Designing an Effective Safety Management System (SMS)

Prof. Nancy Leveson

MIT

I have been working in system safety for the past 40 years and have seen a lot of successful and spectacularly unsuccessful attempts at managing safety. This paper documents what I have learned. Unnecessary loss of life is tragic, but not as tragic as not learning from it. This paper describes how systems thinking can be used to create a new SMS or to improve an existing one.

What is an SMS?

The goal of an SMS should be to proactively control safety in every aspect of the organization. For organizations that produce products, the SMS manages safety in the product development process and in the product itself. For service industries, safety must be designed into and managed in the workplace and the services provided. Because sociotechnical systems are rarely perfect from the beginning and the world changes over time, there must be an effective learning process in place that is continually learning from experience and improving the workplace and the products and services as well as the management system itself. On the other hand, very successful safety management systems can and have degraded over time. There needs to be a way to identify when the SMS is becoming less effective before tragedy strikes.

There is no one design of an effective safety management system. The goal of the design will depend on such factors as the type of organization (product development or service providers), the societal culture in which the organization resides and individual organizational structure and culture, the inherent safety of the product or service, environmental factors, the other goals the organization wants to achieve and the regulatory environment. Some groups, such as the FAA and ICAO, have specified one approach to risk management, but that may not be an appropriate solution for everyone, and risk management is only a subset of what is required in a safety management system. In addition, particular safety management system rules imply underlying cultural values that may not be the best ones for a particular group or organization in achieving high safety levels.

There are, however, some general characteristics of a successful SMS. The systems thinking approach to designing an SMS, described in this paper, is the result of forty years of the author's experience in system safety engineering, in investigating accidents and identifying the factors involved, as well as helping organizations to prevent accidents. Some industries and organizations have lots of accidents, others have few, and some do not have accidents at all (e.g., SUBSAFE, the U.S. Navy nuclear submarine safety program[1]). Certain factors stand out as critical with respect to which of these three categories an organization or even an industry falls. Some of the most effective ways to control safety may not be practical for everyone. Designing an SMS is a system engineering problem: The goal should be to design and operate an SMS that is as effective as is feasible to implement in your organization.

There are several components of the organization that are most important in managing safety: the culture, the safety management (control) structure, and the safety information system. These three components cannot be considered in isolation, though. The systems approach suggests that these must all be part of a coherent and consistent whole in order to be most effective. Culture defines what desirable and acceptable behavior is and how decisions should be made. The management or control

[1] SUBSAFE focuses on two safety goals: (1) watertight integrity of the submarine's hull and (2) operability and integrity of critical systems to control and recover from a flooding emergency.

structure determines how that culture will be implemented in the organization. Finally, the safety information system provides the information necessary to make the management structure successful in achieving the desired safety culture: even the best of intentions will not suffice without the appropriate information to carry them out.

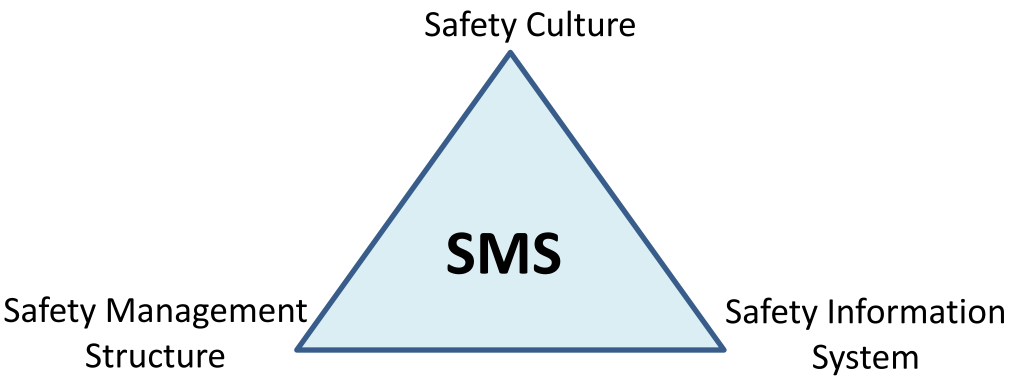

Safety Culture

People define safety culture in different ways. I like Edgar Shein's definition of safety culture:

> *Safety culture is the values and assumptions under which safety-related decisions are made*.

Figure 1 shows Shein's three levels of organizational culture. The bottom level is the essence of the culture. The values and deep cultural assumptions form the basis for creating organizational rules, policies, and practices. These documents, in turn, describe how the surface-level cultural artifacts (e.g., hazard analyses, accident investigation reports, etc.) are produced.

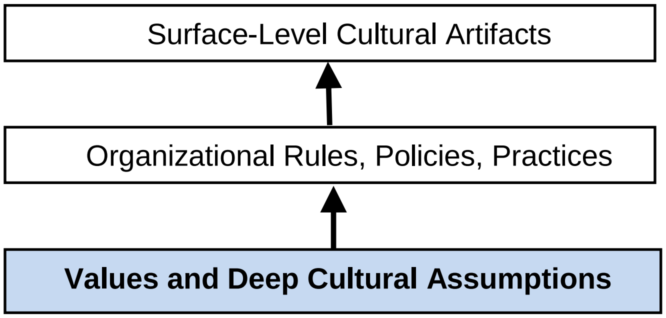
Surface-Level Cultural Artifacts
Organizational Rules, Policies, Practices
Values and Deep Cultural Assumptions

Figure 1: Shein's Three Levels of Organizational Culture

The middle and top levels are not the culture; they merely *reflect* the organizational culture. While attempting to make changes only at the top two levels may temporarily change behavior and even lower risk over the short term, superficial fixes at these levels that do not address the set of shared values and social norms on the bottom level are likely to be undone and become ineffective over time. At the same time, trying to change culture without changing the environment in which it is embedded is also doomed to failure.

While sometimes people talk about "creating" a safety culture, there always exists some sort of safety culture, although the existing one may not be very useful in promoting safety. There is no such thing as a cultureless organization, industry, or society.

Often historical and environmental factors are key in creating the existing safety culture. For example, William Boeing, when building the commercial aviation industry, was faced with the fact that improved safety would be required to sell aircraft and air travel. In 1955, only 20% of the U.S. citizens were willing to fly—aircraft crashes were a relatively common occurrence compared with today. In contrast, the nuclear power industry was initiated by passage of the Price-Anderson Act in 1957 whereby they agreed to be tightly regulated by government in exchange for limited liability in case of accidents. In some industries, the safety culture was the result of strong individual leadership such as Admiral Hyman Rickover in the nuclear navy. Differences in the way safety is handled exist in each of these industries as a result.

While it is difficult to define a "good" safety culture vs. a "poor" one, there are some industries and organizations that have more accidents and major losses than others. The safety culture in organizations having relatively high accident rates tend to have one or more of the following characteristics:

- Culture of Risk Acceptance: This culture is based on the assumptions that accidents are inevitable and that nothing much can be done to prevent them beyond exhorting everyone to be careful. Accidents are considered to be the price of productivity. Often this assumption is accompanied by the belief that everyone should be responsible for safety—their own and others—and that accidents result from a lack of responsible behavior on the part of individuals. The belief is prevalent that if only everyone would act responsibly and safely, accidents would be reduced.
- Culture of Denial: In a culture of denial, risk assessment is often unrealistically low, with credible risks and warnings being dismissed without appropriate investigation: Management only wants to hear good news so that is what they are told. The focus is on showing the system is acceptably safe, not on identifying the ways it might be unsafe.
- Culture of Compliance: Focus here is on complying with government regulations. The underlying cultural belief is that complying with regulations will lead to acceptable results. Because regulatory agencies tend to focus on certifying that a product or service is safe, after the fact assurance is emphasized and often extensive "safety case" arguments are produced with little or no impact on the actual product or process safety.
- Paperwork Culture: This culture rests on the belief that producing lots of documentation and analysis paperwork leads to safe products and services. Piles of paper analyses are produced but they have little real impact on design and operations. The bulk of the safety-related paperwork may be produced by a group that is independent from and have little interaction with those who are designing and operating the products, implementing the processes, or providing the services.
- Culture of "Swagger": Safety is for sissies. Real men thrive on risk.

