazard can be used to determine those barriers with the greatest effect upon the likelihood of occurrence of incidents. The uncertainty associated with estimates of the strength of such barriers should be reduced where possible. The nature and implications of any uncertainties should be carefully documented.

Where barrier strengths are specified with confidence intervals, accurate estimation of the likely mean and distribution of the frequency of occurrence of adverse incidents requires use of statistical simulation techniques.

In order to meet the above requirements, quantitative Consequence Analysis techniques are generally based upon formal mathematical foundations and are supported by a suite of computer based tools.

The typical disadvantages of such techniques should be noted:

- they are generally incapable of addressing temporal variations in data, applying only if barrier strengths remain constant over time; and
- they are generally incapable of addressing interdependencies between barriers.

Since the consequence models are usually generated with the assistance of individual domain experts, they should be subject to peer review in order to enhance confidence in their integrity and correctness.

8.5.3.4 General remarks

It is important in Consequence Analysis to consider the full range of consequences. Do not assume that because a failure is termed a 'Right Side Failure' that it cannot contribute to an accident. Typically, even right side failures lead to alternative, temporary methods of working which increase risks.

8.5.4 Stage 4: Loss Analysis

8.5.4.1 Introduction

Loss Analysis comprises a systematic investigation of the safety losses associated with all incidents and accidents identified through Consequence Analysis.

Loss Analysis involves assessment of the losses associated with the hazards of an undertaking *before* considering risk reduction measures, leaving the consideration of the effect of these measures to later stages.

The losses associated with a system should be aggregated for all hazards of the system. The safety losses experienced by different groups of people (for instance passenger and trackside workers) should be aggregated separately for each group.

Loss Analysis may be carried out qualitatively or quantitatively.

8.5.4.2 Qualitative analysis

Safety losses should be estimated in terms of **Potential Equivalent Fatalities** per annum. In other words, all safety losses should be converted into an equivalent annual fatality figure. The current convention is as follows:

- 1 fatality = 10 major injuries
- 1 major injury = 20 minor injuries

For example, if 1 major injury is estimated as arising from a hazard (over a year), this equates to 0.1 Potential Equivalent Fatalities.

8.5.4.3 Quantitative analysis

In order to convert safety losses into monetary values an indication of what it is reasonably practicable to spend to reduce risk by one fatality is required. Such a figure is often referred to as a **Value of Preventing a Fatality (VPF)**. The VPF is a parameter intended for supporting ALARP decisions only. It is not an estimation of the commercial loss that might follow from such a fatality and so cannot be used for purposes such as arranging insurance cover.

The total Potential Equivalent Fatalities per annum is multiplied by the VPF to yield a monetary loss per annum, for decision making purposes.

VPFs are generally set by railway operators. In *'Reducing Risks, Protecting People'*, HSE suggests that a benchmark of slightly under £1M (at 1998 prices) can be used in some cases. However, a higher figure should be used for risks for which there is high aversion. As risks of major railway accidents fall into this category, the VPFs used in railway decision making are often higher.

Be aware that all benchmarks are only rough reflections of the values held by society at large. If there is significant public concern about a hazard then you should take this into account in your decision making and it may justify precautions that would not be justified otherwise.

8.5.5 Stage 5: Options Analysis

Options Analysis determines options to reduce the associated losses determined during Loss Analysis. These options can typically be divided into:

- those aimed at reducing the rate of occurrence of a hazard;
- those aimed at limiting the consequences of a hazard once it has occurred.

For each option, the costs associated with its implementation should be assessed and recorded. Only costs associated directly with implementation of the option should be estimated. The impact of potential benefits realised by the option should not be included (this will be determined in the next stage).

Demonstration of compliance with the ALARP principle requires that all significant potential risk reduction measures are identified and considered. Unless a comprehensive Options Analysis has been undertaken, therefore, it is not possible to demonstrate that the risk has been reduced ALARP.

Options Analysis is therefore best conducted:

- using empirical and creative processes (for example checklists and brainstorming respectively) in a manner similar to that used in Hazard Identification; it should be noted that a thorough Hazard Identification process may also have identified some potential options;
- through analysis of the results of Causal and Consequence Analysis to guide identification of potential options.

8.5.6 Stage 6: Impact Analysis

Impact Analysis determines the likely effects of each option identified in Options Analysis upon the losses.

Impact Analysis revisits the previous stages, this time allowing for the effects of the option. For each option identified, the following process should be adopted:

1. Determine the impact of the option upon occurrence or escalation of a hazard.
2. On the basis of the revised Causal or Consequence Analysis, revisit the Loss Analysis of the associated hazard to determine the losses to be realised assuming implementation of the option.
3. Calculate the difference between safety losses with and without the implementation of the option. This is the **safety value** of the change.

