SAFETY SCIENCE Monitor

Special Edition 2006

Article 4

VOL 10

AIR TRAFFIC MANAGEMENT

G. MORTON
Safety Regulation Unit, Eurocontrol.

STEVE KINNERSLY
AEA Technology

BOB HUMBERTSON
Federal Aviation Administration/System Engineering and Technical

1.1 OVERVIEW OF RISK AND MANAGEMENT EXPERIENCE (G. MORTON)

1.1.1 Introduction

The ATM [1] system is generally considered to be one of the safest transport systems in the world. It is used by, and is certainly visible to, most of the western world. However, it is being subjected to considerable change due to the expectations of its stakeholders in terms of cost of operation and more recently as it opens its Safety Processes to outside attention.

This paper aims to address how ATM regulation began, what required it to change, how that change is being accommodated, and how to prevent that change from losing the desirable features of the system.

The paper also shows how these changes have been based on the idea of quantified risk.

1.1.2 The history of ATM regulation

ATM regulation followed the model of all regulation since the early railway acts in the United Kingdom[2]. Regulation was a result of a reactive response to significant events –usually accidents and based on derivation of best practices to address these incidents. International Aviation regulation[3] was based on these best practices being transferred to other areas of the world prior to those countries having the same event. We learned from lessons, captured them in regulations and used the regulatory regime to ensure these best practices were implemented correctly.

1.1.3 What were the precursors for change?

The aviation industry took account of the findings from enquiries into accidents in other industries, Piper Alpha, Challenger, the Clapham Junction rail accident, the Canadian Westray mine disaster[4]. All these inquiries indicated that accidents were not the result of single points of failure, but were failures of the system of safety management. In particular the inquiries identified that the role of the regulated organisation in taking responsibility for its own actions had been lost. Safety was something built on afterwards by the regulator, not

---

[1] ATM – Air Traffic Management. The aggregation of ground and aircraft based functions required to ensure the safe and efficient movement of aircraft during all appropriate phases of operations

[2] http://www.woodberry.org/acad/hist/irwww/Transportation/Operation/Railway_Accidents.htm

[3] http://www.icao.org/cgi/goto.pl?icao/en/history.htm

[4] http://www.gov.ns.ca/labr/westray/execsumm.htm

built in by the operating organisation. More specifically, when there was no management of the safety in the organisation safety was given low priority.

Within the aviation sector in ECAC (The European Civil Aviation Conference[5]) the safety objectives were specifically noted as being based on quantified approaches, 'to improve safety levels by ensuring that the numbers of ATM accidents and serious and risk bearing incidents do not increase and, wherever possible, decrease'[6].

With some risks within the ATM system, particularly those associated with engineered systems, it was possible to use mathematical models to simulate some of the risks[7]. From these we could set quantified targets on elements of the system. The concept of quantified risks began to take shape.

That certain landmark legal cases set the idea of liability within organisations, Particularly focussing on the idea of having defined responsibilities to consider safety within organisations[8].

When people with regulatory responsibilities and people with service provision responsibilities exist in the same organisation, there will always be the idea that there is a lack of independence between them. It is becoming increasingly expected that such separation shall exist[9]. This situation becomes especially evident when the ATM service provision, which tends to be a government agency, is removed from government through privatisation or corporatisation.

People notice events, and however misused, they can consider the rates of occurrence of events as a measure of the goodness of a system. The users of a system started to demand quantified risks.

The existing ATM system was shown to be deficient in capacity. New methods of implementing ATM functionality were required to achieve this. Moving and changing functionality meant that pre-existing defences was removed or lost; the same systems were being used for different purposes. So quantified risk evolved to replace the implicit protections.

When replacing or advancing functionality in the system, there is a possibility that the new system will be less safe than the one it replaces. We need a way of comparing systems against a common baseline. So quantified risk achieved this.

When planning for the future, we need to provide a target for assessing performance. Quantification allows this.

1.1.4 The perception of risk

There is a difference however between what the various stakeholders perceive as the risk. ATM regulators are responsible for ATM risk, aircraft regulators are responsible for airworthiness and operational risks, and the users are concerned with gate-to-gate risks.

Who might we consider stakeholders?

- The general public
- The travelling public
- The aviation using general public
- The international aviation community
- The European aviation community
- Governments
- Aviation system suppliers
- Air traffic service providers
- Regulators
- The media

---

[5] http://www.ecac-ceac.org/

[6] http://www.eurocontrol.int/eatmp/library/documents/ATM2000-Vol1en-10a.pdf

[7] http://www.eur-rvsm.com/safety.htm#assessment

[8] http://www.safeware-eng.com/software-safety/orgmanage.shtml

[9] http://www.nats.co.uk/news/(Future%20%20Av%20Reg%20summ)%20future_aviation_regulation_summary.PDF

All these perceive risk differently,

The general public might consider risk on the basis of the weight given to reports in newspapers. Strangely this generally is inverse to risk, the most unusual events (i.e. the least probable and therefore the safest) get reported more. There are also indications that the general public consider risks in relation to how much control they have in the system or its complexity. People accept higher risks in driving their cars, which they control, than in aircraft. People fear nuclear power, which they do not understand, rather than a coal fired power station.

The travelling public when judging methods of travel might take safety into account; reports on the standards applied in other countries transport systems have effect. I.e. people fly to islands after reports of ferry disasters.

The aviation using general public move away from flying after any aviation disaster[10].

The European aviation community, through initiatives at the European level, has moved towards greater harmonisation of standards[11]. While these initiatives have not always been for the purpose of ensuring increased safety, they have required countries to find a common view on the way forward.

The international aviation community judges risk differently, particularly in relation to costs. Many of these issues relate to the economies of developing countries and their ability to afford safety improvements[12].

Governments also perceive risk differently. Or more particularly they perceive their safety improvement focus differently, different countries have different priorities in transport improvement; the UK might be particularly concerned with rail travel, Austria with road tunnel safety.

Aviation equipment suppliers are particularly interested in common markets. They want to be able to sell equipment in multiple markets with a minimum of barriers. They therefore see common standards, with clear requirements as an advantage.

Aviation service suppliers might see a need to limit the liability of their operations. The cost of an aircraft disaster is very high[13], generally far higher than the income of such organisations, and possibly higher than the capitalisation of the companies. How do such organisations make rational decisions on risk cost benefits?

The gate-to-gate concept[14], where the system is considered for the entire flight from start to finish, was first proposed for economic considerations. The basis of the consideration was that it was necessary to consider the system in its entirety in order to ensure that actions in any one part of the system could be prioritised, and to a lesser extant to ensure that actions in one part of the system did not adversely affect the other. An example would include the possibility of increasing the landing rate, while the airport was not capable of moving that number of aircraft around the airport.

With this view of the entire system, it was logical to consider the risks in the same fashion; an aircraft accident to a user is still an accident when it occurs on the ground, or in the air. With the gate-to-gate concept the view of risk was to be taken from a total aviation perspective.

Within ATM we intend to use the Risk methodology to both ensure that the evolution of the system meets the target, but also to ensure that the actual performance meets the requirements.

1.1.5 The relationship between Accidents and incidents

For some time, there has been the realisation that not all failure modes in the system lead to accidents. Particular examples include the experience of other industries. When a serious accident is investigated, it is usually found to be a series of events. There is a chain of events and each stage; something went wrong that allowed the precursor to propagate through to cause the accident. For the sake of this discussion, the most important finding from these reports was that these individual events were found to have occurred on many occasions prior to the serious event. It was the combination of the events that caused the serious accident. But the

---

[10] http://www.newsday.com/news/nytwa97-crash422.story
[11] http://www.jaa.nl/whatisthejaa/jaainfo.html
[12] http://www.ifalpa.org/Press%20Release/01PRL006_ATC_Training_Africa_Oct00.pdf
[13] http://www.airlaw.com/complaint.html
[14] http://www.eurocontrol.int/eatmp/overview/rationale.html

individual events could have been identified if the incidents were evaluated. The model is generally known as the Swiss cheese model or the Reason Model[15].

So the concept of the accident pyramid became standard. In this model there is the idea that there is a relationship between the number of incidents and the number of accidents. More than this there is not just a relationship, but there is causality. In other words, if there are 500 incidents a year, this identifies a number of failure modes in the system, and then accidents occur when these failure modes occur simultaneously. The more incidents, the more identified failure modes, the more accidents[16].

So it became reasonable to not only set a target for the risk of accidents, but based on the causality, to set targets for incidents.

But it was also obvious that there were different classes of incidents. For some time, the prime mechanism for investigating such events was within the remit only of accidents, and to a certain extent, only fatal accidents. When the causal relationship was made between the accident rate and incident rate, there was a case for extending investigations below the levels of accidents.

But incidents were different than accidents, depending on the particular scenario of events, the effect on the user would differ, some incidents have no effect on the user, consider an aircraft wandering onto an empty runway, while others might be noticed by other parts of the system, such as an aircraft wandering onto the runway while an aircraft was at three miles on the approach and the aircraft initiated a go around[17].

There clearly was some difference between incidents; some concept of severity could be formulated taking into consideration the scope of the system to recover from the first event. This recovery of the system related to the idea of systems being composed of levels of defence, the more of these levels that were broken being related to the severity of the incident, and also being related to the degree of investigation necessary. In some way, an incident was an accident that did not happen today. I.e. any incident is the result of a failure of the system, and any incident can propagate to an accident, thus the number of incidents bears a relationship to the number of accidents.

So now risk became two concepts, the severity of the incident, and the rate at which that severity could occur. We not only needed to monitor accidents and consider their rate, but also the severity of incidents and their rate of occurrence.

While there are many requirements to investigate accidents, particularly fatal accidents[18], the limited number of these will be unlikely to reveal the many other factors in aviation, particularly those factors that only grow slowly in importance.

1.1.6 What was the mechanism for change?

Within the Aviation system, like all systems, there are always new functions being evaluated, and new ways of achieving existing functions. The first question is usually, “and how safe should it be?”. To do this we usually look back in the past. If we consider that the present operation is safe enough, then we can evaluate the present risk of the function, and then define this as the risk required in the future. This has been done many times, the prime example being for ‘blind’ landings. In that case, the rate of accidents due to the landing of aircraft in low visibility was set to be the same accident rate for landings in good visibility, based on data available at the time.

Aviation is a very expensive business; it was always considered that many of these costs related to the inclusion in the aviation system of many levels of defence. Some of these levels of defence were more expensive to accommodate than others. There needed to be a way of ensuring that the levels of defence were cost effective in relationship to the risk they addressed.

It’s a well known risk approach that the safest system to address the risk is to remove the risk altogether. The safest aircraft is one that never actually flies. The safest ATM system is maybe the system that only allows one aircraft in the sky at the same time. But there is a cost to such an approach. In fact, the more levels of defence are placed in the system, there is a direct correlation to the usability and hence the cost of the system. We all know

---

[15] http://www.hq.nasa.gov/office/codeq/risk/uva.pdf

[16] http://www.ccohs.ca/hscanada/accident_pyramid.pdf

[17] Go around. A practice where a pilot abandons an approach to landing

[18] http://www.iprr.org/Manuals/Annex13.html

that there is a benefit to the system, the aviation system allows us to take holidays, makes meetings more efficient and allows business and trade. This is all a benefit. For this benefit we are prepared to pay, both in terms of money and terms of risk of death or injury.

Both these ideas were encapsulated in the ALARP principle in the UK and the equivalent in other countries, where the risk is reduced to a level that the benefit of having the system is matched against the cost of the system itself[19].

With new regional groupings, there is a move to have previously different regulatory systems and approaches move towards a common goal. But similarly, we need to respect the approaches taken by different groups when these approaches are equally effective in reaching the 'goal'. To do this we need an abstract concept for which everyone can agree, without too much prescription of the approach to be taken to allow local implementations to take account of cultural differences and differences in the evolution of the regulatory system in that country.

The classic approach to regulation was to identify a failure condition, develop a method to address it and implement this as a best practice. Unfortunately, best practices are not always 'best'. In particular this mode of regulation promotes compartmentalised thinking. Smaller and smaller parts of the system are addressed. There are several problems with this.

Firstly that the solution in one area might cause problems in other areas, secondly, we might miss something at the boundaries between areas, and thirdly we might be focussing too much in one area.

