

# 6 Managing Risk and Safety

The concepts of *risk* and *safety* are not unique to the world of safety-related programmable systems. Many industries have to handle the concept of risk, ie in its broadest terms, the possibility of loss, and some have well developed practices from which one can learn. Similarly, useful techniques have been developed for the definition of security requirements. Two project have examined what could be learned from these parallels: RATIFI in the field of commercial risk analysis, and SPaM from the parallels with security requirements.

The approaches taken by these two projects are quite different. RATIFI investigated the potential for a range of emerging technologies to integrate both quantitative and qualitative risk assessment data, in order to drive a programme of risk reduction. The approach is similar in concept to the DATUM approach to combination of diverse evidence described in section 4.3, though using quite different techniques. SPaM, in contrast, concentrates on early specification of high level ‘safety policies’ and thereafter conformance of subsequent design and development to ensure that those policies are realised in practice. Each approach is claimed to be successful, or at least useful, in their different domains of application. RATIFI in particular, using some new technologies, was able to model ‘soft’ human factors contributions to risk in ways which have not been possible before. Future research might investigate the prospect for a common framework in which to combine these different styles of approach.

## 6.1 Managing risk

The RATIFI project brought together a joint team from Alexander & Alexander UK and the University of Birmingham to classify the wide range of techniques currently applied in different industries to analyse risk. The underlying aims of the project were to develop a normative model of the risk management process, to design and prototype support tools to implement the model, and to define and promote risk management as a safety engineering technique.

### *Current Technologies and Practice*

The project team carried out a study of current risk management practice which demonstrated the great variety of approaches to risk which any would-be generic processing technique should cover.

The team identified five major headings under which practice in different application domains could be contrasted. These were:

### *Goal*

People’s goals in managing risk are varied - some aim for complete elimination of risk (usually at a price), others require risk to be kept within a particular limit, in some cases minimisation is all that is possible, occasionally risk can be transferred to another party, and of course, some people are only ever interested in ‘hedging’ their risk. e.g. an official dealing with nuclear power

risks might be interested in establishing a “tolerable” public risk level, then bringing the technological risk within this limit. Derivatives dealers, e.g. in FOREX, take action to “hedge” (neutralise) clients’ risks.

### *Data Type: qualitative v quantitative*

Most significant risk environments tend increasingly to yield a mixture of quantitative and especially qualitative data. The dynamic contemporary mix of technological, social and political tensions generates complex change drivers whose effects are less and less amenable to classical trend analysis. Yet experts and informed decision makers characteristically supply a view of risk which is qualitative and rich in linguistic content, e.g. “I think A is more likely to happen than B, if C is absent”. Such statements lack the precision that mathematicians look for as they contain qualitative/affective elements which are difficult to model in classical probabilistic formulations. Recently (1980-1995), some new qualitative modelling approaches have emerged, e.g. Fuzzy Set Theory and Causal Networks. To concentrate wholly on either qualitative or quantitative data alone, is to discard information.

### *Representation*

As described above, representations for both the structure of the risk scenario and the uncertainty within it must be chosen. For effective use, objective probability requires large amounts of past data about a situation in which the significance of relevant parameters is constant. Subjective probability calculations require prior probabilities to allow conditionalisation to update one’s degree of belief in a proposition - these are often unavailable. Logical representations carry the overhead of maintaining consistency within a rule base, and obeying restrictive combinatorial laws if certainty factors are attached to rules.

### *Structure*

In practice an initial choice of processing will often dictate what kind of structure is possible, yet better results might be obtained by an initial consideration of what structure would be appropriate. Tree-like structures are the easiest along which to propagate uncertainty measures, as there are no problems with loops and feedback. However, the richness of many problems lies in the vast set of interactions between elements of the risk scenario - forcing a tree-like representation on what is essentially a network-type problem neglects just those factors which are most influential on risk, i.e. complex interactions. For example, it is well known in insurance circles that the vast majority of the cost of a major incident lies in the indirect consequences (often called ‘consequential loss’), not the first-line incident itself. Network representations are variously named causal maps, connectionist models, connection matrices. Tree-structure causal-type models also exist, e.g. flow charts, influence diagrams. An Influence Diagram consists of arcs and nodes in which nodes represent variables or decisions and arcs represent the path by which one node influences another. The path from one node to another is a causal chain and the diagram as a whole represents the set of all acyclic casual chains.

### *Processing*

A common statistical technique is Monte-Carlo simulation. Amongst other models, an Influence Diagram can be analysed by Monte-Carlo simulation methods. Values are inserted into the initial node variables and an output value computed using a random number for the influence outcome. This value is transferred to the successor node where it is combined with all other influences affecting that node. The process is repeated throughout the diagram. The whole

process is then repeated using different random numbers until the output distribution stabilises. Problematically, if there are many inputs to a node ( $> 4$ ), the analysis becomes extremely complex, and the assessment of conditional probabilities impractical. The Bayesian conditionalisation technique for subjective probabilities also faces computational complexity problems when risk scenarios are modelled with many nodes. Both methods are data-intensive in that many probability distributions must be input to the Monte-Carlo methods and many prior probabilities are needed for the Bayesian network. Neural network models escape this problem - their processing relies on matrix multiplication and statistical mechanics - and exhibit properties of graceful degradation (particularly useful where uncertain data is used) and scalability. Ironically, at the level of analysis, Monte-Carlo type techniques are used to consolidate the results generated, but no probabilistic representations are included in the actual model.