The features of a good safety culture are harder to specify, but they often include such things as: openness about safety and the safety goals, a willingness to hear bad news, an emphasis on doing what is necessary to increase safety rather than just complying with government regulations or producing a lot of paperwork, and believing that safety is important in achieving the organization's goals. In an effective safety culture, employees believe that managers can be trusted to hear their concerns about safety and will take appropriate action, managers believe employees are worth listening to and are worthy of respect, and employees feel safe about reporting their concerns and feel their voice is valued. Safety is considered to be a shared responsibility where employees are part of the solution and not just part of the problem. At the same time, responsibility is not simply placed on the workforce to keep themselves and others safe.

How is the organizational safety culture established or changed? The safety culture (the values to be used in decision making) in any organization is set by the top management. A sincere commitment by management to safety is often cited as the most important factor in achieving it. Employees need to feel that they will be supported if they exhibit a reasonable concern for safety in their work and if they put safety ahead of other goals such as schedule and cost.

A manager's open concern for safety in everyday dealings with personnel can have a major impact on the reception given to safety issues. A classic story is about Paul O'Neill when he was hired from the outside as the CEO of Alcoa in 1989. He immediately announced that his primary goal was to make Alcoa the safest company in America and to go for zero injuries. Typical reactions to this announced goal were that the Alcoa board of directors had put a "crazy hippie in charge" and that he was going to destroy the company. In fact, within a year, Alcoa's profits had hit a record high and continued that way until O'Neill retired in the year 2000. At the same time, all that growth occurred while Alcoa became one of the safest companies in the world. O'Neill clearly understood that safety and productivity are not conflicting and, in fact, support each other.

In designing or engineering any system, goal setting and establishing the requirements for achieving those goals are the first step in any successful effort: If you do not know where you are going, you are unlikely to get there. Goals are necessary to guide any successful design process. Management establishes the value system under which decisions are made in an organization. Therefore, the first step in designing or improving the safety culture is for management to establish and communicate what is expected in the way of safety-related decision making and behavior, that is, the bottom rung of Shein's Safety Culture model.

Once the desired values have been identified by top management, the next step is to communicate the basic values of the leaders through a short written safety philosophy and more detailed safety policy to implement the general philosophy. Top management has responsibility to ensure that the written philosophy and policy has wide buy-in from the managers and employees.

Improving the safety culture, and thus safety, starts by the leaders communicating the type of safety culture they want to establish. Most companies have extensive safety policy documents that detail exactly how safety should be handled. While these are important, a shorter statement of the philosophical principles and values is a good way to communicate the expected cultural principles of the organization. This philosophical statement should define the relationship of safety to other organizational goals.

Some general cultural principles are applicable to all organizations, while the applicability of others will depend on whether the organization is developing safety-critical systems, operating them, or providing safety-related services. Some general principles that might be part of the Safety Philosophy statement:

1. All injuries and accidents are preventable.
2. Increasing quality and safety leads to decreasing cost and schedule and, in the long term, to increased profits. Preventing accidents is good business.
3. Safety and productivity go hand in hand. Improving safety management leads to improving other quality and performance factors. Maximum business performance requires safety.
4. Safety has to be built into a product or the design of a service. Adding it later will be less effective and more expensive. After-the-fact assurance cannot guarantee a safe design where safety is not already present. It is better to build safety in than try to ensure it after the fact.
5. The goal of accident/incident causality analysis is to determine why the loss (or near loss) occurred so that appropriate changes can be made rather than to find someone or something to blame.

6. Incidents and accidents are an important window into systems that are not operating safely and should trigger comprehensive causal analysis and improvement actions.
7. Safety information must be surfaced without fear. Safety analysis will be conducted without blame.
8. Safety commitment, openness and honesty are valued and rewarded in the organization
9. Effective communication and the sharing of information are essential to preventing losses.
10. Each employee will be evaluated on his or her performance and contribution to our safety efforts.

Number 4 requires some extra explanation as it violates how many organizations and even industries treat safety today. It is common to find an emphasis on measuring or assessing safety or risk after the fact rather than building safety into products and workplaces from the beginning. Designs are either safe or they are not, and all the arguing and assessing one can do will not change that fact. In addition, assessing or verifying after the fact is not only more expensive than designing safety in originally, it is also not very effective because the possible improvements are usually very limited after the design has been completed.

After-the-fact assessment can only provide confidence that safety already exists or evidence that it does not. Emphasis on after-the-fact assessment, therefore, leads to fudging numbers, often unconsciously, such as emphasizing only positive factors or those factors that will produce positive results (called *confirmation bias*), considering only factors that can be measured and can provide a number while ignoring other factors (such as management and design flaws), making up numbers or manipulating them so that numerical goals are achieved, etc. That does not mean that measurement and assessment are not important, just that at best it can only assure that an effort to create a safe system has been successful or not but it cannot make the system safe if it is not safe already.

An alternative to after-the-fact assessment is to go into the assessment process with a mindset to provide evidence that the design is not safe, rather than providing an argument that it is safe. This approach is more likely to be more effective. However, by the time this assessment occurs, it is usually too expensive to fix any problems that might be found or the fixes will be more costly and less effective than building safety into the design from the beginning. Starting the safety analysis process at the beginning of and during the design process will provide the best results. One new, very effective way to do this is to use a modern hazard analysis technique like STPA and integrate it into the system engineering process.[2] That is, instead of proving safety, the emphasis should be on identifying and eliminating or controlling hazards. Those who emphasize after-the-fact assessment of safety are unlikely to achieve high safety.

In addition, an emphasis on probabilities or likelihood in assessment can lead to overconfidence and unrealistic assessment and to the omission of important factors that are not stochastic or for which probabilities cannot be obtained. In the very successful SUBSAFE program, all evidence used in certifying a system for safety must be Objective Quality Evidence (OQE). OQE is defined as "any statement of fact, quantitative or qualitative, pertaining to the quality of a product or service based on observations, measurements, or tests that <u>can be verified</u>." Probabilistic risk assessment (or even qualitative risk assessment that includes identifying the likelihood of a hazard leading to an accident) is an attempt to predict the future and cannot be verified beyond waiting many years to see if the predictions were valid. Few effective crystal balls exist. Assessing likelihood is only effective when the past is identical to the future so past experience can be applied or it is close enough so that extrapolation from the past can be done. But new products are created to improve on past products in some way, not to duplicate them.

[2] Nancy Leveson, *Engineering A Safer World*, MIT Press, 2012.

And rare is the system (and its environment) that does not change over time. Even if the product itself does not change or degrade, people operating or using the system start to change their behavior, their uses of the system evolve, and the environment changes.

A written statement of the safety philosophy and more detailed policy statements are a start, but they are not enough. Employees quickly identify when the written policy differs from the actual behavior of managers. To be successful, there needs to be real commitment by those at the top, not just sloganeering and going through the motions.

How is commitment demonstrated? It is shown by setting priorities and following through on them; by personal involvement (e.g., top management chairing groups where safety decisions are made); by setting up appropriate organizational structures; by appointing designated, high-ranking leaders to safety-related responsibilities and providing adequate resources for them to be effective; by assigning the best employees to safety-related activities and rewarding them for their efforts; and by responding to initiatives by others. It is also communicated by minimizing blame. Leaders need to demonstrate that their highest priority is to fix the factors leading to losses and not just find someone on which to place blame (usually someone at the lowest level in the organization as possible) and then moving on. Finally, leaders need to engineer the incentive structure in the organization to encourage the behavior desired.

