EUROPEAN ORGANISATION
FOR THE SAFETY OF AIR NAVIGATION

EUROCONTROL

This Document is issued as EATMP Method and Tool. The contents are not mandatory. They provide information and explanation or may indicate best practice.

The Development of a Safety Management Tool within ATM (HERA-SMART)

Edition Number	**:**	**1.0**
Edition Date	**:**	**12.05.2003**
Status	**:**	**Released Issue**
Intended for	**:**	**EATMP Stakeholders**

EUROPEAN AIR TRAFFIC MANAGEMENT PROGRAMME

DOCUMENT CHARACTERISTICS

TITLE

The Development of a Safety Management Tool within ATM (HERA-SMART)

	EATMP Infocentre Reference:	030422-02
Document Identifier	**Edition Number:**	1.0
HRS/HSP-002-REP-08	**Edition Date:**	12.05.2003

Abstract

This report has been developed among a series within the Human Error in ATM (HERA) Project dealing with how error in Air Traffic Management (ATM) can be analysed and evaluated to improve safety and efficiency in European ATM operations. The purpose of this work is to develop an approach, using the Human Error in ATM (HERA-JANUS) classification technique (see EATMP, 2003), for safety management within ATM. The report describes the developments made of the 'Safety Management Assistance and Recording Tool (SMART)' approach over an eighteen-month period, and details the results and recommendations from this process.

Keywords

Human Error	Incident Analysis	Safety Management
Generic Initiator (GI)	Criticality	Safety Architecture (SA)
Recovery	System Safety	Safety Assessment
Safety Principles (SPs)		

Contact Person	Tel	Unit
Michiel WOLDRING, Manager, HRS Human Factors Sub-Programme (HSP)	+32 2 7293566	Human Factors & Manpower Unit (DIS/HUM)

Authors

J. Pariès, C. Bieder, J. Reason and A. Isaac

STATUS, AUDIENCE AND ACCESSIBILITY

Status		Intended for		Accessible via	
Working Draft	☐	General Public	☐	Intranet	☐
Draft	☐	EATMP Stakeholders	☑	Extranet	☐
Proposed Issue	☐	Restricted Audience	☐	Internet (www.eurocontrol.int)	☑
Released Issue	☑	*Printed & electronic copies of the document can be obtained from the EATMP Infocentre (see page iii)*			

ELECTRONIC SOURCE

Path:	G:\Deliverables\HUM Deliverable pdf Library\

Host System	Software	Size
Windows_NT	Microsoft Word 8.0b	

EATMP Infocentre
EUROCONTROL Headquarters
96 Rue de la Fusée
B-1130 BRUSSELS

Tel: +32 (0)2 729 51 51
Fax: +32 (0)2 729 99 84
E-mail: eatmp.infocentre@eurocontrol.int

Open on 08:00 - 15:00 UTC from Monday to Thursday, incl.

DOCUMENT APPROVAL

The following table identifies all management authorities who have successively approved the present issue of this document.

AUTHORITY	NAME AND SIGNATURE	DATE
Please make sure that the EATMP Infocentre Reference is present on page ii.		
HERA Project Leader	A. ISAAC	16.05.2003
Chairman HRT Human Factors Sub-Group (HFSG)	V.S.M. WOLDRING	16.05.'03
Manager EATMP Human Resources Programme (HRS-PM)	M. BARBARINO	16/05/03
Chairman EATMP Human Resources Team (HRT)	A. SKONIEZKI	16/5/03
Senior Director Principal EATMP Directorate (SDE)	W. PHILIPP	25.05.03

DOCUMENT CHANGE RECORD

The following table records the complete history of the successive editions of the present document.

EDITION NUMBER	EDITION DATE	INFOCENTRE REFERENCE	REASON FOR CHANGE	PAGES AFFECTED
0.1	21.11.2001		Working Draft	All
0.2	10.01.2002		Draft	All
0.3	19.02.2002		Second Draft	All
0.4	07.02.2003		Proposed Issue for HRT19	All (document configuration)
1.0	12.05.2003	030422-02	Released Issue	All (final document configuration adjustments)

CONTENTS

EXECUTIVE SUMMARY

Phase 1 of the Human Error in ATM (HERA) Project produced a detailed methodology and technique for analysing and learning from error-related incidents in ATM (see EATMP, 2002a, 2002b, 2003a, 2003b).

The general objective of Phase 2 of the HERA Project (HERA 2) is to investigate several specific areas associated with the prediction, detection and management of human error in Air Traffic Management (ATM), and to develop methods for the implementation of the results of these concepts at various levels of ATM safety management within Europe.

The specific objective of HERA2 is to explore more intensively the potential operational applications of the error analysis technique developed during Phase 1, in relation to four safety-related areas:

- to develop an approach using the HERA-JANUS Technique to investigate how human error can be detected and managed within a real-time simulated ATC environment: HERA-OBSERVE (see EATMP, 2002c, 2002d);

- to investigate the potential of the HERA-JANUS classification as a prospective tool within ATM (error prediction): HERA-PREDICT (report currently under preparation);

- to develop an approach using the HERA-JANUS classification technique for safety management within ATM: HERA-SMART (covered by this report);

- to develop teaching materials on the HERA-JANUS Technique for incident investigators and safety managers within several ECAC States (see EATMP, 2003c).

Page intentionally left blank

1. INTRODUCTION

1.1 The HERA Project

The Human Error in ATM (HERA) Project, Phase 1 (HERA 1), sought to review theories of human error and formulate a practical approach for analysing these errors within the ATM environment. This work arose as a result of increasing automation and the importance of error recovery and error reduction in ATM as the future traffic increases are predicted and as airspace structures are re-aligned to produce maximum traffic flow. The resultant work in this first phase established the rationale for a conceptual framework for this initiative. This conceptual framework outlined a model of human performance and the types of taxonomies that would be required to classify errors and contextual factors relating to ATM incidents. This technique was then used in various validation exercises to establish its robustness, efficacy and usability (see EATMP, 2002a, 2002b, 2003a, 2003b).

Reliability and variations in human performance are an important element in the understanding of aviation safety and in analysis and design of air traffic management systems. The first phase of the project established a framework for understanding human errors in ATM operations and has provided a basis for better categorising ATM incident data. Statistics and trends obtained from applying these concepts have provided a basis for the application to a range of ATM activities, such as incident analysis and to a lesser extent the prediction of human performance with new ATM tools. However, the dearth of similar work indicated that there was a need to extend this activity into another dimension, that of prediction, detection and recovery of human error within the ATM system.

The general objectives of the second phase of the project (HERA 2) are to investigate several specific areas associated with the prediction, detection, and management of human error in ATM, and to develop methods for the implementation of these concepts at various levels of the ATM system, such as safety training, safety management, incident investigation and the application of human error vulnerability within the system.

The specific objectives of HERA 2 are therefore the following:

- to develop an approach to investigate how human error can be detected and managed within a real-time simulated ATC environment: HERA-OBSERVE (see EATMP, 2002c, 2002d);
- to investigate the potential of the HERA-JANUS classification as a prospective tool to predict error-prone conditions within ATM: HERA-PREDICT (report currently under preparation);
- to develop an approach using the HERA-JANUS classification tool for safety management within ATM: HERA-SMART (covered by this report);

- to develop an approach, using the HERA-JANUS classification, for the training of incident investigators which incorporates an understanding of human factors and system safety aspects within the investigation process (see EATMP, 2003c).

1.2 Overall Work Plan and Focus of this Report

The overall work plan for this part of the HERA Project (HERA 2) is divided into four Work Packages (WPs), which reflect the objectives cited in the previous paragraph. Although the four work packages have been / are / will be explored separately, they typically have heavy dependencies. The following figure illustrates the inter-dependencies of each objective and work package, and their link with the HERA 1 work.

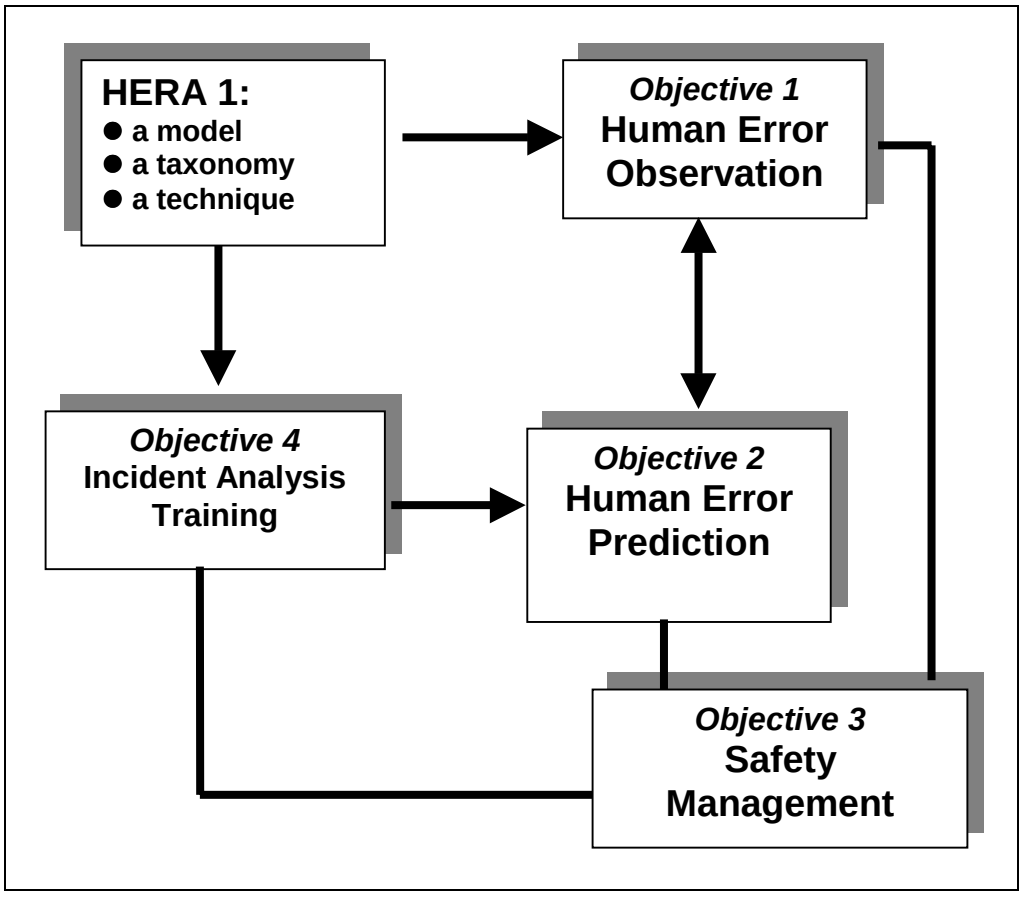

Figure 1: Overall work plan for Phase 2 of the HERA Project (HERA 2)

The present WP3 of HERA 2 describes the development of an approach using the HERA-JANUS classification tool for safety management within ATM. The report describes the approach developed over an eighteen-month period with the support and expertise of ATM incident investigators and safety managers from three ECAC States. Results from this work are detailed and recommendations for future work are discussed.

The remainder of this report contains sections and detailed appendices:

- Section 2, ‘Safety Management’,
- Section 3, ‘The Need for a Safety Management Tool’,
- Section 4, ‘A Complementary Approach - SMART,
- Section 5, ‘Methodology to Develop the Safety Architecture’,
- Section 6, ‘Using SMART with Examples’,
- Section 7, ‘Examples Using SMART’,
- Section 8, ‘Making Decisions in Risky Environments’;
- Section 9, ‘Summary’;
- Annexes: Bibliography, Further Reading, Glossary, Abbreviations and Acronyms, Contributors.
- Appendix 1, ‘Methodology for a Proactive Virtual Exploration of the Effects of Change on Safety’;
- Appendix 2, ‘Consistency between SMART and SOFIA Approaches’;
- Appendix 3, ‘ATM Organisational Safety Assessment’.

Page intentionally left blank

2. SAFETY MANAGEMENT

2.1 Introduction

The majority of accidents in hazardous activities and new technologies are caused by human error. This problem is not new and is discussed by a variety of authors from diverse research backgrounds (Chopra, Bovill, Spierdijk & Koornneef (1992), Hawkins (1987), Perrow (1984), and Reason (1990)). It has also been revealed that the human involvement in major disasters is distributed very widely, both within the organisation and often over several years before the actual event (e.g. the fire at the Three Mile Island nuclear plant, the Exxon Valdez oil tanker spill, the Piper Alpha oil platform fire, the Kegworth air accident, and the Kings Cross and Clapham Junction rail accidents).

A landmark case associated with the factors of an organisational accident was investigated in England in 1987. Mr Justice Sheen's judgement on the causes of the capsize of the 'Herald of Free Enterprise' (a European roll-on, roll-off ferry) went beyond the errors of the Master, the Chief Officer and the Assistant Bosun. He wrote,*....a full investigation into the disaster leads inexorably to the conclusion that the underlying or cardinal faults lay higher up in the company* (Sheen, 1987).

Since this time there have been several other reports into catastrophic system failures which have implicated organisation and management decisions leading to fatal consequences: the Air Ontario crash (Moshansky, 1992) and the Challenger Shuttle disaster (Vaughan, 1990).

In aviation, as in other complex technologies, we are in an era of what Reason calls the age of the organisational accident (1990, 1997). This concept addresses the problem of pre-existing and often long standing latent failures arising in the organisational structure.

In any disaster, as has been mentioned, there will typically be several behavioural pre-conditions which will have originated some years prior to the actual event. In the case of the space shuttle Challenger disaster, the initial faulty booster design decisions were made thirteen years before the disastrous flight. Also, and as a direct consequence of the complexity inherent in modern socio-technical systems, it is often a chain of unanticipated interactions between contributory causes that will lead to a disaster (Perrow, 1984).

The first analysis of the organisational pre-conditions leading to disaster in large technical systems was by Turner (1978). He conducted a detailed analysis of 84 major accidents over a ten-year period in the United Kingdom and concluded that prior to any disaster a number of undesirable events accumulate, often unnoticed or not fully understood. These events gradually develop over a number of years leading to the disaster itself. This

development is brought to a conclusion either by taking preventive action to remove the dangerous conditions where they are noticed, or by an event which might be a final critical error.

As with Reason, Turner discusses the fact that line operators often inherit faulty systems; either as a function of particular equipment, procedures or working practices or more directly as a result of decisions made elsewhere in the organisation. One problem after the event is that the immediate trigger may be confused with the more systemic background causes to the disaster, or may even be taken to be the sole cause.

2.2 Safety Management in ATM

There are no initiatives within air traffic management which are similar to the safety management activities in other hazardous organisations such as the nuclear and offshore petrochemical industries. The present work attempts to take the thinking from these other organisations and develop a new approach to the human factors associated with safety management in air traffic management.

There are many ways in which the use of an error classification and analysis tool such as HERA-JANUS can be envisaged for safety management within ATM. The goal of this work was to develop specifications for a **Safety Management Assistance and Recording Tool (SMART)**, a tool that would act as an interface between individual safety occurrences reports and safety management decisions. SMART is intended to be ultimately a software accessible by anyone in the organisation to retrieve and edit the data. This data would then be used to confirm the robustness of safety assumptions within the ATM organisation and to ultimately increase the strength of the safety management system.

2.3 Working Methodology

In order to develop an operational approach, suited to the needs of the ATS providers, it was decided to involve various national ATM organisations in this work. The following organisations participated in the development of this work:

- LVNL, The Netherlands,
- DFS, Germany,
- CENA, France.

Furthermore, in order to ensure the consistency of various safety management approaches under development within EUROCONTROL, a participation from the Safety, Quality Management and Standardisation (SQS) Unit was also established. The collaborative work was mainly conducted during dedicated technical meetings attended by the Core Development Team, that is the representatives of the Human Factors Team, the SQS Unit and operational ATM controllers in charge of incident analysis within their national organisations.

Eight technical meetings were held by the Core Development Team. During these meetings, the proposed principles[1] and approach was discussed and then examined with several cases from incidents submitted by the different participants from operational ATM organisations.

[1] These principles have been derived and developed from conceptual and applied research conducted by Airbus Industrie to improve the quality of lessons learned from the customer airlines operational experience.

Page intentionally left blank

3. THE NEED FOR A SAFETY MANAGEMENT TOOL

3.1 Safety Management and Feedback from Operational Experience

HERA-JANUS is a technique to analyse errors associated with activity in the ATM environment, and the question which is raised concerning these results is “How can such a technique and the results of incident analysis be used to better manage ATM safety?”.

Analysing incidents to improve safety is a common strategy in virtually all activities, including ATM. Incidents are seen as accident precursors, and therefore the number of incidents is usually considered an indication of the risk level reached by an organisation in its operations. Hence, one common goal of safety management is to minimise the number of incidents, at least the most serious ones, in which only good luck saved the system. Better understanding what caused incidents is expected to generate ideas to amend the design or the operation of the system, in order to make it safer.

Safety management also strives to be as proactive as possible. Consequently, it tries to learn lessons from smaller incidents, from minor failures or deviations, situated some distance from the accident itself. Following this method naturally leads to the addressing of front line operators’ individual errors.

The HERA-JANUS Technique follows this method. The technique defines an error as *any action (or inaction) that potentially or actually results in negative system effects, where more than one possible course of action is available*. Violations are considered *actions that contravene a rule, procedure or operating instruction* (EATMP, 2003a). Once errors and violations involved in an incident have been identified, the HERA-JANUS Technique supports the following analysis processes for each of the errors:

1. The recording of the task involved, the kind of equipment in use and the type of information involved.

2. The identification of the Error and or rule-breaking/violation type.

3. The identification of Error Detail.

4. The identification of the Error Mechanism which failed.

5. The assessment of the Information Processing level involved.

6. The identification of the Contextual Conditions involved.

The goal of this analysis is to understand how errors are generated by the Air Traffic Controllers (ATCOs)’ cognitive processes, and how ATM work contexts and situations, and more generally systemic factors (workplace design,

economic constraints, teamwork, organisational issue), can create or contribute to error-prone conditions.

However, if the ultimate goal is to manage ATM safety, an additional step is needed: we also need to analyse **how errors contributed to, or may contribute to unsafe situations**. This implies a reference to a model of the systemic safety.

A broadly accepted basic perspective on systemic safety is that the operation, if not the existence of a system like aviation, creates dangers that are kept under control by defences and protections. The interaction between dangers and defences, hence the efficiency of the protections, are influenced by both:

- the design of the system itself, including organisational factors; and
- the behaviour of front line operators, themselves under the influence of the organisation.

Reason's (1997) Model demonstrates the integration of these two dimensions:

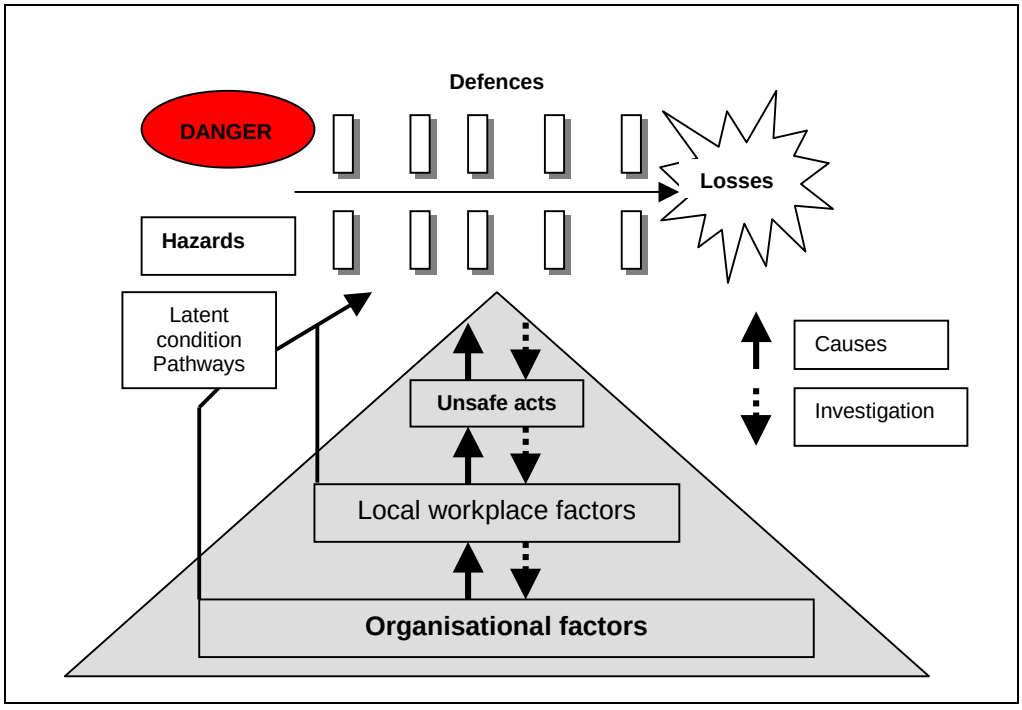

Figure 2: Reason's Model of fallible defences

The main principles of the model are as follows: all activities are exposed to dangers and hazards, and they are protected from these hazards by a series of in-depth defences. As a result of organisational factors, local workplace factors incline front line operators to commit unsafe acts (errors or violations), that weaken the defences, so that dangers can breach the defences. This is

the direct pathway to accidents. There is also an indirect pathway, through which organisational factors create latent conditions for fallible defences.