### *Application areas*

Four very different applications domains were examined and their approaches contrasted using the five main issues raised above. These four application domains were:

#### *Insurance*

Insurers define 'pure risk' as risk of loss alone, whilst 'speculative risk' is a risk which can result in either loss or gain. Thus they allow the magnitude involved to be positive or negative, but distinguish the types of risk in each case. This is because traditionally insurers were interested only in protecting clients from loss - i.e. pure risk, not providing them with a chance of gain. (This definition has blurred of late, when insurers are faced with derivative instruments as risk management tools - these are speculative in nature).

#### *Finance (derivatives-type trading)*

Financiers define risk as 'the volatility of potential outcomes'. They obviously take for granted that potential outcomes can have either a positive or a negative effect, and their terminology emphasises the probabilistic aspects of their processing techniques.

#### *Computer Security*

This relatively recent discipline is a rapidly-moving technical area where new risks arise regularly in complex, highly interactive systems. Little quantitative data is available, and all systems have a 'residual risk', however well-secured they may be. For safety-critical systems, a key risk area, formalised methods are popular to try to prove a risk level conclusively. (The parallel with security is explored more fully in section 6.2.)

#### *Clinical Risk Management*

This area has a statistical flavour when dealing with analyses of patient populations and particular pharmaceuticals, and a very qualitative aspect when applied to clinical practice. Such unquantifiable issues as 'quality of care' and 'quality of life' are prevalent. (The RED project, described in section 14.1 was concerned with safe decision making in the clinical context.)

DISCIPLINE	GOAL	DATA TYPE	REPRESENT-ATION	STRUC-TURE	PROCESSING
INSURANCE	Eliminate or Minimise	Quantitative	Objective Probability	Trees	Statistics
FINANCE	Hedge, i.e. Neutralise	Quantitative	Objective Probability	Trees Nets	Statistics & Stat Mechanics
COMPUTER SECURITY	Eliminate Minimise	Qualitative	Logic	Trees	Rule-Based
CLINICAL RISK MAN.	Minimise	Qualitative	Subjective Probability	Trees	Bayesian Propagation
RATIFI MODEL	Minimise or Optimise	Qualitative	Fuzzy Values	Nets	Stat. Mech Matrix Theory

Table 6.1. A Comparison of Tools Used in Four Established Risk Management Domains with the RATIFI Approach

*The RATIFI approach*

The project developed a normative model which recognised the cyclical nature of risk management and provided a methodological framework for more detailed representation of risk data at application domain level.

The model was tested for general applicability on several application domains, including derivative instruments, IT security, environmental management and human factors, and proved a successful tool in all but the most tightly coupled quantitative risk domains. Design and prototyping of Knowledge-Based support tools proved feasible for the purposes of data gathering, structuring and visualisation. The greatest challenge lay in constructing object-oriented hierarchies representing the risk factors for a new model from the given application domain. Relatively simple inferencing about risk was achieved, and the team developed this work to model the complexities of real risk scenarios, such as the chaining effect of risk events, the unexpected ways in which different risk controls interact and, crucially the effects of qualitative factors. RATIFI work demonstrated that the representation of qualitative data was possible using fuzzy sets, whilst a network approach was used to represent causal relations between risk factors.

*Strengths of the approach: Handling Human Factors*

It was found that the RATIFI approach - a top-down qualitative analysis using a network representation for risk issues, was particularly useful in the area of Human Factors. RATIFI studies suggested that our shortcomings when it comes to handling Human Issues are

unsurprising when one considers the nature of commonly used representations and processing tools within Safety Science. This observation obviously excludes those Human Factors which readily lend themselves to reliable quantitative analysis, e.g. in man-machine environments where there is a finite number of identifiable human error modes and/or a large body of data from which relatively stable human error probabilities can be derived. These fit nicely within the existing framework of objective probabilistic analysis and fixed Boolean causal relationships. However, this omission still leaves a large body of Human Factors which are crucial to safe system operation and/or mission-critical project success. Such 'soft' issues as morale, team dynamics, internal 'politics', and safety culture patently belong to this latter class. In practice, in the absence of specialist analytical skills and resources, many 'harder' potentially measurable issues, e.g. communication, attention, expertise were found to be equally unspecified, particularly in the initial stages of system/project definition.