The major ideas in this section are summarized in the following box:

Tips for how management can improve safety culture

- Set the goals and values to be used in decision making; establish and communicate what is expected in safety-related decision making and behavior.
- Support employees who exhibit reasonable concern for safety in their work
- Create a short written safety philosophy and more detailed safety policy
- Ensure safety philosophy and policy have wide buy-in from managers and employees
- Follow the safety philosophy in your decision making and expect the same from everyone
- Emphasize building safety in and not assurance or after-the-fact assessment
- Perform assessment with the goal of providing evidence that the design is not safe, not providing an argument that it is safe
- Require objective quality evidence for certification or assurance
- Demonstrate commitment to safety by
    - Personal involvement
    - Setting priorities and following through on them
    - Setting up appropriate organizational structures
    - Appointing high-ranking, respected leaders to safety-related roles and responsibilities
    - Providing adequate resources for safety-efforts to be effective
    - Assigning the best employees to safety-related activities, not just those who are nonessential or expendable
    - Rewarding employees for safety efforts
    - Responding to initiatives by others
- Minimize blame; focus on “why” not “who”
- Engineer the incentive structure to encourage desirable safety-related behavior

Safety Management Structure

Designing the safety management structure starts with a statement of the purpose of the effort. The SMS is:

What: A control structure that assists in creating and maintaining safety in an organization

Why: To ensure that hazards are eliminated or, if not possible, controlled (mitigated) and to promote an effective safety culture

How: By establishing management controls and responsibilities to manage hazards and a comprehensive and usable safety information system.

The overall goal is to design a control structure that eliminates or reduces losses. Satisfying this goal requires a clear definition of expectations, responsibilities, authority, and accountability for safety tasks at all levels of the control structure. In addition, to operate effectively the structure requires appropriate feedback and coordination between entities. It should also involve leading indicators to signal when the controls are becoming ineffective because of internal or external changes. Together the entire control structure must enforce the safety constraints on system behavior through physical design, defined processes and procedures, and social interactions and culture.

There will probably be significant differences between a safety management structure for development and one for operations. If you are a regulated industry, government regulatory agencies need to be included in the overall safety management structure. Other external groups may be important to include such as the courts, insurance companies, user groups, etc.

General Safety Management Structure Design Considerations

Some basic considerations when designing or evaluating a safety management structure is how to assign responsibility for safety, the appropriate place(s) in the organization for safety-related activities, communication of information and coordination of activities, managing and controlling for change, designing and encouraging feedback, and determining the design and role of the risk management procedures to be used. In addition, there must be consideration of education and training efforts that will be required and how learning and continual improvement of the management structure itself will be assured. A "checklist" is included in the Appendix titled *Guidelines for Designing and Evaluating an SMS*.

Assigning Responsibility:

Leadership is the key to achieving high levels of safety. This means that leadership of the organizational safety functions should not just be a rotational assignment for training future leaders and managers. Organizations that have few losses appoint leaders of the safety functions who are passionate about safety and the role they play in preventing losses. There should be a career path within the organization that allows those who are committed to preventing losses to rise in the safety management organization.

As with any effective management system, there must be responsibility, authority, and accountability assigned. A belief that "everyone is responsible for their own and others safety as well as the safety of the products" leads to excessive accidents. If everyone is responsible for safety, then nobody is.

Responsibility for safety lies at every level of the organizational structure although appropriate responsibilities will differ at all of them. In some, usually high-accident organizations, a belief is prevalent that responsibility for safety should be pushed down in the control structure. The argument is used that the lower levels have more information about how their parts of the system actually work. The problem with this argument is that the lower levels also lack perspective: Although they have detailed information about their specific parts of the system, they do not have information about other parts of the system and therefore cannot anticipate or prevent unsafe interactions among all the system components. Lower levels also tend to have a short-term focus rather than a longer-term one, while the higher organizational levels have a broader focus on system goals as opposed to the goals of their own subcomponent of the system. The higher organizational levels need to control interactions among the lower level components and ensure that safety constraints are being enforced by those beneath them and that tradeoffs with other goals are being made in a manner consistent with overall organizational goals.

As an example, Figure 2 shows some of the complex interactions involved in a ferry accident that led to 193 passengers and crew being killed in Zeebrugge, Belgium. The details are not important here, but decisions were made by each of the six groups at the bottom of the structure that were individually safe but together led to the tragedy.

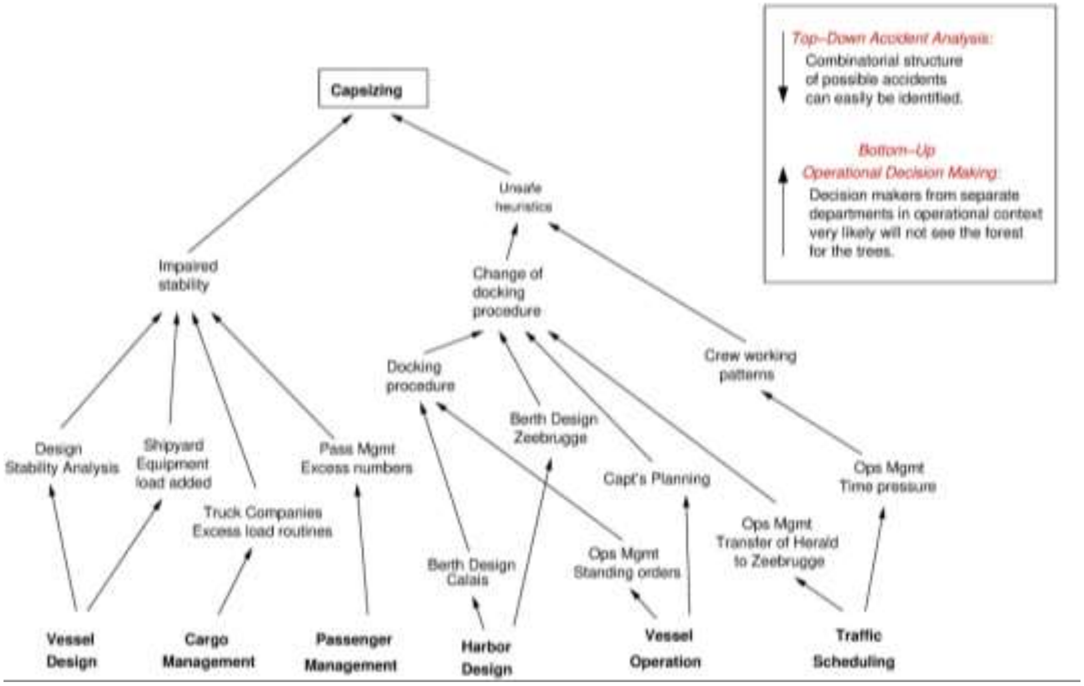
Top-Down Accident Analysis:
Combinatorial structure of possible accidents can easily be identified.
Bottom-Up Operational Decision Making:
Decision makers from separate departments in operational context very likely will not see the forest for the trees.
Capsizing
Unsafe heuristics
Impaired stability
Change of docking procedure
Docking procedure
Crew working patterns
Berth Design Zeebrugge
Design Stability Analysis
Shipyard Equipment load added
Pass Mgmt Excess numbers
Capt's Planning
Ops Mgmt Time pressure
Truck Companies Excess load routines
Ops Mgmt Transfer of Herald to Zeebrugge
Berth Design Calais
Ops Mgmt Standing orders
Vessel Design
Cargo Management
Passenger Management
Harbor Design
Vessel Operation
Traffic Scheduling

Figure 2: The Complex Interactions in the Zeebrugge Accident (adapted from Jens Rasmussen, Risk Management in a Dynamic Society: A Modelling Problem. Safety Science 27(23):183-213, 1997)

Each level must provide oversight of the level below by ensuring they are using appropriate procedures, carrying them out correctly, and that they are effective. There also usually needs to be a management focal point (or points) with responsibility for ensuring that the safety management system is designed and working properly at each level of the management system.