With this model in mind, what could be the link between safety management and the analysis of ATCO's errors?

Safety management is a concept concerned with controlling the level of risk in a system so that it remains within acceptable boundaries. This implies ways of defining what is an acceptable risk, ways of assessing what is the actual risk level, methodologies for action to reduce these, and tools to monitor the effect of actions taken.

The black solid arrows in Figure 2 above show the propagation of cause-effect processes from organisational factors to unsafe acts, through incident-prone work conditions. If we could follow the black broken arrows, going back from the effects - the unsafe acts - to the causes, we would understand how organisational factors influenced unsafe acts, or how latent conditions augmented the potential consequences of unsafe acts. Then we could understand how to change the system design to improve its resilience to unsafe acts, or to reduce the frequency of unsafe acts.

3.2 Limitations of Incident Analysis Strategies

Considering incidents to prevent accidents, and analysing errors and violations to prevent incidents, seems a wise and straightforward thing to do. However, although it may be wise, it is certainly not easy to do it effectively.

The statistical methodology limitations in a domain such as aviation incident investigation are well acknowledged. The conditions for statistical validity of an emerging correlation between safety indicators and behaviour are difficult to meet. Furthermore, what seems uneventful in an investigation may not be reported and therefore it could bias the remaining analysis. In other words, much of the time, we may use analysed data which is not representative of the real system.

Even a thorough analysis of individual incidents can have limitations, they may be based on a process of 'causal attribution', which is always multifaceted. A common and simple definition of the cause of an event is 'something that directly contributed to the occurrence of the event', in other words, something without which the event would not have happened. An example could illustrate this problem:

> A jet airliner suffered a bird strike at FL 80 during the initial approach phase in good visibility, and the windscreen was severely damaged, leading to an emergency landing. The following elements could be attributed as 'causes' of the event:
>
> - the presence of a heavy bird flock at FL 80 on the initial approach path;
> - the decision of the ATCO to maintain the flight at FL 80 for separation purpose;

- the fact that the crew did not look outside the windscreen at that moment.

Each of the above 'causes' can be decomposed into second level causes:

- the presence of the bird flock is due to migration and hunting;
- the controller had to maintain the flight at FL 80 to maintain separation from peak hour departures;
- the crew did not look outside at this very moment because they were busy checking their vertical flight profile from FL 80;

and so on, towards what is generally called the 'root' causes of the event.

With the above example, causes explain **how an event has happened**. However, an incident is an event which is not planned to happen normally. Therefore, what we need to explain is **why something that was not expected to happen actually happened**. And this perspective implies a different, and specific, expectation of the notion of cause in which the 'cause' of an incident is **anything that was supposed to prevent the incident from happening, and failed to do so**. This therefore includes the whole system, including the person, their tasks and the context in which they work. We can again use Reason's (1997) Model to illustrate this idea:

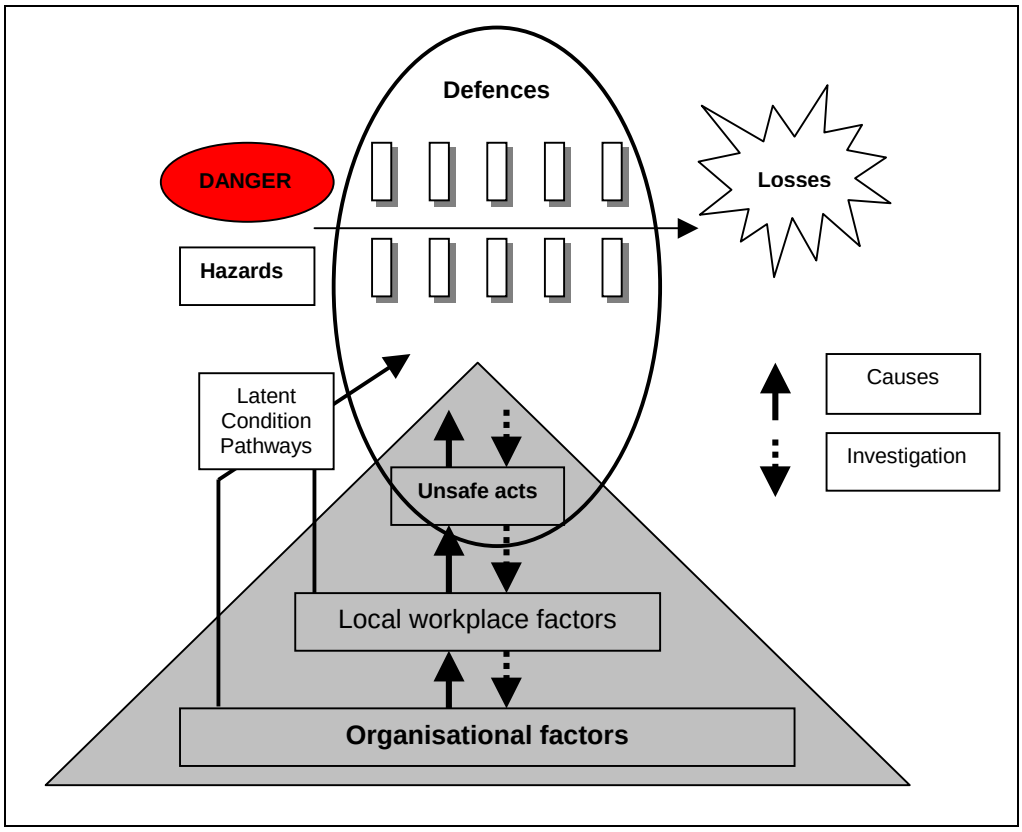

Figure 3: Defences and incident causation

Hence a ‘cause’ is something that failed in the defences within its interaction with the real time dynamics.

Therefore, when we identify the causes of an incident, we necessarily refer to a model of the system’s defences, or more generally, a model of what is supposed to make that system safe. If we take the bird strike example again: was the initial approach path supposed to be clear from birds, or was the controller supposed to make sure a flight level is clear from birds before using it? If the answer to both questions is no, then the presence of the bird flock, or the decision by the controller to maintain FL 80 are not ‘causal’ to the bird strike incident from the current safety model perspective. However, the ATM system requires the safety manager to respond to such incident evidence and this requires precise information in several areas. Safety managers therefore need information about the following:

1. What is the potential damage associated with such an event?
2. What is the frequency of that event (probability of such damage)?
3. What are the protections currently existing against this kind of event?
4. How efficient are they?
5. What protections could be improved, added? At what cost?

Often the answers are not currently provided by the incident reporting and analysis systems, in ATM or in other domains. However, there are many reasons for this:

- As already stated, incident analyses are based on a process of ‘causal’ attribution. Often this causal analysis does not clearly discriminate between the two meanings of ‘cause’ that we have discussed above. The causation model does not clearly refer to what was supposed to keep the system safe.

- Even when causes are identified from a safety model perspective, it is mainly done implicitly. There is generally no explicit description of the safety defences of the system. Consequently, the taxonomy and its attribution of ‘causes’, reflect the analyst and the analyst’s organisational thinking about safety. The problem here is that when one refers to a causation model, implicitly the causation model is taken as a truth, whereas it should be a falsifiable assumption.

- A significant part of the ‘causality’ of an event is context related: it lies in the very specific combination of circumstances, actions, failures and variations in that day. This context will not repeat itself, and lessons can only be learned from an event if some generalisation is achieved. The challenge is therefore to extract safety lessons from the cross-contextual elements that condense the system safety behaviour across all the occurrences.

- The causation models that are used are usually linear (one event or fact produces one effect, and so on). As noted by Rasmussen and Svedung (2000), *one major difficulty in the use of linear causal reasoning is that it is*

unreliable for analysing the behaviour of systems including closed-loop feedback functions. In that case, linear causal reasoning becomes circular.

- The notion of error itself is representative of this difficulty. If error is defined with reference to safety (*actions, or inactions, that can potentially or actually result in unsafe situations which can cause an accident* [EATMP, 2003a]), then the link to risk is a tautology and the link to psychology is not established. If error is defined with reference to cognitive process, then the link to psychological mechanisms are real but the link to safety is not established.

In order to circumnavigate the difficulties discussed above, a complementary approach has been developed, leading to the design of a software-based **'Safety Management Assistance and Recording Tool (SMART)'**. The next section describes the principles of this approach and the following sections will describe the implementation methodology, and the operating procedures.

4. A COMPLEMENTARY APPROACH – SMART

4.1 The Building Blocks

The approach proposed needs to analyse incidents through pre-identified risk management strategies. The core idea is to challenge the cause of incidents directly with the reasons why they happened, or should not have happened given the level of safety in the system. In other words to challenge the causes of incidents which correspond to safety assumptions in order to assess and map the strengths and weaknesses of these safety assumptions. The SMART approach uses several key components in its methodology:

- Firstly, an overall **Safety Architecture** which illustrates a specified ATM system: airport, terminal, en-route. This Safety Architecture can be visualised as a three-part structure concerned with the phases of any occurrence event: the prevention, recovery and mitigation phases (see Figure 6).

- Secondly, the mapping of information on which the Architecture is built. This information is taken from the **incident investigation process** and necessarily includes data regarding errors and contextual conditions (see Figure 5).

- Thirdly, the notion of **Safety Principles** which can be derived from the investigation process above and are found at each of the three phases of an occurrence, i.e. prevention, recovery and mitigation (see Figure 6).

- Lastly, the notion of **Generic Initiators**. A Generic Initiator is any event (or non-event) from which an occurrence would develop, should no specific recovery action be positively taken. These Generic Initiators are derived from the list of reported incident events within the incident investigation process. There are Generic Initiators created from each Safety Principle, and from each of the three phases of an occurrence, i.e. prevention, recovery and mitigation (see Figure 6).

Figure 4 overleaf illustrates these building blocks.

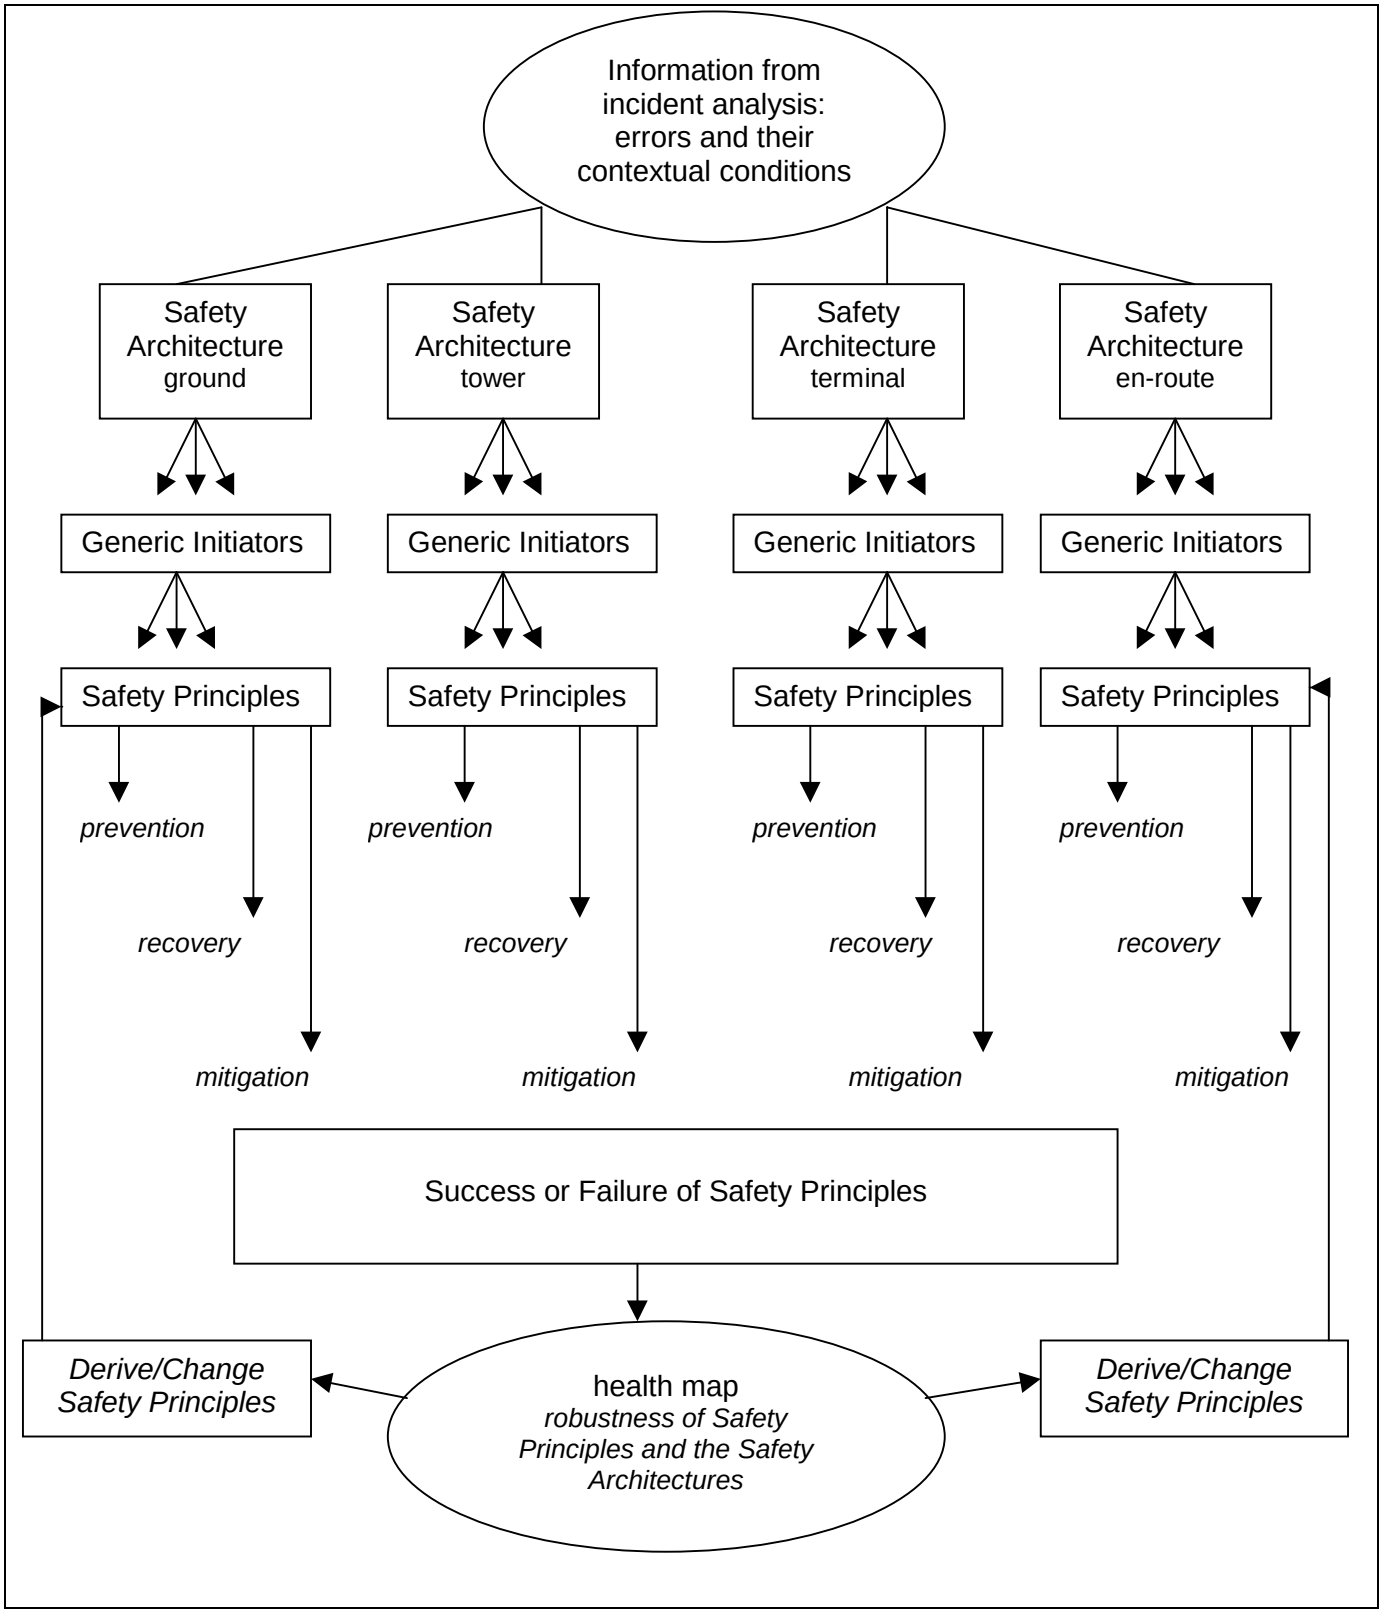

Figure 4: SMART components and structure

4.2 The Complete Model

Each of these components and their interdependencies are described in the following section.

One main goal of the Safety Management Assistance and Recording Tool (SMART) is to explain – from an ATM perspective - the Safety Principles (SPs) and assumptions about the safety of the ATM system, in order to assess/challenge them through feedback from operational experience.

However, the number of SPs meant to ensure the safety of ATM in general is so large that it would be impossible to list all of them for the whole ATM system. It would also be impossible for an analyst to check all of them when analysing an individual incident. Therefore, there is a need for a screening function to identify the relevant subset of SPs associated with a specific incident. Additionally, SPs do not work independently so their synergy must be represented.

Mapping the strengths and weaknesses of the safety assumptions or Safety Principles is illustrated in Figure 5 below.

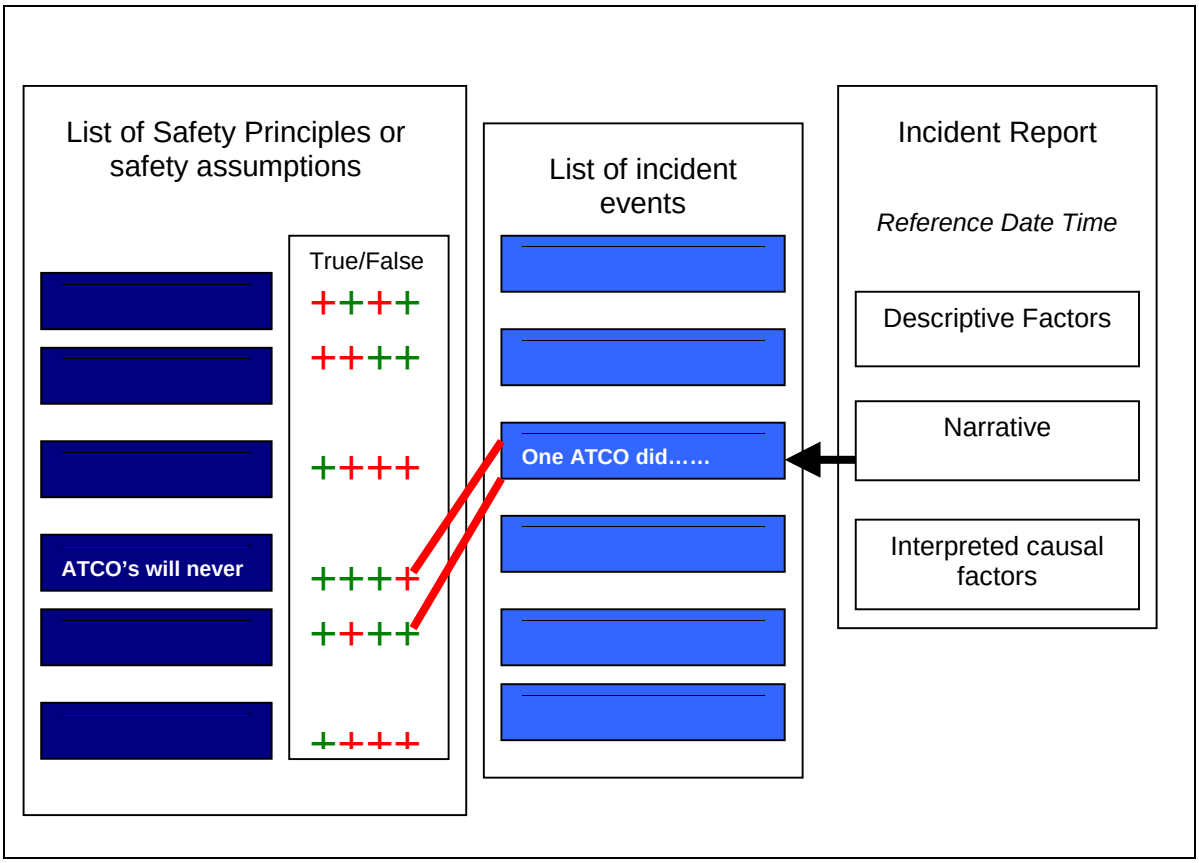

Figure 5: Mapping the safety assumptions or Safety Principles

A significant part of the 'causality' of an event is context related, and this context may not repeat itself in the same way. The challenge to extract a safety lesson is therefore to find cross-contextual elements. The solution suggested here is to seek this information through generic, prototypal paths to accidents. A prototypal path to an accident is a failure in the organised set of protections that are expected to prevent a specific type of accident. The main challenge is then to make the ATM Safety Model explicit in order to test it against reality through actual reported events or incidents.