In some cases, an option may have the potential to mitigate hazards in other railway systems. In that case, you may increase the safety value of the change by the reduction in losses associated with the other system as a result of this option.

Safety values should be determined individually for each affected population, in the same way as for Loss Analysis.

Where more than one risk reduction option has been identified, care should be taken to ensure that the dependencies between these options are properly addressed.

If the previous stages were originally done qualitatively then they should be revisited qualitatively. If they were originally done quantitatively then they should be revisited quantitatively.

Where quantitative analysis is employed, sensitivity parameters may be derived for each of the options through appropriate analysis of the corresponding causal or consequence models. This helps determine the most effective measures for loss reduction.

8.5.7 Stage 7: Demonstration of ALARP and compliance

As explained in section 8.3.2, demonstrating compliance with the ALARP principle involves demonstrating two separate facts:

- 1 that the overall risk is in the tolerability region, that is below the upper limit of tolerability, and
- 2 that risk has been reduced ALARP.

This stage can be divided into two steps, each demonstrating one of these facts.

8.5.7.1 Demonstration of compliance with upper limit of tolerability

The upper limit of tolerability will be defined for any given railway by the railway authority for that railway. Typically, it is defined in terms of the individual risk experienced by a member of an affected group of people.

Upper limits of tolerability may be set for more than one group of people. For instance, Railtrack's Railway Safety Case sets limits for three groups: employees, passengers and the public.

Note that completing this step is not enough to show that you have reduced risk ALARP; to do this you still need to perform the next step – Demonstration of ALARP.

8.5.7.2 Demonstration of compliance (qualitative)

A qualitative argument for compliance with the upper limit of tolerability may be made, on the basis of order of magnitude calculations by showing that the changed railway presents significantly less risk than before, provided that:

- the risk was below the upper limit of tolerability before the change was made;
- the upper limit of tolerability has not since been reduced by a larger factor than the improvement in safety; and
- there has been no significant adjustment of safety targets between railway systems.

Justification should be made that all the above provisos are met.

In general, a qualitative argument of this form can be made by the system supplier alone, using limited, and often publicly available, information on safety performance and policy from the railway authority.

Alternatively, if a likelihood-severity matrix has been constructed for this application, a qualitative argument for compliance with the upper limit of tolerability may be made by showing that:

- the risk of each hazard falls into a Tolerable or Broadly Acceptable category;
- the guidelines associated with the matrix have been followed; and
- the assumptions associated with the matrix hold for application in question.

8.5.7.3 Demonstration of compliance (quantitative)

The quantitative approach to demonstrating compliance with the upper limit of tolerability requires three steps:

- 1 to apportion the upper limit of tolerability between railway systems;
- 2 to derive tolerable hazard rates for the system in question;
- 3 to show that the actual system hazard rates are below the derived upper limits.

The third step is performed by direct comparison with the results of quantitative Causal Analysis.

If the railway authority has already defined tolerable hazard rates for the system (see section 8.4.7), the first two steps can be omitted. Otherwise they may be performed as follows.

To apportion the limit, you will normally employ an existing model of the contribution of safety risk from different railway systems. Typically you will estimate an initial apportionment in line with historical data as follows:

- estimate what fraction of total annual risk of safety loss is attributable to the system;
- multiply the upper limit of tolerability by this fraction.

If upper limits of tolerability are set for multiple groups, then this calculation will be carried out for each group.

The initial apportionment may be adjusted to meet strategic objectives for safety improvement.

Tolerable hazard rates for the system are then set so that the exposed members of each group experience an individual risk from the system below this limit. To confirm that this is the case, you will need to do the following for each group:

- add up the statistical average number of fatalities (F) that would occur for this group if all hazards occurred at their tolerable hazard rates;
- estimate the number of people (n) within this group exposed to the risk; and
- estimate the individual risk (F/n) experienced by an average person who is exposed to the risk and show that this is below the apportioned upper limit of tolerability.

8.5.7.4 Demonstration of ALARP

To show that risk has been reduced ALARP, you have to show that no reasonably practicable options exist which have not been implemented.

A qualitative demonstration may be made relying on informed consensus from a group of experts reviewing the results of Options Analysis that all rejected options are not reasonably practicable. The reasons for this judgement should be articulated and documented.

If a quantitative approach is being followed, Impact Analysis will have calculated, using VPFs supplied by the railway authority, a safety value, that is a monetary value for the improvement in safety arising from each option. Options Analysis will have estimated the net cost of implementing the option. An option may be rejected as not reasonably practicable if the safety value is significantly less than the cost.

Note, that this conclusion can only be made robustly if the difference between the two values is more than the total uncertainty in both of them.