By their very nature, the 'soft' issues do not yield to quantitative analyses, and attempts to 'shoehorn' them into rigorous quantitative paradigms are unsatisfactory as they generally involve indefensible assumptions and results of spurious accuracy. Traditional representations of safety/mission-critical scenarios are characterised by tree-like structures showing fixed dependencies and task sequences. RATIFI studies found that decision-support tools for safety/risk management are usually characterised by one or more of the following:

- record-keeping/administration and structuring of relevant data
- incremental construction of a fixed version of one of the above structures based upon a particular viewpoint on the system to be modelled
- the assignment of probabilities to elements of the structure
- the propagation of probabilities along the model in a specified direction
- scheduling, planning and resource allocation

'Soft' issues don't fit these representations and tools because they are characterised by:

- **Subjectivity.** An individual viewpoint on an issue like communication and its influence on a given man-machine environment cannot be relied upon in isolation. However, an amalgamation of the opinions of several key individuals involved in the operation and management of the MME is a valuable piece of information. Combination of subjective evidence can only be handled by 'mainstream' safety management techniques if *a priori* probabilities can be assigned to the evidence in question and Bayesian probabilistic theory applied.
- **Uncertainty.** Traditional approaches, e.g. FMEA<sup>1</sup>, provide for the representation the uncertainty of a particular failure mode taking place. Yet the relations between the failure modes are fixed within the structure of the model. There is a deeper level of uncertainty within models which include 'soft' Human Factors, that is, the existence of a relation between particular parameters is itself uncertain. A probabilistic representation allows only for the presence or absence of a relation at any given time, and then provides a measure of relative frequency (objective) or degree of belief (subjective) for that relation. Yet 'soft' issues are not suited to the idea of *whether* or not there was a cause and effect between two parameters, but rather *the degree to which* there was a causal relation. It is often difficult enough to put a probabilistic value on the relation between two 'hard' parameters without coming up against the 'cost of information', i.e. the way in which the analytical investment necessary to quantify a particular item increases rapidly with the degree of accuracy required.

---

<sup>1</sup> Failure Modes and Effects Analysis.

It is commonly only possible to make qualitative statements on the relation between 'soft' parameters, e.g. 'Safety culture usually suffers if morale falls dramatically'. This type of information has worth and relevance, particularly if it comes from expert sources - it is just not in a format we are familiar with processing.

- **Feedback.** The relations between 'soft' issues can rarely be assigned a complete set of causal sequences, and they almost always involve feedback. Their structure is far easier to represent with an influence diagram-type structure than any rigidly sequential model. In the past, influence diagrams have been useful for clarification, but they have not facilitated information processing.

This analysis of current approaches to Human Factors in Risk Management highlighted the dilemma that the representations and tools which have become established in the analysis and treatment of safety/mission-critical system data (which, for historical reasons, are quantitative) are by their very nature precluding (or at least obstructing) the representation and processing of 'soft' issues like Human Factors. Yet making safety/risk management decisions without reference to Human Factors is demonstrably bad practice: there are plentiful examples of safety failures which stem from mismanagement or, more often, neglect of the human (and especially the managerial/political) dimension of the system.

### *Case Study*

The following case study concerns a medium sized company involved in the field of IT Security. The company had grown substantially over the past five years, from its origins as a management buy-out and was now suffering the strains of rapid growth and human factor-related problems with their major security projects. The RATIFI methodology was used to perform a high-level analysis of major risk factors for the organisation and recommend appropriate, cost-effective risk control measures.

Figures 6.1, 6.2 and 6.3 illustrate the way in which the methodology builds up a rich picture of the risk factors and the complexity of their interactions. Such an analysis would not be possible using conventional quantitative methods. Initially, major assets were identified (triple-lined border), together with the relevant threats. Then, in 6.2, potential control measures (shaded border) were related to the threats identified. The resulting complex model (triple-bordered assets, shaded border controls) was processed using the RATIFI approach to identify those measures which would most effectively reduce risk for the organisation.