For this purpose the notion of a **'Generic Initiator' (GI)** is introduced.

For any incident, the Generic Initiator (GI) would be identified, and for each of the Generic Initiator GIs, the Safety Architecture or logical combination of Safety Principles (SPs) that it relied on would also be identified:

- SPs intended for preventing the Generic Initiator from happening will be called **Prevention** SPs;
- SPs intended for preventing the Generic Initiator from developing into an accident will be called **Recovery** SPs;
- SPs intended for preventing the Accident from developing into its worse consequences will be called **Accident Consequences Mitigation** SPs.

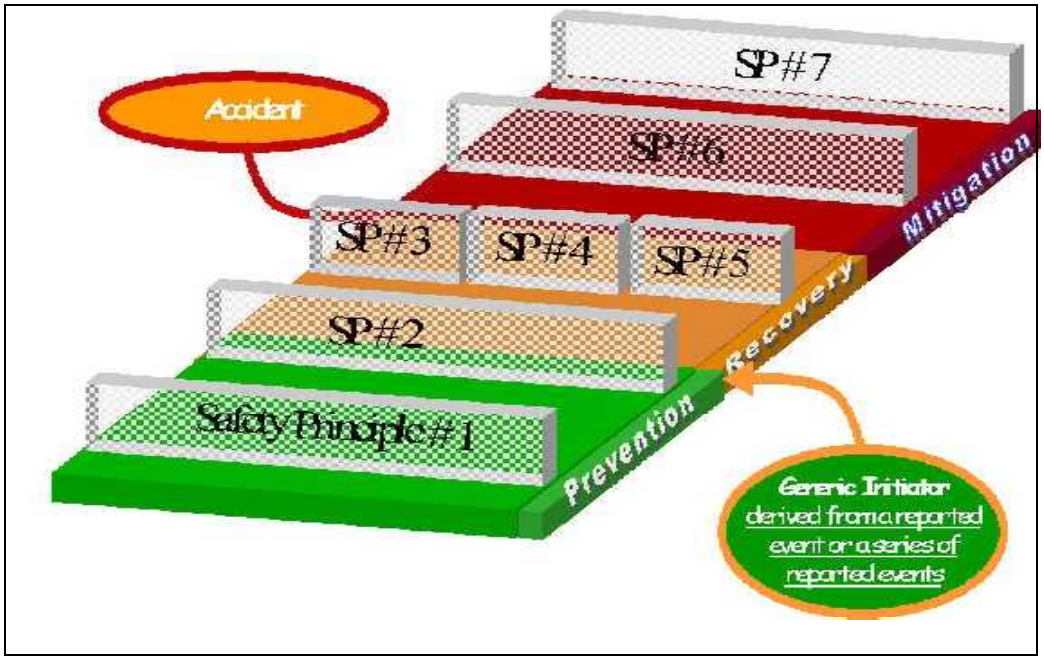

Figure 6: Safety Architecture: Safety Principles associated with a Generic Initiator

The method then challenges each Safety Architecture (SA) with the lessons from operational events. Indeed, the capability of SMART to store relevant information and capitalise on safety lessons is based on a pre-existing model

of SA. For that purpose all the possible Generic Initiators should be identified, and for each of them, the Safety Architecture explained.

Such a task could be achieved through an exhaustive top-down approach based on a functional safety analysis. However, it would take a huge amount of time to complete the work. Therefore, it was decided initially to develop a global ATM *a priori* safety model from a top-down functional safety analysis. In this way one could determine, on a case-by-case basis, the Generic Initiators associated with one reported event (or a series of similar reported events), and then develop the corresponding SAs[2].

[2] The development of the SAs remains totally top-down.

Page intentionally left blank

5. METHODOLOGY TO DEVELOP THE SAFETY ARCHITECTURE

This section describes the methodology to be used to develop the logical combination of Safety Principles (SPs) – within the Safety Architecture (SA). This methodology has been derived from the work performed by the Development Team whilst using several case studies.

5.1 Preliminary Definitions

5.1.1 Boundaries of the system

Any safety management approach focuses on a 'system', be it explicitly or implicitly. The scope of the system taken into account has a direct influence on the scope of the possible recommendations. The system should be large enough to allow a proper understanding of the relevant safety weaknesses, and it should be specific enough to be accessible to modification and improvement. For example, focusing on the individual ATCO using their radar screen and radio/telephone would at best lead to recommendations concerning these human-machine interactions, and leave teamwork and organisational aspects out of the scope of potential recommendations. Conversely, addressing an extremely wide system, such as the whole European ATM system, will only reflect shared problems and lead to very generic recommendations.

In this work it was decided that the **system** would include **operational ATM organisations the size of an ACC, as well as the corresponding ATM functions implemented on the aircraft flight deck**.

5.1.2 Accident or incident

It is necessary to define what is meant by an accident or incident in the suggested safety management approach. An accident or incident is no more than an event that is determined to be unacceptable and depending on the domain, it can be either a hull loss, a separation infringement or the loss of a certain amount of money.

In the approach suggested here an **accident or incident** is defined as the physical occurrence, i.e. runway excursion, collision: **the accident or incident is considered at the level of the whole aviation system rather than at the level of the ATM system**.

5.1.3 A reported event or incident

A reported event recounts an occurrence that will probably never occur again in the same way. Therefore, in order to make the lessons learned more generic, a reported event should be matched to a more generic story covering

a wide range of individual occurrences, all of which challenge safety in a similar way. A key concept to allow identification of these generic occurrences is the concept of a Generic Initiator (GI). An Initiator is an event at the ATM system's level from which an accident would develop, should no specific recovery action be positively taken. A GI is a high-level initiator encapsulating a set of initiators corresponding in a similar way to manage/impair safety.

5.1.4 Generic Initiators

The underlying vision of safety in this development is a dynamic vision. Following Weick's (1987) words, safety is seen as a *dynamic non event*. This means that safety is not seen as an absence of unsafe events (e.g. errors, violations), but as the result of the system being under control, in a dynamically stable, intrinsically safe state[3]. The different states of the system can then be represented metaphorically as shown by Figure 7:

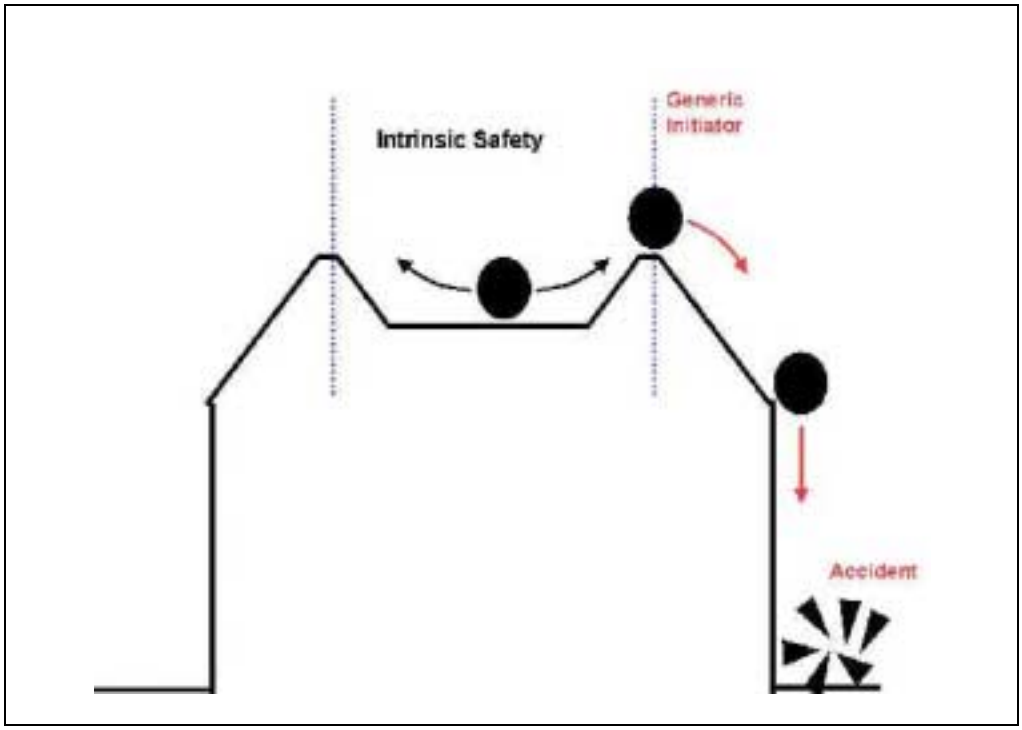

Figure. 7: A basic model of safety

[3] The similarities between this approach and the approach used in 'Sequentially Outlining and Follow-up Integrated Analysis (SOFIA)', which is the method developed by EUROCONTROL SQS Unit, can be found in the Appendices.

The ideas conveyed by the concept of Generic Initiator (GI) are twofold:

- it is **generic,** which means that it is independent from a particular instance or circumstance - a specific layout or piece of equipment or organisation;

- it is an **initiator***,* which means that at one point the system switches from a stable, controlled, intrinsically safe state to an unstable, uncontrolled, intrinsically unsafe state.

The methodology to identify the GI associated with a reported event therefore includes two steps:

1. Identification of the Initiator in the specific context of the reported event.
2. Generalisation of the initiator to derive the associated Generic Initiator.

Both steps are presented in detail below and are illustrated by a case study based on the fictitious incident.

Case Study

The incident occurred at Anytown Airport and involved the Ground Controller, Tower Controller and two flights, Flight ABC, and Flight XYZ. It was night with visibility of 1500m, ceiling overcast at 300 ft, and no wind.

23 01 10: Flight XYZ, a B767 from Airline B, is cleared by Ground Control to taxi to runway 36 and required to report at the holding point.

23 02 10: Flight ABC, a Beech 1900, is cleared by Ground Control to taxi and required to report “holding point runway 36”.

Note: Runway 36, [3900m long] had many access taxiways. Flight ABC, as usual, gets to its normal holding point, Delta 36. Delta 36 was set on a high-speed exit for runway 18, located at approximately a third down runway 36.

23 05 15: Flight XYZ reports “holding Alpha 36” to Ground Control and is cleared to contact Tower.

Note: Taxiway Alpha is the first taxiway at the threshold of runway 36.

23 05 20: Flight XYZ contacts Tower and is cleared to line up and take-off.

23 05 30: Flight ABC reports “holding runway 36” to Ground Control and is cleared to contact Tower.

23 05 35: Flight XYZ initiates take-off and reports at take-off to Tower;

23 05 36: Flight ABC contacts Tower and is cleared to line up and hold.

Note: Both the Ground controller and the Tower controller were unaware of the actual Flight ABC holding position, and assumed it was behind Flight XYZ at holding point Alpha 36.

23 05 42: Flight ABC moves on to the runway to line up.

Note: Flight ABC crews cannot see Flight XYZ at take-off, due to the acute angle between the high-speed exit centreline and the runway centreline.

23 06 05: Flight XYZ aborts take-off after the crew have seen Flight ABC moving on to the runway to line up.

Step 1: Identification of the Initiator

The objective here is to identify how, in the reported event the man-machine system became unstable and switched from an intrinsically safe state to an intrinsically unsafe state.

Questions to be asked in this phase are:

- *What overall ATM safety function failed?*

If we say that the system switched from an intrinsically safe state to an intrinsically unsafe state, it means that something failed in at least one of the basic safety functions. One then needs to refer to the three basic roles of ATM as far as safety is concerned:

– to ensure separation between aircraft,
– to inform,
– to alert.

The first thing to do is to identify which one(s) failed in the incident. We should therefore discriminate between three families of initiators, depending on the function concerned: **separation, information, alert**[4].

In the case study incident, the initiator belongs to the **'separation'** family.

Additionally, ATM services are commonly provided through four main types of control: **En-Route / Approach / Tower / Ground**. These four types of control actually correspond to four different types of work, with different goals and methods. Consequently, we should further discriminate between the four sub-families of initiators, depending on the type of control concerned: En-Route, Approach, Tower or Ground.

In the case study incident, the initiator belongs to the **'ground'** sub-family. **Consequently, the initiator is now characterised as follows:** ***separation/ground/…***

- *When, during the incident, would the natural course of things have led to an accident if no specific recovery action had been taken?*

An important challenge is to ensure that what is identified as the critical event actually is an Initiator, i.e. an event that actually represented the switch between stable and unstable safety states. Consequently, we need to identify

[4] By convention, we will consider 'alert' as a main ATM safety function when related with an accident not caused by ATM, and as an accident consequences mitigation safety principle when related to an accident caused by ATM.

at what moment, within the course of the event, the natural course of things would have led to an accident if no specific recovery action, had been taken.

> In the case study incident, the critical event is when **Flight ABC moves from the holding point to line up on the runway.**

- *What triggered the unstable safety state?*

The answer to the previous question should be helpful to identify what exactly triggered the switch between stable and unstable safety states. In other words, what happened that suddenly led to the loss of one of the main ATM safety functions.

This requires reasoning at the man-machine system level (rather than at an individual or equipment level) to establish how the functional safety objective is achieved in the real world.

For example, separation is maintained by ensuring that protection safety spaces around aircraft are not compromised. Therefore, identification of the Initiator through the following questions can be established:

- What were the protection 'space(s)' of the aircraft involved? (e.g. runway strip if aircraft is on the runway; safety space for cruising aircraft; initial/final approach fix to runway for an approaching aircraft with no radar system).

- In what way were the protection safety spaces broken?

In most cases the Initiator will be the infringement to the separation rule itself.

> In the case study incident the protection safety space associated with the aircraft at take-off is the part of the runway strip ahead of it. The Initiator could be worded as follows: **"Abnormal penetration of the runway strip by an aircraft during take-off operations"**.
>
> **Consequently, the initiator is now characterised as follows:**
>
> ***Separation / Ground / Abnormal penetration of the runway strip by an aircraft during take-off operations.***

Step 2: Identification of the associated Generic Initiator

The issue now is to 'generalise' the initiator identified above. The main characteristic of a Generic Initiator (GI) is that it describes a typical, but detailed, way in which safety can be impaired and managed (prevented, recovered or mitigated). For example, for an incident concerning radar control, the GI will not encompass non-radar environments, because radar control implies specific safety strategies.

On the other hand, a GI should be generic enough to include, incorporate, and represent all the events that are assumed to be managed in a similar way. It should be as independent as possible from the specific equipment or organisation, as long as these elements do not have a major influence on the safety management. It should also be independent from the configuration of the runways and taxiways, or from the position of the tower, unless these specific features play a major role in the way safety is managed.

In brief, a GI should encompass all the initiators that would impair safety, or would be handled in a similar way. Therefore, the objective of this second step is to abstract the initiating event as identified from the particular incident analysed.

The methodology used in this phase is known as a **substitution test**. The substitution test consists in reviewing the environmental or contextual conditions involved in the incident, in order to assess if they could be changed without changing the safety management to prevent or recover from the Initiator, or mitigate the consequences of the accident. The basic question to be asked in this phase is:

- *What if the environmental or contextual conditions had been different?*

The following aspects should be reviewed:

- flight rules,
- phase of flight,
- ATS airspace types,
- ATS airspaces/routes,
- restricted areas,
- types of airspace in relation with the applicable vertical separation,
- types of separation vertical/horizontal,
- runway configuration,
- taxiway configuration,
- meteorological conditions.

If the global *a priori* strategy associated with the reported incident is strong when substituting one aspect, then the GI is independent from that considered aspect. In this case the wording should be modified if needed.

> In the case study incident, the Initiator was worded as follows: **"Abnormal penetration of the runway strip by an aircraft during take-off operations".**
>
> This wording is independent from flight rules, ATS airspace types/routes, runway configuration, meteorological conditions, restricted areas, types of airspace in relation with the applicable vertical separation, and types of vertical/horizontal separation.
>
> ... / ...

However, the wording is dependent on the phase of flight, as it considers a runway incursion while a take-off is performed. However, the same basic threat to safety remains if a landing, instead of a take-off, is considered. Additionally, the same safety problem may remain if, instead of an aircraft incursion, the take-off/landing operation took place while any obstacle was present on the runway strip. Conversely, penetrating the runway strip behind a take-off/landing may be a normal situation.

Consequently, the Generic Initiator characterisation should be changed as follows:

Separation / Ground / Abnormal start of a take-off/landing on an occupied runway

5.1.5 Safety Principles

The Safety Architecture (SA) associated with a Generic Initiator (GI) is the comprehensive logical combination of Safety Principles (SPs) (i.e. their associated argument in terms of 'and' and 'or' statements) that is expected to prevent the GI occurrence, or prevent it from developing into an accident, or mitigate the accident consequences.

The SPs are the assumptions made by the designers (in a broad sense) and the operators of the ATM system about the behaviour of the ATM system within the environment in which it is embedded.

These assumptions are meant to ensure that any traffic within the boundaries addressed by design are safely handled by ATM. These SPs are in the first place defined from the ATM system point of view.

5.1.5.1 Identifying Safety Principles - general

As already mentioned, three main families or categories of Safety Principles (SPs) can be listed:

- **Prevention Safety Principles** are meant to prevent the GI from occurring;
- **Recovery Safety Principles** are meant to prevent a GI from developing into an accident;
- **Accident Consequences Mitigation Safety Principles** are meant to mitigate the consequences of an accident, once an accident has occurred.

The method to identify the SPs proceeds along these three categories.

The identification of SPs must be based on a functional approach. In other words, the reasoning to identify Safety Principles is mainly based on a '**how**' questioning process as presented hereafter. A reasoning based on '**there is'**

statements on existing equipment or features (there is a radar, there are two ATCOs) is not recommended as this approach leads to too much detail on certain aspects, and complete omission of others.

In addition, the method starts with the 'high-level' Safety Principles, then each SP is decomposed into a logical combination of lower level SPs, and so on. Consequently, a Safety Architecture (SA) looks like a genealogy tree.

The method goes from the most abstract (strategy) to the most concrete (expected behaviour of the ATM system, its components and interactions). It may be helpful to see the successive levels of SPs as levels in a means-ends abstraction hierarchy. Thus the highest level describes the overall protection strategy, the objectives to be achieved in the real world. The next level describes how the ATM system manages comply with this strategy through design, operations, and training. A further level addresses the functions to be performed to reach the objectives. A next level tackles the processes needed to implement the functions. The last level deals with the resources needed (people, skills, time, equipment). In practice, it may be too difficult to follow such a means-ends abstraction hierarchy and, therefore, it should only by used as guidance.

Some general rules can help speed up the Safety Principles identification process such as:

Whenever an SP refers to something the air traffic controllers are expected to do (or not do), it can be decomposed into three sub-Safety Principles linked together with an OR/AND relationship as follows:

- ATCOs know they are expected to do it (or know the rule) OR/AND
- ATCOs are able to do it (or to follow the rule) OR/AND
- ATCOs are willing to do it (or to follow the rule) OR

5.1.5.2 Identification of Prevention Safety Principles

How is the ATM system *a priori* protected against the occurrence of the GI? The answer to this question will determine the top-level prevention SP that conveys the overall prevention strategy (in a non-radar-equipped context the overall protective strategy for the approach phase is: No crossing of initial approach fix unless preceding aircraft has terminated its approach (landing/go-around)).

- *How does the ATM system manage to comply with this Safety Principle?* - How is the crossing of initial approach fix made 'impossible' unless preceding aircraft has terminated its approach?

- *What are the design, operations, training philosophies, supporting the compliance with this Safety Principle?* - Is there a rule saying so? Does equipment design make its implementation possible?

- *How does this translate at the level of detailed design, operations, training,… specifications?*

- *How is the system, its various components and their interactions (with one another and with the environment) supposed/expected to behave in practice to achieve this?* - What are the corresponding expected behaviours?

The answer to these questions allows the identification of the SP at a level of detail which is clear enough to model the Safety Architecture and build up the physical model associated with the prevention of the identified GI.

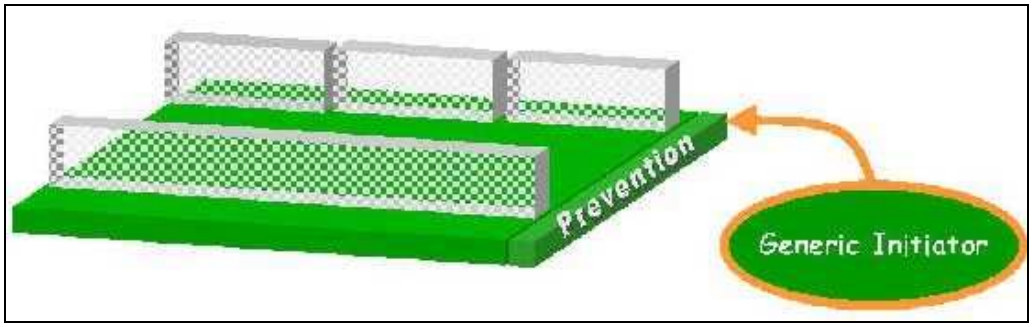

Figure 8: Prevention Safety Principles

5.1.5.3 Identification of Recovery Safety Principles

How is the ATM system *a priori* protected against the development of the GI into an accident? The answer to this question will determine the top-level recovery SP that conveys the recovery strategy: 'Discontinuation of approach':

- *How does the ATM system manage to comply with this Safety Principle?* - How is the discontinuation of approach made possible?