8.6 Related guidance

Chapter 7 provides guidance on defining the boundaries of a system as a pre-requisite to risk assessment.

Chapter 9 explains how risk assessment is used to set safety requirements in general and safety integrity levels in particular.

Chapter 13 describes the maintenance of a Hazard Log, which will act as a repository for risk assessment data.

Appendix C provides supporting checklists.

Appendix E describes some relevant techniques.

This page left intentionally blank

Chapter 9

Safety Requirements

Your organisation must set safety requirements for any change, to reduce the risk associated with the change to an acceptable level.

9.1 Guidance from volume 1

Safety requirements are requirements that should be met to make sure that the safety risk presented by a change is reduced to an acceptable level. Safety requirements may specify:

- features or functions of the change, including any which help people avoid dangerous mistakes,
- what the change must not do to ensure safety,
- environmental conditions under which the change must operate to stay safe,
- targets for carrying out a function reliably, or reliably avoiding a dangerous state,
- features of the design and build processes, and
- operational procedures and restrictions.

You will set some safety requirements to meet regulations or standards. Others may arise when you identify hazards and assess and reduce risk.

9.2 Background

A project carrying out safety-related work should identify the hazards and accidents that may result from the work, assess the risk associated with these, reduce the risk ALARP and set Safety Requirements to ensure this level of risk is met. There is a legal requirement to assess the risks involved in safety-related work. Safety Requirements should also be consistent with the operator's stated targets.

Safety Requirements may be quantitative or qualitative. Good engineering practice for meeting integrity requirements for components susceptible to systematic failure is to use Safety Integrity Levels (SILs). SILs are described in section 9.6 below.

The activity of establishing safety requirements follows and builds on the work described in the previous chapter. If you have not already done so, you should read the Background section of the previous chapter as it also provides important background for this chapter.

The Safety Requirements Specification consolidates information provided by these activities into specific requirements, which form the basis against which the safety of the system is tested and assessed.

The activity of establishing Safety Requirements is iterative to reflect the iterative nature of safety analysis.

This chapter is written for people writing or reviewing Safety Requirements.

9.3 Reducing risk

The following is a widely accepted order of precedence for reducing risk

- 1 Respecify or redesign to eliminate hazards or reduce their likelihood.
- 2 Reduce risk in the design, by adding safety features.
- 3 Reduce risk by adding warning devices.
- 4 Reduce risk through procedures and training.
- 5 Reduce risk by adding warning signs and notices

For any given hazard you should first seek to set Safety Requirements to eliminate it. Only where this is not possible should you proceed to set Safety Requirements on the design of the system. And only when all reasonably practicable risk reduction has been accomplished on the design should you consider procedures and training as risk reduction options.

9.4 Overview of process

Setting safety targets is normally done by working from a fault tree (or similar representation of cause and effect logic) and the event probabilities to:

- a) derive numerical accident targets which conform to the ALARP principle, that is either they are tolerable and further risk reduction is not reasonably practicable or they are negligible;
- b) derive hazard occurrence rate and/or unavailability targets which are consistent with (a);
- c) (if systematic failure modes exist) relate hazards to system functions and 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, systematic targets should not be set below sub-system function level. Refer to IEC 61508 [F.14] or ENV 50129:1998 [F.13] for further guidance on this decomposition.

Any functional requirements on the system or equipment that are necessary to reduce risk to an acceptable level should be incorporated as qualitative safety requirements.

The analyst may set other qualitative safety requirements such as conformance to external standards and should do so whenever:

- such conformance is assumed in the calculation of safety targets; or
- such conformance is otherwise required to reduce risks as low as reasonably practicable.

If the seven-step process described in chapter 8 is being used, then some requirements will arise from the fifth step, Options Analysis. However requirements may also arise from relevant regulations, standards and codes of practice.

9.5 Apportionment of random failure targets

It is not generally necessary to descend the fault tree fully, that is to set targets for base events. The analyst should set targets at a level coincident with the hierarchical breakdown of the system being developed.

9.6 Assignment of Safety Integrity Levels

There are well-established techniques for assessing and controlling random failures but practice is not as advanced in the treatment of systematic failures. Current best practice is to define a number of Safety Integrity Levels (SILs) representing different levels of rigour in the development process and to relate these to approximate probability targets.

Five levels are defined. There are four safety-related SILs, ranging from SIL 4, the most stringent, to SIL 1, the least stringent. Functions which are not relied upon at all to control risk may be described as having SIL 0. Each level is populated with increasingly stringent processes and techniques.

Systematic failures are of particular concern in software-based systems and have generally been applied to these sort of systems, and, in particular to the software within them. However hardware can exhibit systematic failures and SILs are applicable to them as well.