The other problem with 'prescription' is that it tends to set the system in concrete. If we only address existing risks within our own frame of reference we never allow the system to change, i.e. we might address a problem with ATC understanding radio communications with better headphones rather than considering digital communications.

Prescription also reduces the desire of the organisation to take responsibility for its own actions. It begins to slavishly follow the rules.

1.1.7 What was the attempt at defining acceptable risk [20]

Generally, risk as perceived by users is in terms of the harm that can affect them. However, there is no set relationship in ATM between an accident and the number of injuries. This is for several reasons. Firstly ATM accommodates many classes of aircraft from general aviation to the largest commercial flights. An accident can occur to any such aircraft so one aircraft might have two occupants, another 300. Secondly, there is nothing in the ATM system that can affect the survivability of the occupants. One accident at low speed and low altitude results in possibly many survivors but an accident at high speed or high altitude might result in no survivors.

It was felt therefore that by defining ATM risk in terms of fatalities or injuries would add extra layers of risk decomposition that would merely make the risk approach more difficult to apply. In this example risk decomposition is any method that uses a risk defined in one term (such as fatality) and decomposes this to a risk in another term (such as loss of a radio).

Generally, the only risks that can be addressed by any organisation are those for which it has legal competence to consider. We are ATM regulators so concerning ourselves with the flammability of cabin furnishings is not something we can address. Therefore our definition of risk is limited to our scope of regulation.

Within the risk approach, it is our preference to be able to apportion risk from the total aviation scenario to the ATM element. There are several reasons for this preference. Firstly, to ensure that we are not over-engineering the ATM system in relation to the aircraft, in other words, that we are not expecting the accident rate from ATM hazards to be grossly different from Airworthiness or operational hazards. Secondly, is the overall safety of the aviation system adequate? To show this we would have to have two things, an overall target, and a method of apportionment. At the present there is no overall safety objective for aviation and the different regulatory regimes within the aviation system do not allow an apportionment mechanism. Note however that future advances in Aviation regulation in Europe; in particular EASA might provide this framework.

---

[19] http://www.iee.org.uk/Policy/Areas/Health/hsc36.doc

[20] http://www.eurocontrol.be/src/documents/deliverables/srcdoc1ri.pdf

There are many examples of similar methodologies appearing in different areas of regulation. For example, computer systems appear in ATM, avionics etc. The failure categories are the same in each domain I.e. loss of information, corruption of data, the failure modes are the same, eventual loss or corruption of function and the end hazard is the same, accident, incident etc. Therefore each domain should be able to come to similar approaches to mitigate the failures. In particular, we wish to prevent the situation that two 'boxes', existing in separate functional parts of the aircraft (ATM and Airworthiness for example), are built to different standards of design rigour but can fail and have similar consequences.

We recognise that the ATM system is but one part of the aviation system. There are also, at least, airworthiness, airline operational parts and airports too. It is therefore desirable to be able to compare our risks with these other parts; we therefore prefer to choose metrics that support this comparison.

Whenever we talk about risk, we have the problem of identifying meaningful units. This is because each person in the system has a different view of the situation. The first comparison is between what units can we use to define a flight. A passenger might gauge the risk in terms of flights; they take maybe one flight a year, the length of the flight is not material to them, they just want to get somewhere without harm. A manufacturer of landing systems is only interested in the risk associated with landings; there is only one landing per flight, irrespective of the length of the flight. An ATM provider might consider the flight in terms of the risk of exposure, the longer the flight is in their control, the more possibilities for accidents, so flight hours is a consideration.

However some of these metrics are dependent on the design of the system. One of our principles is for goal setting regulation, setting the *what* not the *how*. So our choice of metrics must be, as far as possible, independent of design.

In the ATM domain, we have chosen two metrics of importance to recognise these requirements. Firstly we define risk in terms of accidents with ATM contribution per flight hour. Secondly we define risk in terms of accidents with ATM contribution per flight.

We have found that a risk objective has three essential parts;

The thing we are trying to prevent, in this case accident but could be environmental

The scope of the events that could lead to the event, we here are only considering ATM events.

The measure, here we choose flights and flight hours being considered most appropriate to our regime.

However any system is composed of people equipment and procedures, which have failure modes that have some undefined relationship with this top effect. Just what is the relationship between the loss of a radio transmitter and the risk of an accident? We therefore need a method of being able to decompose the risk to lower levels of the system.

To include some examples we have heard risks being defined in terms of phase of flight, in terms of ATC sector[21] operation hour, in terms of volume of airspace, in terms of airspace years. All of these represent how the service provider wishes to use the top-level safety objective, and by using some part of the architectural consideration of their system, makes the risk more appropriate to their system design.

For certain stakeholders, the risks have been considered as follows:

Aircraft fleets: based on one catastrophic loss in a fleet over its lifetime.

Airline personal: radiation exposure based on solar radiation over a lifetime equivalent to x-ray accumulated dosage.

Societal risk: comparing risk of loss of life between lifetime usages in different transport domains.

There is also some use of a risk benefit methodology. In these there is the idea that we are prepared to risk harm for a benefit. We live far away from a workplace because we have a better living space, but we put up with the risk of travel. Most investments are now made on the basis of a cost benefit study, and the same idea has been

[21] ATC Sector. An area of airspace in which Air Traffic Control is carried out by one set of ATC Officers. The basic unit of En route ATC.

applied to risk. However the approach generally fails in application because it sets an explicit value on harm. It has however been successfully used as a method of prioritising road improvements[22].

There have been examples of industries making risk comparisons. These entail making a decision to spend money on safety improvements where they will have the most effect in reducing injuries and deaths. The classic example is the US FAA decision not to mandate child seats because the increased cost would mean more children carried by car, where they have a greater chance of injury[23].

1.1.8 Data

Within the ECAC ATM framework, we were set with the strategy that the overall safety objective shall be that the absolute number of accidents should not increase. This led to the need to assess data to determine the present number of accidents per year.

The first problem identified is that most data collected is collected for a particular purpose, and that purpose restricts the data. So when you try to reuse the data you find various features that indicate that the data set might not be as complete as you wish. These limitations were found to include;

Geographical differences; some data is collected across a country, or another artificial grouping. For example, we wished to consider ECAC states, but the number of ECAC states kept changing. So year on year data did not compare like with like. Other databases recorded statistics based on worldwide groupings.

Scope of events; some databases recorded accidents to different groups of users. Some applied to all aircraft, to all commercial aircraft or just to all commercial aircraft above 5070 Kg.

The direct contribution: what part of the accidents was ATM related? Most databases were built up from accident data, however the classifications of the causes of the accidents varied. Therefore it was not always possible to identify whether ATM was a contributing part of the accident and/or how much it contributed to the accident. Most databases provided only the analysis of the incident and did not allow the raw event to be reprocessed using any different scheme of severity or causes.

The different categories of events recorded in databases also caused problems. Most incidents and accidents are classified in some manner, either in terms of damage, violation of separation standards, or risk of collision etc. These classifications are not directly comparable.

Then there was the simple aspect of how much assurance could be provided that the data sets were complete. There were many queries about possible under reporting, whether due to cultural differences or maturity of reporting schemes. There was also the issue of whether there was sufficient input from ATM experts in Aircraft accident investigations.

In our approach, the different data sets were used for comparison purposes to validate the findings. It was not a simple process of averaging the findings.

In fact, complicated statistic treatments were not necessarily desirable. Such treatments tended to remove the association of people understandings of the data and the answer. In other words taking the average of two data sets might make the answer more statistically correct, but in some minds the answer just took two dubious numbers and produced a third dubious number.

It was further felt that the use of statistics was sometimes beyond the evolution of the practitioners, with certain terminology meaning little to the people actually using the data.

The important aspect was to ensure that while data was used to derive the figures, it was essential to use expert judgement to validate it. In this usage validation is addressing whether an expert would agree that the number agrees with their experience. Such validation process also ensured ownership of the final result. In itself the validation process could only supply limited actual validation, but the more people who considered the data and either endorsed it or pointed out discrepancies certainly added to the credibility of the data.

---

[22] http://www.icbc.com/Road_Safety/roadsafety_support_improve.html

[23] http://www.ntsb.gov/speeches/S960801.htm

The other importance was to ensure that the practitioners also realised the difficulties with data collection and reuse, this has the useful side effect of highlighting the importance of the data collection process and should ensure future analysis are more reliable.

Another point to note is the absence of data points. Thankfully, there are few accidents, even fewer with an ATM contribution. However this makes the assessment even more difficult, with the possibility of having missed an event making a significant difference to the result. Missing one event in a hundred is not significant, but missing one in 3 is. In general this is mitigated in one of two ways, firstly, the most significant events are the least likely to be lost, and secondly, any accident pyramid will help assure that an accident missed is accounted for (in some way) in an incident.

The last adjustment we made to the target was to account for traffic growth. Our target was to ensure *"that the number of ATM induced accidents and risk bearing incidents do not increase and, where possible, decrease"*. So this set an absolute number of accidents as a limit. But as we knew that the traffic would increase, we had to modify this based on traffic growth. As the traffic grows, so the allowable rate of accidents based on flight hours or number of flights has to reduce.

1.1.9 Future data monitoring

The purpose of setting goal-based targets is also to ensure that these goals are met. Therefore we need to have in place a monitoring system that collects the necessary data to identify the trends. This guidance is in place and takes due account of the process for collecting data and the typical failure modes that can occur in data collection[24].

The data collection is built up off a set of data that substantiates the top number, but provides the hierarchy of sub events that can allow the identification of trends in causes[25].

The main point of collecting data is to provide assurance that the goal setting requirements are achieved in practise. When they are not, there are requirements to identify a safety improvement plan that will correct the deviations. The data is then collected in such a manner that the top level goal requirements can be shown to be met, but to also allow the development of a safety improvement plan where appropriate.

1.1.10 What not to lose when starting the risk approach

There are certain intrinsic features of the ATM regulatory regime that must not be lost with the move to risk analysis. The first is the history of the system. While we can expect new failure modes with new systems we will very likely always have a frame of reference in which it can be considered. For example Pilot – ATC communication is still communication whether it is digital or analogue or by flags. We can learn much from the present system: functions, the functional failure modes, the severity of hazards and the present rate of failures. The second is the process for regulation. This is built on the idea of focussing and capturing experience. Changes are evaluated by an ever more fine set of filters representing the regulatory structure. While this might be slow and imprecise, it has useful features (review stages) that should be captured.

Aviation regulation has an interesting feature in that the higher management are well represented by persons who developed from within the organisation. There has been a history of higher levels of management being ex controllers, pilots or engineers. This means that decisions at the top are made on a more balanced (though some may say more reactionary) rational as the concepts are better understood.

Do not focus on the average numbers; each event must be assessed for lessons. Another important feature of the ATM system is its ability to review activities. Within the system are a large number of voluntary schemes, both formal and informal that provide the ability to assess system failures and to provide lessons learned and propose changes. Many of these mechanisms are built into the regulatory structure, officially or unofficially and should not be lost.

The regulatory partnership: most ATM organisations have traditionally been regulated from within. This allowed the regulator to be a partner with the service provider and provided a degree of trust and a sense of common purpose with the service provider. This allowed an easier transfer of issues across the boundary, which promoted safety. Separating the functions risks losing this free flow of information.

---

[24] http://www.eurocontrol.int/src/documents/deliverables/esarr2_awareness_package/esarr2e20ri.pdf

[25] http://www.eurocontrol.int/safety/GuidanceMaterials_HeidiTaxonomy.htm

The ATM system is build up of defined groups of professionals with recognised qualifications, whether they be the ATC operators, the Engineer designers and Maintainers or the Pilots users. Because of their common experience and qualification, these groups have developed a culture that promotes the dissemination of information. This covers problems coming out, and solutions going in. Within ATM the culture can be exhibited by a willingness to learn, willingness to share.

Hazard and risk identification is only one part of the Safety Management Process. There are plenty of non-quantifiable risks left over, and lots of out of scope risks. It is unlikely that humans can be treated in the same way as equipment in terms of design; we cannot yet determine the probabilistic failure rate of a person in certain circumstances. We will therefore always have an element of non-deterministic risk within the system to address. While we can promote new approaches we need to ensure that the mechanisms presently used to address human risk (simulations etc) are retained.