- *What are the design, operations, training philosophies, supporting the compliance with this Safety Principle?* – Are go-around clearances prescribed by operations?

- *How does this translate at the level of detailed design, operations, training,… specifications?*

- *How is the system, its various components and their interactions with one another and with the environment, supposed/expected to behave in practice to achieve this?* - What are the corresponding expected behaviours?

The answer to these questions allows the identification of the SP at a level detailed enough to model the Safety Architecture and build up the physical model associated with the recovery of the identified GI.

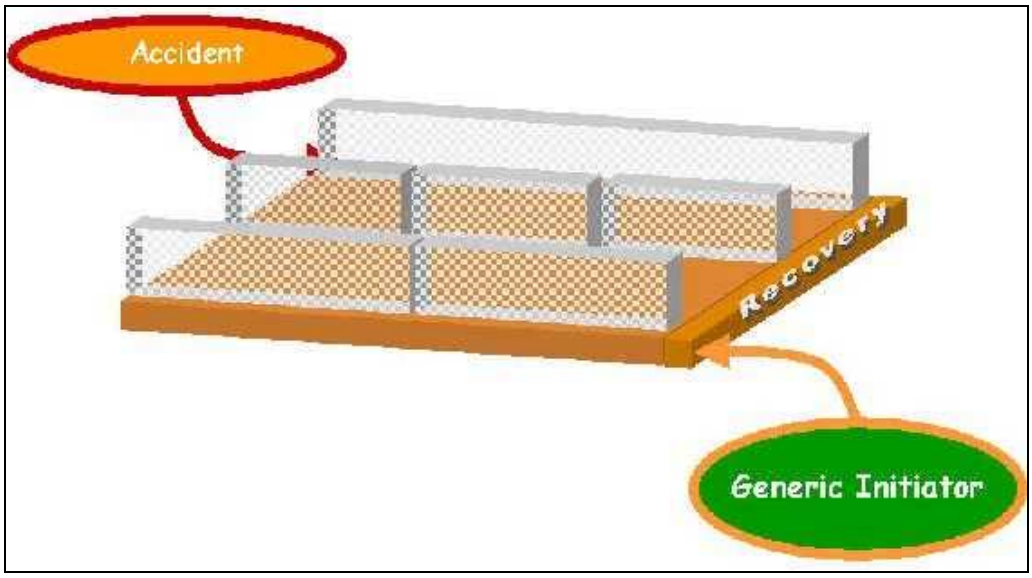

Figure 9: Recovery Safety Principles

5.1.5.4 Identification of Mitigation Safety Principles

Assuming that the accident occurred, how is the ATM system *a priori* protected against the occurrence of the worst consequences further to the accident? The answer to this question will determine the top-level mitigation SP that conveys the mitigation strategy: accident alert, fire and rescue guidance.

- *How does the ATM system manage to comply with this Safety Principle?*
- *What are the design, operations, maintenance philosophies, supporting the compliance with this Safety Principle?*
- *How does this translate at the level of detailed design, operations, training,… specifications?* - Are accurate maps of the airport are available?
- *How is the system, its various components and their interactions, with one another and with the environment, supposed/expected to behave in practice to achieve this?* - What are the corresponding expected behaviours?

The answer to these questions allows the identification of the SP at a level detailed enough to model the SA and build up the physical model associated with the consequences mitigation of the identified GI.

5.1.5.5 Level of detail of Safety Principles

Each Safety Principle (SP) can in turn be broken down into a combination of more detailed SPs. This process has no definable end, therefore, there is a

need for a decomposition stop-rule. There is no absolute rule to define the appropriate level of detail at which Safety Principles should be expressed. However, an intuitive stop rule is that they should be expressed at a level of detail consistent with the information available in occurrence reports, or expected to be gathered through the investigation process. Indeed, since the SMART philosophy is associated with putting SPs to the test of reality through real events, tuning the level of detail of SPs in accordance with the level of detail of operations feedback information would seem the most reasonable option. Finally, it is worth noting that nothing in the design or the use of SMART requires that all Safety Principles are decomposed at the same level of detail for all Generic Initiators (GIs).

However, it must be said that such level of detail may turn out to be insufficient when it comes to deriving operational recommendations. While the conformation of Safety Principles with real events, at this level of detail, should allow the identification of where problems lie, the same level of detail may not allow the understanding of where these problems originate.

This is the role of the next step: exploring the **'failure modes'** of the Safety Principles. Once an SP or a group of SPs have been shown to be failing repetitively, the incidents in which these failures occurred can be selected and a search for a potential explanation for such a failure can be undertaken. Looking into the associated incident reports, or investigating further, may lead to the factors that were shared by all the events: common contexts, common tasks, common type of error, and so on. Here the outcomes of analyses based on the HERA-JANUS Technique can be used.

Page intentionally left blank

6. USING SMART WITH EXAMPLES

The concept of a Generic Initiator (GI) is a key concept to screen the relevant Safety Principles for a specific reported event. Indeed, matching a reported event with one or more GIs allows the direct linking with the Safety Principles that were likely to be involved in this event.

6.1 Matching the Reported Event with a Generic Initiator

Because GIs within the Safety Architecture (SA) may be matched with several similar events the first question which should be asked is: **"Can the reported event be matched with any of the Generic Initiators already determined?"**

A GI describes a way of impairing and managing safety and may be restricted to certain environmental conditions, should these influence the way safety is dealt with. Therefore, determining whether the reported event can be matched to an available Generic Initiator consists of:

- going through the list of existing GIs;
- determining whether the specific initiator that occurred during the event is one instance of any of these GIs;
- checking if the conditions in which the specific initiator occurred are compatible with the conditions (if any) restricting the scope of the GI[5].

If a Generic Initiator can be found in the list, then the incident can be processed with that GI.

If no Generic Initiator (including potential environmental conditions) can be found in the existing list, then a new GI has to be defined according to the process described previously, and the associated SA has to be developed according to the process described in the previous sections.

Once the incident is matched with a Generic Initiator, the Safety Principles to be considered are simply those involved in the safety architecture associated with the GI identified as illustrated in Figure 10.

[5] In order to make the report more easily readable, GI will refer to the type of separation infringement plus the environmental conditions possibly characterising it. In practice, these conditions will be mentioned in the label of the GI, e.g. separation infringement en-route in a radar-equipped ATCC.

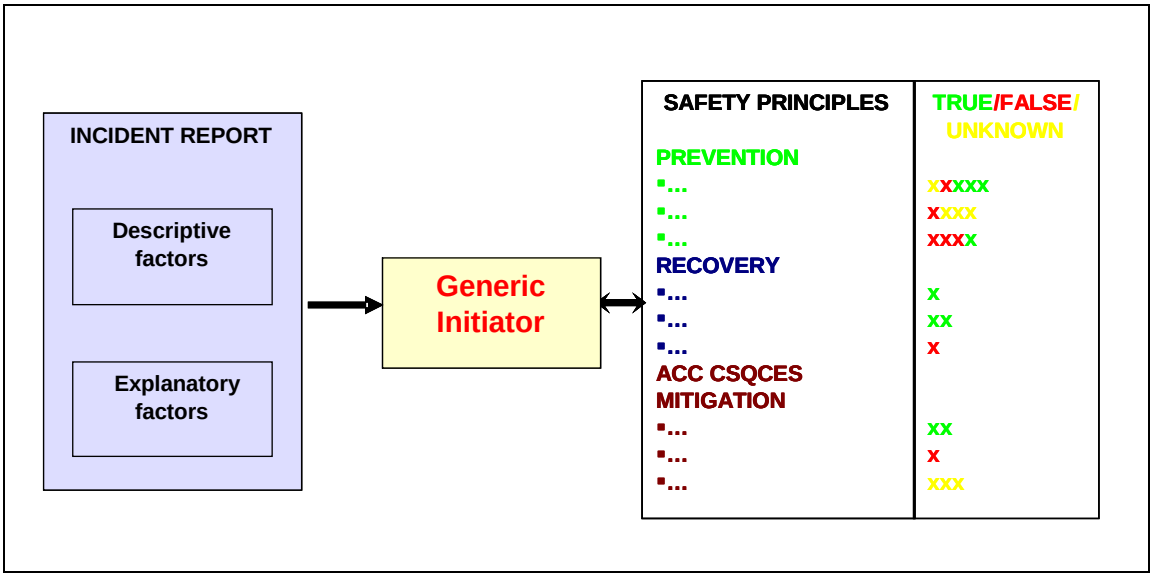

Figure 10: Assessing Safety Principles through an incident report

6.2 Amending the Safety Architecture

A first lesson learned while matching an incident report in detail with the Safety Architecture is that the Safety Architecture may be wrong. The analyst may find out that something in the Safety Architecture is inconsistent with the real world, as cited in the event. It may happen that, according to the Safety Architecture logic, the reported event should not have happened, whereas in reality it did happen. Conversely, it may happen that, according to the Safety Architecture logic, the reported event should have led to an accident, whereas in reality it did not.

In brief, if the reported event conveys information on the actual safety system that turns out be inconsistent with the Safety Architecture, the Safety Architecture has to be modified accordingly. This will be addressed more specifically in a future section.

6.3 Assessing the Empirical Robustness of the Safety Principles

6.3.1 Assessing the behaviour of the Safety Principles involved in an event

Once the reported event has been related to a Generic Initiator (GI), the Safety Principles (SPs) potentially involved in the event are those included in the associated Safety Architecture (SA). Assessing these SPs then consists of going through the following list:

- which of the SPs failed;
- which of the SPs were successful;

- which of the SPs were called upon but had an unknown outcome;
- which of the SPs were clearly not prompted. This recording will generate data about the prompting frequency of an SP.

The behaviour qualification (**success**, **failure**, **unknown outcome**, or **not prompted**) of the Safety Principles involved in one event translates into colour codes attributed to the link established between the corresponding SP and the event. This is illustrated in Table 1:

Table 1: Behaviour qualification of the Safety Principles

SP behaviour qualification	Meaning	Link colour code
Success / True	The SP was called upon during the event and actually contributed to safety as expected. In other words the underlying safety assumption took the logical value 'True' during that event	Green
Failure / False	The SP was called upon during the event and failed to contribute to safety as expected. In other words the underlying safety assumption took the logical value 'False' during that event	Red
Unknown outcome	The SP was called upon during the event, but its actual contribution to safety cannot be determined (e.g. lack of data). In other words the logical value ('True' or 'False') of the underlying safety assumption during that event cannot be determined	Amber
Not prompted	The SP was clearly not called upon during the event	White

6.3.2 Assessing the strength of Safety Principles

In practice, a Safety Architecture (SA) is a construction in which each SP is decomposed to a certain level of detail. Therefore, the question arises in relation to the level of detail that should be chosen when assessing the robustness of SPs against an incident event? If the analyst remains at the top level, the lessons learnt could be rather limited as the top-level SPs are often specific to an SA. Indeed, since at each level, the SPs are decomposed into a logical combination of SPs with AND and OR relations, the information of the weakness/robustness of SPs propagates logically in a bottom-up direction.

The following examples illustrate both bottom-up and top-down processing of the SP's assessment.

Let us consider a top-level SP - SP1, that can be decomposed into SP2 AND SP3 OR SP4. The SP1 decomposition can be represented as follows.

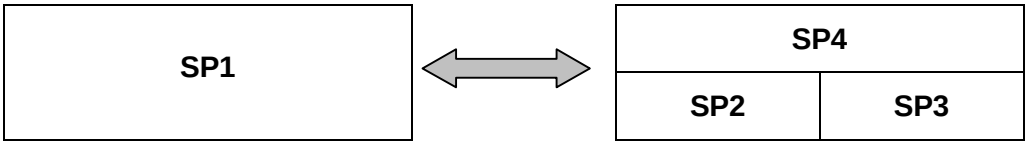

Illustration of a bottom-up process - assessing the robustness of the most detailed SPs

- If during the incident analysed, SP2 **AND** SP4 turned out to be false, logically, it is possible to affirm that SP1 was false - see following sequence.

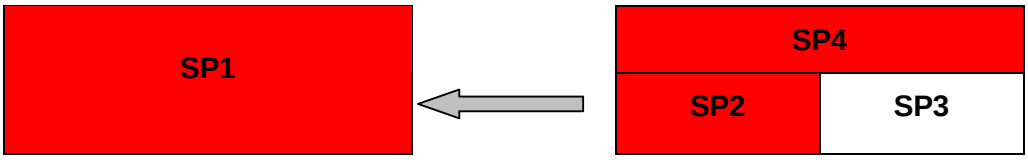

- In the same way, if both SP3 **AND** SP4 turned out to be false, logically, it is possible to affirm that SP1 was false - see following sequence.

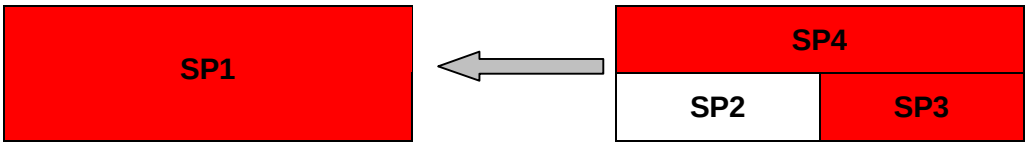

- If SP4 was true during the event, or if both SP2 and SP3 were true during the event, then, SP1 was true during the event.

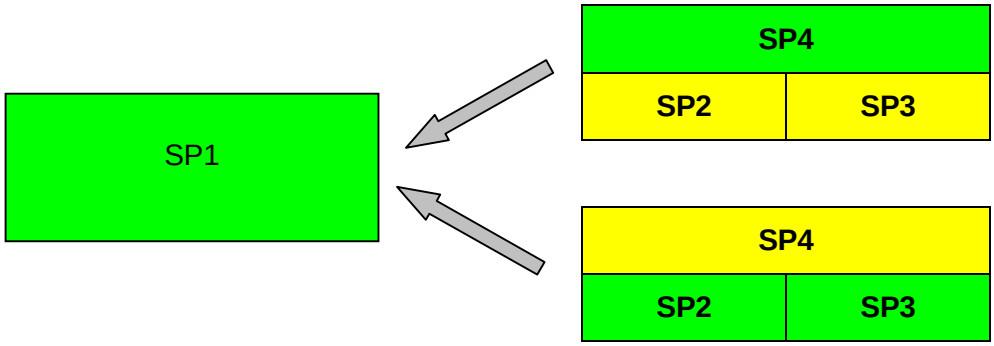

Illustration of a top-down process - assessing the robustness of the most general SPs

- If SP1 was true during the event, the only thing that can be said is either both SP2 and SP3 were true during the event and/or SP4 was true during the event. In other words, nothing can really be said on the robustness of SP2, SP3 and SP4 apart from conditional assumptions, which will not be helpful when it comes to decision-making.

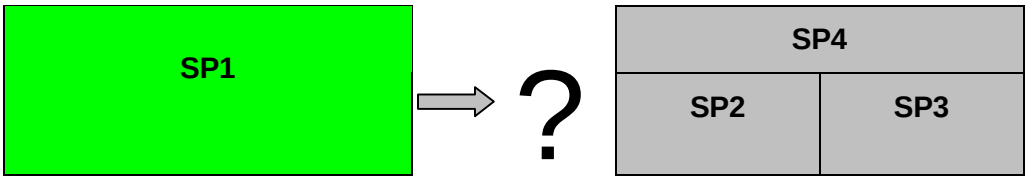

- If SP1 was false during the event, the only thing that can be said is either SP2 or SP3 was false during the event **AND** SP4 was false during the event. In other words, nothing can really be said on the robustness of SP2, SP3.

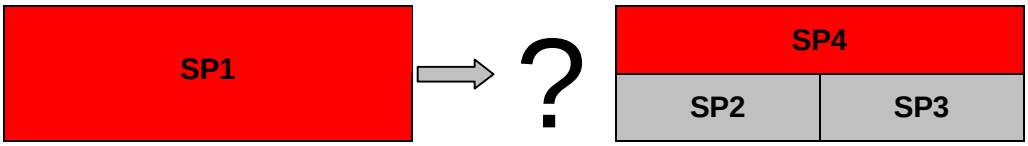

It is strongly recommended to start from the most detailed Safety Principles to assess their behaviour during the event. Less detailed Safety Principles should only be addressed if the information available in the report does not allow the handling of more detailed SP levels.

6.3.3 Completing the Safety Principle health map

While the process of assessing SP behaviours is repeated, incident report after incident report, the SMART system starts to accumulate the coloured links declared for each SP (be it success, failure, or called upon with unknown outcome). At this stage, SAs are no longer discriminated. If an SP is shared by several SAs, and if links to event reports have been declared for this SP through several GIs (hence several SAs), then all these links will be gathered. The series of links associated to one SP will then form what could be called the trustworthiness or 'health map' of that Safety Principle, as illustrated in Figure 11 below.

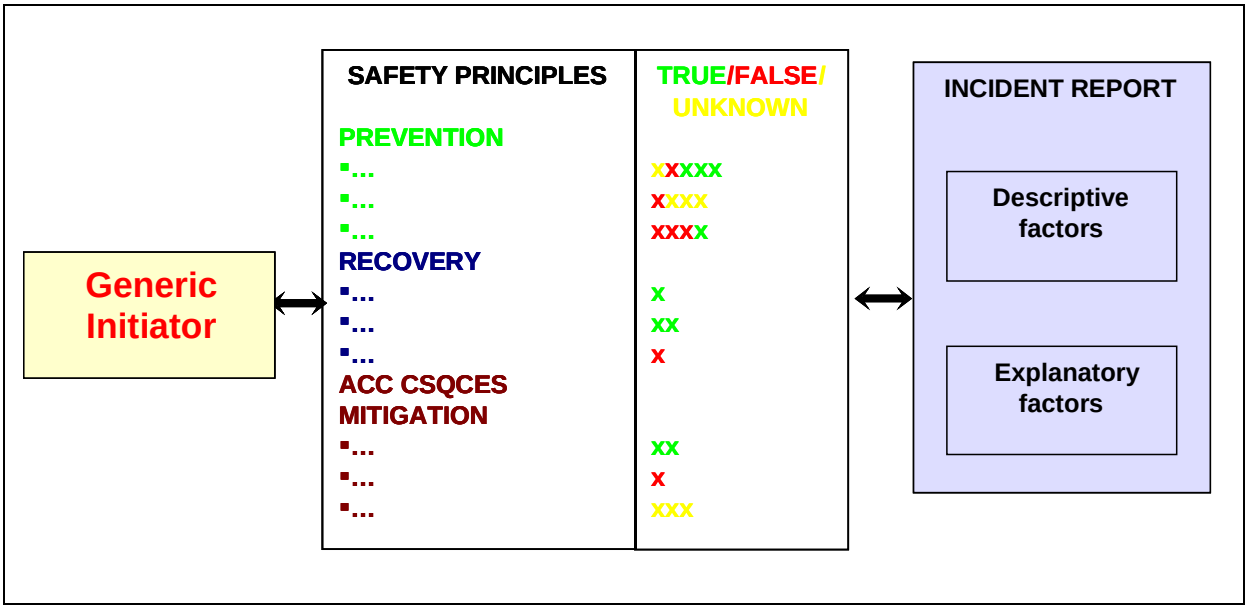

Figure 11: Building up the Safety Principles health map

6.3.4 Assessing the empirical robustness of Safety Principles

The next step is to interpret the health map of each Safety Principle to assess its empirical robustness, in other words, the level of trust that can be placed in it on the grounds of the knowledge gained through an interpretation of the available feedback from experience.

At this stage a display of the health map of the SP can be explained by Figure 12 below.