Each integrity level is associated with a target probability of failure. One widely accepted association is shown in Table 9-1, which is derived from IEC 61508 [F.14]. The Low Demand column should be used if demands are expected to occur:

- no more than once per year, and
- no more than twice as often as the system is checked out.

Otherwise use the Continuous/High Demand column.

Low Demand Mode of Operation (probability of failure on demand)	Continuous / High Demand mode of operation (Dangerous failure rate per year)	Safety Integrity Level
$\geq 10^{-5}$ to 10^{-4}	$\geq 10^{-5}$ to 10^{-4}	4
$\geq 10^{-4}$ to 10^{-3}	$\geq 10^{-4}$ to 10^{-3}	3
$\geq 10^{-3}$ to 10^{-2}	$\geq 10^{-3}$ to 10^{-2}	2
$\geq 10^{-2}$ to 10^{-1}	$\geq 10^{-2}$ to 10^{-1}	1

Table 9-1 - Safety Integrity Levels

Target probabilities of failure for systematic functions should be set to achieve an acceptable level of risk for the overall system.

Each sub-system within the overall system will generally take the maximum SIL of all the functions that it implements. The components within that sub-system may then be allocated SILs according to the guidance given in section 9.7.

However, if it can be clearly demonstrated that a sub-system's functions are wholly independent of each other (that is, the immediate effects of a function's failure are restricted to that function), then these functions (or groups of functions) may be considered as sub-systems in themselves and assigned SILs accordingly. In this way, the apportionment of SILs need not be confined to physically separate units.

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. If possible, functions with differing SILs should be segregated either physically or logically.

Once the SIL for a sub-system has been established then appropriate techniques to develop the sub-system to that level can be established by reference to tables in standards including ENV 50129:1998 [F.13], IEC 61508 [F.14] and DEF-STANs 00-54 [F.15] and 00-55 [F.16].

9.7 Apportionment of Safety Integrity Level

Having set a SIL for a function to achieve the necessary probability target, the analyst may need 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.

However, it is possible to use a redundant architecture to build high SIL systems from sub-systems of lower SIL by building in back-up or protection functions.

If the architecture ensures that a top-level function can only fail if both a *main* and *backup* function fail and the two functions are independent, then the SIL of the top function may sometimes be higher than that of the main or backup function.

In some cases, there may also be a *combinator* function (for instance a voting scheme) which combines the results of the main and backup functions.

Table 9-2 (derived from DEF-STAN 00-56 [F.4]) 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.

Note that the combinator always inherits the top level SIL.

The table should not be repeatedly applied to allow a SIL 4 system, say, to be made of many SIL 1 systems.

Top Level SIL	SIL of Lower Level Function		Combinator (if present)
	Main	Other	
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 2
SIL 1	SIL 1	None	None

Table 9-2 - Apportionment of Safety Integrity Levels

9.8 Software Safety Requirements

This guidance is applicable to any railway system containing software, including embedded systems such as programmable logic controllers.

For programmable systems, it is normal to derive a Software Requirements Specification (although other titles may be used). This should define the functions that the software must perform which, taken together with the capabilities of the hardware components, will allow the overall system to meet its requirements.

In just the same way as safety requirements are set at the system level and form part of the overall system requirements, it is usual to establish a Software Safety Requirements Specification, either as a subset of the Software Requirements Specification or as a separate document.

The software safety requirements will normally include requirements for features which can tolerate faults as well as requirements for dependability of the software.

prEN 50128 provides guidance on fault-tolerant features.

Dependability should be treated by specifying the SIL of the software. This will be the same as the SIL for the system unless it has been explicitly apportioned as described in the previous section.

Guidance on the development of software for safety-related railway applications can be found in prEN 50128 [F.17] which also describes techniques appropriate to each SIL.

Evidence of validation of the software against its requirements should be produced. If prEN 50128 is used then this is documented in a software assessment report and a software validation report. This evidence will form an important part of the overall system Safety Case.

9.9 The Safety Requirements Specification

The following structure is recommended for a Safety Requirements Specification:

- Introduction.
- Background. A summary of the system and project, including configuration information where appropriate.
- Statement of Safety Requirements. A list of all Safety Requirements.
- Justification of Safety Requirements. The assumptions and calculations supporting the statement of Safety Requirements, including a record of the techniques employed, the manner in which they were applied.
- Reference to safety documentation. References to all documents used together with version numbers.

Other effective formats are in common use. The Safety Requirements Specification does not need to be a separate document. and is sometimes combined with other documents. A Safety Requirements Specification will, however, normally include at least as much information as provided in the structure above.

The Safety Requirements Specification should be submitted to the Safety Authority for endorsement.

9.10 Related Guidance

Chapter 8 provides guidance on the safety analysis processes which should be carried out before setting safety requirements.