We similarly have a number of possible out of scope risks still left to consider. While we try to set the regulatory boundaries, there is always a boundary between regulators and the possibility that something will fall outside the boundaries. ATM risks for example do not cover Aeronautical Information Services, and while these are presently addressed as a strategic risk addressed by tactical mitigation (check the data when it becomes more critical to safe flight), this situation may not always exist.

1.1.11 Risk assessment scheme

In ATM we aim to define the tolerable risk in terms of a criticality of the event and the tolerable rate at which that event can occur. This reflects the concept above derived from the accident pyramid[26].

Every event that happens in the ATM system has an end effect on the Safety of aircraft. The degree to which safety of aircraft is affected is termed the severity. In the scheme chosen for ATM there are five severity categories. In normal usage this relates to the seriousness of the event. Five severity levels were chosen as this has been found to be desirable in other parts of the aviation regulation scheme (JAR 25.1309). However, while the highest severity (accident) and the lowest (no effect) are essential categories, the number of categories between them is largely arbitrary. General practice however in aviation classification schemes used for other purposes (AIRPROX, accident investigation) shows that at least 2, but less than 4 intervening categories are needed.

The five severity classifications are defined as follows;

Severity Class	1 [Most Severe]	2	3	4	5 No safety effect [Least Severe]
Effect on Operations*)	Accidents	Serious incidents	Major incidents	Significant incidents	No immediate effect on safety
Examples of effects on operations Include*):	❑ one or more catastrophic accidents, ❑ one or more mid-air collisions ❑ one or more collisions on the ground between two aircraft ❑ one or more Controlled Flight Into Terrain ❑ total loss of flight control. No independent source of recovery mechanism, such as surveillance or ATC and/or flight crew procedures can reasonably be expected to prevent the accident(s).	❑ large reduction in separation (e.g., a separation of less than half the separation minima), without crew or ATC fully controlling the situation or able to recover from the situation. ❑ one or more aircraft deviating from their intended clearance, so that abrupt manoeuvre is required to avoid collision with another aircraft or with terrain (or when an avoidance action would be appropriate).	❑ large reduction (e.g., a separation of less than half the separation minima) in separation with crew or ATC controlling the situation and able to recover from the situation. ❑ minor reduction (e.g., a separation of more than half the separation minima) in separation without crew or ATC fully controlling the situation, hence jeopardising the ability to recover from the situation (without the use of collision or terrain avoidance manoeuvres).	❑ increasing workload of the air traffic controller or aircraft flight crew, or slightly degrading the functional capability of the enabling CNS system. ❑ minor reduction (e.g., a separation of more than half the separation minima) in separation with crew or ATC controlling the situation and fully able to recover from the situation.	No hazardous condition i.e. no immediate direct or indirect impact on the operations.

Within the ATM scheme it is possible to decompose these definitions. This is desirable as it is not always possible to easily define a relationship with an event in the ATM domain with these definitions. The scheme therefore allows for a subset of definitions to be produced, typical sub-classification that might be needed include:

For each degree of severity, we define a tolerable rate of occurrence. The tolerable rate of occurrence is based on the overall safety objective of the system. I.e. the target rate of ATM related aircraft accidents.

[26] http://www.eurocontrol.int/src/documents/deliverables/esarr4v1.pdf

Severity Class	1	2	3	4	5
Maximum tolerable probability (of ATM direct contribution)	$1,55.10^{-8}$ Per Flight/Hour	To be included in a future revision of ESARR 4, once enough safety data have been collected according to ESARR 2♦.	To be included in a future revision of ESARR 4, once enough safety data have been collected according to ESARR 2♦.	To be included in a future revision of ESARR 4, once enough safety data have been collected according to ESARR 2♦.	To be included in a future revision of ESARR 4, once enough safety data have been collected according to ESARR 2♦.

♦ *To be determined at national level based on past evidence on numbers of ATM related incidents.*

The target figure has only been defined for the accident rate; this figure is derived from the ECAC safety strategy and is derived in SRU Policy doc 1.

The other figures are to be determined, and a data collection system has been put in place to collect the data required to populate this graph and the assumptions and to validate the tolerable risk matrix.

This model however only defines the tolerable risk at the ECAC level. To allow this to be used to design systems, we need a method of apportioning these targets to elements of the system.

The first level of decomposition is to the national level. This level is expected to be defined through the data driven approach, similar to the method used to define the ECAC level.

Further decomposition beyond this point is then the responsibility of national regulators. Decomposition is necessary as there may be any number of hazards within the ATM system in any one country, depending on the ATM system chosen. Each of these hazards will have a different severity classification. It is the total number of these hazards that must meet the national safety objective. For example if there are 50 hazards of severity level 1, then the tolerable rate of each of them must be (at the most) $1/50^{th}$ of the allowable rate for severity 1 occurrences.

It is how to do this proportionment that is presently taxing the knowledge of the industry. There are however the following approaches in use.

One approach is to simply define that there are no more than a certain number of independent hazards of a certain severity in the system. An approach that has been used in the design of aircraft systems (JAR25.1309) for some years. Then if there are 100 independent hazards, then the tolerable rate for each system level hazard is 1/100 the tolerable rate at the national level.

It is possible to separate the ATM system by operational functions, for example, the landing function and the separation function. Providing these functions are truly independent, it is possible to simply divide the risk between them. If the systems are constructed in such a way that there is a clear connection and independence of function at system level and operational function, then the risk can be allocated.

Within international standards, particularly ICAO, some numerical safety requirements already exist. It is possible that these can be used to provide a framework for interpreting the ones that are missing.

If the system under consideration is merely replacing functionality directly, it is possible to determine the safety requirements from the performance of the existing system (subject to the existing system being adequately safe).

One of the major reasons for using the risk-based approach was to enable reuse of systems without re-qualification. Therefore it is possible to take a system that has been assessed in one part of the world, and subject to revalidation of the operational usage, that system can be assumed to be as safe in the new situation.

Within other standard bodies there are frameworks available for the consistent treatment of system design that might be usable within the ATM service[27].

One of the purposes of the risk-based methodology is to allow design qualification. The idea being that it is possible to define sets of design, specification and testing methodologies that allow the performance of any system designed using these methodologies to be assumed to have a certain performance. Various international bodies provide frameworks for the development of such qualification standards[28].

There are various areas of development that will be necessary before the scheme will be fully in use. Though it is important to note that these developments are not necessary to allow the approach to be used, organisations have applied the scheme successfully; development is necessary to get consensus that the approaches are workable. These developments include.

The abstraction of the hazard definitions to the ATM system level is difficult for people who have a normal scope within the function of one system.

The absence of reliable data makes the setting of the tolerable rate of occurrence for hazard severity's 2 through 5 difficult.

The method of allocating the rate of occurrences down through the systems is difficult.

At some point in the system, a safety requirement is allocated to a person, and at this point the performance element of the requirement loses its meaning. The method of relating human related safety requirements is presently unclear.

1.1.12 Risk assessment is only one part of SMS

We should not however only focus on risk assessment; within the ATM regulatory scheme risk assessment is only one part of the overall process.

Partly in order to preserve the intrinsic safety elements of the existing ATM system, and partly from experience in other industries, we have mandated the use of Safety Management Systems [29]

1.2 ANALYTIC ISSUES, DECISION PRINCIPLES AND POLICIES: TARGET LEVELS OF SAFETY (STEVE KINNERSLY)

This paragraph addresses the issue of how to demonstrate that Target Levels of Safety (TLS) are met. This is important not just for technical and regulatory reasons, but also regarding the public acceptability of the risk from potentially hazardous industries. Three specific aspects are considered:

- Measuring compliance with very high TLS (i.e. very low risk)
- Consistency in the international arena
- Coherence among TLS

Compliance with a very high TLS can be difficult to demonstrate. It typically requires data for many years in order to have confidence that the TLS has been met. However, these are historical data, so how can we be sure that the TLS is still being met or will be met next year? Similarly, a single accident might violate the TLS for that year. However, the TLS might still be achieved over the longer term. Lower-level targets provide a possible solution. Issues raised by this approach are considered.

For an international industry such as aviation, consistency in the international arena is important. Passengers expect the same level of safety wherever they are. Consistent monitoring of compliance with TLS in different countries is therefore important. Issues considered here include international consistency in the understanding of risk concepts, incident monitoring, interpretation and recording. Examples are given from air traffic management.

---

[27] Examples include IEC61508

[28] UL 1998

[29] http://www.eurocontrol.int/src/documents/deliverables/esarr3v10ri.pdf

Aviation, in common with some other industries, uses many TLS. This situation has arisen partly for historical reasons, partly for technical reasons. New TLS are still being introduced. Coherence among these TLS and their compliance criteria is needed. Implications and ways in which they can be addressed are considered.

The aviation industry is used here as an example. However, the principles and lessons apply generically.

1.2.1 Introduction

A Target Level of Safety (TLS) is intended to be achieved. Whether or not a TLS has been met for a particular industry, plant or activity is important for many of the stakeholders involved. For instance:

- A Regulator needs to know whether the terms of a licence have been satisfied.
- An owner or operator of a plant needs to know whether their operations have been acceptably safe
- The public needs to know that they have not been exposed to undue risk

How to know that a TLS has actually been met is therefore an important, fundamental question. This paper addresses three specific and important aspects:

- Measuring compliance with very high TLS (i.e. very low risk)
- Consistency in the international arena
- Coherence among TLS

1.2.2 Measuring Compliance With Very High levels of safety

Problems With Small Numbers

Determining compliance against some TLS is conceptually straightforward, whether or not it is easy in practice. For example, fatal road traffic accidents are, unfortunately, fairly common in industrialised countries. A TLS of 'No more than X fatal accidents per year' is clear and unambiguous. It is straightforward to demonstrate compliance. Simply count the number of fatal accidents each year and compare with X. If it is less, then the TLS has been achieved. If it is more, then it has not been achieved.

The main reasons why it is straightforward (if not easy) to demonstrate compliance are:

- The events that must be measured (fatal accidents) are clear and unambiguous
- There are enough events each year that the total can be clearly and unambiguously compared with the TLS

For very safe (i.e. low risk) industries, however, one or possibly both of these reasons is absent. Typically, the TLS is for very few events each year. Often, the TLS is a small fraction of the event each year. Clear and unambiguous counting of events each year and comparison with the TLS is often difficult or impossible. Similarly, the type of events that are used for the TLS may now be less straightforward to detect or measure. Perhaps more typically, residual ambiguity in some events that would not be significant if there were many events per year becomes important if there are very few.

Consider, for example, an idyllic time in the future when driverless control systems for road traffic have reduced the number of accidents per year to a low level, safety devices mean that immediate death is very unlikely even if there is an accident and medical treatment can give even the very severely injured several years of good quality life. There is still a TLS for fatal accidents, but it is now very small: less than 10^{-1} per year (for an entire country!). Demonstrating unambiguous compliance is now much more difficult:

- Direct fatalities can still be measured, but there are very few (perhaps none). So should accidents resulting in delayed deaths (people given a few more years life my advanced medical treatment) be counted? Should they be given the same weighting as accidents causing direct fatalities?
- Fatal accidents come in units of one, but the TLS is less than one per year. So does one fatal accident per year not comply but no fatal accident does comply? And what if there is an accident where medical treatment has given a seriously injured person a few more years of life?

These difficulties are not new. They are present in principle now, when fatal accidents are measured in thousands per year. Delayed death is real (and well known to lawyers) and the possibility of exceeding the TLS by one fatal accident is possible in principle. What is different is that such issues are not significant now. They become significant only when the target level of safety is very much higher (i.e. the target risk is very much lower). This is summarised diagrammatically in Figure 1.