The box below summarizes some of the important factors in assigning responsibility for safety.

Tips for assigning responsibility

- Appoint leaders who are passionate about safety
- Create career paths for those committed to preventing losses
- Provide a clear definition of expectations, responsibilities, authority and accountability at all levels of the safety management structure
- Do not make everyone responsible for safety
- Assign safety responsibilities throughout and at all levels of the control structure
- Assign someone to be responsible for ensuring the SMS is designed and working properly

<u>Place in the organization</u>:

Some basic principles need to be considered in deciding where to place system safety activities in the organization and how to design and manage feedback and control change.

First, there needs to be a high-level group with enterprise-level responsibilities such as ensuring that required activities are taking place and that they are effective. This group also provides leadership and coordination. Although safety activities will permeate every part of the development and operation of large organizations, a common methodology and approach will strengthen the individual activities. The leader of this group needs to have the ability to communicate directly with top management and provide input to all types of management decision making. This implies that the person in this position must report to someone with influence and be seen as having the support of senior management.

Below this top-level safety management, there will be activities at all levels of the organization, with the higher levels having broader responsibilities and each successive lower level having more focused responsibilities appropriate to the level at which they operate, for example, at the business unit, program and project level of a development organization. A common mistake is to make safety a staff function that has little impact on line operations.

There is no one or even several correct designs for a safety management structure. An effective design will depend on the industry, organizational management style, etc. Some responsibilities that should be included in the safety management structure for development and operations are the Appendix.

Second, and very important, separating system safety engineering from system engineering is a mistake in product development organizations; it can lead to higher costs in achieving safety goals because safety concerns must be addressed starting early in the development process. Integrating safety analysis and decision making into system engineering decision making is not only critical for ensuring safety in the final system but also ensuring the cost of developing an acceptably safe system is minimized.

Safety engineering efforts should not be focused on after-the-fact assurance. Service organizations have different concerns, but there should be some safety engineering component within all departments or groups where decision making about safety takes place. At the same time, there are reasons for also having independent oversight and decision-making authority outside the primary decision-making group. For all types of organizations, safety needs to be designed into the workplace and decisions need to be made about the appropriate place for this responsibility in the organization.

Some general principles for allocating safety responsibilities in the organizational structure are:

- Decision makers need direct links to those who can provide safety information. If critical information has to float up a chain of command, it may be lost or modified either deliberately, usually because of schedule or budget pressures, or inadvertently. Direct communication channels provide more chance that information is received in a timely manner and without being changed by being filtered through groups with potentially conflicting interests. Some decision makers may also need fast access to information.
- Direct communication channels to most parts of the organization are required. Safety may be involved in almost all organizational activities.
- Safety must have influence on decision making, which means that decision makers have access to necessary safety information at the time safety-related decisions need to be made.

The SUBSAFE program for safety in U.S. nuclear submarines uses a unique design that satisfies many of these requirements. They describe their structure as a separation of powers or a “three-legged stool” (Figure 3). Managers can select only from a set of acceptable options (with respect to safety) derived by the Independent Technical Authority (ITA). Technical Authority is defined as a process that establishes and assures adherence to technical standards and policy. The ITA provides a range of technically acceptable alternatives with risk[3] and value assessments. Responsibilities (and accountability) include:

- Setting and enforcing technical standards
- Maintaining subject matter expertise
- Assuring safe and reliable operations
- Ensuring effective and efficient systems engineering
- Making unbiased independent technical decisions
- Providing stewardship of technical and engineering capabilities.

The third leg of the stool is the compliance verification organization. It is equal in authority to the program managers and the ITA.

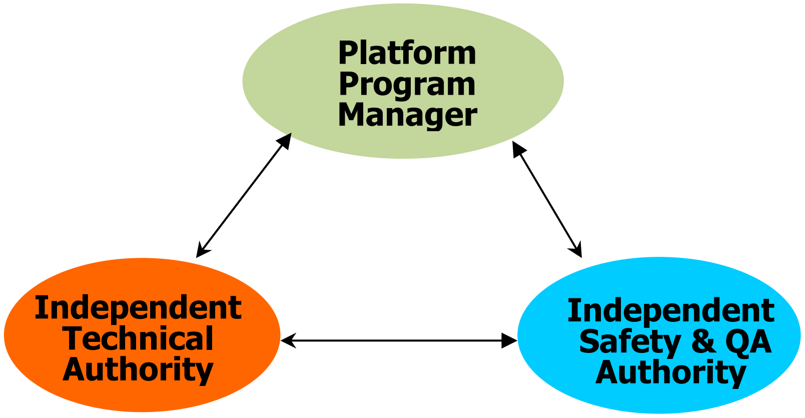
Platform Program Manager
Independent Technical Authority
Independent Safety & QA Authority

Figure 3: The SUBSAFE 3-Legged Stool Concept

[3] Risk may be specified quantitatively or qualitatively but, as explained elsewhere in this chapter, in SUBSAFE it must use objective quality evidence and therefore probabilistic risk assessment or likelihood estimates that cannot be tested or verified are not allowed.

Tips for where to locate activities in the control structure

- Create a high-level group with enterprise level responsibilities, such as ensuring required activities are taking place and they are effective and providing leadership and coordination.
  - This group ensures a common methodology and approach is being used by everyone
  - They must have a direct path to top management and provide input to management decision making
  - They must report to someone with influence and be seen as having the support of senior management
- Assign responsibilities so safety is not just a staff function but has direct impact and influence on line operations
- Do not separate system safety and system engineering in product development organizations. Integrate the safety function into the engineering function.
- Include safety engineering components in all departments or groups where safety-related decision making occurs.
- Provide independent oversight in addition to integrated functions
- Ensure decision makers have direct links to those who can provide safety analysis and information.
- Create direct communication channels between any parts of the organization where there are safety-related activities and safety decisions need to be made.
- Consider implementing the SUBSAFE separation of powers structure. At the least, ensure that technical decisions are independent of programmatic ones.

Communication and Coordination:

Feedback and dissemination of information to those who need it to perform their safety management roles is an important factor to consider in the design of your safety management structure. It is not just a matter of collecting lots (sometimes enormous amounts) of information. In fact, collecting too much information may degrade the information system when it overwhelms the ability to identify the information each person actually needs to make effective safety decisions. In addition, required information for improved safety-related decision making must be handy, that is, available when needed. Because it is so important in reducing losses, the safety information system is discussed separately later in this paper.

Communication is also important in coordination of activities and responses to events. People with overlapping responsibilities need communication channels and ways to coordinate their activities to ensure that safety constraints are enforced. For example, safety-motivated changes in one subsystem may affect another subsystem and the system as a whole. Safety activities must not end up fragmented and uncoordinated. Interactions must be defined not just between hierarchical components but also between different parts or types of systems at the same level, such as between development and manufacturing or between development and operations. Figure 4 shows an example of the types of safety-related information that needs to be communicated between system developers and system operators.