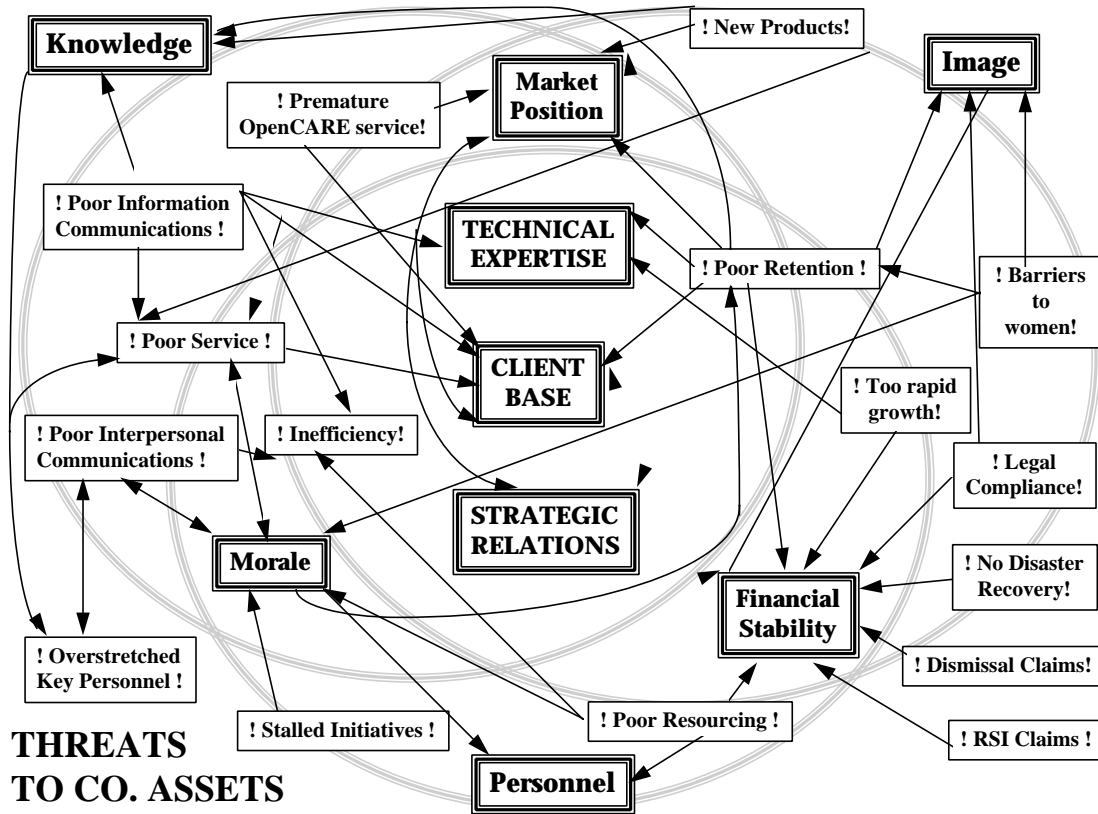
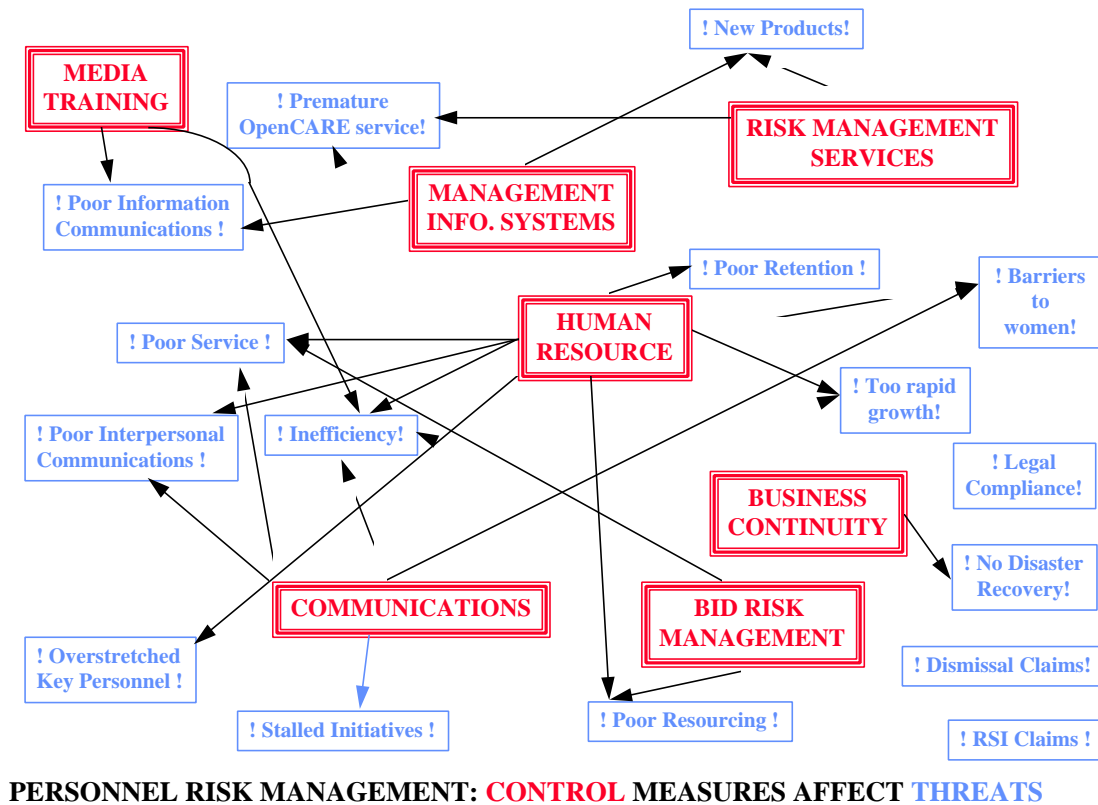