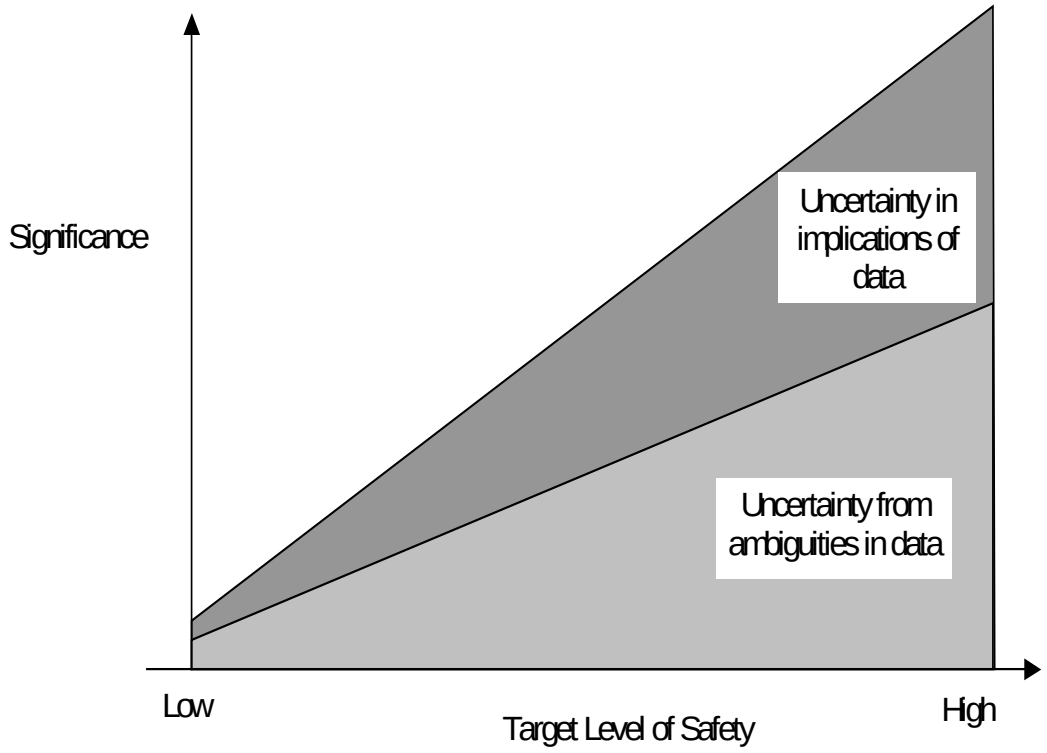

Figure 1 : Significance of Data Uncertainties for Compliance with TLS

What Does Compliance Mean?

Considerations such as the above make it important to question what complying with a TLS means. The simplest, most direct and intuitive meaning can be summarised as:

1. The relevant events (or risk) have been counted or measured.
2. The number has been compared with the TLS.
3. The number is less than the TLS.

This meaning is clearly applicable to the case when there are a lot of events, such as current fatal road accidents. It is more difficult to apply when the TLS is a very small number, such as the hypothetical future fatal road accident scenario given above. The reason is because of the implicit assumptions made at each step. For instance:

1. The relevant events (or risk) have been counted or measured.

   [Assumptions: They are unambiguously defined and counted or measured; Data are complete;]

2. The number has been compared with the TLS.

   [Assumptions: Comparison is possible; Comparison is unambiguous;]

3. The number is less than the TLS.

   [Assumption: It is possible to say this with certainty;]

The assumptions are not necessarily true.

However, it is also possible to view compliance with a TLS as a form of hypothesis testing rather than a direct comparison of numbers. The hypothesis is 'There is compliance with the TLS'. The meaning of compliance with the TLS is then:

1. The number of events (or risk) has been counted or measured.
2. The hypothesis has been assessed in the light of this number and any other relevant evidence.
3. The evidence supports the hypothesis to a sufficient degree of belief.

Compliance with the TLS is expressed in terms of a degree of belief in the hypothesis of compliance with the TLS supported by relevant evidence. In the 'large number' case, direct comparison of numbers can be strong or compelling evidence for (or against) the hypothesis and the hypothesis and 'direct comparison of numbers' meanings are equivalent in practice. Hypothesis testing, however, is not restricted to a direct comparison of numbers. Sophisticated statistical tests are routinely used for hypothesis testing in, for example, medical trials and epidemiology.

Expressing compliance with a TLS in terms of a degree of belief that the TLS has been met is not in itself unambiguous. Degree of belief can be interpreted in terms of classical confidence levels and long-run probabilities. However, it can also be interpreted in terms of subjective probabilities i.e. a Bayesian view of compliance. In this case, the degree of belief is the subjective likelihood that that the hypothesis of compliance with the TLS is true. Consistent use of evidence to support the hypothesis, and the updating of the likelihood due to new evidence, is then ensured by the use of Bayes' Theorem.

There has been much debate over the years about the pros and cons of the Bayesian interpretation of probability and its use in hypothesis testing. A readable account of the application of the Bayesian interpretation and methods is given in (Loredo 1990). The motivation in this case for the use of the Bayesian interpretation and methods is astrophysics. Events such as supernova and gamma ray bursts are rare, so data are few and often uncertain. A Bayesian approach rather than long-run probabilities appears to be the most natural.

Whether or not a Bayesian interpretation is taken, the above discussion has highlighted a number of important points regarding a coherent approach to monitoring compliance with TLS for very safe industries:

- Events should be defined and identified as unambiguously as possible
- There should be at least one clear, coherent and justifiable method for assessing compliance with a TLS
- Compliance may have attached to it the appropriate degree of belief

The last point is interesting because it would permit compliance with a TLS to be shaded according to the strength of evidence. For example, suppose a TLS is 10^{-1} fatal accidents per year. None occurred last year, but there were a few non-fatal accidents. A straight comparison of the number of fatal accidents says that the TLS was achieved. However, the few non-fatal accidents suggest that avoidance of a fatal accident involved a certain amount of luck. Thus, the degree of belief in compliance with TLS is assessed to be (say) 90%. Next year there are still no fatal accidents, but number of non-fatal accidents increases markedly. Again, a straight comparison of fatal accidents says that the TLS is achieved. However, the increase in the number of non-fatal accidents suggests that safety is actually being degraded. The degree of belief in compliance with the 10^{-1} TLS is therefore reduced to (say) 70%.

It might be argued that public acceptance and understanding of a degree of belief in achieving a safety target would not be forthcoming. Nevertheless, the public (at least in the UK) is now quite used to weather forecasters saying that there is an 80% chance of rain and opinion pollsters saying that their results are only accurate to +/- 3%. They are also used to the idea that English criminal law requires proof of guilt beyond reasonable doubt while civil law requires proof only on balance of probabilities. Perhaps 'beyond reasonable doubt' or 'on balance of probabilities' could be applied to compliance with very safe TLS.

1.2.3 Measuring Trends

Detection of trends in safety performance is an important part of managing safety. Simply knowing that a TLS was met last year is not usually sufficient on its own. There needs to be confidence that the TLS will be met this year and in the future. If that confidence is not there, then something must be done to ensure that the TLS will be met. Thus, detecting trends in measured data is important even though the data may currently show compliance with the TLS.

Trends can arise for a number of reasons. Examples from ATM include:

- More miles flown
- More flights
- New ATM technologies
- New aircraft types
- Procedural changes
- Changes to airspace structure
- New safety management regimes

An imperative behind much of the current developments in ATM in Europe is to increase the capacity to cope with increased demand while having no increase in the number of ATM-related accidents each year. A TLS has been set that takes into account the currently achieved level of safety and projected growth in air traffic (ESARR 4). Complying with that TLS in future implies a decreasing trend in the accident rate expressed in terms of accidents per unit of distance flown per year (or flights per year). How can we know that there is such a trend? Might we believe that there is the right trend, only to find out (when it is too late to do anything about it) that our belief was unfounded and the TLS is actually breached?

A number of factors make detection of trends a difficult problem for very safe industries such as ATM. For example:

- Technology changes – what does the new technology really do?
- Operational noise – things that cannot be controlled such as weather - becomes significant
- You have dealt with what you know about, so what is left is inherently ill-understood

The net result is that there the small number of residual accidents can appear as a random sequence of unique events. Indeed, it has been argued that for such industries what is seen now (in terms of safety) gives little idea of the future (Amalberti, 1999). Detection of trends, let alone extrapolating into the future, is then difficult and uncertain.

As an example of the difficulties of detecting trends, consider Figure 2. This shows hypothetical accident rates (accidents per year) for a 20-year period (Years –9 to +10). What trend (if any) is apparent and what (if anything) does it say about the future? The accident rate appears to be fairly uniformly distributed about a mean of about 20 throughout the 20-year period. An unexpected increase to above 40 in Year +7 is followed by a comforting decrease to about 20. Thus, the future is likely to be like the past: accident rates scattered about a mean of about 20.

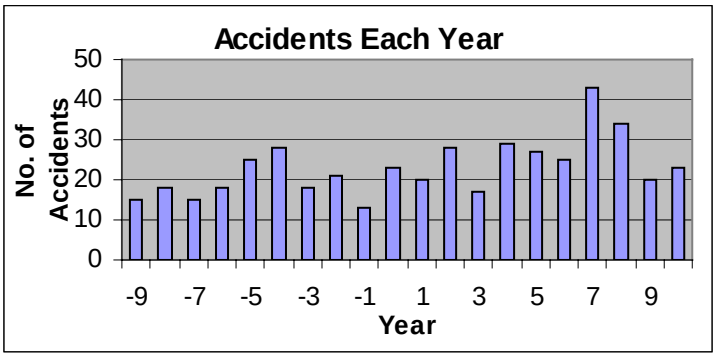

Figure 2 : Hypothetical Accident Rates for a 20 Year Period – Example 1

Consider now the hypothetical accident rates in Figure 3. These are the same as in Figure 2 for the first 10 years but differ over the last 10 years. The last 10 years shows a gradually increasing accident rate to over 50 at

the end of the period. It would be reasonable to expect the increase to continue over the next few years, the accident rate reaching 60 or beyond. This is quite different to the picture of the future given by Figure 2.

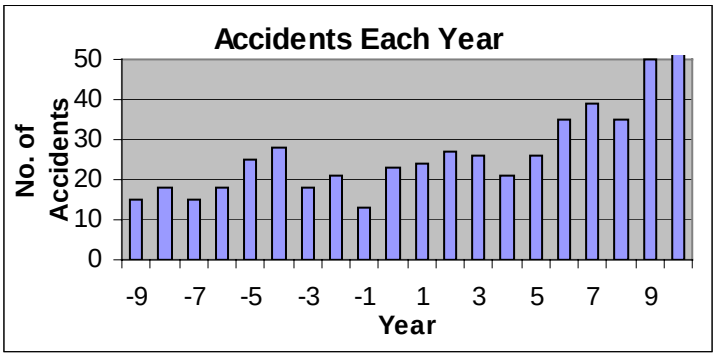

Figure 3 : Hypothetical Accident Rates for a 20 Year Period – Example 2

In fact, both Figures are examples of exactly the same underlying trend. The first 10 years (the same in both cases) are an example of real accident data from the aviation industry. They have a mean of 19.4 accidents per year and a standard deviation of 4.79. The last 10 years are modelled in both cases as random numbers from a normal distribution with mean and standard deviation that start with Years –9 to 0 values and each doubling linearly over the final 10 years (i.e. linear increase in mean from 19.4 to 38.8 accidents per year, and in standard deviation from 4.79 to 9.58). The difference between the two sets of accident rates is entirely due to the throw of the dice. The underlying trend is actually the same in both cases.

Now, suppose there is a TLS of 60 accidents per year (this is purely hypothetical and for illustration only). The data in Figure 3 shows an increasing trend that threatens to exceed the TLS in a few years. However, an increasing trend is discernable, albeit not necessarily accurately quantifiable, from the data up to Year +7. There is warning of an underlying problem and something can be done before the TLS is exceeded. Contrast this with Figure 2. Move forward enough years and the comforting picture of 'no change' given by Figure 2 will undoubtedly change. The accident rate will increase. However, there has been no advanced warning so no preventive measure are likely. Exceeding the TLS is likely to come as a surprise.

These examples show how difficult it can be to determine underlying trends in accident data (and perhaps even incident data) from very safe industries. They also show how important it is to detect those trends. The next section considers an often-used approach for detection.

Lower-Level Targets

A common approach when data for determining compliance with a TLS are scarce is to define lower-level targets. These are TLS for events that are correlated in some way with the higher-level target events and occur more often. Thus, they suffer less from the problems of small numbers.

For instance, the number of accidents due to ATM is very small indeed. However, safety-related events such as loss of minimum separation between aircraft are monitored and TLS can be set for them (such events are usually called incidents). Loss of minimum separation is correlated with accidents since a collision between two aircraft is always preceded by loss of safe separation. Of course, the converse is not true. Loss of minimum separation is not usually followed by a collision. Another example of a lower-level safety-related event for ATM is level bust. A level bust involves a descending or ascending aircraft descending or ascending past its cleared flight level. Level bust can clearly lead to loss of minimum separation and thence to a collision.