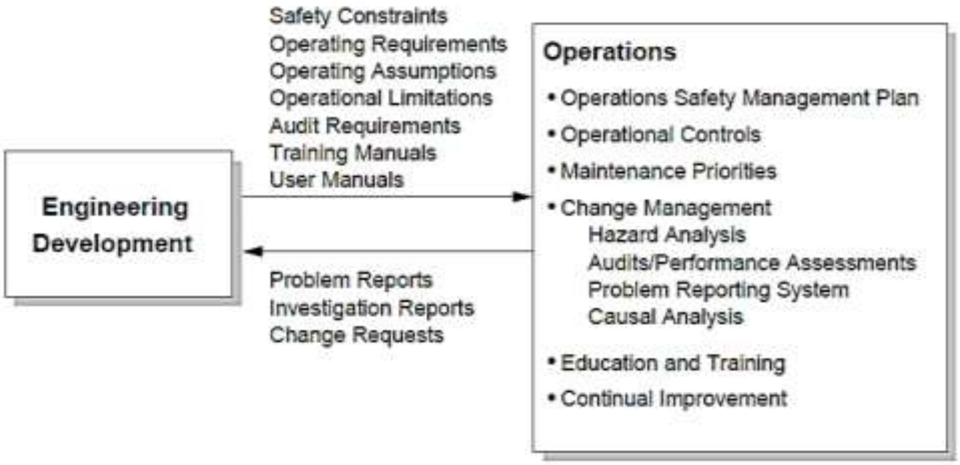

Figure 4. Necessary information channels between development and operations to support safety activities in both.

Similar descriptions of the information needing to be communicated among all the components of the safety management structure need to be identified and documented.

One very effective means for communication and coordination devised by the defense department is *working groups*. Defense department projects can span many years or even decades, be extremely large and complex, and often involve a large number of participants who are geographically and organizationally distributed. Coordination and communication can become a major problem in such projects. Part of the solution is provided by a hierarchical structure where the high-level management of the project may lie in the Defense Department, but there is a Prime Contractor that provides the system engineering and coordination among the subcontractors. For such a structure to work, communication becomes critical. Working groups have been successful in such coordination and communication efforts and can be adapted for less complex projects.

A safety working group provides an interface between two hierarchical components of the safety management structure or between two or more components at the same level. Members of these groups are responsible for coordinating efforts, reporting the status of unresolved safety issues, and sharing information about independent safety efforts. There may be working groups at each level of the control structure, with their members and responsibilities depending on the level at which they operate. A corporate working group may be composed of the safety managers for different divisions or programs while at the lower levels working groups may be composed of representatives from groups designing different parts of a product. When there are special types of safety problems that require new solutions, a safety group may be created to share experiences and approaches, such as a software safety working group. A government agency might create a safety working group for a large project composed of representatives from the agency offices and the prime contractor. A prime contractor may create a safety working group composed of the safety managers for all the subcontractors.

Tips for designing communication and coordination

- Make sure information necessary for decision making involving safety is available to decision makers.
- Provide communication channels and a way to coordinate activities among those with overlapping safety responsibilities
- Make sure safety activities are not fragmented and uncoordinated.
- Define required interactions and ensure that necessary information flow among them is defined and used
- Identify and document necessary safety-related communication channels among all components of the management structure.
- Ensure feedback and coordination channels exist and are working
- Consider establishing safety working groups

Managing and controlling change

Most accidents occur after some type of change.[4] This fact is not surprising as continuing to do what one has done without any changes in behavior or the environment should theoretically result in the same consequences. Adaptation and change is an inherent part of any system and is required for an organization to thrive. The problem is not change, but *unsafe* change. The SMS, therefore, must have carefully designed controls to prevent unsafe changes and to detect them if they occur despite the efforts to prevent them. The goal is to allow change as long as it does not violate the safety constraints. Two types of changes need to be considered: planned changes and unplanned changes.

Planned Changes: The SMS must continue to be effective in the face of planned changes, including changes in the organization, in human behavior and processes, in operations (including changes to maintenance phases and back to operational phases), or in the environment. Most companies have (and, if not, they should have) Management of Change (MOC) policies that require all planned changes to be evaluated with respect to their impact on safety. The cost of such evaluations will depend on the scope of the change, the quality of documentation, and how the original hazard analysis was performed. Traceability should be provided from the identified hazards to their causal scenarios and to the design features that are used to prevent them and vice versa. Without such traceability, the cost of reanalysis after a change may be impractical and therefore skipped.

In fact, MOC procedures are often skipped in practice, resulting in losses that need never have occurred. The safety management structure must have assigned responsibility for enforcement of MOC procedures and feedback channels to determine whether they are being followed and, if not, why. If the cause of skipping them is in the MOC change procedures themselves (too onerous, too difficult, too time consuming, etc.), then the MOC procedures may need to be changed.

The cost of implementing the MOC procedures will depend on the quality of the system documentation and how the original hazard analysis was done. A design goal for STPA was minimization of the cost in evaluating changes.

Unplanned changes: Unplanned changes present more difficult challenges. There needs to be (1) a way to identify potentially unsafe changes and (2) ways to respond to these changes. Leading indicators, should be created from assumptions made in the system (including the SMS) design and safety analysis

[4] See Nancy G. Leveson, *Safeware*, Addison-Wesley Publishers for references.

process.[5] Identifying leading indicators requires recording assumptions during design and development of the organizational structure, the SMS, and the products and workplace. A sophisticated and comprehensive hazard analysis method, such as STPA, should provide the information necessary to create effective leading indicators. If the causal scenarios cannot be eliminated, information from the STPA analysis can be used to create system design constraints as well as to identify what to check in audits and performance assessments.

A common reason for unsafe changes to occur is risk re-evaluation. After a period of few or no losses, people begin to re-evaluate their view of the risks downward. Nothing may have changed so risk has not really decreased, but people may believe it has. Under pressure, they start to violate their own rules and justify it by arguing that safety will not be affected. The processes in the safety management structure need to interrupt this risk re-evaluation process before safety margins are eroded. One requirement is the communication of appropriate information about the actual level of risk. Management needs to be aware of the state of risk in the processes they are controlling. That requires appropriate feedback about the state of the designed safety controls.

In addition, there needs to be an alerting function to a person with responsibility when behavior is contrary to the true level of risk.

The key to preventing this risk re-evaluation phenomenon is to allow flexibility in how safety goals are achieved and to provide information that allows accurate risk assessment by decision makers.

An effective SMS is characterized by continual vigilance against degradation of the system over time. The SUBSAFE program, as an example, puts a lot of effort into combating what they describe are their three most difficult challenges:

1. Ignorance: Not knowing;
2. Arrogance: Pride, self-importance, conceit, or assumption of intellectual superiority and presumption of knowledge that is not supported by fact;
3. Complacency: Satisfaction with one’s accomplishments accompanied by a lack of awareness of actual dangers or deficiencies.

Much of the design and focus of the SUBSAFE program is aimed at combating these three challenges, and the design of any SMS should include steps to provide vigilance against them.

[5] Nancy Leveson, A systems approach to risk management through leading safety indicators. Reliability Engineering and System Safety, 136:17‑34. 2015

Tips for managing and controlling change

- Design controls and MOC policy to prevent unsafe changes and detect if they occur
- Evaluate all planned changes, including temporary ones, for their potential impact on safety
- Assign responsibility for ensuring MOC procedures are enforced and are being followed. If they are not, find out why and fix the problems.
- Create documentation and procedures that minimize the cost of performing the MOC procedures.
- Create ways to identify unplanned changes that could be unsafe ones and to respond to these changes.
    - Devise assumption-based leading indicators and create a risk management program that effectively monitors and responds when potentially unsafe changes are identified.
    - Record assumptions during design and development of the organizational structure, the SMS, the products, and the workplace.
    - Create shaping and hedging actions and a leading indicator checking program including audits and performance checking as well as signposts
    - Implement leading indicators to signal when controls are becoming ineffective
- Ensure that decision makers have information about the current level of risk and state of the designed safety controls
- Assign responsibility to respond when feedback shows that behavior does not match the true level of risk.
- Remain vigilant against the degradation of the safety management structure and the safety culture over time and any increase in complacency.