Chapter 10

Safety Evidence and Authorising Change

Your organisation must convince itself that risk associated with a change has been reduced to an acceptable level. It must support its arguments with objective evidence, including evidence that it has met all safety requirements.

No change can be authorised until all necessary safety approvals have been given.

10.1 Guidance from volume 1

10.1.1 Evidence of safety

You should normally put these arguments together in a safety case to show that:

- you have adequately assessed the risk;
- you have set adequate safety requirements and met them;
- you have carried out the safety plan; and
- all safety-related work has been done by people with the proper skills and experience.

If other people must take action before a change is safe, the safety case should describe these actions and show that the other people have accepted responsibility for carrying out these actions.

You may include relevant in-service experience and safety approvals as supporting evidence.

If you are working on signalling systems or equipment, CENELEC standard ENV 50129:1998, '*Railway Applications – Safety Related Electronic Systems for Signalling*' [F.13] is relevant. It places requirements on safety cases.

10.1.2 Safety approval

You must get safety approval from the necessary safety authorities. You will usually need approval from both the railway authority (such as Railtrack and London Underground Limited) and the regulatory authority (HMRI in the UK). Safety approval will normally be based on accepting the safety case.

The approving authority will normally produce a certificate, setting out any restrictions on how the work is used.

The approving authority will usually give safety approval at the end of a project, when the change is about to go into service. Some projects make staged changes to the railway in which case each stage will need safety approval. Large or complicated projects may need additional approval before they change the railway, for example for a safety plan or for safety requirements.

10.2 Background

The Safety Case is a document that provides an argument for the safety of a change to the railway. It provides assurance that risk has been reduced to an acceptable level to the project itself and to the Safety Authorities who will approve the change to the railway.

The main sources of evidence called up by the Safety Case are the records that have been kept and the checks that have been made by independent engineers.

The Safety Case can also be presented as an incremental document which will include ESM data as it becomes available.

The Safety Case required by ESM is an *engineering* Safety Case and should not be confused with the railway safety case which the '*Railways (Safety Case) Regulations 1994*' require all train and station operators to produce. The two are linked however: an operator's railway safety case may rely in part on the engineering safety cases for the operator's major systems.

The Safety Case provides much of the evidence for safety that the Safety Authority requires in order to grant safety approval for the change to proceed.

Note that the phrase 'safety approval' is used by some people to describe a process during which the safety authority accepts liability for the railway change. The phrase 'safety acceptance' is used to describe an endorsement without acceptance of liability. In this volume 'safety approval' is used to describe any process by which a Safety Authority grants its approval for a proposed change to the railway to proceed, regardless of the implications for legal liability.

This chapter is written for:

- anyone preparing a Safety Case, and
- anyone reviewing a Safety Case.

10.3 Application

The size of the Safety Case will depend on the risks and complexity of the project. For example, the Safety Case for a simple and low-risk project should be a concise document with brief arguments justifying that the risk is acceptable. A Safety Case for a high-risk or complex project will require a comprehensive Safety Case with comprehensive safety arguments.

10.4 Responsibilities

The Project Manager is responsible for ensuring that a Safety Case is prepared, maintained, and submitted to the Safety Authorities. He may delegate the preparation to a Project Safety Manager but should retain overall responsibility.

The relevant Safety Authorities are responsible for endorsing the Safety Case.

10.5 Submission

The Safety Case should be submitted to the relevant Safety Authorities for endorsement.

In the UK, the Safety Case may need to be submitted to HMRI as part of type approval as required by the '*Railways and Other Transport Systems (Approval of Works, Plant and Equipment) Regulations 1994*'.

Interim versions of the Safety Case may be submitted as the project proceeds. For instance, for the introduction of rolling stock onto the Railtrack railway network it is usual to make safety submissions after design, before track test, before interim operations and before full fleet operation.

The points at which versions of the Safety Case will be submitted should be agreed with the Safety Authorities and documented in the Safety Plan.

A complete version of the Safety Case should be submitted and endorsed before any change is introduced to the railway. If the project is making staged changes then several versions may need to be submitted and endorsed, each covering one or more stages.

The Safety Case should be modified during operation if the system is changed or if further safety-related information is obtained.

Note: The Safety Case forms part of the evidence of safety, submitted to the Safety Authority. It will typically be accompanied by other documents, including an Independent Safety Assessor's report. Also, other ESM documents such as the Safety Plan and Safety Requirements Specification may have been submitted for endorsement beforehand. The Safety Authority's review of these documents can constitute a Safety Review in the terminology used by prEN 50126 [F.8].

10.6 General guidance

The Safety Case should demonstrate that the system complies with its Safety Requirements and that risk has been reduced to an acceptable level.