Figure 6.1. Case study example



**PERSONNEL RISK MANAGEMENT: CONTROL MEASURES AFFECT THREATS**

Figure 6.2. Case study example

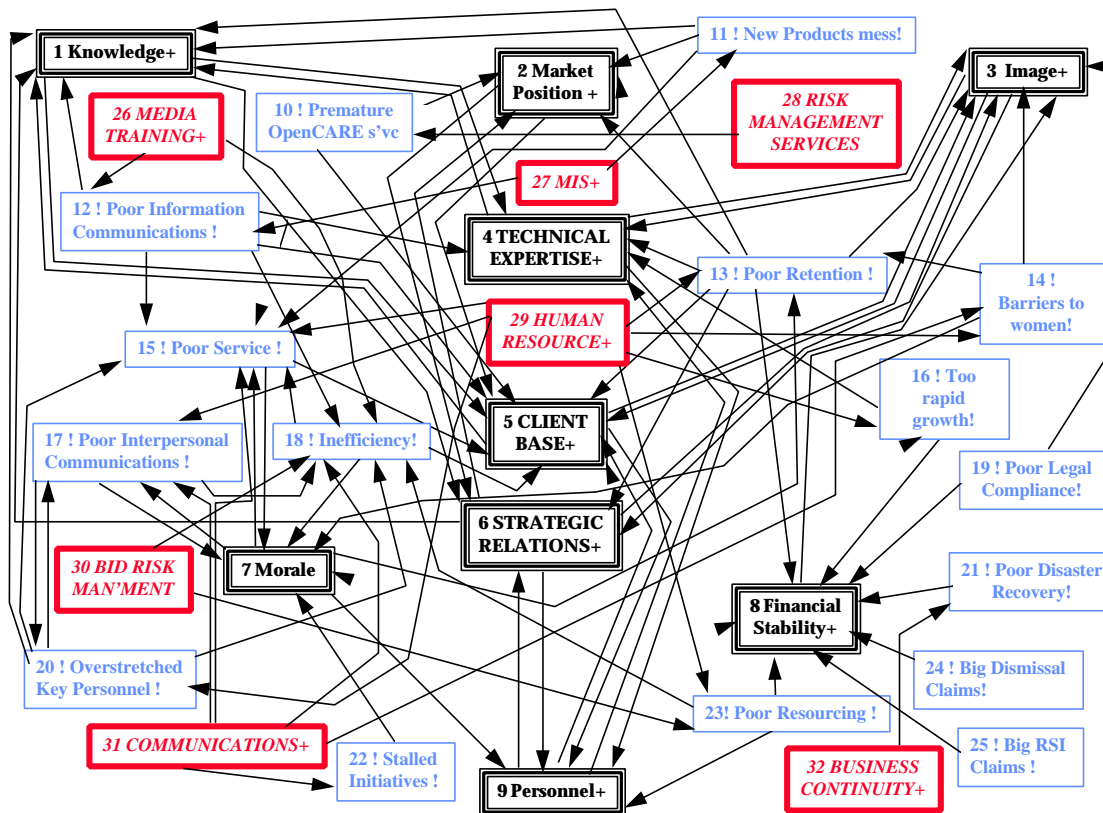


Figure 6.3. Case study example

The RATIFI project recommended that:

- The wider Safety Community (which embraces mission-critical systems, risk management and large project management) would benefit from the more widespread use of structured qualitative analysis techniques which can represent and process ‘soft’ factors. To exploit them fully, a broader view of the range of tools and techniques which can be applied to safety problems is required. This could be communicated by changing the emphasis of the relevant academic and professional training to equip practitioners with an understanding of a variety of qualitative and quantitative techniques and their strengths and limitations in relation to different types of safety/risk scenarios. Such a change would lessen the emphasis on the details of specific quantitative methods simply because they are well established and understood.
- This more balanced approach would promote the use of a ‘top down’ process to address safety problems, where an initial qualitative analysis which included Human Factors would provide the general picture, and subsequent quantitative analyses of specific issues would complete the model.
- Just as effective quantitative techniques spawn support tools for the job, more professional support tools for qualitative analysis are required. Only when it is natural and commonplace to use representations and tools which include ‘soft’ issues in building and communicating a safety/risk management programme will we be able to reason about, argue over, and justify our decisions on Human Factors.

### *Inferencing in complex situations*

The RATIFI project demonstrated that there exists a (perforce) more abstracted level of representation for complex causal scenarios (e.g. Safety) which, at a first pass, can represent both the qualitative and quantitative, subjective and objective aspects of a system. Moreover, *it is possible to perform some degree of inferencing using such representations*. Neural and Fuzzy System technologies have made this possible. Neural networks thrive on the feedback relations so common in ‘influence diagram’-type models of ‘soft’ issues. They provide a numerical means of processing qualitative data. Fuzzy Set Theory provides a consistent method for combining subjective opinion, and a non-probabilistic representation of uncertainty. The inferencing performed by such systems is naturally qualitative in nature, but in so far as it sheds light upon the complex interactions of human factors and the rest of the system (which may be far more tractable to traditional analysis) it is likely to be extremely valuable. It should be regarded as the initial stage in a top-down analysis, where all those concerned with system development and/or maintenance can be involved, including those who manage ‘soft’ issues like Human Factors. It should not replace in any way further analysis of a more traditional type, either on the system as a whole or on selected subsystems.

The project concluded that new technologies can provide the means to construct and process qualitative models in ways which were hitherto impossible. They can complement existing techniques and widen the range of tools available to safety/risk engineers. As systems become increasingly complex, and ever more tightly coupled, the work carried out in RATIFI, and its successor STREAM (Support Technologies in Risk Evaluation And Management) aims to make full use of their potential.