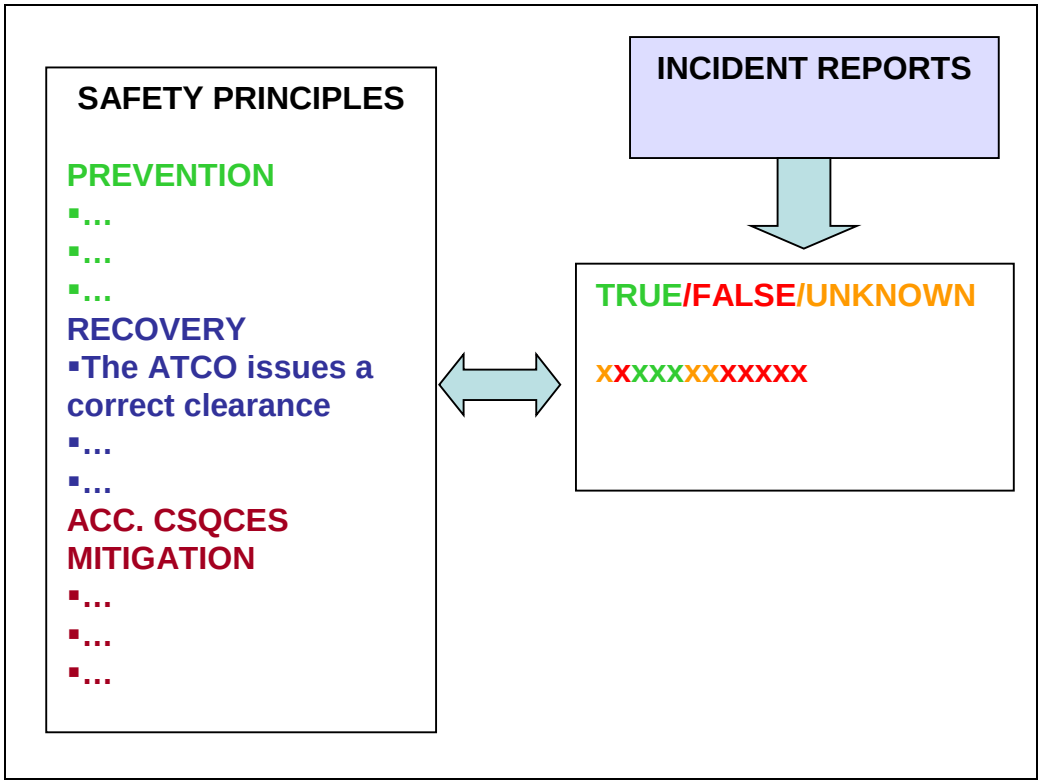

Figure 12: Safety Principle health map

This health map provides a sample of the actual behaviour (success / failure / unknown outcome) of the SP during reported events.

Whenever the health map is robust enough (i.e. allows for decisions to be made in relation to the appropriate level of trust that can be placed in an SP), a four-level rating scale of the SPs reliability will be used. A colour coding of the ratings will be implemented in the SAs displays - SPs will appear under their corresponding rating colour to better visualise the consequences of a rating on the Safety Architecture.

Table 2: The assessment

SP robustness rating	Meaning	SP colour code[6]
Reliable	The SP can be relied on: its reliability is up to that assumed by the *a priori* safety model	Green
Unreliable	The SP can no longer be relied on: its reliability is significantly less than assumed by the *a priori* safety model	Red
Unsure	The SP has been called upon, but available information prevents assessing its reliability (a lot of uncertainty still affects the assessment).	Amber
No rating	No assessment could be done because of lack of data, or excessive uncertainty	White

These ratings are expected to be the outcome of an expert judgement, possibly by collective consensus. The health map is only expected to support that judgement process through a colour coded visualisation of the available experience.

However, the notion of probability is not really intuitive, and it sometimes leads to misinterpretations. When empirical reference to actual events occurrence or frequency is involved, it might be more difficult to understand. For instance, the probability for a coin thrown in the air to come down tails is 0.5 if the coin is not fixed. If in one, two of even one hundred attempts it comes down heads, the probability for the same coin coming down tails in the next attempt is not increased by the experience of feedback. For most people the occurrence of an incident, not to mention an accident, is evidence that the system is unsafe and that more accident/incidents will occur. However, this is not necessarily the case.

In order to avoid any inaccurate interpretation of the health map, one should consider the following:

- a ‘red link’ associated with a Safety Principle (or even more than one red link) does not mean that the actual reliability of the Safety Principle is lower than that expected;

- in the same way, exclusively green links for a Safety Principle does not mean that the Safety Principle can be trusted all the time;

[6] The same colour coding has been selected for: 1. The links between SPs and events, 2. The rating of the robustness of an SP. This does not mean that they have the same meaning. The link qualification pertains to the local behaviour of one SP during one event. The SP robustness qualification pertains to a level of trust found in the SP across all SAs.

- 'no link' should not be interpreted as a green link by default. On the contrary, the absence of experience feedback on an SP, depending on its role in the overall safety model, should eventually lead to further investigation on its reliability;

- a series of links of the same colour – be they red or green - may not tell the absolute (un)reliability of an SP. They may simply result from the transparency of this SP towards the two alternative behaviours (success/failure).

  For example, an SP such as 'The ATCO issues a correct clearance' may only become important if challenged. It will probably not be reported in any event in which it was successful. It means that little data will be available about its success, but it does not mean that it is not reliable.

6.3.5 Criticality of Safety Principles

The reliability of an SP, as introduced previously, is an important feature to assist safety reasoning. But it is not sufficient. The safety lesson further depends on the role played by the SP, mainly from two main perspectives:

- How **critical** the SP is: What happens to the safety of the system if that SP fails? What is the amount of risk involved?

- How **common** the SP is: How far is the overall ATM safety affected if that SP fails?

The fact that in the SMART approach a Generic Initiator is substituted for each reported event has in itself the effect to extract safety lessons from the 'local' conditions associated with a specific event. But the SMART approach allows much more than that. Because the same SPs can be shared across different Generic Initiators, hence Safety Architectures, a lesson learned locally can be used and its validity can be checked, in different contexts. Because the Safety Architectures are a rational representation of the protections against the risks, (the GIs), the criticality of information can be assessed.

The next sections introduce the notion of **local**, **extended** and **generalised criticality** of a Safety Principle.

6.3.5.1 Assessing the 'local' criticality of Safety Principles

The 'local' criticality of a Safety Principle is the strength of the remaining protections against an accident, should that SP fail in the situation of a specific reported event.

Within each category – prevention, recovery, accident consequence mitigation - the SP criticality depends on two parameters:

- the number of backups or remaining protection layers available, should that SP fail;

- the empirical robustness of the remaining layers, that is of the weakest SP in each layer.

Across all SP categories, the SP criticality is higher when:

- The associated accident(s) are more severe.
- The impaired category – prevention, recovery, accident consequence mitigation - is the one emphasised by the global safety strategy (e.g. Prevention Dominant Strategy). Indeed, depending on the Generic Initiator and the associated accident type, the protection strategy may emphasise prevention (air collision), recovery, or consequence mitigation.
- The lost protection is prevention, rather than recovery, or mitigation. The prevention category determines how far from a Generic Initiator the situation went during the event. The recovery category determines how far from an accident the situation would have developed during the event, should the Generic Initiator have occurred in addition. The accident consequence mitigation category determines how far from the worst consequences the situation developed, should the Initiator have occurred, and also developed into an accident.

An SP is said to be **Absolutely Critical in a Local context** if for that context, within its category – prevention, recovery, accident consequences mitigation - it is a minimal cut set on its own (its impairment is **necessary and sufficient** to make the whole protection collapse). This happens when an SP is common to all protection layers, or part of a single protection layer.

An SP is said to be **Relatively Critical in a Local context** if for that context, within its category – prevention, recovery, accident consequences mitigation - the robustness of the protection layers in which it does not take part is so poor that its actual role in safety is far more crucial than foreseen.

6.3.5.2 Assessing the 'extended' criticality of Safety Principles

The 'extended criticality' of a Safety Principle is the strength of the remaining protections against an accident, should this Safety Principle fail, and had the configuration/context been different from that of the reported event.

It corresponds to a substitution test or in other words to a **'what if'** exploration. Indeed, when it comes to assessing the potential severity of the invalidation of a Safety Principle, it is important to examine the story as it was, but also as it could have developed. Different conditions can be envisaged through a question list:

What if:

- the aircraft type had been different?
- the airport had been different?
- meteorological conditions had been different?
- a different system had been affected?

An SP is said to be **Absolutely Critical in an Extended Contex***t* if, within a category – prevention, recovery, accident consequences mitigation - there is a context in which there are no remaining protection layers beyond those collapsed by the impairment of the SP.

An SP is said to be **Relatively Critical in an Extended Context** if, within a category – prevention, recovery, mitigation - there is a context in which the robustness of the remaining protections layers (beyond those collapsed by the impairment of the SP) is so poor that the SP's actual role in safety is far more crucial than foreseen.

6.3.5.3 Assessing the 'general' criticality of Safety Principles

The 'general criticality' of a specific Safety Principle is the strength of the protections against an accident still in place, should that Safety Principle fail, across all (identified) Generic Initiators (Safety Architectures) in which that Safety Principle is participating.

In order to assess the generalised criticality, the analyst has to identify all the protection layers across all the Safety Architectures in which the SP is involved. For each Generic Initiator involving the SP, the 'local' criticality determination approach should be used.

Once the analysis has been carried out for all the (available) Generic Initiators potentially affected by the challenge of the SP, the generalised criticality can be inferred, taking into account the severity of the associated potential accidents.

An SP is said to be **Absolutely Critical in a Generalised Context** if within a category –prevention, recovery, accident consequences mitigation - it is a minimal cut set on its own (its impairment is **necessary and sufficient** to make the whole protection collapse) for all the Safety Architectures it belongs to.

An SP is said to be **Relatively Critical in a Generalised Context** if within a category –prevention, recovery, accident consequences mitigation - the robustness of the protection layers in which it does not take part is so poor that its actual role in safety is far more crucial than foreseen, for all the Safety Architectures it belongs to.

Finally a Safety Principle is defined as:

- **Absolutely Critical** if it is Absolutely Critical locally, in an extended context or in a generalised context;
- **Relatively Critical** if it is not Absolutely Critical, but has been said Relatively Critical at least once in a local, extended or generalised context.

Page intentionally left blank

7. EXAMPLES USING SMART

Identification of the GI derived from the two similar incidents and the development of the associated Safety Architecture up to three levels for some Safety Principles.

Level 1

SP4	Accident alert	Safeguard ongoing flight operations	SP5
	Accident: ground collision		
SP3	Conflicting runway occupation is removed in due time		
SP2	An avoiding maneuver will be successfully performed in due time		
	Generic Initiator: Abnormal start of a Take-off / Landing on an occupied runway		
SP1	The penetration/location of any obstacle (moving and controlled by ATM) on the runway strip is not conflicting with ongoing TO/L		

Level 2

SP1 decomposition / level 2

SP1.1	No TO/L operations will start unless the runway strip is cleared from clonflicting obstacles	No runway strip penetration of conflicting obstacle will take place during TO/L operations	SP1.2

SP2 decomposition / level 2

Pilots recognise the presence of conflicting obtacle in due time	Pilots react successfully
The ATM system recognizes the problem in due time	The ATM system issues an abortion clearance (or other avoiding manoeuvre) in due time

SP3 decomposition / level 2

The 'obstacle' recognises the conflict situation in due time	The 'obstacle' takes evasion action in due time		
The ATM system recognises conflicting rwy occupation in due time	The ATCOs issue an evasion action instruction in due time	The 'obstacle' reacts to the instruction correctly in due time	SP3.3

SP4 decomposition / level 2

The ATM system recognises there is an accident/incident	The ATM system locates the accident correctly	The ATM system informs the relevant party	The interaction between ATM and other parties works correctly

SP5 decomposition / level 2

The ATM system correctly assesses the consequences of the accident on airport operations	The ATM system correctly manages the operations according to the anticipated consequences of the accident/incident

Level 3

SP1.1 decomposition / level 3

Crews will detect and recover erroneous TO/L clearances before actual start of TO/L		
ATC will detect and recover erroneous TO/L clearances before actual start of TO/L		
No TO/L clearance will be delivered unless the runway strip is anticipated free from conflicting obstacles	The ATCOs anticipation of the absence of obstacle on the runway is correct	No A/C will TO/L unless crews have received and understood a clearance to do so from ATC

SP1.2 decomposition / level 3

Ground crews will detect and recover erroneous TO/L clearances before actual start of TO/L		
ATC will detect and recover erroneous TO/L clearances before actual start of TO/L		
Clearance to penetrate the runway strip will be delivered only if anticipated not to lead to a conflict with ongoing TO/L operations	The ATCOs anticipation of the absence of ongoing TO/L operations is correct	No penetration of a controlled obstacle on the runway strip without a clearance to do so from ATC

SP3.3 decomposition / level 3

The ‘obstacle’ receives the instruction correctly	The ‘obstacle’ understands the instruction correctly	The ‘obstacle’ is willing to comply with the instruction	The ‘obstacle’ is able to comply with the instruction

Analysis of Incident 1
> recording the success/failure/unknown outcome of the detailed SPs
> propagating these assessments to higher level SPs
> checking the realism of the Safety Architecture

Level 3

SP1.1 decomposition / level 3

Crews will detect and recover erroneous TO/L clearances before actual start of TO/L		
ATC will detect and recover erroneous TO/L clearances before actual start of TO/L		
No TO/L clearance will be delivered unless the runway strip is anticipated free from conflicting obstacles	The ATCOs anticipation of the absence of obstacle on the runway is correct	No A/C will TO/L unless crews have received and understood a clearance to do so from ATC

Comment: from an outsider point of view

SP1.2 decomposition / level 3

Ground crews will detect and recover erroneous TO/L clearances before actual start of TO/L		
ATC will detect and recover erroneous TO/L clearances before actual start of TO/L		
Clearance to penetrate the runway strip will be delivered only if anticipated not to lead to a conflict with ongoing TO/L operations	The ATCOs anticipation of the absence of ongoing TO/L operations is correct	No penetration of a controlled obstacle on the runway strip without a clearance to do so from ATC

SP3.3 decomposition / level 3

The 'obstacle' receives the instruction correctly	The 'obstacle' understands the instruction correctly	The 'obstacle' is willing to comply with the instruction	The 'obstacle' is able to comply with the instruction

Comment: none of these SP is related to this incident ➔ remains white

Level 2

SP1 decomposition / level 2

SP1.1			SP1.2
	No TO/L operations will start unless the runway strip is cleared from clonflicting obstacles	No runway strip penetration of conflicting obstacle will take place during TO/L operations	

SP2 decomposition / level 2

Pilots recognise the presence of conflicting obtacle in due time	Pilots react successfully
The ATM system recognises the problem in due time	The ATM system issues an abortion clearance (or other avoiding maneuver) in due time

SP3 decomposition / level 2

The 'obstacle' recognises the conflict situation in due time	The 'obstacle' takes evasion action in due time		
The ATM system recognises conflicting rwy occupation in due time	The ATCOs issue an evasion action instruction in due time	The 'obstacle' reacts to the instruction correctly in due time	SP3.3 *Comment: remains white since not related to the incident*

SP4 decomposition / level 2

The ATM system recognises there is an accident/ incident	The ATM system locates the accident correctly	The ATM system informs the relevant party	The interaction between ATM and other parties works correctly	*Comment to the red link: they didn't diagnose the problem correctly. They thought it was a technical failure*

SP5 decomposition / level 2

The ATM system correctly assesses the consequences of the accident on airport operations	The ATM system correctly manages the operations according to the anticipated consequences of the accident/incident

Level 1

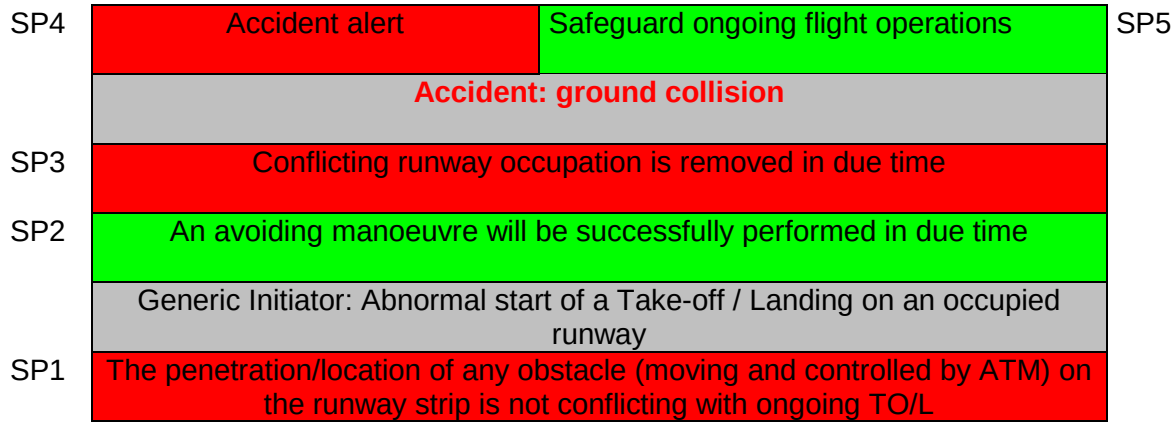
SP4
Accident alert
Safeguard ongoing flight operations
SP5
Accident: ground collision
SP3
Conflicting runway occupation is removed in due time
SP2
An avoiding manoeuvre will be successfully performed in due time
Generic Initiator: Abnormal start of a Take-off / Landing on an occupied runway
SP1
The penetration/location of any obstacle (moving and controlled by ATM) on the runway strip is not conflicting with ongoing TO/L

Analysis of Incident 2
> recording the success/failure/unknown outcome of the detailed SPs
> propagating these assessments to higher level SPs
> checking the realism of the Safety Architecture

Level 3

SP1.1 decomposition / level 3

Crews will detect and recover erroneous TO/L clearances before actual start of TO/L		
ATC will detect and recover erroneous TO/L clearances before actual start of TO/L		
No TO/L clearance will be delivered unless the runway strip is anticipated free from conflicting obstacles	The ATCOs anticipation of the absence of obstacle on the runway is correct	No A/C will TO/L unless crews have received and understood a clearance to do so from ATC

SP1.2 decomposition / level 3

Ground crews will detect and recover erroneous TO/L clearances before actual start of TO/L		
ATC will detect and recover erroneous TO/L clearances before actual start of TO/L		
Clearance to penetrate the runway strip will be delivered only if anticipated not to lead to a conflict with ongoing TO/L operations	The ATCOs anticipation of the absence of ongoing TO/L operations is correct	No penetration of a controlled obstacle on the runway strip without a clearance to do so from ATC

SP3.3 decomposition / level 3

The 'obstacle' receives the instruction correctly	The 'obstacle' understands the instruction correctly	The 'obstacle' is willing to comply with the instruction	The 'obstacle' is able to comply with the instruction

Level 2

SP1 decomposition / level 2

SP1.1			SP1.2
	No TO/L operations will start unless the runway strip is cleared from clonflicting obstacles	No runway strip penetration of conflicting obstacle will take place during TO/L operations	

SP2 decomposition / level 2

Pilots recognise the presence of conflicting obtacle in due time	Pilots react successfully
The ATM system recognises the problem in due time	The ATM system issues an abortion clearance (or other avoiding maneuver) in due time

SP3 decomposition / level 2

The ‘obstacle’ recognises the conflict situation in due time	The ‘obstacle’ takes evasion action in due time		
The ATM system recognises conflicting rwy occupation in due time	The ATCOs issue an evasion action instruction in due time	The ‘obstacle’ reacts to the instruction correctly in due time	SP3.3 *Comment: remains white since not related to the incident*

SP4 decomposition / level 2

The ATM system recognises there is an accident/incident	The ATM system locates the accident correctly	The ATM system informs the relevant party	The interaction between ATM and other parties works correctly

SP5 decomposition / level 2

The ATM system correctly assesses the consequences of the accident on airport operations	The ATM system correctly manages the operations according to the anticipated consequences of the accident/incident

Level 1

SP4	Accident alert	Safeguard ongoing flight operations	SP5
	Accident: ground collision		
SP3	Conflicting runway occupation is removed in due time		
SP2	An avoiding manoeuvre will be successfully performed in due time		
	Generic Initiator: Abnormal start of a Take-off / Landing on an occupied runway		
SP1	The penetration/location of any obstacle (moving and controlled by ATM) on the runway strip is not conflicting with ongoing TO/L		

Page intentionally left blank

8. MAKING DECISIONS IN RISKY ENVIRONMENTS

8.1 Using Feedback from Operational Experience to Manage Safety

Firstly, it is worth considering the implications of the use of feedback from experience and how it can be used in safety management:

- Without any specific feedback from experience, we rely on an *a priori* model of the system's safety: the reasons for the system to be acceptably safe can be expressed in terms of a logical combination of assumptions (Safety Principles), with various levels of trust in each of them.

- These trust levels are actually the outcome of a probabilistic approach (be it a real quantitative assessment or a more subjective judgement). A global assumption that the system is acceptably safe is permanently derived from the combination of these levels of trust.

- A decision to modify the system is made when this global assumption, that the system is acceptably safe, can no longer be made. This can happen if it is realised, by further reasoning or by the availability of new scientific evidence, that some safety assumptions are not as trustworthy as expected.

- With each specific feedback from experience this a priori model is refined and modified by the knowledge gained through interpretation of that specific experience.

- In other words, *a priori* probabilistic reasoning is complemented/amended by conditional probabilistic reasoning; for instance, the *a priori* probability that the ATCO issues a wrong clearance is replaced by the probability that the ATCO issues a wrong clearance, knowing that this has happened several times, in different circumstances.