Designing and Encouraging Feedback

Accurate risk assessment requires information flow and properly functioning feedback channels. Cultural problems are a common reason for problems in feedback. There are three general ways to implement feedback channels: audits and performance assessments, accident/incident causal analysis, and reporting systems.

Audits and Performance Assessments: There are many types of audits and performance assessments, but those whose goal is to evaluate the state of safety start from the safety constraints and the assumptions in the design of the safety controls. Audits should involve not just the products and processes but also the safety management system itself and the effectiveness of the controls designed to ensure that losses are prevented. At least some part of a safety audit and performance assessment should be focused on the operation of the safety management structure as it was designed and assumed it would operate.

The entire safety management structure must be audited and not just the lower levels. In the SUBSAFE program, even the top Admirals are subject to a SUBSAFE audit. Not only does this ensure that the program is operating as assumed, but it also provides a positive cultural component to the audit in that all employees see that even the leaders are expected to follow the SUBSAFE rules and that top management are willing to accept and resolve audit findings just like any other member of the SUBSAFE community. People at lower levels of the SUBSAFE safety management structure participate in the performance assessment of those above them, providing a visible sign of the commitment of the entire program to safety and the importance of everyone in preventing losses. Accepting being audited and implementing improvements as a result, that is, leading by example, is a powerful way for leaders to convey their commitment to safety and to its improvement.

Participatory and non-punitive audits can be very effective. The goal of such an audit is that it be a constructive learning experience and not a judgmental process. It should be viewed as a chance to improve safety rather than a way to evaluate employees. Instead of the usual observation-only audit process by outside audit companies, experts from other parts of the organization not directly being audited should make up the audit team. Various stakeholders may play a role in the audit and even individuals from the group being audited may be on the audit team. The goal should be to create an attitude that this is a chance to improve our practices and provide a learning activity for everyone involved, including the auditors.

The audit itself should involve continuous communication with those being audited so as to obtain full understanding of any identified problems and potential solutions. Unlike the usual rules for outside audits, in a participatory audit immediate feedback should be provided and solutions discussed. Doing this will reinforce the understanding that the goal is to improve safety and not to punish or evaluate those involved. It also provides an opportunity to solve problems when a knowledgeable team is on the spot and not after the usual written report is provided months after the audit actually occurs.

An important goal of safety audits and performance assessments is to measure the level of safety knowledge and training that actually exists, not what managers think exists or what exists in the training programs and user manuals. The audits can provide important feedback about potential improvement to the training and education activities. In keeping with the non-punitive audit philosophy, knowledge assessments should not be used in a negative way that is viewed as punishment by those being assessed. Instead, the goal should be to improve training and education efforts.

Incident and Accident Investigation: Incident and accident investigation provides clear evidence that the SMS is not working as designed or expected. The investigation procedures must be embedded in an organizational structure that allows exploitation of the results. *All* systemic factors involved must be identified and not just the symptoms, technical factors, or low-level employees. Assigning blame should not be the goal; rather the goal should be to find out why the safety management structure was not effective in preventing the loss or near loss. That means that the entire safety management structure must be investigated to identify each component's potential contribution to the adverse events. My CAST (Causal Analysis based on Systems Theory) accident causal analysis process was designed to achieve these goals.

Managers should not be responsible for investigating incidents that occur in their chain of command: investigators and causal analysts must be managerial and financially independent from those in the immediate management structure involved. Using trained teams with independent budgets and with high-level and independent management should be considered.

Investigation requires more than just writing a report. There must be assignment of responsibility for ensuring that appropriate measures are taken to strengthen the aspects of the safety management structure that contributed to the events. Then there should be follow-up to ensure that the fixes were effective. Too often, fixes are made but there is no attempt to determine whether the fixes were successful in improving safety management. Finally, the findings should be used as input to future audits and performance assessments. If there is a reoccurrence of the same factors that led to incidents and accidents in the past, there needs to be an investigation of why those factors were never corrected or why they reoccurred even if they were removed for a while. If fixes are not effective in removing the causes of incidents, then an investigation of the process of creating recommendations and responding to them is warranted to identify any weaknesses in these processes in the organization and to improve them. That is, not only must the factors involved in the incident be corrected but also the process that led to inadequate fixes being implemented after previous incidents or accidents.

Reporting Systems: Reporting systems are critical. Often, after an accident, it is found that the same events occurred multiple times in the past but were never reported or, if reported, were never corrected. These events may involve near misses. For example, several aircraft may not fly an approach correctly, but because it did not result in an accident they may not report it, even if a reporting system exists. Not until an accident occurs is action taken.

A common finding in accident reports is that a reporting system existed but was not used. These reports then usually include recommendations to train people on how to use the reporting system and to require that they use it. This recommendation assumes that the problem is with the potential reporter and not with the design of the reporting system itself.

Examination of the reporting system may find that it is difficult or awkward to use and that information reported appears to go into a black hole. People may believe that there is no point in going through the official reporting system because the organization will not do anything anyway about the factors reported. Often, those finding problems bypass the reporting system and simply go directly to a person or group they believe may be able to solve the problems. While this might be effective and efficient in the short term, it may lead to the same problems occurring in the future because the systemic causes are not eliminated. In some cases, reporting is inhibited by a fear that the information reported will be used against the reporter. Anonymous reporting systems can be helpful here.

Another factor to consider is that often events are not reported by front-line operators because they identified the problem before it created a perception of risk or they may have perceived it as only their own error. Most people are not trained to recognize risk when it is created as part of a normal job process. People accept the flaws in design as “normal” and perhaps already known so rather than reporting them, they just work with or around them. A related factor is that people will generally not report hazardous events when those events do not meet the criteria for required reporting. Near misses may fall in the latter category.

In general, reporting needs to be encouraged. This can be accomplished by maximizing accessibility, minimizing anxiety, and acting on the information obtained. Reporting forms or channels should be easily and ubiquitously available and not cumbersome to fill in or use. To minimize anxiety, there should be a written policy on what the reporting process is; the consequences of reporting; and the rights, privileges, protections, and obligations of those doing both the reporting and those following up on the reports. Without a written policy, ambiguity exists and people will disclose less. Alternatively, ambiguity about who is responsible for following up may lead to everyone assuming that someone else will take care of it.

Finally, encouraging reporting involves providing feedback. Immediately after a report is created, the person who provides the information should be informed that the report was received, assured that it will be investigated, and thanked for their input. A second crucial component is providing feedback later about the results of any investigation and any steps that were taken as a result to prevent a reoccurrence. Reporters should not feel like their concerns are being ignored.

Tips for designing and encouraging feedback

- Design and ensure the continued efficacy of audits, performance assessments, and reporting systems.
- Audit the safety management structure itself (including all levels) and the effectiveness of the designed controls.
- Design audits so they are constructive learning experiences and not a judgmental process.
    - Include as participants members of the groups that are being audited.
    - Use audits as a way to improve safety and as a learning activity and not to evaluate employees.
- Take advantage of audits to evaluate the effectiveness of training and education activities and use them to provide feedback (knowledge assessment) to be used for improving training activities.
- Create effective system-level accident/incident causal analysis procedures that focus on *why* and not *who***.**
- Create incident and investigation procedures based on systems thinking (such as CAST) that identify systemic factors and not just the symptoms of the deeper problems.
- Embed the investigation procedures in an organizational structure that allows exploitation of the results. Assigning blame or finding a "root cause" should not be the goal.
- Accident investigation should be managerially and financially from those in the immediate management structure involved. Consider using highly trained teams with independent budgets and high-level management. Follow up on recommendations to determine whether they were effective and, if not, then why.
- Ensure reporting systems are easy to use and available and anonymous reporting channels exist.
    - Encourage reporting and train people to know when it should be used.
    - Provide a written policy
    - Maximize accessibility, minimize anxiety, and act on information obtained.
    - Provide feedback to those using the reporting channels. Reporters need to feel like their concerns are not being ignored.
-

Risk Management

Sometimes the risk management system is called the "safety management system," but they are not the same. The SMS is a larger concept, that may (and usually does) involve risk management activities. Risk management is the set of activities associated with identifying hazards, analyzing them, and using this information to reduce losses through the design of products, processes, services, and workplaces. These activities will be embedded somewhere in the safety management structure. While risk management is comprised of the technical activities involved in preventing hazards, the SMS also contains organizational, managerial, and cultural aspects of safety.