## **6.2 Safety policies and models**

Considerable progress has been made in the technology of computer security. In particular the use of security policy models aids the definition and understanding of fundamental security requirements, and provides an objective measure of the effectiveness of any mechanisms intended to implement security. The SPaM project has explored how these concepts and practices, which have proved powerful in the security field, might be paralleled in the field of system safety.

One key security technique involves using the concept of *security policy* to enhance the definition of security requirements in the form of security behaviours or properties. A security policy is used to define the essential properties necessary to assure the intended secure behaviour. The SPaM project has highlighted this notion of policy as a likely means of improving the definition of safety requirements, one of the key weaknesses within the safety system development life cycle.

In order to provide a systematic and rigorous approach, the idea of a security policy model became the starting point for the development of the *safety policy model*. Thus the project is proposing safety policy models as the vehicle for expressing and validating safety requirements, i.e. the intended safety behaviour described as safety properties.

The project argues that using safety policy models would improve the way in which safety systems are verified against safety requirements. This would require an associated method for applying the safety policy model to the safety life cycle. In addition, an infrastructure would be required to support the development and use of safety policy models. This led to other research activities, such as the impact that the safety policy model would have on the definition of safety integrity, and the notion of a safety taxonomy.

### *Safety Policy*

The term *safety policy* is derived from its counterpart in the computer security domain where a *security policy* essentially defines the security intent of a “system” within a defined operational environment. The level of “system” may vary from the whole organisation involving people, procedures, equipment and projects to the technical components of systems such as software.

SPaM’s use of the concept of a safety policy is to define safety objectives, free from implementation considerations. The project treats safety policies and safety objectives as synonymous. Policies or objectives must be specified, verified and validated in a demonstrable manner throughout development, and they form part of a broader safety requirements document which can then be used as a part of (or be referenced from) a safety case document.

Safety objectives relate to the safe behaviour of the system, i.e. if these objectives are always satisfied during the system operation then the system will always be in a safe state. The concept is that these behaviour objectives can be defined, verified and validated during requirements phases, and verified and demonstrated during design and subsequent phases up to delivery. The role of these objectives and the ensuing safety policy models is to provide key technical criteria for all areas of system development, encompassing the activities involved with design, safety, assessment and management.

### *Defining safe behaviour*

A safety policy model defines the required safe behaviour of a system component. It states how the procurer of safety-critical system would like their system to operate if it is to remain safe. For example, an interlock for a train door should be set when the train is moving and released when the train is stationary. Exceptions to these basic rules can then be established and added so that an effective model may be produced. Assurance that the model is both adequate and complete can then be gained by applying various verification and validation techniques.

Figure 6.4 shows the component parts of a safety policy model.

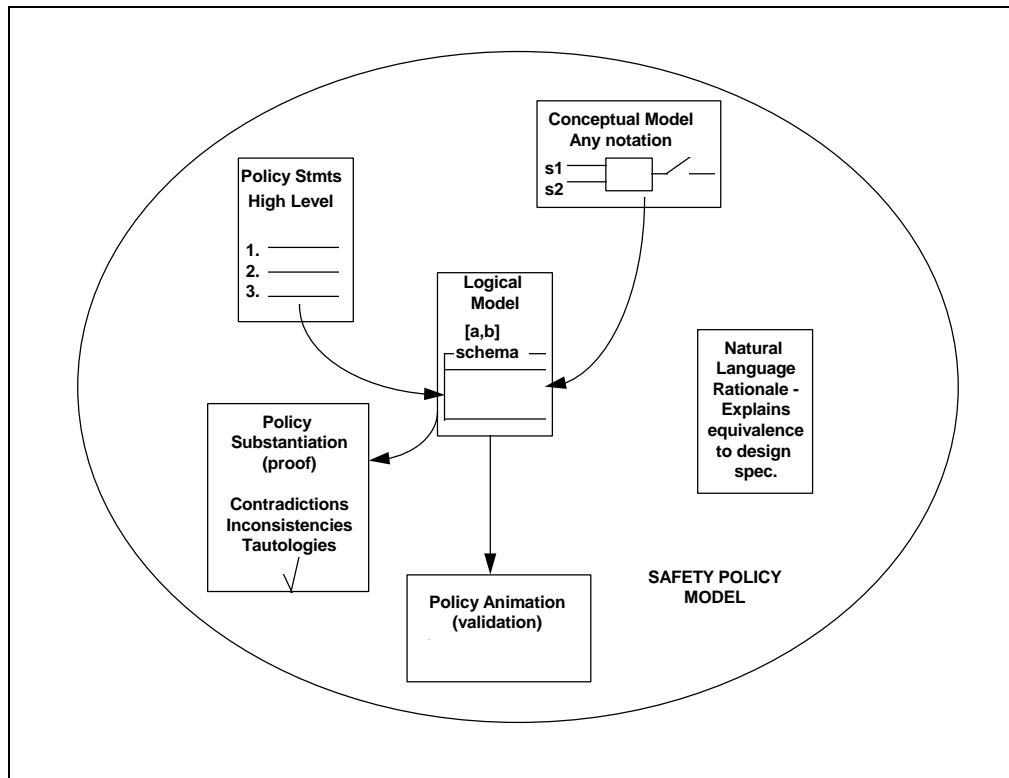


Figure 6.4. Components of the safety policy model

The safety policy model has six component parts:

- Policy Statements - creating a set of safety objectives based on the results of a preliminary hazard analysis stated as a set of high level policy statements.
- Conceptual Model - an informal model is created which may (or may not) satisfy these objectives and which provides concise rules governing the safe operation of the modelled system or component.
- Logical Model - the model and policy statements are incorporated in a well-defined notation (the logical model).
- Policy Substantiation - argues whether the safety model satisfies the safety objectives.
- Policy Animation - an executable prototype of the safety model is used to provide sufficient confidence through demonstration and test that the safety objectives are valid and complete.
- Rationale - builds up an argument that presents the verification and validation evidence to be used in a safety case context.

The key role of each of these parts of a safety policy model is to provide a sound understanding of what safety means for a given system.

### *SPaM Methodology*

The SPaM Methodology has three components:

- *Safety Policy Model Development*
- *Safety Policy Model Application*
- *Safety Policy Model Infrastructure.*

The Development Method produces a definition of safety for systems and components, and rigorously verifies and validates the safety requirements in a pragmatic and organised manner.

The Application Method contains guidance for the developer on how to build and use the safety policy model for a safety-critical system within the framework of a standard product development. It places the model and its development in the system life cycle, and allows one to see how the components of the model influence design.

Supporting the above two methods is the Safety Policy Model Infrastructure: a set of concepts, techniques and activities that make using the Policy Model easier and more effective. Example infrastructure components include a taxonomy of safety problems, using the Policy Model to specify safety integrity levels and a techniques matrix.

One issue surrounding these methods concerns balancing the cost effectiveness of creating any model of safety with the desire to be practical yet extremely rigorous. The approach being taken is that any modelling technique may be beneficially used to apply the SPaM philosophy and methods, subject to compliance with agreed integrity level requirements for the Policy Model.

#### *Safety Policy Model Development Method*

The Safety Policy Model Development Method consists of six stages, each designed to produce a part of the safety policy model. These are:

- D1 Policy creation.
- D2 Conceptual model creation.
- D3 Creating the logical model.
- D4 Policy substantiation.
- D5 Policy animation.
- D6 Produce rationale.

Development begins in stage D1 by creating a set of safety objectives which are referenced to a set of high level system hazards (e.g. returning incorrect data from a database or the incorrect biasing of a power transistor). Linking the system hazards to the safety objectives helps to demonstrate the 'safety problem' that the model is addressing, and will establish the model's scope. It provides the evidence that fundamental hazards are known and have been addressed accordingly.

After the safety objectives have been proposed, the rest of the Development Method concentrates purely on ensuring that the safety objectives are the correct ones for the system, and that they are complete. The degree of rigour is dependent on the safety integrity level of the system: the higher the integrity level, the more the assurance that is required. This assurance is gained by increasing the rigour of the techniques employed in the modelling stages of model development.

The benefits to the developer do not end purely with a correct set of safety objectives. By following the Safety Policy Model Development Method, major components of the system's design are specified (to be safe), and verification and validation evidence is created to support safety case requirements.

Stage D2 of safety policy model development concerns the building of the ‘mental’ conceptual model, e.g. the model of the part of the system the model is addressing. This should reflect the entities and events identified by the safety objectives and provide an idea to the developer on how the safety problems are being solved. This can be expressed in any convenient form or notation.

Stage D3 creates a logical model that allows some reasoning about the safety properties of the system. This model may be expressed in a mathematical formal language such as Z or using a semi-formal structured method such as Yourdon. Whatever modelling technique is used, it is necessary to develop and formalise the conceptual model in a notation that allows an adequate degree of reasoning about the correctness of the model. For example, by systematic argument using Yourdon diagrams or by formally verifying Z.

Stage D4 creates an argument that the safety objectives are consistent, mutually achievable and effective (correct). The argument is based on parts of the logical model. For example, this might involve the translation of a safety objective into a hypothesis (conclusion) to be tested, the translation of component parts of the logical model into a series of conditions (premises), and the demonstration that the hypothesis can be argued or proved to be true if the conditions are also true.

Stage D5 concerns the development of an animatable prototype from the logical model. From this prototype, and from the results of tests run on the prototype, it is possible to demonstrate sufficient confidence that that the safety objectives are valid and complete.

The final stage, D6, builds up a rationale that presents the verification and validation evidence to be used in a safety case context.

### *Safety Policy Model Application Method*

The application method allows the developer or procurer of a safety-critical system to use the safety policy model and the Model Development Method with their standard development life cycle. Figure 6.5 gives an overview of the Application Method and the interaction between stages.

The four boxes on the right-hand side of the diagram represent the conventional project documentation to which the Method contributes. The verification and validation evidence is that associated with the safety aspects and hence is ‘safety’ V&V, ie checking the safety requirements are valid and verifying that they have been met. The fifth box, the Safety Target, is a new concept representing the safety-related system to be evaluated or assessed, including its safety objectives and associated safety policy model.