8.2 Using SMART as a Support to Risk-informed Decision-making

At the end of the event analysis process described in the previous part of this report, the analysts may come out with an updated health map of SPs, for all the SPs called for during the reported events processed. These health maps of impacted SPs are then used to support an assessment of the empirical robustness of the SPs. In parallel, the criticality of SPs can be defined.

In this section we will describe how the empirical assessment of the SPs robustness and their criticality can be used as an input to the decision-making process related to safety management.

The objective of SMART is not to develop a complete decision-making rationale, but to provide decision-makers with safety-related information to

help them take safety issues into account in their decision. In other words, whereas a real decision encompasses multiple aspects such as economical, historical, social, technological and safety-related considerations, SMART focuses on the safety perspective to be integrated into the decision process.

SMART can support and assist safety-related decision-making from several perspectives that will most probably be revealed by the analysts themselves during their operational use of the system. However, it can be anticipated that SMART will assist two kinds of safety-related decisions:

- decisions concerning the administration of the SMART tool itself, e.g. modification of the robustness rating of an SP, modification of a Safety Architecture;
- decisions related to the management of the ATM system safety, e.g. monitoring of the safety implications of a change in a Safety Architecture, propagation of implications to other Safety Architectures, decisions related to the modification of the ATM system to improve its safety level (safety recommendations).

The following sections elaborate on these aspects.

8.3 Administrating SMART

Administrating SMART means keeping the way it represents or models ATM safety as consistent as possible with the realities of the system itself. This implies decisions and actions, at least whenever one or the two following situations occur:

- one Safety Principle must be added or modified;
- one Generic Initiator and its associated Safety Architecture is missing;
- one Safety Architecture turns out to be too optimistic - some possible failure paths are missing;
- one of the Safety Principles robustness ratings must be changed, e.g. from ‘reliable’ to ‘unreliable’ - on the ground of empirical evidence.

The first two situations have already been addressed under the development of the SMART tool. Therefore the next sections will address the latter two issues.

8.3.1 Unrealistic Safety Architectures

The methodology proposed to identify the Safety Architectures associated with a Generic Initiator is neither infallible nor exhaustive. Additionally, people may have wrong assumptions about their system’s safety. Consequently, while matching a Safety Architecture in detail with an incident report, the analyst

may find out that some elements are not properly combined (e.g. 'AND' instead of 'OR'), are missing, or conversely are erroneously mentioned in the combination of SPs.

Optimistic Safety Architectures will be detected when, according to the Safety Architecture logic, the reported event should not have happened, whereas in reality it did happen. Particularly, it may be the case that an SP has been omitted in a 'AND' position in one or more lines of defence.

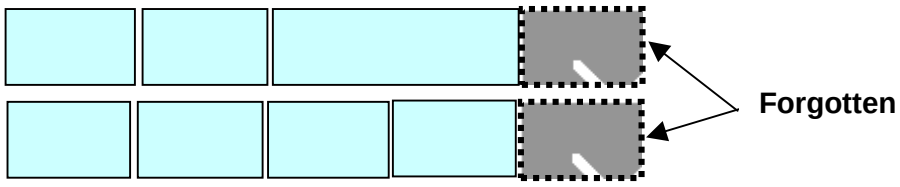

This easily happens with actions or processes that critically contribute to safety, while they look so obvious, so natural, that they become virtually transparent to an observer. The following example illustrates this point:

> *A clearance to line up "behind the 737 at landing". The restriction "behind the 737 at landing" is critical for safety, and it totally relies on the assumption that the crew will actually identify the right aircraft, or report they cannot see it. If this assumption is not fulfilled, then the whole line of defence associated with the assurance of a line up clearance will be invalidated, even if each of the components in the line of defence is true (e.g. relevant clearance, perfect phraseology, etc.). Consequently, if that 'proper aircraft identification' condition is not included in the Safety Architecture, a whole potential failure path has been omitted.*

A similar situation occurs when a failure mode of a sub-system has been omitted in the architecture. If that sub-system plays a role in several lines of defence (e.g. the ATCO him/herself), then forgetting one of its failure modes (e.g. incapacitation) leads to the omission of a complete common mode failure. An optimistic Safety Architecture can also result from a series of SPs being considered redundant ('OR' links) whereas they actually are complementary ('AND' links).

Conversely, it may happen that, according to the Safety Architecture logic, the reported event should have led to an accident, whereas in reality it did not. This indicates that the Safety Architecture is pessimistic: it lacks at least a whole layer of protection (i.e. an additional 'OR' line of Safety Principles in the recovery part). Therefore, the Safety Architecture should be amended to include that real layer of protection. However, since the safety model aims at representing what is thought to positively ensure safety, aspects such as good luck or exceptional reaction of either a pilot or an ATCO, should not be included. Indeed, passengers, or safety managers, do not want to rely on these variables.

In brief, if SPs seem to be missing or conversely badly defined in a Safety Architecture, this Safety Architecture should be corrected.

8.3.2 Change of Safety Principle robustness

Changing the empirical robustness rating of a Safety Principle may have crucial consequences on the resulting strength of the Safety Architectures in which it is involved. Such decisions will be made recurrently, on the basis of the SP health map data, and with reference to the criticality of the SP.

The following table is a reminder of the meaning of the different ratings:

Table 3: Safety Principle robustness rating

SP robustness rating	Meaning	SP colour code
Reliable	The SP can be relied on: its reliability is up to, or better than, assumed by the *a priori* safety model	Green
Unreliable	The SP can no longer be relied on: its reliability is significantly less than assumed by the *a priori* safety model	Red
Unsure	The SPs reliability can not be rated, although it has been called upon in real events. A lot of uncertainty still affects the assessment.	Amber
No rating	No assessment could be done, by lack of data, or excessive uncertainty	White

The information available on a Safety Principle is:

- its potential absolute or relative criticality, and
- its health map.

There is no general rule, or threshold, to derive from that information when an SP should start to be considered reliable or unreliable. Considering an SP unreliable as soon as it has been invalidated during one event would be too conservative. Moreover, since ‘reliable’ means ‘up to assumed reliability’, and all SPs do not imply the same reliability assumption, it would not be relevant.

While no generic rule can be established, for certain situations and categories of SP, some principles can be defined, as presented hereafter. The following sections discuss the decision cases indicated in the change matrix below (see Table 4).

Table 4: Matrix of changes in Safety Principle robustness rating

SP robustness rating **From:**	**To:**			
	Reliable	Unreliable	Unsure	No rating
Reliable				
Unreliable				
Unsure				
No rating	✓	✓	✓	

8.3.2.1 Changing a Safety Principle rating from 'no rating' to 'unreliable'

- Absolutely and Relatively Critical Safety Principles

  For both Absolutely and Relatively Critical Safety Principles, the first 'red link' should make the Safety Principle turn from white to red. The absence of 'link' should call for close monitoring since no information is available either on its reliability or on its unreliability.

- Safety Principles with 'red links' only

  A further investigation is needed regarding SPs that have always been invalidated when called for in reported events. Indeed, two extreme situations can be envisaged: it may turn out that whenever they function properly, they are never mentioned in the report, or that whenever they are actually called for, they are invalidated. In the latter case, the Safety Principle should be turned to red.

8.3.2.2 Changing a Safety Principle rating from 'no rating' to 'reliable'

Operational experience may reveal that an SP is always 'true' (successful) in the reported events. It is essential to check whether such information is influenced by some bias in the database (e.g. non-compliance with such an SP is never reported by analysts). When no reasons are found to be doubtful of the representative nature of the figures, the SP rating should turned to 'green'. Note that such a decision is an important one. This SP will then be relied on in the safety model, at least until experience leads to review this rating. This can have an influence on the urgency of a decision related to other SP involved in the same Safety Architecture.

8.3.2.3 Changing a Safety Principle rating from 'no rating' to 'unsure'

Some event reports include no information about the actual behaviour of an SP. When such an SP is critical, either locally or generally, it may be wise to change its rating to 'unsure'. Such a change can help visualise the potential consequences of an unreliable status. It can also help focus an investigation, or have an influence on the urgency of a decision related to other Safety Principles involved in the same Safety Architecture.

8.4 Using SMART to Manage ATM System Safety

8.4.1 The model of ATM safety has missed a failure mode

In the case of such an event, the first appropriate decision should be to further investigate and determine whether the unforeseen failure or combination of failures should be considered in the Safety Model or not.

A combination of failures may be unlikely enough (even if it occurred once) to be acceptable from a strict safety standpoint. This investigation should not only focus on the Generic Initiator concerned by the reported event, but also a review of the consequences of the omission on the whole Safety Model across all Generic Initiators.

If a failure mode of a system has been omitted, this omission has repercussions on all the Safety Principles referring to the functioning of this system, whatever the Generic Initiator. If any of these Safety Principles is Absolutely Critical or Relatively Critical, the unforeseen failure can be considered as unacceptable. Apart from these 'obvious' cases, the acceptability of the problems in the Safety Model relies on expert judgement.

If it finally turns out that the event resulted from acceptable 'bad luck', i.e. that the unforeseen failure can be neglected in the Safety Model, then the wise decision from a safety viewpoint would be to do nothing. If conversely, it turns out that the unforeseen failure or combination of failures is unacceptably likely to occur again, modifications should be agreed. An SP proving to be less reliable than expected does not necessarily mean that this particular SP should be changed or made more reliable. A systemic view may suggest an upstream modification changing the relative role of this SP, and leading to a much more efficient overall Safety Architecture.

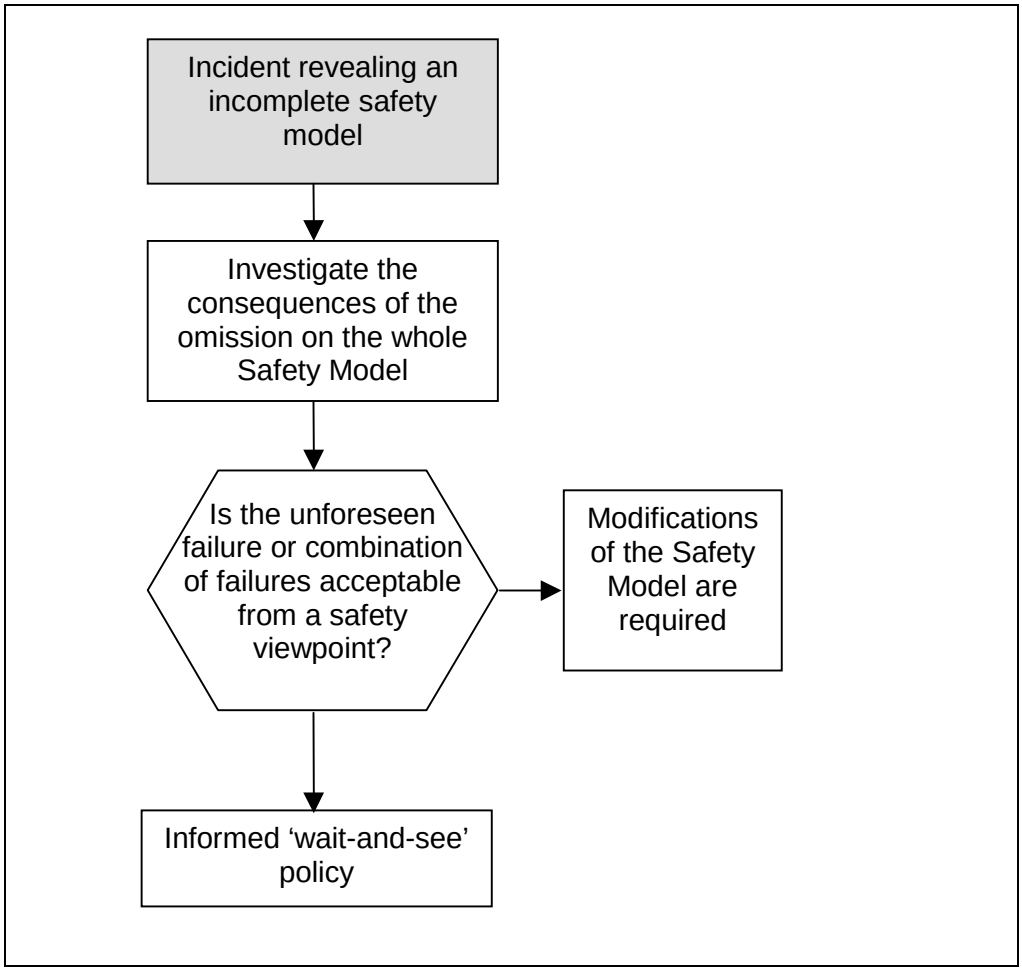

Flowchart 1

8.4.2 A Safety Principle rating has been changed to 'unreliable'

When the decision has been made to change the empirical robustness rating of an SP to the 'unreliable' status, this decision will cause changes within the SMART Model in two directions:

- Vertically: In any specific Safety Architecture in which the SP is a component at a given level of decomposition, the new SP status will lead to a 'false' value of the logical variable describing the SP behaviour. In other words, the SP will then appear 'red' in the corresponding layer of defence. The SMART Model will then compute the logical consequences of this situation for all the higher decomposition levels, using the corresponding combinations of 'ANDs' and 'ORs'. Ultimately, the system will generate a warning if the top level of the safety management strategy is affected.

- Horizontally: The new SP status will be transmitted to all the Safety Architectures involved. Hence they will all be affected by the vertical

propagation. All the logically derived consequences on all the Safety Architectures concerned will be visualised every time they affect the integrity of the higher level of protection (prevention, recovery, mitigation).

With this assistance, a safety manager will be able to visualise all the consequences induced by a decision to consider an SP 'unreliable'. The reaction to the switch of an SP to a 'red status' is an expert judgement that mainly depends on:

- the families in which the SP is involved – Prevention, Recovery, Accident Consequences Mitigation;
- the roles these families play in the Safety Strategy associated with the Generic Initiators;
- the criticality of the SP.

In case of a Critical SP, a more specific decision can be made, as discussed below.

8.4.2.1 Absolutely Critical Safety Principles

In case of a Critical SP, the decision to change its status to 'red' should be considered at the first occurrence of failure, in other words at the first 'red link'. The recommended reaction to the first 'red link' of an Absolutely Critical SP then depends on:

- the families with which the SP is involved – Prevention, Recovery, Accident Consequences Mitigation;
- the roles these families play in the Safety Strategy associated with the Generic Initiators for which the SP is Absolutely Critical.

The suggested decisions in these different situations are indicated by the following table:

Table 5: Criticality of Safety Principles

Safety Strategy is dominated by:	**Safety Principle is critical for:**		
	Prevention	**Recovery**	**Accident Consequence Mitigation**
Prevention	modification of the ATM system	close monitoring policy	close monitoring policy
Recovery	close monitoring policy unless, further to an investigation, it turns out that a more efficient strategy could be based on prevention	modification of the ATM system	close monitoring policy
Accident Consequence Mitigation	close monitoring policy unless, further to an investigation, it turns out that a more efficient strategy could be based on prevention or recovery or both.	close monitoring policy unless, further to an investigation, it turns out that a more efficient strategy could be based on prevention or recovery or both.	modification of the ATM system

If the Safety Strategy is distributed among two or more categories – Prevention, Recovery, Accident Consequences Mitigation, then the first 'red link' calls for further investigation. This will determine the health of the remaining category of Safety Principles, and the cost in terms of safety of the loss of one category; for instance, if the potentially lost category is Prevention, the investigation should also consider the consequences in terms of safety of the occurrence of the associated initiator. The following flowchart summarises this decision strategy.

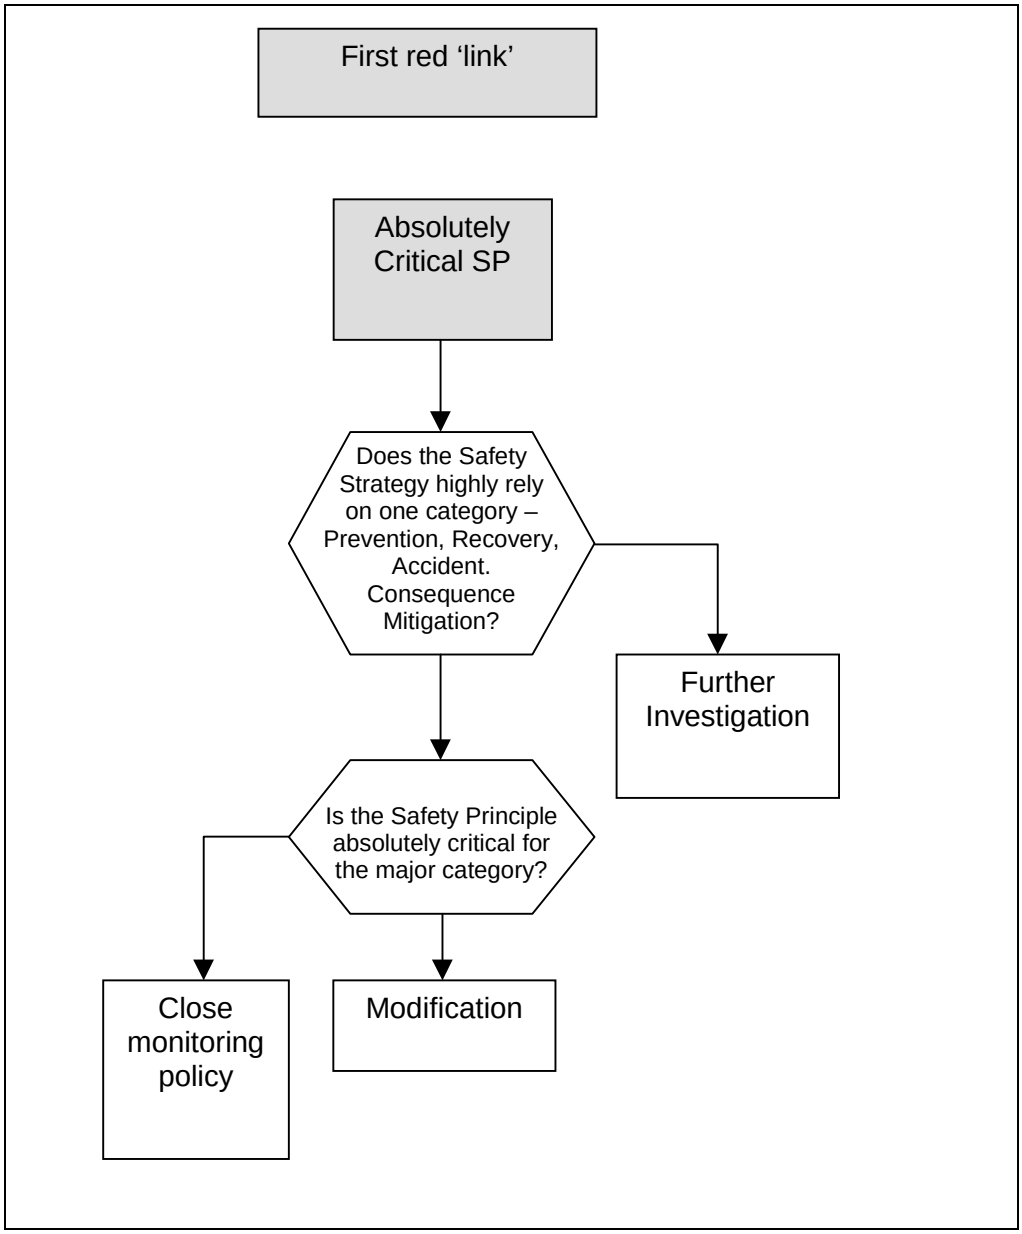

Flowchart 2

8.4.2.2 Relatively Critical Safety Principle

The first 'red link' should call for a first investigation on the possibility of a simultaneous failure of this SP and the surrounding 'red' ones. If it turns out to be incompatible, then no action is required from a safety viewpoint. If it turns out to be compatible, a modification is needed, especially if the SP is relatively critical for a major category in the Safety Strategy.

The following flowchart summarises the decision strategy:

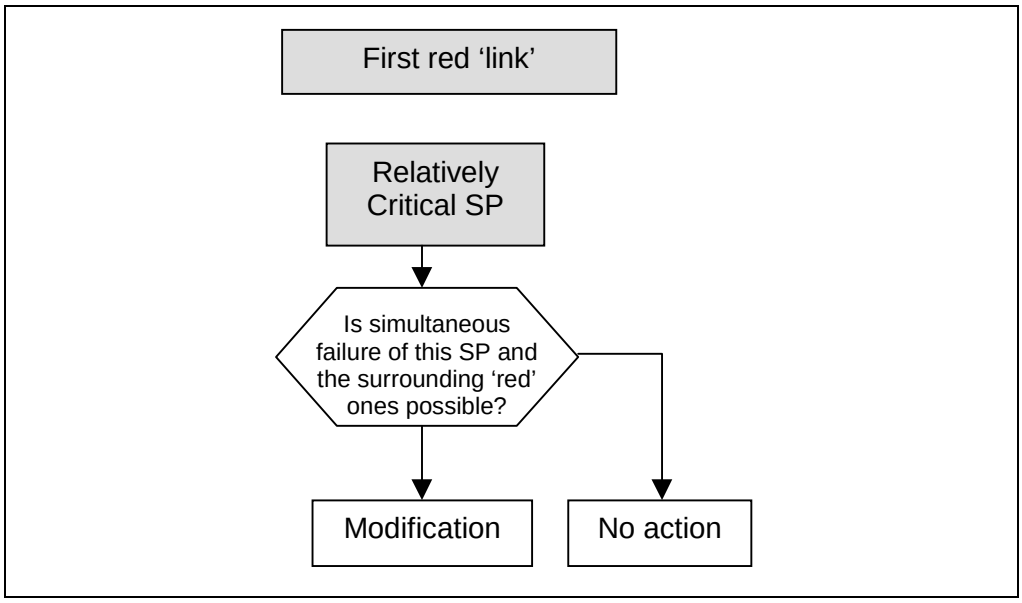

Flowchart 3

8.4.2.3 A Critical Safety Principle which still has an 'unknown behaviour rating'

An absence of 'link' after a while doesn't mean that the Safety Principle is 100% reliable, but rather that reality has not given a clue to its actual reliability. Therefore, an appropriate decision would be to carry out a dedicated investigation on this SP.