The larger number of lower level events (incidents) permits trends to be identified and action taken if necessary. A recent good example of this is level bust. Concern by the UK regulator about increasing numbers of level busts lead to action and a significant reduction in the number of level busts. That is clearly good for safety.

More generally, TLS for ATM incidents are being set in the EUROCONTROL area. ESARR 4 defines four levels of incident severity below 'accident'. Maximum tolerable probabilities (i.e. TLS) will be determined once enough data have been collected and validated.

Use of lower-level TLS is clearly beneficial for safety. However, the extent to which they can help to support claims for compliance with the top-level TLS is less certain. A common argument is based on the

'accident pyramid'. This claims (and is supported by evidence) that there is a correlation between the rate of incidents and the accident rate, with corresponding correlations between more severe and less severe incidents. Accidents are more frequent than severe incidents, which in turn are more frequent than less severe incidents. Furthermore, the ratio of accidents to incidents, and of severe to less severe incidents, can be determined from accident and incident data. Thus, by setting TLS for incidents and accidents that are consistent with these ratios then monitoring compliance with TLS for incidents, compliance with the TLS for accidents can be demonstrated.

This argument is valid as long as the data are valid and sufficient. In general, the argument that compliance with lower-level TLS implies compliance with the top-level TLS can be made as long as the data apply to the actual system under consideration. Unfortunately, the argument is weakened in one situation where compliance with the top-level TLS is of particular concern: when the system has recently changed. Confirming compliance with the top-level TLS after a change is obviously important. However, the change may have altered the characteristics of the system such that the original ratios of accidents and incidents no longer apply. Using the old ratios may be misleading. In particular, it may hide degradation in safety that is not yet apparent at the top-level because of the small numbers issues.

The long-term solution is, of course, to collect more data until new ratios can be determined. However, this takes time and it may take too long to collect the data for it to be useful – indeed, with the current pace of change in ATM another change might be introduced before enough data have been collected. This is not unlike the situation faced by developers and users of safety-related software. Operational data may be essential to justify the safety integrity of the software. However, new versions may need to be introduced regularly for operational reasons. Each new version sets the clock back to zero; previous data are no longer valid.

The Bayesian 'degree of belief' approach to compliance with TLS offers a possible resolution to this problem. Before the change to the system, compliance with the top-level TLS might be justified to a high degree of belief – say 95%. After the change, uncertainty about the ratios of accidents and incidents weakens the evidence compliance so the degree of belief in compliance is less. Analysis of the change might suggest a significant, but currently unknown, effect on the ratios corresponding to a significant reduction in degree of belief in compliance – to say 75%. The degree of belief could, however, be increased if the analysis were to provide a quantitative estimate of the change in the ratios and the new incident data were consistent with the new ratios. Eventually, incident data from the changed system would re-establish the ratios and allow the high degree of belief in compliance with the top-level TLS to return.

1.2.4 Consistency

Consistency in determining compliance with TLS is important, particularly in an international industry such as aviation. Inconsistent monitoring for compliance year on year can result in trends being missed or spuriously detected. Inconsistent monitoring by different organisations or in different countries can produce a distorted picture of safety that may be dangerously misleading. Action may not be taken when it should, or unnecessary action taken that diverts money and resources away from other areas that are actually more important.

Such issues are, of course, not just the concern of the safety community. They arise whenever monitoring for compliance is important. Examples include compliance with arms control agreements, carbon dioxide emission targets and fishery quotas. Each has its own special technical, sociological and political issues while sharing some generic common ground.

The following sections address three important areas:

- Completeness of reporting
- Definitions and meanings
- Use of computer tools

Completeness of Reporting

Credible demonstration of compliance with a TLS requires that reporting of incidents and accidents is as complete as reasonably possible. Incomplete reporting results in uncertainty in, or overestimation of, the level of safety achieved. It also makes it difficult or impossible to make realistic comparisons of safety and potentially dangerous trends may not be recognised sufficiently early. These issues are particularly significant for a very safe industry where the number of incidents or accidents is very small. Omitting or misreporting one accident when the TLS is 1,000 per year is not likely to be significant for compliance. Omitting or misreporting one accident when the TLS is 5 per year may be very significant.

Completeness depends on two factors: recognition that something has occurred that should be reported and then reporting it. When a TLS is expressed in terms of easily detected accidents (e.g. aircraft crash) or clearly defined specific incidents (e.g. loss of separation), the monitoring system can be designed to look for these specific events. In principle, they can all be detected and reported. However, recognition that something has occurred that should be reported can be impaired if the definition of what is a reportable incident is not prescriptive. A non-prescriptive definition is nevertheless desirable from the point of view of needing to capture all relevant incidents but places a burden of judgement on the person carrying out the monitoring.

An example of both prescriptive and non-prescriptive accident and incident reporting requirements occurs in ESARR 2 'Reporting and Assessment of Safety Occurrences in ATM' issued by the EUROCONTROL Safety Regulation Commission. This gives a prescriptive list of accidents and incidents that is the minimum that must be reported and assessed (Appendix A of ESARR 2). However, there is also a non-prescriptive requirement '… for any person or organisation in the aviation industry to report any such occurrence or situation in which he or she was involved, or witnessed, and which he or she believes posed a potential threat to flight safety or compromised the ability to provide safe ATM services. Such provisions shall not be restricted to the reporting of aircraft accidents or serious incidents, since other types of occurrences could reveal the same types of hazards as accidents or serious incidents'. Judgement is required to decide whether an occurrence or situation posed a potential threat to safety. Different people may interpret a situation in different ways or have different thresholds for what is a threat. Completeness of reporting requires that suitable steps are taken (e.g. training) to ensure that all occurrences and situations that are a potential threat to safety are, in fact, recognised as such by all relevant people.

Completeness can also be compromised if the occurrence of an incident is not noticed even though it would have been reported if it had been noticed. For instance, a busy human being may not notice an incident that is a brief transient (e.g. separation only just lost and for a short period). Computer tools are therefore being developed to address this problem by the automatic detection of incidents (see section on 'Computer-based Tools' for an ATM example).

Definitions and Meanings

Unambiguous reporting and investigation of incidents and accidents for compliance with TLS always requires careful definition of words and terms. This is exacerbated in the aviation industry by the need for consistency across different languages and cultures. There is therefore a strong emphasis on precise, unambiguous terminology. This has resulted in the development of detailed taxonomies for incident and accident reporting.

A recent example of a taxonomy for ATM is HEIDI. The HEIDI Taxonomy is a set of fields and definitions together with a classification scheme supporting the reporting and investigation of ATM accident, incidents and occurrences as defined into the Gate-to-ate Concept. This has been developed by EUROCONTROL as part of the harmonisation of incident and accident reporting in European ATM. HEIDI covers background terms, event types, descriptive factors, explanatory factors, classification scheme and safety recommendations. Standardisation of reporting based on HEIDI will contribute to consistency of safety monitoring. The HEIDI taxonomy is publicly accessible via the EUROCONTROL website [EUROCONTROL, 2001].

Level	Term Title	Reference	Definition	Explanations	Inputs
3	Short/medium term ATC "Planning"	HEIDI	A situation where a conflict was detected and a plan for a course of action elaborated for some reason was not/ could not be implemented..	Planned not implemented, not planned too late etc…etc… Forgotten, not physically implementatble (e.g. frequency congestion)	Planned And Not Implemented/Too Late/Physically Not Implementable

Table 1: HEIDI Example – From 'Descriptive Factor'

Level	Term Title	Reference	Definition	Explanations	Inputs
2	Accident	ICAO	An occurrence associated with the operation of an aircraft which takes place between the time any person boards the aircraft with the intention of flight until such time as such persons have disembark, in which: - a person is fatally or seriously injured (as a result of..) - the aircraft sustains damage of structural failure (which..) - the aircraft is missing or is completely inaccessible	ICAO Annex 13: - mid-air collisions between aircraft or between aircraft and other objects - collisions on the ground including Controlled Flight Into Terrain or collisions on the ground between aircraft or between aircraft and other objects	Y/N

Table 2 HEIDI Example – From 'Descriptive Factor'

HEIDI has been produced in English, which is the common language in aviation. However, different languages and cultures can have an important impact on the consistency of reporting. Meanings must be consistent in all the languages that may be used by the people who are required to report incidents or accidents. Concepts must be consistent as well as words. For instance, 'safety' and 'security' might be translatable into another language, but the difference between them may be understood as something quite different. And does 'mishap' mean the same on both sides of the Atlantic? Consistency in determining compliance against TLS in the international arena requires careful attention to definitions and meanings.

Computer Tools

Computer tools can facilitate consistent reporting and investigation of incidents and accidents. In particular, for the multi-national, multi-cultural aviation industry, the use of a consistent set of computer tools is a means of ensuring that the data from many countries is reported and can be combined consistently and coherently.

Computer tools can be particularly valuable in two areas: automatic detection and reporting of events and collection and collation of data. Two examples of developments by EUROCONTROL show the type of developments that are taking place in ATM and raise general issues regarding the use of such tools.

Firstly, the Automatic Safety Monitoring Tool (ASMT) is being developed for automatic monitoring of specific safety-related events. Automatic monitoring has the advantages of being objective and rigorous in detecting and recording. However, it can only provide facts, and then only those facts that have been chosen for recording. It cannot provide reasons, opinions or judgement. Automatic monitoring also raises human issues relating to 'the spy in the machine' and the use that is made of the data. Furthermore, automatic monitoring raises the possibility of systematic, repeated reporting errors due to deficiencies in the specification, design, development or implementation of the tool. While automatic monitoring can bring many benefits, these other issues should be addressed whenever introduction of an automatic monitoring system is being considered.

The second example is the TOKAI toolkit (TOKAI = Tool Kit for ATM occurrence Investigation). Occurrences may be recorded initially either automatically or by a human-based reporting system. Whatever the source of the incident report, TOKAI provides functionality for:

- Notification of safety occurrences in standard format
- Investigation activities, including data gathering and input, safety occurrence reconstruction, analysis and classification
- Safety recommendations
- Exchange and reporting according to specified requirements

TOKAI uses the HEIDI taxonomy and internationally agreed standard formats to facilitate consistent reporting and data exchange. Thus, both the strengths and weaknesses (if any) of HEIDI are inherited by TOKAI. This is a characteristic of automatic systems for collecting, collating and reporting on incident or accident reports (or, indeed, other types of report). Automation pushes towards reporting by pre-determined categories and

definitions. In doing so, care needs to be taken that the loss of the flexibility and nuances possible with free-text reporting is not to the detriment of the quality and completeness of the reports.

1.2.5 Coherence Among TLS

Coherent TLS – Why?

People are protected by a multitude of TLS. As an aircraft comes in to land at a UK airport, its passengers are protected by TLS for the aircraft itself and for the air traffic management system that is guiding the aircraft towards the runway. The latter includes TLS for both the air traffic control system and the automatic landing system. Meanwhile, people under the flight path are protected by TLS incorporated into airport Public Safety Zone planning requirements.

This simple example illustrates some of the reasons why coherence among TLS is an important issue. If a person arrives in an aircraft, are they better protected from the hazards of ATM or from the hazards of the aircraft and its on-board systems? Why? Is their exposure to risk from the automatic landing system greater or less than that from the air traffic control that is ensuring safe separation between aircraft during their approach and descent? Why? And if that person then returns to their home in the Public Safety Zone, are they better protected while landing in the aircraft or while at home as another aircraft flies overhead on its descent to the runway? Whatever the answer, why should it be like that? Are the TLS for the airport Public Safety Zone and for the descending aircraft consistent? And how do both compare with the risks on the road as the passenger drives through the heavy traffic from the airport to their home?

Passengers and local residents are not the only people with an interest in the coherence of the TLS in this example. Achieving high levels of safety costs money. Is the ATM service provider paying an excessive cost for protecting the passenger compared with the aircraft manufacturer or airline? Must the on-board electronics for the automatic landing system be designed and produced to a higher level of integrity than the electronics that control the engines? Why?

Coherent TLS – Meaning?