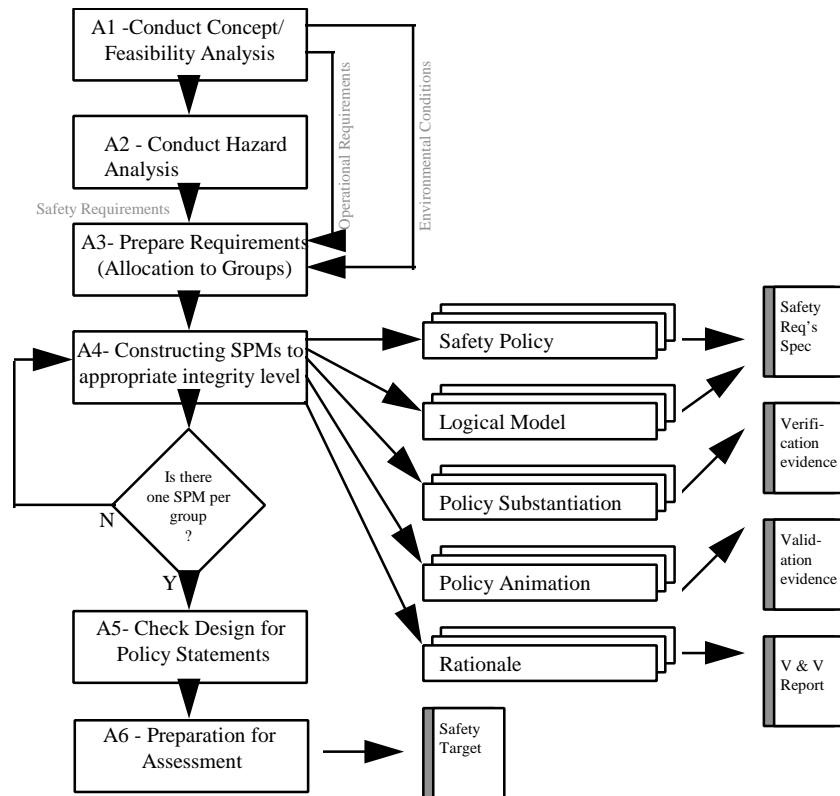


Figure 6.5. An Overview of Safety Policy Model Application Method

The four stages are as follows:

**A1: Conduct Concept/Feasibility Analysis.** The initial investigation of the system is conducted and information about the system gathered. This includes: system entities, information about the environment, operational characteristics and the system scope.

**A2: Conduct Hazard Analysis.** An attempt is made in understanding the safety implications of the system. This includes: hazard identification, events leading to hazards, probability and risk assignment, risk severity levels and safety integrity levels.

**A3: Prepare Requirements.** This is where all the information (e.g. functions) about the system is allocated to perceived components (e.g. assemblies, hardware, software).

**A4: Construction of Safety Policy Models.** The relevant information for safety policy model development becomes available during this stage allowing the models (for systems, components and sub-components) to be built. As the models are constructed, the outputs from this development are used to build up: the Safety Requirements Specification, Verification Evidence, Validation Evidence and the V & V Report.

**A5: Checks During System Development.** The system being developed is checked for conformance with the Safety Policy (high level system safety requirements as objectives).

**A6: Preparation of Deliverables for Assessment.** This stage collects the necessary information with which to build the safety target used as a basis of assessment.

*Security and safety parallels.*

In their comparison of security and safety issues, the SPaM project identified the parallels shown in Table 6.2 and proposed the terms in the third column of the table for use (at least) within the project.

<b>Security Term</b>	<b>Safety Term</b>	<b>SPaM Term</b>
Security Policy Model	Not Defined	Safety Policy Model
Security Objectives	Safety Goals	Safety Objectives
Security Enforcing Functions	Safety Functions	Safety Functions
Evaluation Level	Safety Integrity Level	Safety Integrity Level
Security Target	Overall Safety Requirements	Safety Target
Target of Evaluation	Not Defined	Target of Assessment
Evaluation	Assessment	Assessment
Evaluation Framework	Not Defined	Assessment Scheme
Assurance (Correctness & Effectiveness)	Integrity (Correctness)	Assurance (Correctness, Effectiveness & Management)
Commercially Licensed Evaluation Facility	Independent Safety Auditor	Assessor
Certification Body	Regulatory Body	Regulatory Body
Threats	Causes of Hazards	Safety Problems
Vulnerability	Safety Weakness	Safety Weakness
Mechanisms	Not Defined	Off the Shelf Algorithms
Corporate Security Policy (CSP)	Company Safety Policy	Corporate Safety Policy
System Security Policy	Overall Safety Requirements Specification	System Safety Policy
Technical Security Policy	Safety Functional Requirements Specification	Technical Safety Policy

Table 6.2. Parallel terms

Further information about the project can be found in the references on the project page in Appendix A or can be obtained from the project contacts on the same page.