8.4.3 From 'patch' solutions to changes in the Safety Management System

Failure of an SP, or missing parts in a Safety Architecture, may be an opportunity to question the design of the ATM system.

- Failure of a Safety Principle may indicate a poor match between reality and the expected behaviour (man-machine behaviour) on which design relies. Since resources are limited, a 'patch' solution could then be envisaged. Such an option would consist of adding a patch on a protection layer, i.e. a fix to increase the reliability of the failed SP. This option is probably the least demanding in terms of both time and money in most cases. Typical examples of this are HMI design modifications intended to reduce the frequency of ATCOs errors.

- The next step would be to amend the Safety Architecture within the same SP family (Prevention) by adding an additional protection layer in order to increase the redundancy and improve the 'defences in depth'.

- However, when a Safety Architecture appears to be questionable, the overall safety strategy, i.e. the balance between Prevention, Recovery and

Accident Consequence Mitigation should be questioned. For example, Recovery may prevail in the *a priori* strategy, while the actual 'cost' associated with the Initiator's occurrence is felt to be unacceptable by some of the operators. Then, most probably, these operators will invent their own way of reinforcing Prevention, and will deviate from expected behaviour. Conversely, if the strategy is Prevention dominant, while the Recovery component is felt to be efficient, deviations will probably occur to alleviate the prevention protections.

No systematic rule can be used to decide the most adequate level of change. The challenge is to determine whether the potential mismatch remains at the practice level or is due to a deeper incompatibility between functions, policy or even philosophy on one hand, and reality on the other. In other words, the challenge is to assess the extent of the inadequacy of the Safety Model.

Any decision is acceptable as long as it is worked out in the full knowledge of the facts. The most important aspect is that the underlying reasons for the decision are correctly traced. The SMART structure can obviously assist in this activity.

8.4.4 Event Reports, Incident Analysis and other data

The evidence that some critical SPs are unreliable calls for an in-depth analysis of the actual weakness of the Safety Architecture. This detailed analysis requires returning to events to better understand the contexts and the conditions surrounding the failure of SPs.

The reported events relevant for each SP are accessible through the links associated with the SP health maps. Some features may turn out to be recurrent in the corresponding contexts, thereby suggesting an explanation of the SP behaviour and indicating the need for a context modification.

For example, it can turn out that in similar events, one SP only fails with flights of the same airline, or at the same time of the day, or with the same aircraft type, or the same ATCO team. Such an investigation should help focus on the relevant problems or trends. In the case of a localised failure of some Safety Principle, the whole Safety Architectures does not have to be modified.

Some of the Safety Principles critical for Safety and unacceptably unreliable will inevitably refer to human factors aspects. Often the failure of an SP will translate into the commission of errors or violations by the front line operators. Understanding why these safety-related behavioural assumptions turn out to be wrong implies a better understanding of the contexts and the conditions surrounding these 'errors' or 'violations'. The analysts will then need to gather more data, and go back to human factors science to establish more realistic, hence reliable, HF-related Safety Principles. The HERA-JANUS Technique, in the investigation of errors and violations, will then become very useful.

Additionally, while SMART is designed to allow the assessment of SP reliability through feedback from operational experience, it does not dismiss any complementary approach. For instance, a safety manager may wish to assess the reliability of an SP through more proactive methods, such as audits, questionnaires, experiments or event scientific research.

8.4.4.1 Organisational factors

This can particularly be applied to organisational factors that are widely recognised as the key issue for safety improvement strategies. For example, the following issues have been identified (Reason, 1997) as latent organisational factors relevant for safety:

- Administrative policies: All facilities have administrative policies that regulate acceptable and required behaviours (e.g. safety requirements). However, the policies in place can interfere with successful accomplishment of work, either directly or indirectly.

- Communications: All work involves communications; between team members, with other groups, and between the team and its management. Effective communications require appropriate communications tools (e.g. adequate numbers and quality of telephones and radios) and pathways.

- Equipment condition: Poor equipment is a burden on operators who must constantly overcome repeated (and often multiple) failures, develop and apply 'band-aid' fixes, and sometimes cope with unsafe work environments.

- Interfaces with other groups and departments: Any work group must interact with others in order to accomplish their tasks. However problems with interfaces with other groups can cause significant problems in accomplishing these tasks.

- Man-machine interface: This issue refers to the usability of all interfaces with information and control systems.

- Roles and responsibilities: It is important for the effective performance within teams and groups that the roles and responsibilities of the individuals in the teams are understood and accepted to be appropriate and reasonable.

- Scheduling, time pressure and shift scheduling: Work within teams can be degraded significantly when inappropriate product scheduling and high time-pressure requirements are applied, or when shift scheduling results in worker fatigue.

- Task structure and design: The way tasks are designed can influence strongly the ability of people to accomplish their tasks; for example, poorly

designed tasks can encourage operators to violate procedures because they believe the work can be performed more efficiently.

- Tools and equipment: The use of deficient or inappropriate tools and equipment can cause the tasks to be performed inadequately, cause damage to the processes being performed, or cause harm to others.

- Work control documents: Work control documents (e.g. procedures and operating rules) provide the instructions to operators on specific tasks and activities that must be performed. Poor work control documents can mislead operators into performing tasks incorrectly.

- Work environment: Poor work environments (e.g. too hot/cold, poor lighting or inadequate workspace) are well recognised as problems that can create work performance problems.

These organisational factors can be addressed, as discussed above, through operational feedback. When incidents are analysed, the influence of organisational factors is inferred through the investigation process. In the HERA-JANUS taxonomy the link to organisational factors that create or influence error-prone work contexts is addressed through the notion of Contextual Conditions.

They can also be monitored directly through 'organisational markers' that are indicators of the organisation health from a safety perspective. An example of such markers could be the data from checklists or questionnaires given to a representative selection (from top management to the front line operators) of personnel within an organisation. The data can then be reviewed with the perspective of assessing their implication on various SP reliabilities. An example of such an Organisational Safety Assessment Questionnaire can be found in Appendix 3.

The results of these approaches should then be combined with that of incident analysis when rating the SP empirical robustness.

8.4.5 Virtually exploring new system options

Once a reasonable number of Safety Architectures are introduced into the SMART system, it is possible to 'play' with these in a proactive way and examine the consequences of a contemplated change. This can be simply achieved by replacing a part of a current safety structure by those envisioned.

Whatever the change decided, unforeseen repercussions are to be expected. However, the SMART system as it is designed, allows for a relatively global view of the role of an SP within the overall Safety Architecture. Therefore, when an SP is finally designated to be replaced by another, the robustness of all the protection layers it is involved with should be compared to what they were before the change. In other words, the global view of the role of an SP to be changed in the Safety Architecture helps specify the requirements the new SPs should comply with, in order to maintain safety. Appendix 1 illustrates in a

more detailed way the rationale that can be followed to virtually explore new options.

8.4.6 Monitoring the effects of decisions

It would not be reasonable to expect that all the effects of any modification to an existing system can be completely and accurately foreseen even through a virtual exploration as described in the previous section. Therefore, whenever a decision is made to perform a modification, specific attention should be paid to unfamiliar or unexpected phenomena even if they seem disconnected with the implemented change. The HERA-PREDICT Technique (report currently under preparation) can be used to assist in the decisions in terms of possible man-machine error problems.

Page intentionally left blank

9. SUMMARY

9.1 Introduction

The **Safety Management Assistance and Recording Tool (SMART)** concept presented in the current report relies on the assumption that the most efficient way to derive valuable information from reported events is to use them to put the *a priori* Safety Principles that presided over design and operational choices.

These Safety Principles are organised into a series of logical, hierarchical, structures, called Safety Architectures. Each of these Safety Architectures represents the assumed protection – prevention, recovery, accident consequences mitigation - against a prototypal incident characterised by a Generic Initiator. In other words, they represent the rational proof that satisfies the designers and users of a system that each Generic Initiator is acceptably unlikely to happen, or degenerate into an accident, or reach the worst accidental consequences.

Matching a reported event to a Generic Initiator, then reading the event through the Safety Architecture corresponding to that Generic Initiator, allows the assessment of the local success or failure of the Safety Principles involved in that event. Building up these local assessments through a series of events ultimately allows for assessing the 'health' or robustness of the concerned Safety Principles. The robustness of all corresponding Safety Architectures can then be derived through logical computing.

Through this process, an analyst can examine the consequences of the failure of some Safety Principles not only in an occurrence situation, but also in slightly different contexts as well as in totally different contexts in which these same Safety Principles are supposed to play a role in Safety.

9.2 Developing SMART

The development of the SMART tool includes two main parts: the development of the methodologies and the development of a software.

9.2.1 Development of methodologies

One key component of the SMART approach is the notion of Generic Initiator. Any reported event will be first matched to a Generic Initiator. If no matching Generic Initiator can be found, a new one must be identified. The Generic Initiators are the insights to incident scenarios. They are identified through a methodology starting from reported incident scenarios. In a first step, Initiators - incidents likely to develop into an accident, should no protection be activated

- are derived from the incident scenario. In a second step, the Initiator is 'generalised', i.e. freed from local contexts and circumstances.

For each Generic Initiator Safety Principles are then identified and categorised into three main families: Prevention (of the Generic Initiator occurrence), Recovery (from the Generic Initiator leading to an accident) and Accident Consequences Mitigation.

The identification of Safety Principles is based on a functional approach, in other words, on a *'*how*'* questioning process. In addition, the method starts with the 'high-level' safety principles. Then each SP is decomposed into a logical combination of lower level SPs, and so on.

The method goes from the most abstract (goals, strategy) to the most concrete principles (expected behaviour of the ATM system, its components and interactions) in a means-ends abstraction hierarchy.

9.2.2 Development of a software tool

The software tool which supports the SMART approach is in the developmental stages but the following options must be included:

- capturing and editing an event report;
- capturing and editing a Generic Initiator;
- capturing and editing a Safety Principle;
- capturing, editing and visualising a Safety Architecture;
- recording and visualising SP behaviour assessments linked to event reports;
- editing and recording SP robustness status;
- simulating the consequences of a change in an SP status for all concerned Safety Architectures.

9.3 Using SMART

Generic Initiators play an important role, not only in the preliminary development phase but also in the operational use of the tool. Indeed, matching a reported event with one or more relevant Generic Initiator(s) allows the screening of Safety Principles in order to focus on those called upon in the reported event.

The behaviour – success or failure - of the Safety Principles involved in the reported event is then assessed in the situation described in the event report.

This information is stored, event after event, into the SMART database. It builds up and allows for the assessment of the health – the empirical robustness - of Safety Principles through experience.

In parallel, the logical combination of Safety Principles to compose Safety Architectures allows the exploration of the criticality of the failure of a Safety Principle. This can be done locally – with reference to a Generic Initiator and a specific event context – or extended to all contexts and then generalised to all Generic Initiators involving this Safety Principle.

Considering all Safety Principles, and the combination of their empirical robustness, their criticality and the role they play in the overall safety strategy – prevention, recovery, accident consequences mitigation - can then assist a safety manager in their safety-related decision-making process.

9.4 Conclusions

The approach suggested in the Safety Management Assistance and Recording Tool (SMART) fundamentally differs from traditional incident analysis approaches in that it tends to describe why the ATM system is supposed to be safe, instead of trying to understand why it failed. It seeks to learn safety lessons through a permanent comparison of expected and actual behaviour of the ATM safety system. Beyond the ‘negative’ experience analysis, it also affords to build up information on what functioned in the way it was expected to. Additionally it allows the building of experience on all protection layers which is seldom examined by traditional approaches.

Through the notion of Generic Initiators, the SMART approach attempts to go beyond the specific circumstances of an incident, but records the reasons why they could breach the systems protections. This systematic recording, enhanced with the exploration across other situations, allows greater insights not only into the protection layers that have been affected in events, but also on protection layers that were never called upon. Such information is usually never considered in other approaches.

All this information is enriched as reported events are analysed, through a growing and living database. Thus, the decision-making process proposed in the SMART approach benefits from a complete aggregation of experience rather than from a collection of single event experiences.

Page intentionally left blank

REFERENCES

Chopra, V., Bovill, J.G., Spierdijk, J. & Koornneef, F. (1992). Reported significant observations during anaesthesia: A prospective analysis over an 18-month period. British Journal of Anaesthesia. 68:13-17.

EATMP Human Resources Team (2002a). *Technical Review of Human Performance Models and Taxonomies of Human Error in ATM (HERA)*. HRS/HSP-002-REP-01. Ed. 1.0. Released Issue. Brussels: EUROCONTROL.

EATMP Human Resources Team (2002b). *Short Report on Human Performance Models and Taxonomies of Human Error in ATM (HERA)*. HRS/HSP-002-REP-02. Ed. 1.0. Released Issue. Brussels: EUROCONTROL.

EATMP Human Resources Team (2002c). *The Investigation of Human Error in ATM Simulation.* HRS/HSP-002-REP-05. Ed. 1.0. Released Issue. Brussels: EUROCONTROL.

EATMP Human Resources Team (2002d). *The Investigation of Human Error in ATM Simulation – The Toolkit.* HRS/HSP-002-REP-06. Ed. 1.0. Released Issue. Brussels: EUROCONTROL.

EATMP Human Resources Team (2003a). *The Human Error in ATM Technique (HERA-JANUS)*. HRS/HSP-002-REP-03. Ed. 1.0. Released Issue. Brussels: EUROCONTROL.

EATMP Human Resources Team (2003b). *Validation of the Human Error in ATM (HERA-JANUS) Technique*. HRS/HSP-002-REP-04. Ed. 1.0. Released Issue. Brussels: EUROCONTROL.

EATMP Human Resources Team (2003c). *HERA-JANUS Teaching Materials.* HRS/HSP-002-REP-09. Ed. 1.0. Released Issue. Brussels: EUROCONTROL.

Hawkins, F.H. (1987). *Human Factors in Flight*. Aldershot: Gower Publishing Company.

Moshansky, Mr. Justice (1992). Commission of inquiry into the Air Ontario crash at Drydon, Ontario. Ottawa: Ministry of Supply and Services, Canada.

Perrow, C. (1984). *Normal accidents: Living with High-Risk Technologies*. USA: Basic Books.

Rasmussen, J. & Svedung, I. (2000). *Proactive Risk Management in a Dynamic Society.* Karlstad, Sweden: Swedish Rescue Services Agency.

Reason, J. (1990). *Human Error*. Cambridge: Cambridge University Press.

Reason, J. (1997). *Managing the Risks of Organisational Accidents*. Aldershot, England: Ashgate Publishing Limited.

Sheen, Mr Justice (1987). MV Herald of Free Enterprise. Report of Court No. 8074. Formal Investigation. London: Department of Transport.

Turner, B.A. (1978). *Man-made Disasters*. London: Wykeham.

Vaughan, D. (1990). Autonomy, interdependence and social control: NASA and the Space Shuttle Challenger. Administrative Science Quarterly. 35:225-257.

Weick, K.E. (1987). Organisational culture as a source of high reliability. California Management Review, 19: 112-127.

FURTHER READING

Amalberti, R. & Barriquault, C. (1999). Fondements et limites du retour d'expérience. Annales des Ponts et Chaussées 91 (Sept.), 67-75.

Bourrier, M. & Laroche, H. (2001). Risque de défaillance : les approches organisationnelles. In: R. Amalberti, C. Fuchs & C. Gilbert (Eds). *Actes de la première séance du séminaire. Le risque de défaillance et son contrôle par les individus et le organisations dans les activités à haut risque*. CNRS, MSH-Alpes.

Decker, S.W.A. (2001a). The disembodiment of data in the analysis of human factors accidents. Communication to OSU Symposium.

Decker, S.W.A. (2001b). *The Field Guide to Human Error Investigations*. Ashgate: Cranfield University Press.

Hale, A. (2000). Le risque de défaillance et son contrôle par les individus et les organisations dans les activités à hauts risques. Communication to CNRS Seminar. Gif sur Yvette, France.

Hollnagel, E. (1998). *Cognitive Reliability and Error Analysis method*. Elsevier, Oxford, England.

Kahneman, D. & Tversky, A. (1979). Prospect theory: an analysis of decision under risk. Econometrica, 47, 263-291.

Koornneef, F. (2000). *Organised Learning from Small-scale Incidents*. Delft, The Netherlands: Delft University Press.

Pariès, J., Merritt, A. & Schmidlin, M. (1999). Development of a Methodology for Operational Incident Reporting and Analysis Systems. Final Report – Convention DGAC 96/01, Paris: DEDALE.

Rasmussen, J. (1997). Risk Management in a Dynamic Society: a modelling problem. Safety Science Vol. 27, N° 2/3 pp. 183-213.

Reason, J. (2000). Events, Individuals and Organisations. In: I. Swedung & G.M. Cojazzi. *Risk Management and Human Reliability in Social Context, Proceedings of the 18th ESReDA Seminar*, European Communities, Luxembourg.

Wioland, L. & Amalberti, R. (1998). Human error management: towards an ecological safety model: a case study in an air traffic control microworld. ECCE 1998', Limmerick, Ireland.

Woods, D., Johannesen, L., Cook, R. & Sarter, N. (1994). *Behind Human Error: Cognitive systems, computers and hindsight*. Columbus, Ohio: CSERIAC.

Page intentionally left blank

GLOSSARY

For the purposes of this document, the following definitions shall apply:

Accident Consequence Mitigation (Safety Principles): Safety Principles meant to prevent an accident from developing into its worst potential consequences.

Behaviour (of a Safety Principle): The qualification (success, failure, unknown outcome, or not prompted) of a Safety Principle involved in one event.

Criticality (of a Safety Principle):

- The 'local' criticality of a Safety Principle is the strength of the remaining protections against an accident, should that Safety Principle fail in the situation of a specific reported event.
- The 'extended criticality' of a Safety Principle is the strength of the remaining protections against an accident, should this Safety Principle fail, and had the configuration/context been different from that of the reported event.
- The 'general criticality' of a specific Safety Principle is the strength of the protections against an accident still in place, should that Safety Principle fail, across all (identified) Generic Initiators (Safety Architectures) in which that safety Principle is participating.

Generic Initiator (GI): A macro-initiator, encapsulating a set of initiators corresponding to a similar way to manage/impair safety.

HERA-JANUS: A EUROCONTROL/FAA technique to investigate the human factors associated with ATM occurrences.

HERA-PREDICT: A technique to predict the human factors issues and errors associated in ATM adaptations and changes.

Human Error (HERA definition): Any action (or inaction) that potentially or actually results in negative system effects, where more than one possible course of action is available.

Initiator: An event at the ATM system's level from which an accident would develop, should no specific recovery action be positively taken.

Prevention (Safety Principles): Safety Principles meant to prevent the occurrence of a Generic Initiator.

Recovery (Safety Principles): Safety Principles meant to prevent the Generic Initiator from developing into an accident.

Safety Architecture (SA): The logical, hierarchical combination of Safety Principles that compose the safety protections associated with a Generic Initiator.

Safety Principle (SP): Any assumption about what is supposed to make the ATM system safe in the *a priori* safety model.

Safety Management Assistance and Recording Tool (SMART): A software-based tool that acts as an interface between individual safety occurrences reports and safety management decisions.

Status (of a Safety Principle): The rating (Reliable, Unreliable, Unsure, No rating) of the empirical robustness of a Safety Principle across all behaviour qualifications in individual events.

System (ATM): Includes operational organisations ATM organisations of the size of an ACC, as well as the corresponding ATM functions implemented in the A/C cockpits.

Violations (HERA definition): Actions that contravene a rule, procedure or operating instruction.