So what does coherence among TLS mean? One can think of coherence as meaning that differences among TLS are both reasonable and rational. ‘Reasonable’, as used here, means that the differences are compatible with the levels of safety that are acceptable to the stakeholders. ‘Rational’ means that the differences are supported by argument and evidence. Of course, ‘Reasonable’ and ‘Rational’ are not easily achieved. Stakeholders do not always agree on what is an acceptable level of safety. ‘Compatible with …’ may require careful compromise. Alternative logical arguments may imply different TLS. Evidence may be incomplete or ambiguous. It is difficult to avoid the conclusion that coherence among TLS is a subjective and sociological issue as well as a technical issue.

Achieving Coherent TLS

There is no generally accepted methodology for achieving coherence among TLS. The ASTER (Aviation Safety Targets for Effective Regulation) study recently carried out for the European Commission (Joyce, 2001) identified 51 TLS for the aviation industry alone. They range from accident rates for in-flight collisions to probabilities for aeronautical data corruption at airports (which could lead to unsafe flight and landing). The TLS have been developed over many years and by many organisations and groups of people. A wide range of derivation processes has been used. These include:

- Historical data
- Historical data modified to reduce future risk
- Consensus among experts
- Using an existing TLS for another system

Some degree of coherence among some TLS is achieved by basing them on a common generic TLS, for instance the JAR 25.1309 TLS of 10^{-9} per flight hour for failure of a single aircraft system. However, in many cases coherence is far from apparent. Even for those TLS based on 10^{-9} per flight hour for failure of a single aircraft system, the degree of coherence may be limited for technical reasons. For instance, the increasing number of safety-related aircraft systems means that in some cases a TLS of less that 10^{-9} per flight hour may be more appropriate in order to avoid a gradual decrease in the level of safety achieved for the aircraft as a whole.

Similarly, increasing inter-dependence between systems means that 10^{-9} per single system may be inappropriate if failure of one system can compromise other systems as well.

How, then, to achieve coherence among TLS? As noted previously, coherence involves subjective and sociological issues as well as technical issues. In addition, it must address both coherence among unrelated sources of risk (e.g. air travel and global warming), among related systems (e.g. aircraft engines and air navigation systems) and among different levels is a risk hierarchy (e.g. aircraft engines and total aircraft risk).

A reasonable starting point is high-level generic risks, such as the risks from different industries (e.g. nuclear power; chemical process;), transport (e.g. air transport; road;), life essentials (e.g. food; water; air;) and external sources (e.g. global warming; earthquake;). There is no a-priori reason why TLS for any of these should be related. Coherence depends on whether people think they should be related. For instance, it might be considered reasonable for TLS for different modes of transport to be related but for air transport and earthquake TLS to be quite unrelated. Thus, before even considering the actual values for TLS it is reasonable to establish the degree of relatedness among them. This necessarily involves strong subjective and sociological elements. As a result, the degree of relatedness cannot be unique. Nevertheless, something like a Risk Relatedness matrix (see Figure 4) or its corresponding network provides a starting point for coherent development of TLS.

	Air	Rail	Road	Nuclear Power	Global Warming	Food
Air						
Rail						
Road						
Nuclear Power						
Global Warming						
Food						

Entries: S Strongly Related; M Moderately Related; U Unrelated

Figure 4 : Risk Relatedness Matrix

‘Related’ does not, of course, imply a deterministic, quantitative link between TLS. However, one would expect there to be a reasonable basis for the pattern of TLS for Strongly Related areas. Similarly, one might expect at least some plausibility for the pattern TLS in Moderately Related areas. Unrelated areas, however, need have no basis for differences between their TLS.

Having determined degrees of relatedness, a coherent chain of quantified high level TLS can be established. Since ‘Related’ does not imply a quantitative link between TLS, how can it be manifest in the quantified TLS? Figure 5 illustrates a plausible scheme. A value for TLS X implies a reasonable range of values for Strongly Related TLS Y. The range might be higher or lower than TLS X, as long as there is a reasonable basis for it. TLS Y can be anywhere in the range. However, once TLS Y is given a value, it sets a range for the Moderately Related TLS Z based on a plausibility argument. The range is larger than that for TLS Y because Y and Z are only Moderately Related rather than Strongly Related. Again, TLS Z can be anywhere within the range and remain part of the coherent set.

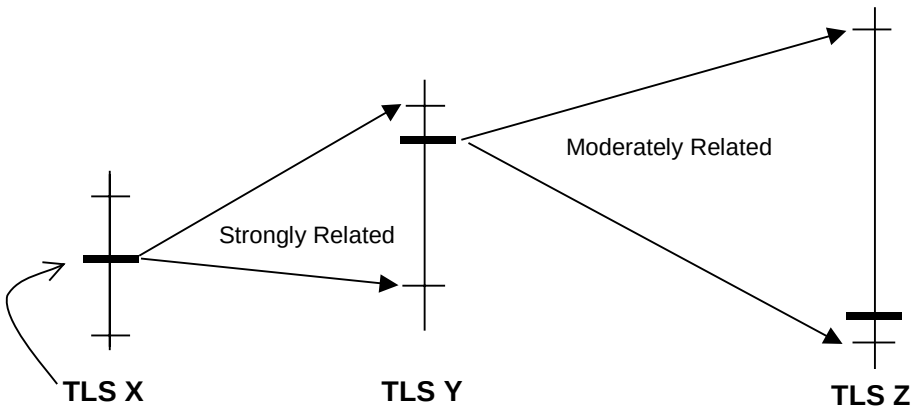

Figure 5 : A Coherent Chain of TLS

Having determined a coherent set of quantitative high-level TLS, coherent sets of lower level TLS can be established. This is essentially a matter of rational partitioning of the high-level TLS and allocation to lower-level entities. Coherence now applies both vertically under a high-level TLS (between different levels), horizontally under a high-level TLS (for instance, between different systems or components) and horizontally between appropriate levels in different areas. This last might appear somewhat surprising. However, it might not be unreasonable to expect, for example, that the level of safety against faults in an aircraft and a car were related.

The overall picture is given in Figure 6. The triangles represent rational partitioning and allocation of high-level TLS to lower levels. Horizontal arrows within a triangle represent relationships between different entities at the same level. Horizontal arrows between triangles represent relationships between appropriate levels in different areas. Of course, the degree of relationship may be 'no relationship' (as one might expect between a TLS for a car and a TLS for sea level change due to global warming).

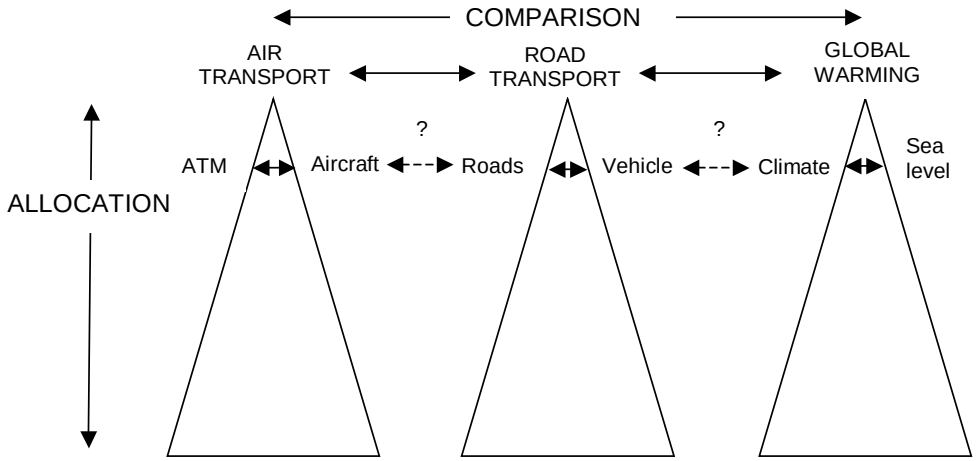

Figure 6 : Vertical and Horizontal Coherence

Again, subjective and sociological factors are present. For instance, local circumstances may well lead to different countries partitioning an overall TLS for global warming into TLS for climate change and TLS for rise in sea level in different ways. Economic factors also come into play as well as technical factors.

Various approaches to partitioning of high-level TLS exist. The ASTER project developed a method to enable safety targets to be set and optimised for each of the actors of the air transport system to achieve the optimum level of safety for the system as a whole. The method allows an assessment of the safety benefit of any change (including changes in legislation and rulemaking) in relation to the costs of those changes. It is based on the use of historical data, representative accidents and a Bayesian network to determining a coherent set of TLS.

However, as noted in the ASTER project, no single process for establishing TLS is best in all situations. The process should be appropriate for the commercial, legal and political situations of the industry concerned and should take into account technical issues and timescales. It should be acceptable to at least the most important

stakeholders; it should be clearly defined and should result in a TLS that is both achievable and not open to significant dispute.

1.3 COMMERCIAL OFF-THE-SHELF PRODUCTS (BOB HUMBERTSON)

It has been said that process accounts for 80 percent of all problems, while people account for the remaining 20 percent. With risk defined as any situation or circumstance that creates uncertainty in achieving an objective, analyzing an organization's processes is essential to identify and mitigate any risks inherent in the processes in themselves. One method is to identify via the Safety Risk Management (SRM) Decision Model the work products that are exchanged between processes, or subprocesses, in order to identify risks to reach your goal. These intensive workflow areas are more susceptible to risk and thus deserve greater oversight and tailoring for efficiency. This section focuses on the commercial off-the-shelf (COTS) acquisition strategy and its relationship to TLS. For purposes of this document, COTS describes a product or service that has been developed for sale, lease, or license to the general public and is currently available at a fair market value.

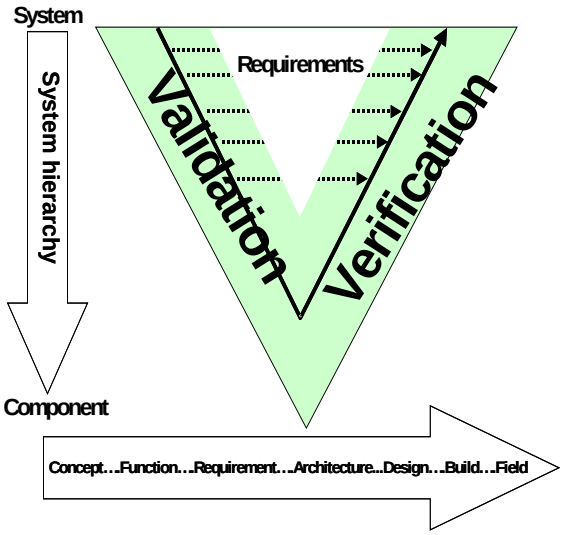

Figure 7: Validation and Verification

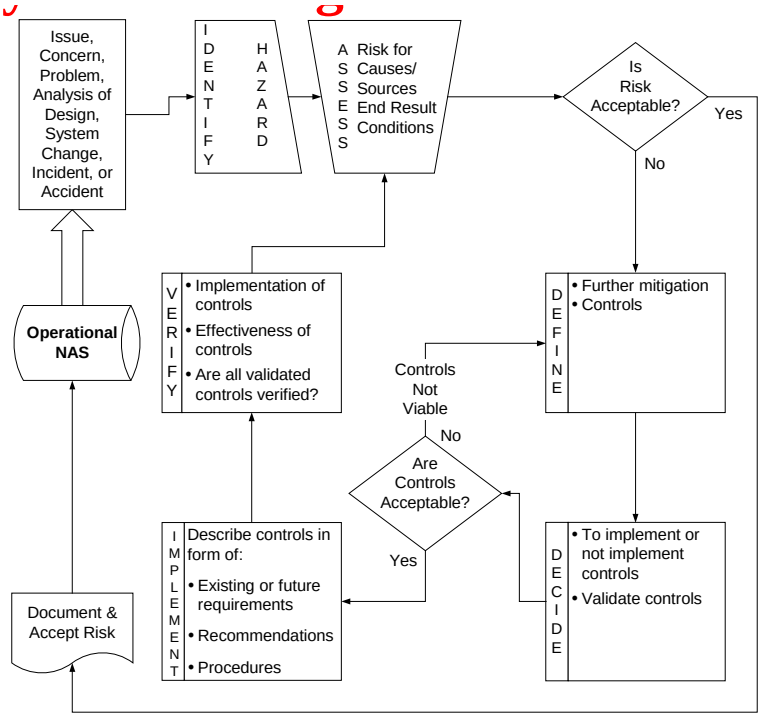

Figure 8: Safety Risk Decision Model

1.3.1 Introduction

Among the many factors that must be considered during programmatic risk management planning is the acquisition strategy. For example, the Federal Aviation Administration's life-cycle acquisition management system "stresses preference for commercial and non-developmental solutions to mission needs." However, acquisition strategies can range from custom-developed systems containing little or no COTS products to COTS-based systems containing mostly or exclusively COTS products. Since custom-development programs have traditionally not developed system components that are readily available on the market (processors, displays, power supplies, disc drives, application software, etc.), most system acquisitions fall into the hybrid systems category. These systems contain varying mixes of COTS, modified COTS, nondevelopmental items, glue ware, middleware, custom interfaces, and so on. Therefore, the COTS risk-mitigation strategies described in this paper apply to most new system acquisition strategies as well as systems already fielded.