Risk has traditionally been defined as the severity and likelihood of hazards or accidents occurring. An alternative is:

> *Risk is defined in terms of the effectiveness of the controls used to enforce safe system behavior, i.e., the design and operation of the safety management structure.*

This definition does not require the determination of the likelihood of the events occurring but rather an evaluation of the effectiveness of the controls being used to prevent them.

Creating a risk management system involves designing the procedures to be used in performing the technical risk management activities and assigning responsibility for implementing these procedures to components of the SMS. Risk management may also involve creating leading indicators to identify when risk is increasing.

In fact, an important factor in the effectiveness of the risk management system lies in being able to identify when risk is increasing before a major loss occurs. All safety efforts involve assumptions about how the components of the system will behave and about the environment in which they operate. Violations of these assumptions will undermine the original risk identification and management assumptions. Virtually all accidents have precursors that were unrecognized or for which the responses were non-existent or ineffective. Precursors to accidents are incidents or conditions in which losses do not occur but could have under other circumstances. Leading indicators are characteristics of an organization or an organization's operation that indicate the safety management system is not operating as was assumed when it was designed.

Another way of looking at this is that leading indicators are simply evidence before losses occur that the assumptions under which safety was assured were flawed or are no longer true and that risk may be increasing. As an example, before the Shell Moerdijk chemical plant explosion in the Netherlands, there had been two instances where the same conditions led to explosions in other Shell plants. These explosions were never thoroughly explored or at least the results of the analyses were not used in later design activities, such as the design of the Shell Moerdijk plant where the assumption was made that the conditions were impossible and could not lead to an explosion. Surprisingly, probabilistic risk assessments are often not re-evaluated even after concrete evidence is gathered from actual use of the system that the assumptions made in the analysis are not true. Apparently, the power of a calculated number often trumps evidence that the number is untrue. Generating evidence about the truth of risk assessments and the assumptions that underlie them should be an important part of any safety management system.

While it is unrealistic to expect people to recognize precursors that are only evident in hindsight, too often the precursors could and should have been identified and led to further investigation. *An effective SMS is characterized by continual vigilance against degradation of the safety management structure over time*.

A common paradox is that the fewer accidents an organizational has, the more likely it is that there will be more in the future. If the safety management system is effective, then losses are few. Unfortunately, people tend over time to judge the inherent amount of risk with respect to the current accident rate. If there are few accidents, namely, the safety management system is very effective, over time inherent risk is judged as low and safety management starts to degrade and become less effective. Feedback and information is required to keep the mental models of decision makers up-to-date with the actual level of risk at any time and to prevent this degradation process. Other types of controls can also be designed to prevent this phenomenon.

Tips for risk management

- Create procedures for identifying hazards, analyzing them, and then using the resulting information in product development and operations. Assign responsibility to ensure this is taking place.
- Emphasize qualitative procedures that provide the information necessary for improvement and not just quantitative assessment. Don't make up quantitative assessments for likelihood that is unknowable.
- Reevaluate risk assessment results as data is obtained and when changes occur over time.
- Provide ways to identify when risk is increasing over what was originally assumed.
- When procedures are not being followed, don't just try to enforce them. Instead evaluate why they are not being followed and redesign them. Try to understand why gaps between what is specified and what is done are occurring.

Education and training

Everyone in the safety management structure, not just the lower-level controllers of the physical systems, must understand their roles and responsibilities with respect to safety and why the system—including the organizational aspects of the SMS—was designed the way it was. If employees understand the intent of the SMS and commit to it, they are more likely to comply with that intention rather than simply follow rules when it is convenient to do so. Training is not enough; education is required.

Education must include not only information about the hazards and safety constraints enforced by the controls, but also about priorities and how decisions are to be made. The safety philosophy statement, discussed earlier, provides information about the safety values to be used in decision making. In addition, everyone needs to know the risks they are taking in the decisions they make. Often, poor decision making arises from having an incorrect assessment of the risk being assumed. Using the non-probabilistic definition of risk provided above, this means that the decision makers must know how their decisions will impact the designed controls in the safety management structure.

Telling managers and employees to be “mindful of weak signals” (a common suggestion in the High Reliability Organization or HRO literature) simply creates a pretext for blame after a loss event occurs and hindsight provides the clarity that changes a weak signal (noise) into a strong signal. Instead, everyone must be trained on the hazards associated with the operation of a system and how to recognize them if we expect them to recognize the precursors to an accident. People need to know what to look for, not just be told to look for an undefined something.

Training should also include “why” as well as “what.” Understanding the rationale behind the safety rules they are asked to follow will help reduce complacency, what appears to be reckless behavior (but to the person made perfect sense), and unintended changes leading to hazards. The rationale includes understanding why previous accidents occurred and what changes were made to try to prevent a reoccurrence. With increasing automation and people interacting with it, understanding the controlled hazards involved may require more training than was necessary in the past when the hazards were more obvious and intuitive. People who interact with complex systems need to learn more than just the procedures to follow: they must also have an in-depth understanding of the controlled physical process as well as the logic used in any automated controller (software) they may be supervising or with which they may be interacting. A list of what controllers—at all levels of the safety management structure—need to know was included in *Engineering a Safer World* and is repeated here:

- The system hazards and the reasons behind safety-critical procedures and operational rules.

- The potential result of removing or overriding controls, changing prescribed procedures; and inattention to safety-critical features and operations: Past accidents and their causes should be reviewed and understood.
- How to interpret feedback: Training needs to include different combinations of alerts and sequences of events, not just single events.
- How to think flexibly when solving problems: Controllers need to be provided with the opportunity to practice problem solving involving safety.
- General strategies rather than specific responses: Controllers need to develop skills for dealing with unanticipated events.
- How to test hypotheses in an appropriate way: To update mental models, human controllers often use hypothesis testing to understand the current system states and to update their process models. Such hypothesis testing is common with computers and automated systems where documentation is usually so poor and hard to use that experimentation is often the only way to understand the automation design and behavior. Hypothesis testing, however, can lead to losses. Designers need to provide operators with the ability to test hypotheses safely and controllers must be educated on how to do so.
- Emergency procedures must be overlearned and continually practiced. Controllers need to be taught about operating limits and specific actions to take in case they are exceeded. Requiring operators to make good decisions under stress and without full information simply ensures they will be blamed after a loss occurs.

Training should not be a one-time event for employees but should be continual throughout their employment, if only as a reminder of their responsibilities and the system hazards. Learning about recent events and trends should be part of this training. Assessing for the effectiveness of the training, perhaps through regular audits and performance assessments, can be useful in implementing an effective improvement and learning process. Incident and accident investigation results are an important source of information about training effectiveness.

Some companies have created a practice where managers provide safety training. Training experts help manage group dynamics and curriculum development, but the training is provided by project or group leaders. By learning to teach the materials, supervisors and managers are more likely to absorb and practice the key principles. In addition, it has been found that employees pay more attention to a message delivered by their boss than by a trainer or safety expert.