The Safety Case should identify and justify any unresolved hazards and any non-conformances with the Safety Requirements Specification and Safety Plan.

The Safety Case should consider safety relating to the entire system as it consists of a combination of hardware, software, procedures and people interacting to achieve the defined objective.

The Safety Case should present information at a high-level and reference detail in other project documentation, such as the Hazard Log. Any referenced documentation should be uniquely identified and traceable. References should be accurate and comprehensive.

The Safety Case should present or reference evidence to support its reasoning. Evidence may come from many sources, although the Safety Case is likely to depend heavily on entries in the Hazard Log and the results of Safety Assessments and Safety Audits.

The Safety Case should accurately reflect information obtained from other project documentation.

Although the Safety Case is primarily used to satisfy the project and Safety Authorities of the safety of the system or equipment, the Safety Case may have a wider readership, including Safety Auditors and Assessors, and this should be taken into account when preparing the Safety Case.

10.7 CENELEC standard ENV 50129:1998

This pre-standard defines the conditions which should be met to accept a safety-related electronic railway signalling system.

The principal normative contents of the standard are:

- requirements on Safety Case structure and content (clauses 5.1 through 5.4 and appendix B);
- requirements on safety acceptance and approval (including types of safety cases) (clause 5.5);
- requirements on the establishment of Safety Integrity Levels (appendix A); and
- requirements on the identification of hardware component failure modes (appendix C).

Note though that this standard is not intended to provide comprehensive guidance on writing a Safety Case. It provides a structured framework for demonstrating safety but requires interpretation to deliver a convincing demonstration of safety for a railway change.

Volume 2 of the Yellow Book in general, and this chapter in particular, have been written to allow the reader to comply with ENV 50129:1998 [F.13] while following the guidance provided, and to help with interpreting this standard effectively.

10.8 Types of Safety Case

Three different types of Safety Case can be considered, see ENV 50129:1998 [F.13]

- A *generic product Safety Case* provides evidence that a generic product is safe in a variety of applications.
- A *generic application Safety Case* provides evidence that a generic product is safe in a specific class of applications.
- A *specific application Safety Case* is relevant to one specific application.

These may be used to allow efficient re-use of safety evidence. For instance a specific application Safety Case for a resignalling scheme may refer to a generic application Safety Case for the use of a points machine in a particular type of junction which may in turn refer to a generic product Safety Case for that points machine.

NB ENV 50129:1998 [F.13] requires that a specific application Safety Case be split into two Safety Cases: *application design* and *physical implementation*. This publication however does not recommend splitting the Safety Case in that way for all applications.

10.9 Content of the Safety Case

The structure described in Figure 10-1, which is consistent with ENV 50129:1998 [F.13], is recommended for a Safety Case.

Alternative structures may be appropriate in some cases but they should cover the same topics as this structure.

Note though that above all, a Safety Case should deliver a convincing and comprehensive argument for safety. This cannot be provided just by complying with any given structure but should arise from an effective programme of ESM activities.

1. **Executive Summary**
2. **Introduction**
3. **Definition of System**
4. **Quality Management Report**
5. **Safety Management Report**
 - Introduction
 - Roles and responsibilities
 - Safety lifecycle
 - Safety analysis
 - Safety requirements
 - Safety standards
 - Safety audit and assessment
 - Supplier management
 - Safety controls
 - Configuration management
 - Project safety training
6. **Technical Safety Report**
 - Introduction
 - Assurance of correct functional operation
 - Effects of faults
 - Operation with external influences
 - Safety-related application conditions
 - Safety qualification tests
 - Other outstanding safety issues
7. **Related Safety Cases**
8. **Conclusion**

Figure 10-1 – Recommended Safety Case Structure

While this section provides a framework for *structuring* the Safety Case, the ESM activities should drive the *content* – any activity which was necessary to achieve acceptable risk should contribute some content to the Safety Case.

10.10 Safety Case: Executive summary

The executive summary should summarise the key information contained in the Safety Case. It should contain the following:

- a brief description of the change, its purpose, functionality and location;
- a summary of the safety design and development process undertaken;
- a summary of the assessment and audit processes undertaken;
- a summary of the test and operational experience; and
- a summary of the current safety status in terms of evidence obtained and unresolved hazards.

10.11 Safety Case: Introduction

This section should describe the aim, purpose, scope and structure of the Safety Case.

10.12 Safety Case: Definition of system

This section should provide an overview of the change in order to provide an understanding of the safety issues raised. It should cover, or reference, documentation dealing with the purpose, functionality, architecture, design, operation and support of items under review.

It should include:

- a description of the system including its physical location;
- definition of system boundaries and interfaces, including assumptions about other systems, services and facilities; and
- identification of constituent sub-systems, and if appropriate, a reference to sub-systems' Safety Cases.