1.3.2 Acquisition Strategy and COTS Risks

For any system, safety objectives and requirements must be validated and verified:

- Prior to investment decision;
- Prior to implementation;
- During transition to operational service;
- During system lifetime;
- During transition to deactivation.

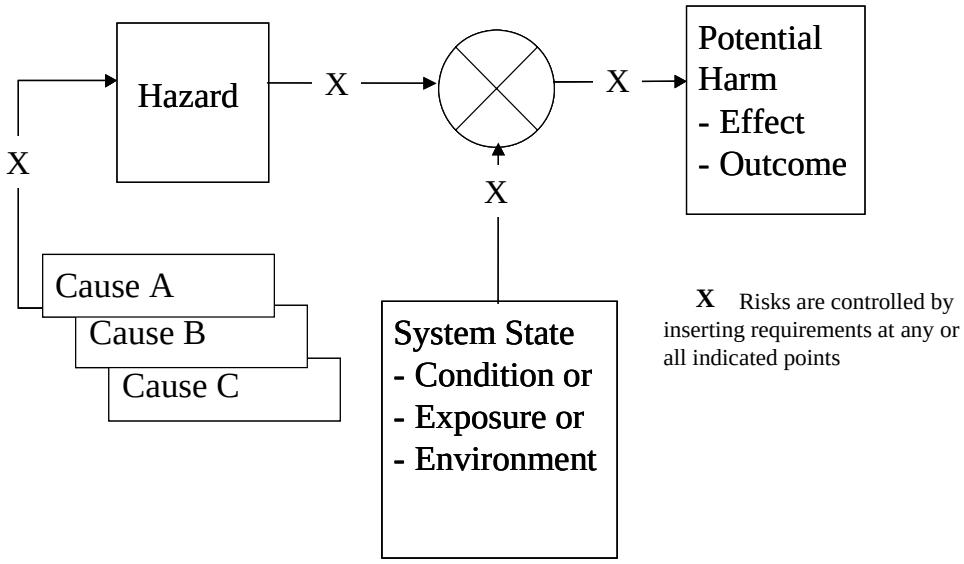

Figure 9: Controlling risk

An acquisition strategy that uses a high percentage of COTS products can provide the following potential benefits:

- Avoidance of the risks typical of the custom system approach;
- Reduction of the time and resources required to acquire and deploy systems and automated tools capable of meeting user and mission needs;
- Reduction of front-end development cost "spikes" in the budget;
- Allowance for a more rapid infusion of current technology and open standard interfaces to support modernization and sustainment; and
- Expansion of product competition across a broader market/vendor base.

The first benefit listed suggests a COTS-based acquisition strategy that seeks to manage risk by reducing or eliminating potentially severe risks and resultant adverse effects typical of custom-developed systems. However, while using COTS products helps to deal with these "custom acquisition" risks, such use introduces other types of risk that stem directly from the unique characteristics of COTS products.

The increased use of COTS products creates a new acquisition operations and support environment that requires development of a standard approach for identifying and managing (i.e., mitigating) the unique risks of COTS products and reaching the desired target level of safety. This paper suggests that any acquiring activity with a standard methodology for acquiring and supporting COTS products and is specifically structured to:

- Identify and understand the risks associated with using COTS products;
- Describe COTS-specific risk-mitigation strategies; and
- Define the relationship of risk-mitigation strategies to system engineering and programmatic risk management processes.

The following paragraphs discuss understanding the risks associated with using COTS products, their mitigation strategies, and how they may affect program management decisionmaking.

1.3.3 COTS Product Risk Factors

The need for risk mitigation of COTS products as depicted below stems from the products' unique risk factors (or characteristics). These risk factors, as identified in Table 1, are based on an extensive analysis of common U.S. Government/industry lessons learned described in numerous technical documents.

Table 1. COTS Product Risk Factors

COTS products can exhibit rapid and asynchronous changes.
COTS product obsolescence can affect systems in different ways.
COTS products are typically documented with proprietary data.
Low initial costs of COTS products can be offset by higher life-cycle costs.
Functionally equivalent COTS products/systems can have multiple configurations.
Different COTS product manufacturers have different quality practices.
A COTS product's form, fit, and function are sold "as is."
COTS products are developed to commercial standards.
COTS products typically have time-limited manufacturer support.
COTS product interoperability can introduce information security susceptibility.

These risk factors, which are at the heart of the behavior of COTS products and the challenges of managing their use, can be distilled into three related market-driven factors:

- In using COTS products, the acquiring activity loses market control. In the custom-development acquisition mode, the acquiring activity defines the market as long as firms are willing to develop and supply products that the activity can afford.
- In using COTS products, purchasers must search the market to find offerings at least closely compatible with their needs. This market-driven situation forces a more flexible and proactive technology evolution planning approach.
- In using COTS products, the acquiring activity must deal with a "high- speed" market. The market, technology, and products change at a substantially higher rate than the historic speed found with use of the custom-development acquisition approach. This increased and accelerating rate of change in the

market forces the acquiring activity to deal with rapid obsolescence (i.e., diminishing levels of product support) for COTS products.

Following is a description of the 10 COTS product risk factors or characteristics.

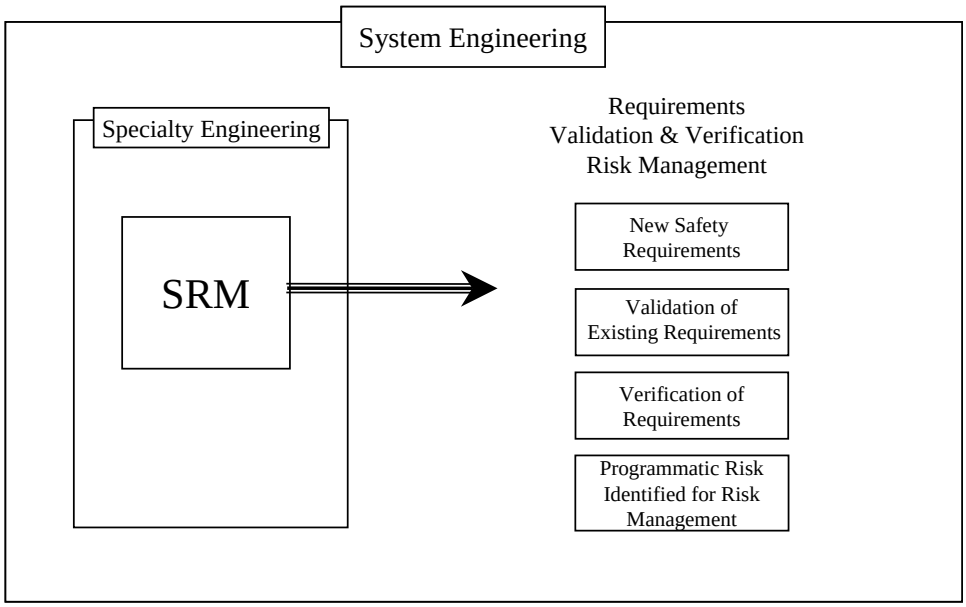

Figure 10: Process Flow to Control Risk

COTS Risk Factor No. 1: COTS products can exhibit rapid and synchronous changes.

Commercial markets are driven by competition for larger profits and, therefore, for expanded market share.

In the information technology sector in particular, the ongoing and rapid sequence of technological advances (such as greater power or speed, miniaturization, capacity, and bandwidth of the underlying components) has both permitted and stimulated a correspondingly fast-paced development and introduction to the market of increasingly more capable COTS products. This competitive environment and rapid advances in underlying technologies both drive and allow COTS product manufacturers to anticipate customer demands and to quickly develop and market their COTS products.

The recent explosion of information technologies has created families of COTS products to satisfy information management needs. The rapid rate of change in technologies and products, a direct consequence of the competition within the commercial market, means that new commercial products are released at a pace based on the speed of market and technology evolution, not necessarily on their continued usefulness to the acquiring activity. As a result, these products become obsolete (i.e., diminishing levels of product support). To compound the challenge of managing rapid change, various types of products (e.g., processors, displays, power supplies, memory tools) have different market cycles. The products tend to be introduced at different times with varied service lives and are therefore out of phase or asynchronous with each other.

COTS Risk Factor No. 2: COTS product obsolescence can affect systems in different ways.

When a COTS product is projected to be nearing end-of-service-life (EOSL) (i.e., out of production or no longer supported by the manufacturer), the effects of these projected changes of state on the product and on systems using the product must be examined to determine what action if any is needed. It is not a foregone conclusion that all products declared to be EOSL must be replaced immediately by newer versions. Effects may range from no impact to high impact.

The obsolescence support options that are available to address these impacts can range from taking no action to making a major system redesign. The categories of impact due to obsolescence are defined as follows:

- **No impact.** This applies when a COTS product is considered reliable and sufficient spares exist to support the projected failure-driven demand over a predetermined timeframe. In this case, the product's

projected status has no impact on the product or on any system using that product. Although ongoing monitoring must continue, no action needs to be taken.

- **Low impact.** This applies when a COTS product status is projected to reach EOSL and the product must be eventually replaced. A low-impact situation exists if the manufacturer's next-generation product is form, fit, and function (F_3)-compatible (i.e., interchangeable); if other manufacturers' products are F_3-compatible; and if no there are associated changes to interfacing products within the system. This situation typically requires compatibility testing for the new product and a documentation change to identify the new product as a suitable replacement when the old product fails.

- **Medium impact.** This category, like the low-impact category, also applies when a COTS product is projected to reach EOSL, and the product must be replaced. A medium- impact situation exists if the manufacturer's next-generation product is only fit/function or form/function (F_2)-compatible; if other manufacturer sources have only F_2 products available; if minor software changes are required; and/or if related changes to interfacing products are required. This situation can be addressed by using several options, including a lifetime buy of that product or product spares; technology refreshment; purchasing the manufacturer's data rights; extending a maintenance or Service Level agreement; or including the change in a major redesign or integrated system change.

- **Major impact.** This category, like both the low- and medium-impact categories, applies when a COTS product is projected to change status, and the product must be replaced. However, a major-impact situation exists when no F_3- or F_2- compatible replacement products or technologies are available on the market. This situation typically calls for a major redesign or an integrated system change. The situation can be addressed by using several options, including a lifetime buy of that product or product spares; technology refreshment; purchasing the manufacturer's data rights; extending a maintenance or Service Level agreement; or including the change in a redesign or next buy.

COTS Risk Factor No. 3: COTS products are typically documented with proprietary data.

A COTS product manufacturer remains in business because it owns and controls the research and manufacturing processes needed to meet market demands. The information a COTS product manufacturer typically labels as proprietary (i.e., not usually for sale) includes software/firmware source code, specific manufacturing processes, detailed specifications, schematics, and drawings. Such information is routinely required and delivered as part of a government custom-development program. However, with COTS products, data are limited to specification sheets and commercial-style operations and maintenance documentation. As a result, the purchaser must view the COTS product as a "black box" with defined interface and performance characteristics, but which usually does not allow purchaser insight into the product's internal composition.

The maintenance concept for systems using COTS products must also change accordingly. Historically, this involves a maintenance concept that adopts a circuit card or replaceable unit swapping procedure.

COTS Risk Factor No. 4: Low initial cost of COTS products can be offset by higher life-cycle costs.

Accelerating the introduction of COTS products into industry, government and military information management systems is usually advertised as touting the products as a "faster, better, cheaper" way of meeting requirements. Using COTS products is usually a "faster" way of meeting a requirement because the products are readily available or at least assembly lines are ready for immediate production. Marketed as capable of meeting a large and diverse demand, COTS products are characterized as being "better" able to meet general information management needs than custom-developed solutions. Since the marketplace forces competition among the product manufacturers, users of COTS products can usually obtain a "cheaper" acquisition alternative to custom development. The fact that the manufacturer has already assumed COTS product development costs lowers the front-end development costs of system acquisition for the acquiring activity.

