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:

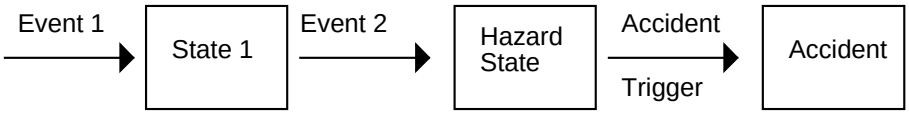

Figure 8-1 – Accident sequences

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.

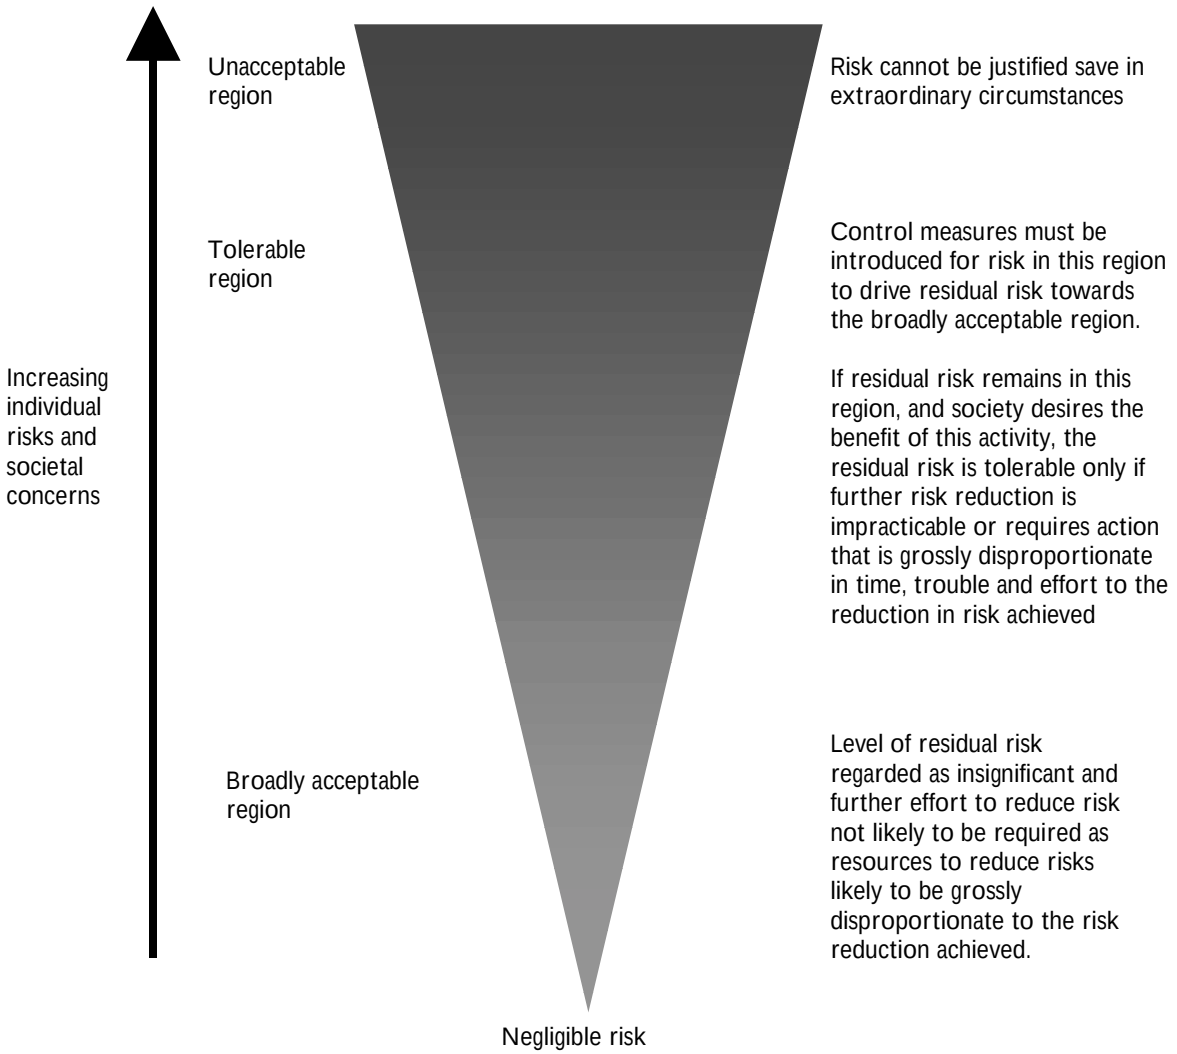

Figure 8-2 – The ALARP Principle

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.