The configuration of the system to which the Safety Case applies should be explicitly identified. This section should demonstrate that the system is subject to effective configuration management and change control, referring to any standards called up in the Safety Plan.

10.13 Safety Case: Quality management report

A pre-requisite for an effective Safety Case is that the quality of the work is and has been, controlled by an effective quality management system (QMS). This section should summarise the QMS activities and justify their appropriateness to the project. Large volumes of detailed evidence and supporting documentation need not be included provided precise references are given to a description of the relevant QMS.

10.14 Safety Case: Safety management report

10.14.1 Introduction

This section should describe and discuss how ESM aspects of the project were carried out. It should summarise and refer to the activities described in the Safety Plan and provide or refer to evidence to show that the activities were carried out as planned and justify that these activities proved to be appropriate and adequate.

The Hazard Log will be the primary source of evidence that hazards have been controlled.

The following ESM issues should be addressed:

- Roles and responsibilities;
- Safety lifecycle;
- Safety analysis
- Safety requirements;
- Safety standards;
- Safety audit and assessment;
- Supplier management;
- Safety controls;
- Configuration management; and
- Project safety training.

Each issue is treated in a separate section, below.

10.14.2 Roles and responsibilities

This section should provide evidence to show that the key safety personnel on the project carried out the roles defined in the Safety Plan.

It should justify the appointment of the key safety personnel by referring to competence and experience.

10.14.3 Safety lifecycle

This section should justify the project and safety lifecycles followed during the project, particularly if they differed significantly from those defined in the Safety Plan.

10.14.4 Safety analysis

This section should present a detailed discussion of the safety analysis process used on project. It should provide assurance that all foreseeable hazards have been identified, that intolerable risks have been eliminated and that other risks have been reduced ALARP.

This section should show that the safety analyses have taken into account the scope of the system and its normal and abnormal operation. System and component failure and malfunction, procedural failures, human error and environmental conditions should be considered.

The following should be provided:

- A list of the analysis methods used and their application on the project.
- Identification of the design documents referenced during the analysis work, clearly indicating the configuration and status of the design for each analysis.
- Evidence that the safety analysis process is capable of addressing the safety of future system changes.

This section should review all the incidents that have occurred. It should state the cause, potential and actual effects, and the actions required to prevent the re-occurrence of all incidents that have occurred during operational experience which could have compromised safety during operation in-service. The review should refer to the Hazard Log.

This section should also present a review of reliability data based on data obtained from operating experience, including the Hazard Log. The data should be used to quantify and justify the safety analysis evidence.

This section should discuss the approach used to demonstrate that risk has been reduced ALARP and demonstrate that the approach follows good practice.

This section should record any elements of the safety policy set by the railway authority which are relevant to the analysis. These may include safety targets set by the railway authority and latitude allowed to their contractors and suppliers to change aspects of the railway environment in which the system or equipment will run.

10.14.5 Safety requirements

This section may either restate the safety requirements for the system or equipment or summarise them and refer to the Safety Requirements Specification.

A discussion of the safety implications of the requirements, indicating how each requirement affected the project, should be included.

Any assumptions made should be stated and justified.

Evidence for compliance with the Safety Requirements is addressed in section 10.15.3 below.

10.14.6 Safety standards

This section should provide evidence that the procedures and standards called up by the Safety Plan were followed, and justify any non-conformances.

10.14.7 Safety audits and safety assessments

Evidence for the implementation of the Safety Audit and Assessment programme is a key element of the Safety Case. The findings of these audits and assessments are normally presented in separate documents. This section should present the following:

- a description and justification of timing of the audits and assessments should be described;
- a justification that the auditors and assessors had sufficient competence and independence;
- a justification of any decision not to take action in response to a finding or recommendation.

10.14.8 Supplier management

This section should show that the work of contractors and suppliers has been carried out to the safety standards expected for the SIL applicable, and as specified in the supplier's Safety Plan.

10.14.9 Safety controls

This section should provide evidence that the safety controls identified in the Safety Plan have been applied.

10.14.10 Configuration management

This section should justify the configuration management system employed and show that it has been implemented correctly.

Evidence that all safety-related project items are under configuration management should be provided.

10.14.11 Project safety training

This section should show that the personnel carrying out the safety-related activities were adequately trained by providing evidence for the implementation of defined training plans.

10.15 Safety Case: Technical safety report

The Technical Safety Report should explain the technical principles which assure the safety of the design including (or giving references to) all supporting evidence (for example, design principles and calculations, test specifications and results, and safety analyses). Large volumes of detailed evidence and supporting documentation need not be included, provided precise references are given to such documents. The following gives a guideline for the structure of the Technical Safety Report,

- 1 Introduction;
- 2 Assurance of correct functional operation;
- 3 Effects of faults;
- 4 Operation with external influences;
- 5 Safety-related application conditions;
- 6 Safety qualification tests;
- 7 Other outstanding safety issues.