ABBREVIATIONS AND ACRONYMS

For the purposes of this document, the following abbreviations and acronyms shall apply:

A/C	Aircraft
ACC	Area Control Centre
ATCC	Air Traffic Control Centre
ATCO	Air Traffic Controller / Air Traffic Control Officer *(US/UK)*
ATM	Air Traffic Management
ATS	Air Traffic Services
CENA	Centre d'Etudes de la Navigation Aérienne *(France)*
DFS	Deutsche Flugsicherung GmbH *(Germany)*
DIS	Director(ate) Infrastructure, ATC Systems & Support *(EUROCONTROL Headquarters, SDE)*
DIS/HUM	See 'HUM (Unit)'
DSA	Director(ate) Safety Airspace, Airports & Information Services *(EUROCONTROL Headquarters, SDE)*
EATCHIP	European Air Traffic Control Harmonisation and Integration Programme *(now EATMP)*
EATMP	European Air Traffic Management Programme *(formerly EATCHIP)*
GI	Generic Indicator
HERA (Project)	Human Error in ATM (Project) *(EATMP, HUM, HRS, HSP)*
HRS	Human Resources Programme *(EATMP, HUM)*
HRT	Human Resources Team *(EATCHIP/EATMP, HUM)*
HSP	Human Factors Sub-Programme *(EATMP, HUM, HRS)*
HUM	Human Resources (Domain) *(EATCHIP/EATMP)*

HUM (Unit)	Human Factors and Manpower Unit *(EUROCONTROL Headquarters, SDE, DIS; also known as DIS/HUM)*
IANS	Institute of Air Navigation Services *(EUROCONTROL, Luxembourg)*
LVNL	Luchtverkeersleiding Nederland *(ATC The Netherlands)*
REP	Report *(EATCHIP/EATMP)*
SA	Safety Architecture
SMART	Safety Management Assistance and Recording Tool *(EATMP, HUM, HRS, HSP, HERA)*
SOFIA	Sequentially Outlining and Follow-up Integrated Analysis *(EUROCONTROL Headquarters, EATMP, DSA, SQS)*
SP	Safety Principle
SQS (Unit)	Safety Quality Management and Standardisation Unit *(EUROCONTROL Headquarters, EATMP, DSA)*
Take-off / Landing	TO/L
WP	Work Package

CONTRIBUTORS

NAME	ORGANISATION / STATE
REVIEW GROUP	
Capt. M. O'LEARY	British Airways, UK
B. CONSIDINE	EUROCONTROL IANS
A. GUERRA	NAV Portugal
P. MANA	EUROCONTROL Headquarters
WORKING TEAMS	
Safety and Investigation	LVNL, The Netherlands
Safety and Investigation	DFS, Germany
Investigation	CENA, France
Document Configuration	
C. HELLINCKX *(External contractor)*	EUROCONTROL Headquarters

Page intentionally left blank

APPENDIX 1: METHODOLOGY FOR A PROACTIVE VIRTUAL EXPLORATION OF THE EFFECTS OF CHANGE ON SAFETY

The analysis of a reported event may highlight that for one Initiator, a Safety Principle is a minimal cut set for the occurrence of the Initiator[7] as in the following figure.

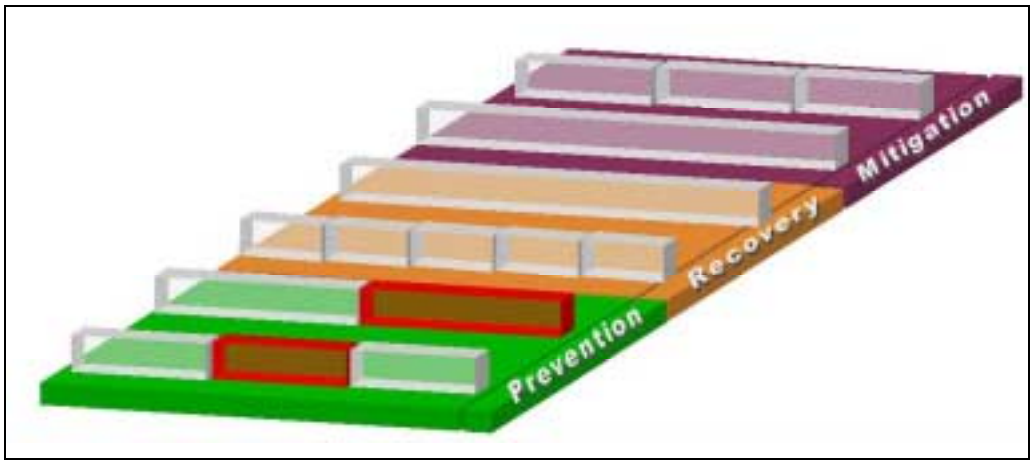

Figure 13: Safety architecture associated with Initiator X for which the Safety Principle (coloured in dark red challenged in the event under study is a minimal cut set for the occurrence of the Initiator.

In other words, should the Safety Principle be impaired in this situation, the occurrence of the associated Initiator is unavoidable. If the 'cost' of Initiator X is unacceptable to the ATM system, a change is to be made in the Safety Architecture associated with this particular Initiator. However, if the decision aims at replacing this Safety Principle by a more robust one (or several ones), it is essential to make sure that the new one(s) are compatible with the whole safety architecture, not only associated with this Initiator, X, but also with any other Initiator.

[7] The Initiator mentioned here is not necessarily that associated to the event under study, otherwise it means that the Initiator actually occurred; it may simply be for example an Initiator associated to another phase of flight.

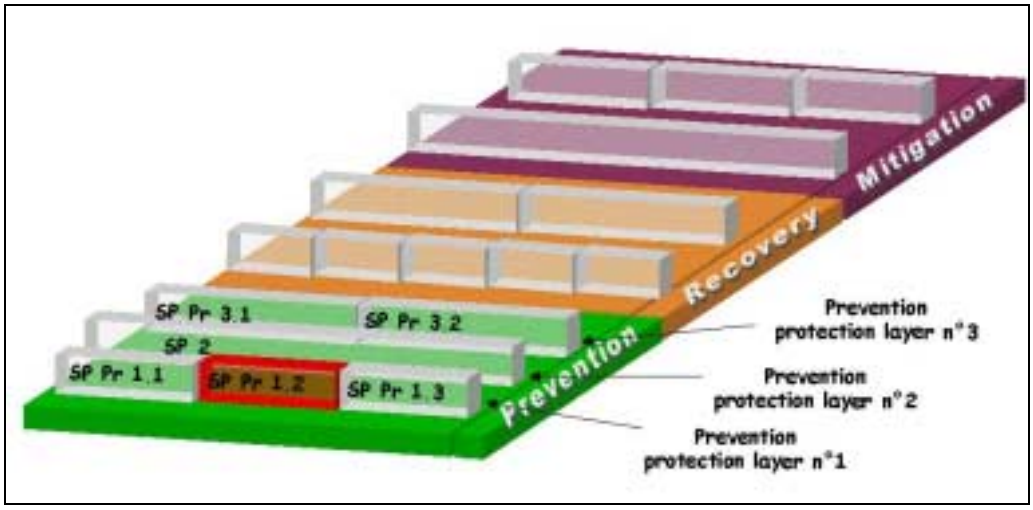

Figure 14: Safety architecture of Initiator Y the challenged Safety Principle is involved in, though not minimal cut set

The challenged Safety Principle also participates in the protection of Initiators Y. It is not a minimal cut set. However, replacing Safety Principle(s) should not make SP Pr 2 or SP Pr 3.1 or Pr 3.2 inoperative – if the cost of Initiator Y is unacceptable or if the recovery and accident consequences mitigation protection layers are not robust enough, or be such that a common failure mode exists that challenges at the same time prevention protection layers 1, 2 and 3.

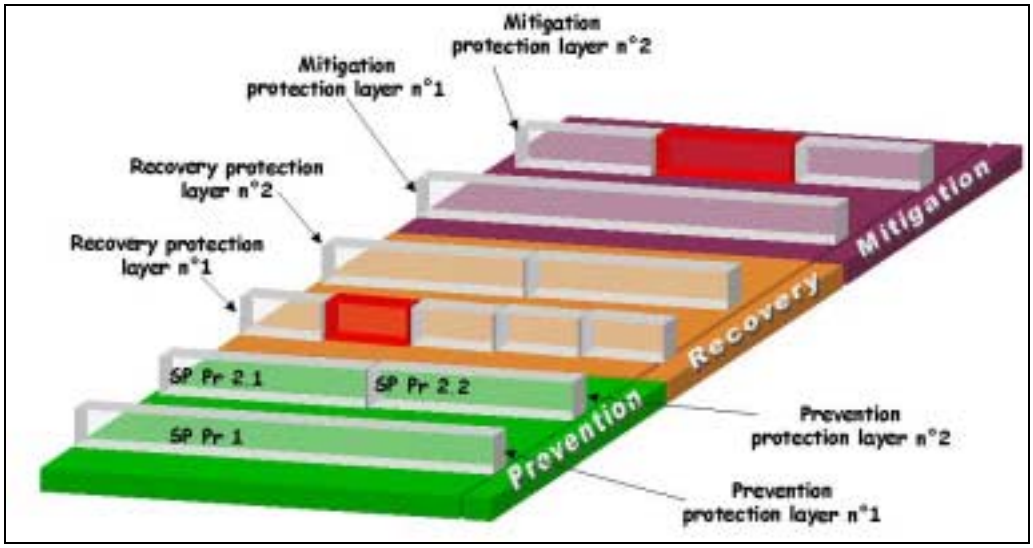

Figure 15: Safety architecture of Initiator Z the challenged Safety Principle is involved in, although not minimal cut set

In the same way, the challenged Safety Principle participates in the protections of Initiator Z, though not as a minimal cut set either. However, if the safety strategy to cope with Initiator Z heavily relies on recovery, it is essential to make sure that the option proposed as an alternative to the failing Safety Principle (coloured in red) does not have negative repercussions in terms of robustness on recovery layer 2. Side effects on other categories should also be checked. For example, it is important to check whether the new alternative modifies something at the prevention stage and if so, modifies it equally or better.

When the analysis concludes that a broader change is needed for the inadequacy of the safety model lies at another level than that of Practice (i.e. at the level of top-level Safety Principles describing strategies rather than practices), the same rationale applies. It is no longer at the level of a Safety Principle to be replaced or reinforced, but at the level of several ones, possibly involved in a range of protection layers.

This process requires a particularly careful attention when it comes to modifying something in the system that impacts the ATCO's activity (change in the interface, in the procedures, in the ATCC organisation). In such cases the consistency of the ATCOS' task has to be ensured within each situation and also across situations.

Page intentionally left blank

APPENDIX 2: CONSISTENCY BETWEEN SMART AND SOFIA APPROACHES

System excursion out of the safety envelope – Generic Initiator

The safety envelope referred to in SOFIA is a similar concept to that of flight envelope for aircraft, i.e. within that envelope, the aviation system is considered controllable whereas it is not outside of this envelope. In other words, the envelope constitutes the frontier between normal and abnormal flight operations. A system excursion out of the safety envelope is an event or a non-event that leads the aviation system to abnormal flight operations.

SMART introduces the concept of Generic Initiator (GI). A GI is defined as any event (or non event) from which an accident would develop, should no specific recovery action be taken. In that sense a GI is an excursion out of the safety envelope defined in SOFIA.

If both concepts describe the same phenomenon of transition between an intrinsically safe state to an intrinsically unsafe state, the level of detail of system excursions out of the safety envelope and of GIs is different. A system excursion out of the safety envelope as currently defined encompasses several GIs. Indeed, given the use of the safety model suggested in SMART, there is a concern to keep the GIs generic enough to derive general safety lessons from singular events, but not too generic so that the model is usable when it comes to analysing an event. Therefore, a GI corresponds to an excursion out of the safety envelope that is handled through a consistent strategy.

The definition of system excursions out of the safety envelope, e.g. ‘separation infringement between two aircraft’, could correspond to various very different strategies to handle safety, for example overseas traffic vs overland traffic, or about class A airspace vs class F airspace. A Generic Initiator is a system excursion out of the safety envelope restricted to some environmental conditions to the point that there is a consistent strategy to deal with safety under these conditions.

SOFIA introduces an intermediate concept, ‘Critical event’, that draws the frontier/limit between desired flight operations and undesired flight operations. This limit depends on subjective judgement and on the current situation (equipment available, layout of the ATCC). A critical event is a generic precursor of a system excursion out of the safety envelope. It can be either ‘two aircraft on crossing paths but with no separation infringement’, or ‘entering a runway strip which is occupied’. Depending on the situation, a critical event has a fluctuating limit, and it may not even be that critical. SMART does not use such a concept in its modelling. The modelling attempts to be as robust as possible with various situations, and as such this concept is not needed to develop the *a priori* safety model.

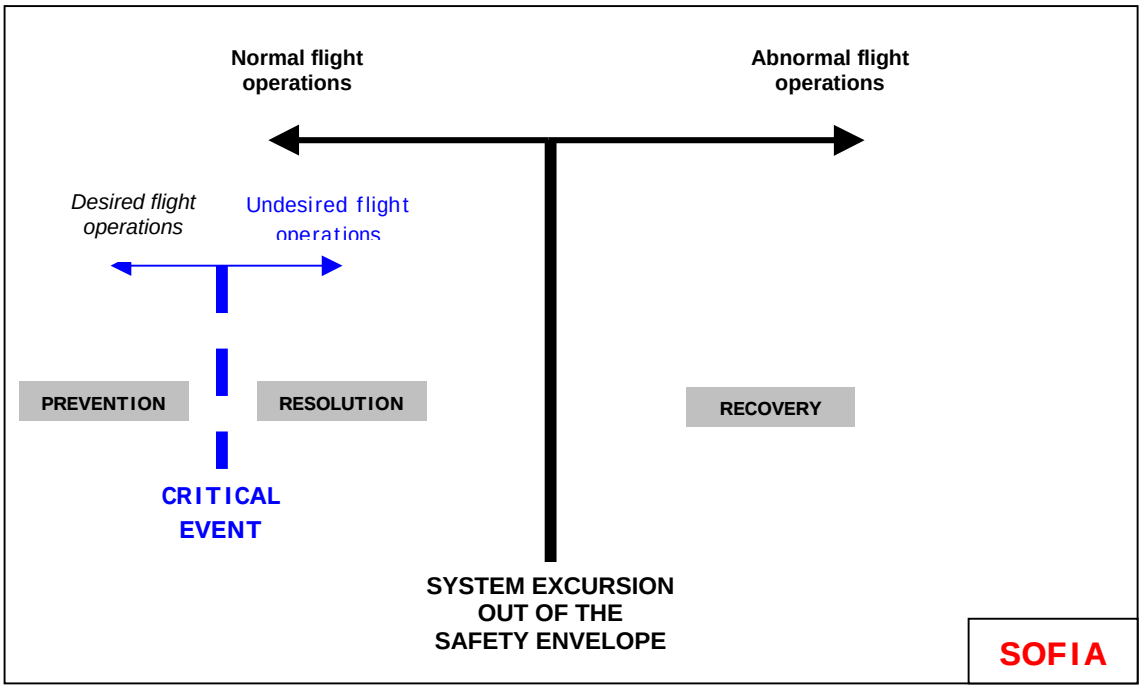

Figure 16: SOFIA Approach

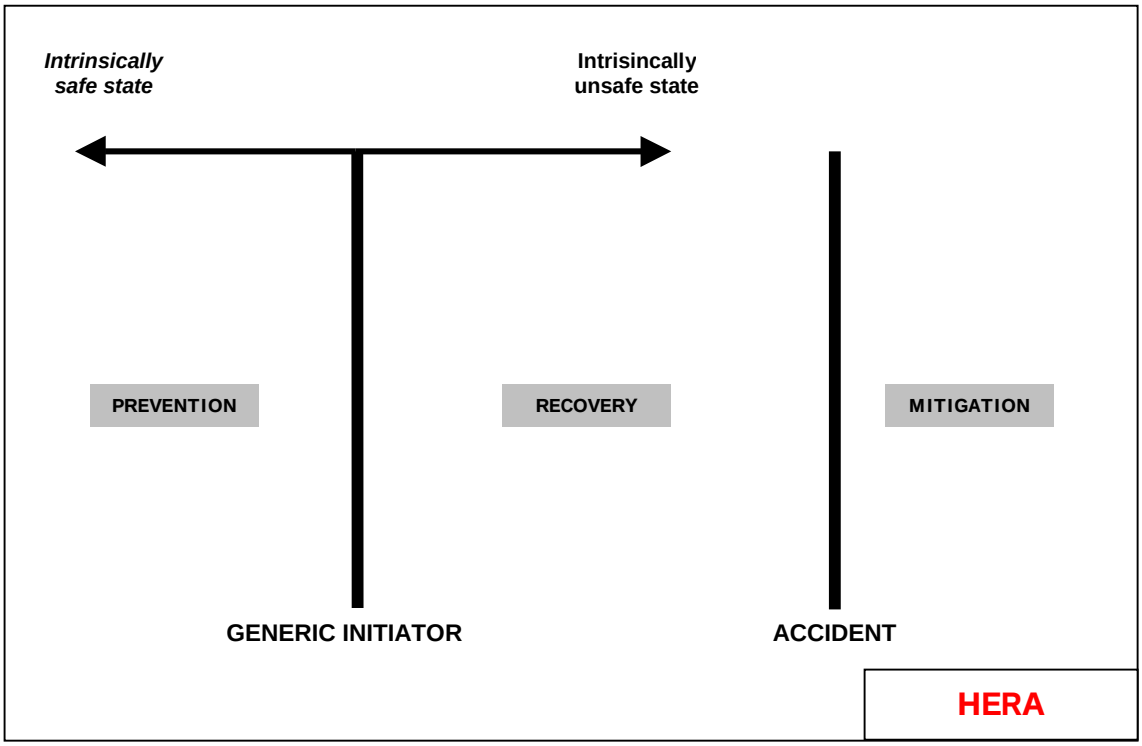

Figure 17: HERA Approach

APPENDIX 3: ATM ORGANISATIONAL SAFETY ASSESSMENT

N°	STATEMENT	YES / NO / DON'T KNOW
1.	Air Traffic Management safety is recognised as being everyone's responsibility, not just that of the safety and licensing specialists.	
2.	Supervisory staff anticipate that all staff will inevitably make errors, and train them to detect and recover them.	
3.	All managers are genuinely committed to the goals of flight safety and provide adequate resources to serve this end.	
4.	Safety-related issues are considered at high-level meetings on a regular basis and not just after an incident.	
5.	Serious incidents are thoroughly reviewed at high-level meetings and the lessons learned are implemented as global reforms rather than local repairs.	
6.	After an incident, the primary aim of the incident investigation is to identify the failed system defences and improve them, rather than seeking to pin blame on specific individuals in the operation room.	
7.	Senior management adopts a proactive stance towards safety. That is, it does some or all of the following: takes steps to identify recurrent 'error traps' and remove them; works to eliminate the work environment and organisational factors likely to provoke errors; 'brainstorms' new failure scenarios and conducts regular 'health checks' on organisational processes known to contribute to problems.	
8.	Senior management recognises that fixing error-provoking organisational factors (e.g. understaffing, inadequate equipment, inexperience, inadequate supervision, poor HNI design) is easier than stopping isolated psychological problems such as distractions, inattention and forgetfulness.	
9.	It is understood that the effective management of safety, just like any other management process, depends critically on the collection, analysis and dissemination of relevant information.	
10.	Management recognises the necessity of combining data from reactive outcomes (near miss and the incident reporting) with proactive process information (regular sampling of working practices such as rostering, workload, procedures, training etc) to identify which of these issues is in most need of attention, and then carrying out remedial actions.	
11.	Meetings related to ATM safety are attended by staff from a wide variety of areas and levels.	

N°	STATEMENT	YES / NO / DON'T KNOW
12.	It is appreciated that commercial goals, financial constraints and flight safety issues can come into conflict and that mechanisms exist to identify and resolve such conflicts in an efficient and transparent manner.	
13.	Policies are in place to encourage everyone to raise ATM safety issues.	
14.	The organisation recognises the critical dependence of a safety management system on the trust of the controllers – particularly in regard to reporting systems.	
15.	The position of Safety Manager is seen as an important position and attracts the appropriate status and salary.	
16.	There is a consistent policy for reporting and responding to incidents in all areas of the operations environment.	
17.	Disciplinary policies are based on an agreed (negotiated) distinction between acceptable and unacceptable behaviour. It is recognised that a small proportion of unsafe actions are undoubtedly reckless and deserve punishment, but the large majority of unsafe actions should not attract punishment.	
18.	Supervisors train their personnel in non-technical (TRM) as well as technical skills necessary to achieve safe and effective performance.	
19.	The organisation has in place rapid, useful and intelligent feedback channels to communicate the lessons learned from both the reactive and proactive safety information systems. Throughout, the emphasis is upon generalising these lessons to the system at large rather than merely localising failures and weaknesses.	
20.	The organisation has the will and the resources to acknowledge its errors, to address them, and to reassure the controllers that lessons learned from such incidents will help to prevent their recurrence	

YES - ***This is definitely the case in my organisation (scores 1)***
NO - ***This is definitely not the case in my organisation (scores zero)***
DON'T KNOW - ***Don't know, maybe or could be partially true (scores 0.5)***

INTERPRETING YOUR SCORE

16 –20	So healthy as to barely credible
11 –15	You're in good shape. But don't forget to be uneasy
6 –10	Not at all bad, but there's still a long way to go
1 – 5	You are very vulnerable
0	You will not survive