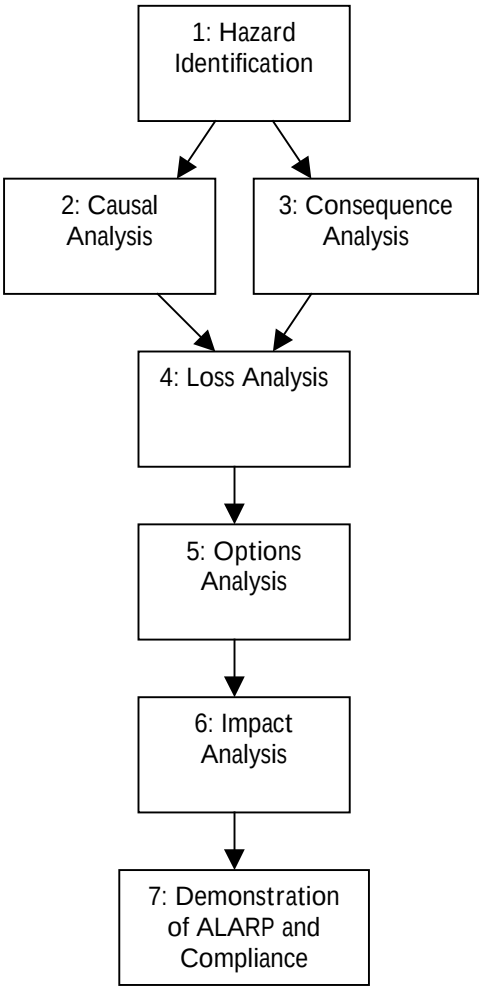

Figure 8-3 – Risk assessment stages

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.

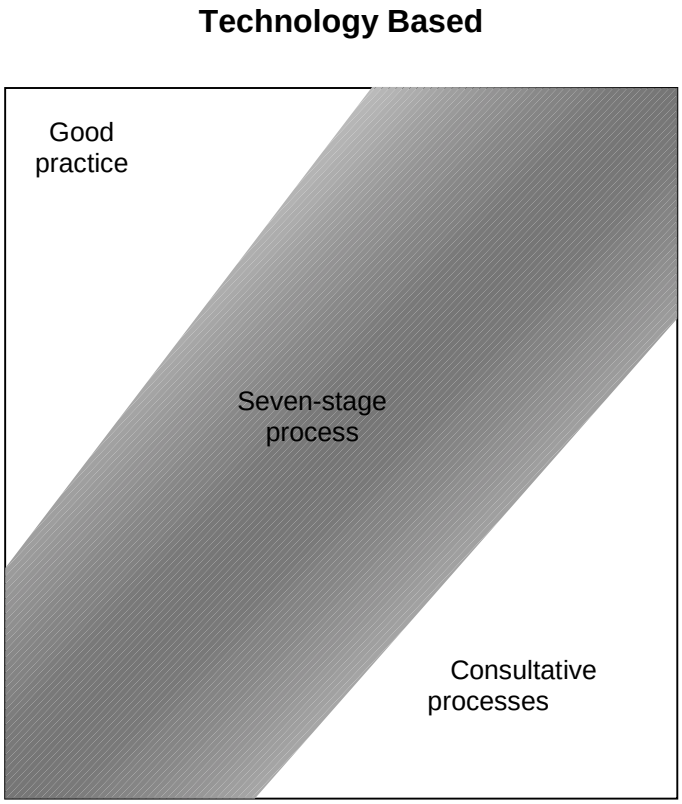

Figure 8-4 – Approaches to different risk decisions

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 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 measure 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