However, unless predetermined safety goals have been established and include a risk management program that incorporates proactive mitigation strategies specifically oriented toward COTS-unique risks, the initial cost benefit can be offset by the often more costly fixes of the risks that were not effectively managed.

Examples of the cost considerations for a COTS-based acquisition strategy that should be included as part of a total cost of ownership analysis include:

- **Inadequate planning costs.** Probably the major life-cycle cost-driver associated with using COTS products is the lack of effective COTS-specific planning and budgeting. When a program fails to apply

COTS risk-mitigation strategies, the program loses the advantage of proactive planning and becomes increasingly reactive to emerging COTS-driven obsolescence situations. These situations limit management options and force programs to adopt sub-optimal and consequently more costly solutions.

- **Test and integration costs.** Although development costs are reduced compared to development costs for a custom approach, the effort required to successively test, integrate, and deploy multiple COTS products into a system can be substantially greater over a system's operational life than is the case with custom solutions. In addition to the actual costs of the test facilities needed to support the possibility of multiple system configurations, different COTS products with varying characteristics typically require that "glue code" be developed to allow the products to interact effectively. Each product must be tested and the results validated against the safety goals and requirements for safety compliance, to performance requirements, conformance to open system standards, and compatibility with the system with which it will be integrated.

- **Modification costs**. In some cases, a COTS product must be modified to meet a particular or unique requirement to meet or exceed the TLS. To even reach the acceptable TLS, there may also be cost to actually modify the COTS product itself. There is also a cost to assume life-cycle management responsibility for that specific product because modifying a COTS product typically voids any warranty, and the vendor may no longer provide support, unless it is at a greatly increased cost. This forces the acquiring activity to assume life-cycle support costs for that product. Costs for documentation, maintenance, training, and spares costs will increase and must be planned for in the life cycle budgeting for the modified product.

- **Configuration management costs.** A consequence of using rapidly changing COTS products within a given system is the strong likelihood that an acquisition of multiple copies of that system will include more than one configuration. This situation not only demands a rigorous application of configuration management (CM) processes to document and manage system baselines, but also requires that test facilities have the ability to replicate all fielded configuration baselines to ensure performance factors are met. Documenting product and system changes and instituting strong CM processes ensure that the acquiring activity can determine the impact of product changes to all affected configurations.

- **Continuous system engineering costs**. COTS-based systems are dynamic in nature. Therefore, continuous systems engineering activity is needed to perform market surveillance, research, and investigation; analyze obsolescence projections; determine the available options to limit obsolescence impacts; and integrate the resulting information with new requirements and field data as part of the overall integrated program planning.

  A continuous system safety engineering approach also requires a continuous systems engineering test environment in order to perform conformance (to commercial standards), compliance (with specified requirements) and compatibility (F3 interchangeability) testing of new products and technologies. An advantage of using a COTS approach is that the acquiring activity can package smaller, more evenly planned changes using technology refreshment and other change management options. This COTS-driven continuous system safety engineering effort should be planned as an additional cost to a program during its life cycle.

- **Obsolescence management costs.** The continuous system engineering activities needed to manage obsolescence can result in more frequent engineering changes to the system. Development, deployment, and configuration management of these changes can add costs that must be included in all COTS-based system program planning. These costs are initially developed as part of the obsolescence management strategy chosen for a program early in the acquisition planning cycle and are then continually refined as system product obsolescence information is gathered, analyzed, and acted upon.

COTS Risk Factor No. 5: Functionally equivalent COTS products/systems can have multiple configurations.

COTS product manufacturers are subjected to constantly changing market availability of components (i.e., microchips, diodes, resistors, capacitors, etc.) and subassemblies (i.e., disk drive, memory device, display, etc.). For example, one production lot can be functionally equivalent to the next lot but contain different components and subassemblies. If a product contains firmware, or if it is a software product, revisions can be made to subsequent product releases to correct deficiencies or to add unique features to enhance product marketability. A COTS product manufacturer may or may not elect to identify these configuration changes to its customers. Similarly, during the integration of COTS products into a system, the models of COTS products evolve over time.

A manufacturer's claim for new COTS product compatibility requires system testing to verify that claim. When the new product is substituted into the system, the physical configuration has changed, but the functional baseline remains the same. Depending on the number of systems to be fielded and the length of time it takes to manufacture and deploy them, the number of configurations could be significant.

COTS Risk Factor No. 6: Different COTS product manufacturers have different quality practices.

While individual COTS products from different manufacturers might satisfy a particular set of functional requirements, marked differences can exist from one product to the next. Differences in the components manufacturers choose to use, safety requirements, quality assurance practices, manufacturing processes, labor force composition, market share, product support, upward/downward compatibility, corporate longevity, and other factors can all affect the quality and desirability of the products that are offered.

COTS Risk Factor No. 7: A COTS product's form, fit and function is sold "as is."

Until recently, governments drove technology development for military applications with large infusions of research and development (R&D) funding for custom-developed systems. Governments could afford to specify exactly what was desired and, therefore, promoted a "buyers" market of companies interested in meeting this demand. However, the recent military downsizing throughout the world and, more importantly, the rapid increase of consumer demand for information processing technologies have fostered a "sellers" marketplace that is no longer driven by government R&D, but by a much larger commercial customer base. This means that the products made available on the open market are manufactured to meet more general consumer demands, instead of being configured to meet specific purchaser requirements. If a system's requirements are stated in absolute (i.e., inflexible) terms and no COTS products exist that meet those very specific requirements, then the choices will be limited to custom development or COTS product modification.

COTS Risk Factor No. 8: COTS products are developed to commercial standards.

COTS products are typically designed and built to a variety of commercial standards that provide high-level guidance on such product characteristics as performance, quality, and interoperability. Whereas different manufacturers can develop products featuring similar or even identical performance characteristics, these products can be limited in their ability to operate with each other (i.e., their interoperability) due to the use of proprietary interfaces. The rapid obsolescence of COTS products (and their resultant replacement) requires stable interfaces that are designed to be "open" (i.e., to allow flexibility and adaptability) to the use of many products from different sources. Such "open systems" interfaces are provided with COTS products that have been developed using such interface standards as those promoted by the International Organization for Standardization and the Institution of Electrical and Electronics Engineers. These standards allow for product improvements in such areas as quality and functionality while maintaining interface stability through use of consistent interface design standards.

COTS Risk Factor No. 9: COTS products typically have time-limited manufacturer support.

As succeeding generations of COTS products are introduced into the commercial market, the manufacturer must determine at what point it is no longer profitable or desirable to support the older-generation products. The manufacturer must make a tradeoff between selling its newer product line while at the same time not alienating the older-generation product consumer base. Manufacturers for both hardware and software COTS products strive for upward/downward product compatibility and typically support two to three previous generations of products before declaring EOSL.

Successive generations of COTS products are rapidly introduced on the market to meet (and/or perhaps to stimulate) consumer demand. Therefore, it is not in the best interests of the manufacturer to stockpile or warehouse large quantities of an existing product or repair parts that may be superseded by a next-generation product the manufacturer wishes to sell. To avoid both costly warehousing expenses and unmarketable inventory, the manufacturer minimizes its product stock to meet current consumer demand and limits

the support period for that product by using a "just-in-time" parts-ordering strategy.

COTS Risk Factor No. 10: COTS product interoperability can introduce information security susceptibility.

When governments develop their own custom systems, there can specify and develop system information security characteristics very precisely. But the use of COTS products developed to commercial standards may introduce potentially significant information security risks. First, the increased interoperability among different

products that meet commercial standards raises the chances that unauthorized access can be gained. Second, using commercial standards allows a greater number of people to be familiar with the software protocols used to manage information. This knowledge can be used to access or disrupt information flow. The "openness" of a particular architecture, the degree to which it links with other external COTS-based systems, and the nature of the security measures in place will determine the extent to which the products and systems using them are susceptible to unauthorized access.

Identifying and understanding COTS risks is the first step to ensure that the acquiring activity can achieve the benefits of using COTS products. The next step is to manage the risks by implementing proactive risk-mitigation strategies.

1.3.4 COTS Risk-Mitigation Strategies

To effectively address COTS risks, the acquiring activity must implement a set of interrelated risk-mitigation strategies that reflect government and/or industry-recognized lessons learned. Risk-mitigation strategies must be implemented early in the acquisition process and continue through the products life cycle. The strategies are:

- Forecast product obsolescence changes and select product support options for obsolescence;
- Determine if the level of documentation offered by a COTS product manufacturer is sufficient to meet product integration and testing, maintenance, training, and/or reprocurement needs;
- Reduce life-cycle costs by proactively planning for obsolescence situations before support options become limited and more costly;
- Recognize market change cycles for various COTS product types when

  formulating technology evolution strategies and budgeting profiles;
- Identify changed or substitute alternatives for compatibility testing;
- Assess a COTS product manufacturer's capability to satisfy the

  functional/performance requirements as well as to develop and apply the nontechnical COTS selection criteria;
- Select the most appropriate combination of commercial standards for inclusion in the functional/performance requirements documents;
- Assess impacts to the system resulting from projected COTS product EOSL situations; and
- Determine the support and information security characteristics of potential substitute products.

Implementing a risk-mitigation strategy to reach your desired goals begins with identifying the COTS knowledgeable resources available to the acquiring activity. A COTS-knowledgeable individual understands the interrelationships among commercial market forces, market research, technology trends, commercial standards, COTS product risks, and risk-mitigation strategies. Individuals become "COTS-knowledgeable" through:

- General experience gained with other COTS-based acquisitions;
- Specific COTS experience gained within their subject-matter expertise; or
- Training specifically focused on understanding and managing COTS risks.

Depending on the phase of the acquisition and risk-mitigation activities that need to be implemented as part of the total system engineering process, the acquiring activity can solicit COTS-knowledgeable individuals from:

- Within the acquiring organization;
- Within the parent organization;
- Program support contractors;
- Consultants;

- Consortiums;
- Centers of excellence; and
- Systems integration agents.

The organization performing the analyses should ensure that COTS-knowledgeable individuals from the applicable systems engineering disciplines are included in all phases of the system's life cycle. This activity should be reflected in the program's Work Breakdown Structure to ensure that the necessary resources are identified. This ensures a cohesive COTS-oriented approach to the analysis that maximizes the benefits for controlled levels of the risk of using COTS.

1.3.5 COTS Obsolescence Planning

Managing COTS product obsolescence entails initial use of a system-level strategy and subsequent use of product-level support options. A system-level strategy for obsolescence management must be formulated early in a COTS-based system's acquisition cycle. It provides a life-cycle system evolution path that integrates such activities as preplanned product improvements and new requirements changes with projected obsolescence-induced system upgrades. The strategy also provides the basis for system budget projections and risk management. The system-level strategy must be reviewed periodically and adjusted as needed.

As the system architecture is defined and the COTS product composition becomes known, the system-level assumptions and resultant planning can be refined to reflect EOSL data gathered through market research activities. When a product's manufacturer projects EOSL dates, one can determine which mix of several product support options should be implemented to support the overall system evolution/obsolescence management strategy. Implementing a mix of options over time provides management cost, schedule, and risk flexibility to address market-driven COTS variances.

1.3.6 COTS Obsolescence Management Strategies

At the system level, a broad range of strategies exists for managing COTS obsolescence in order to mitigate risks. Possible management strategies include:

- Continuously refreshing all COTS products to maintain currency of

  manufacturer support;
- Freezing the hardware/software baseline during development and then using product obsolescence support options to sustain the system over a defined period; and
- Freezing the hardware/software baseline for a defined period and then refreshing as required.

Each obsolescence management strategy exercises a different level of control over market-driven product obsolescence and consequently invokes a different level of program risk. For example, adopting a continuous refresh strategy is beneficial from a product support standpoint, but it is high risk due to the loss of programmatic control over the impact and cost of frequent successive vendor-driven product refreshes to system evolution.

The SRM develops the set of requirements to attain the specified level of safety in the product. These requirements are the controls that ensure that the target level of safety is achieved. Through Validation and Verification processes, the system requirements are confirmed as correct and satisfied.