Items 2 to 7 inclusive are each treated in the sections below.

10.15.1 Assurance of correct functional operation, effects of faults, operation with external influences

These three sections should describe and discuss the activities carried out in each phase of the project in order to satisfy the Safety Requirements. They should summarise and refer to the activities described in the Safety Plan and provide or refer to evidence to show that the activities were carried out and that these activities proved to be appropriate and adequate for the defined SIL.

These activities should be provided under three headings:

- **Assurance of correct functional operation**

Demonstrating that the system will contribute acceptable risk in the absence of faults and external influences. Routine maintenance should be considered as well as normal operation.

- **Effects of faults**

Demonstrating that the system will contribute acceptable risk in the presence of foreseeable internal faults. Relevant safety features, fall-back modes and alternative operating procedures should be described.

- **Operation with external influences**

Demonstrating that the system will contribute acceptable risk in the presence of foreseeable external influences, such as weather, electromagnetic interference and vandalism, Relevant safety features, fall-back modes and alternative operating procedures should be described.

The Hazard Log should be used as the primary source of evidence.

These sections should show that the approach adopted reduced risk ALARP.

These sections should show that the Safety Requirements have been met. They should include the following, where relevant:

- a) Evidence that the Safety Requirements were defined according to good practice (see chapter 12);
- b) Identification and justification of major changes made to the Safety Requirements throughout the project;
- c) A summary and reference to all analyses carried out during requirements definition;
- d) Evidence that high-level allocation of safety requirements to sub-systems has been carried out;
- e) An explanation and justification of all use of sub-systems, pre-fabricated sections, dependencies on other systems and so on, which have been produced outside the direct control of the project;
- f) Safety evidence acquired from verification and validation of the system summarised, including the strategy and method employed and the results and evidence obtained;
- g) Evidence that further work or improvements identified as a result of validation and verification activities have been carried out;
- h) Evidence that the system has been integrated with existing systems and procedures in a safe and controlled manner;
- i) Evidence that commissioning activities have been examined for commissioning-specific hazards;
- j) Evidence that, for systems with extensive or complex hardware, software or human-factors considerations, a formal hazard identification and analysis activity has been carried out;
- k) Practical experience of operating the system, including testing, integration, commissioning and any in-service experience summarised;
- l) Evidence that the hazards associated with system operation will be adequately controlled under both normal and abnormal conditions and for all modes of operation;
- m) Evidence that all aspects effecting safe operation and maintenance, including staffing levels, training requirements, operational management and interfaces to other systems, have been addressed;
- n) Evidence that response time requirements and other analogue issues, such as non-overlapping tolerances have been considered.

10.15.2 Safety-related application conditions

This section should specify (or reference) the rules, conditions and constraints which should be observed in the application of the system. This should include the application conditions contained in the Safety Case of any related sub-system or item of equipment.

10.15.3 Safety qualification tests

This section should provide evidence of test activities which demonstrate that each Safety Requirement has been met.

If this evidence is not available or sufficient, this section should present analytical evidence that each Safety Requirement has been met, or adequately justify any that have not been met. Such a justification should include an assessment of the residual risk presented by the non-compliance.

This section should review the fault history and status of the system or equipment as recorded in the Data Recording And Corrective Action System (DRACAS) and justify the conclusion that risk has been reduced to an acceptable level in the light of this evidence.

10.15.4 Other outstanding safety issues

All outstanding safety issues not covered by documented Safety Requirements should be discussed here, whenever they would have a bearing on operational safety.

10.16 Safety Case: Related safety cases

This section should contain references to any other Safety Cases upon which this Safety Case depends, together with a demonstration that any assumptions, limitations or restrictions in the related Safety Cases are either fulfilled or carried forward into this Safety Case.

10.17 Safety Case: Conclusion

This section should make a statement on the acceptability of the system in terms of the safety requirements. This statement should include:

- a list of assumptions made in the safety case especially those made about the safety requirements;
- a statement of the residual risk presented by the system;
- a statement of system deficiencies;
- identification of all unresolved hazards and other outstanding issues;
- operating restrictions or procedures imposed for safety reasons; and
- recommendations for or identification of further work to be carried out.

The conclusions section should document any caveats on which the conclusion is based including assumptions and limitations and restrictions on use. The Safety Authority may carry these forward as conditions of safety approval.

10.18 Related guidance

Chapter 12 provides guidance on establishing safety requirements.

This page left intentionally blank