Tips for education and training

- Educate, don’t just train
- Everyone should understand their roles and responsibilities, why the system was designed the way it was, information about hazards and safety constraints enforced by the controls, and the risks they are taking in the decisions they make.
- Training should include “why” and not just “what.”
- Everyone should learn:
    - The system hazards and the reasons behind safety-critical procedures and operational rules that are related to their jobs.
    - The potential result of removing or overriding controls, changing prescribed procedures, and inattention to safety-critical features and operations.
- Make sure the causes of past accidents are well disseminated and understood.
- Give people the opportunity to practice problem solving involving safety
- Teach general strategies rather than specific responses so controllers can develop skills for dealing with unanticipated events.
- Teach operators how to test hypotheses in an appropriate way and provide ways for them to learn in this way.
- Emergency procedures should be overlearned and continually practiced.
- Contingency procedures should be reviewed before any safety-critical activity.
- Training should be continuous and not just a one-time event when hired.
- Teach about recent events and trends
- Consider having managers provide safety training.

Learning and continual improvement:

Because SMS designs are rarely perfect from the beginning and the world changes over time, there must be an effective process in place to ensure that the organization is continually learning from safety incidents and improving the SMS itself as well as the workplace, the products and the services.

Experimentation is an important part of the learning process, and trying new ideas and approaches to improving safety should be encouraged but also evaluated carefully to ensure that improvement actually results.

The Safety Information System (SIS)

A comprehensive and usable safety information system is key to having a successful SMS. It provides a source of information about the effectiveness of the SMS and controls, i.e., risk.

The models of controllers and decision makers of the current level of risk must be kept accurate and coordinated. In essence, the SIS acts as a shared model and a source for updating individual models. Therefore, accurate and timely feedback and data are important. The SIS can provide the information necessary to detect trends, changes and other precursors to an accident; to evaluate the effectiveness of the safety controls; to compare models and risk assessments with actual behavior; and to learn from events and improve the SMS. After major accidents, it is often found that the information to prevent the loss existed but was not used or was not available to those involved. In general, lots of information is collected only because it is required for government reports and not necessarily because it is a necessity for the operation of an effective SMS.

To determine what should be in the SIS, the responsibilities of those in the SMS and the feedback they need to fulfill those responsibilities can be used, that is, the information required to make appropriate decisions. In general, the SIS contains, at a minimum:

- The safety management plan (for both development and operations),
- The status of all safety-related activities for the system,
- The safety constraints and assumptions underlying the design, including operational limitations,
- The results of the hazard analyses (hazard logs) along with tracking and status information on all known hazards,
- The results of performance audits and assessments,
- Incident and accident investigation reports and corrective actions taken,
- Lessons learned and historical information, and
- Results of trend analysis.

Using the SMS design to identify what each person in it requires for safe decision making will help to design a usable SIS where each person is able to find the information they need.

Information may be collected by companies or by industries. No matter how the information is collected, understanding its limitations is important. To be most useful, the information must be accurate and timely and it must be disseminated to the appropriate people in a useful form. Three factors are involved: collection, analysis, and dissemination.

Collection: Data may be distorted by the way it is collected. Common problems are systematic filtering, suppression, and unreliability. Data collected for accident reports tend to focus on proximal events and actors but not the systemic factors involved such as management problems or organizational deficiencies. Studies[6] have shown that data from near-miss reporting by operators are filtered and primarily point to operator error as the cause of the events. Reports by management are similarly filtered and also point to operator error as the cause. Even general safety inspections tend to identify limited categories of conditions, especially when checklists are used.

Data collection tends to be more reliable for accidents that are similar to those that have occurred in the past than for new types of systems where past experience on hazards and causal factors is more limited. Software errors and computer problems are often omitted or inadequately described because of lack of knowledge, lack of accepted and consistent categorizations for such errors, or simply not considering them seriously as a causal factor.

Experience has shown that it is difficult to maintain reliable reporting of incidents over an extended period, and therefore this type of data is not very useful for estimating probabilities of events or their potential consequences. There is data to suggest that in some industries, up to three quarters of accidents may be unreported. Sometimes, information is involved that a company wants to keep secret or that those making the reports may worry will be used against them in job performance evaluations. Potential legal liability may also be involved.

Discussions about operational problems that have occurred should be documented in the SIS and analyzed for their hazard potential. While it is common, as previously discussed, for concerns to be brought directly to those who can address them, bypassing documentation in the SIS, very often these discussions or emails may contain safety issues that are not recognized as such by the participants. In addition, documenting these discussions can be beneficial as they contain valuable corporate knowledge that would otherwise be lost as people leave positions or retire or due to email retention standards. A

[6] Many scientific references for this section can be found in Nancy G. Leveson, *Safeware*, Addison Wesley Publishers, 1995, Chapter 11. I wanted to keep this paper from reading like an academic paper.

requirement to retain and document these communications should be implemented or a corporate communication channel that automatically retains such information should be developed.

Some measures can be used to improve the comprehensiveness and reliability of data collection such as special training of data collectors, feedback of results to the data collectors, fixed routines, and anonymous reporting. Training on techniques such as CAST can encourage the inclusion of information that is often omitted. Automated monitoring systems are now commonplace, but a problem is that they can and usually do collect so much information that it is difficult to analyze the results and detect problems. A sophisticated and systemic hazard analysis method, such as STPA, might be a useful way to help determine what types of information need to be collected and to identify leading indicators in the large amounts of data collected, as suggested earlier.

Analysis: Systematizing and consolidating a large mass of data into a form useful for learning is difficult. Raw quantitative data can be misleading, especially if statistical significance is not included: Data is not the same as information. Again, STPA results can be used not only to identify what data needs to be collected, but to provide guidance on the importance of the events that are occurring. In addition, data to validate the STPA causal analysis results and to identify factors that were thought to be eliminated or mitigated should be part of the SIS collection and analysis process.

The biggest problem in analysis is simply that while it is easy to design data collection channels, finding the time and manpower to analyze all the data that results may be difficult or impractical. Tools like STPA can help to identify the most important data to collect and the most useful analysis processes.

Dissemination: Disseminating information (not data) in a useful form may be the most difficult part of an SIS. Information needs to be presented in a form that people can learn from, apply to their daily jobs, and use throughout the life cycle of projects, not just in the conceptual design stage. And it must be updated in a timely manner: accidents have resulted from changes during operations or due to insufficient updates to the hazard analysis when engineering modifications are made. Usable documentation that includes rationale and intent as well as traceability to assist in the change process is needed. Also, the assigned safety-related responsibilities to individuals and groups in the safety management structure will provide the information necessary to tailor information to the needs of those that receive it. Providing too much information to individuals can be as dangerous as providing too little. The problem of providing too much information cannot be overemphasized. There must be an easy way for decision makers to get the information they need when they need it. Managers have told me they simply delete the frequent emails they get about risk assessments because they do not provide any useful information.

The method for presenting information should be adaptable to the cognitive styles of the users and should be integrated into the environment in which safety-related decisions are made. Again, the safety SMS design process will provide a great deal of useful input about what information is needed, when it is needed, and how it will be used by individual controllers in the SMS structure.

Tips for designing a safety information system

- Accurate and timely feedback and data are important
- The SIS should provide the information necessary to detect trends, changes and other precursors to an accident; to evaluate the effectiveness of the safety controls; to compare models and risk assessments with actual behavior; and to learn from events and improve the SMS.
- Use the defined responsibilities of those in the safety management structure to identify the information they need to keep their process models accurate enough for good decision making.
- Understand the limitations of your collected data.
- To be most useful, information must be accurate and timely and it must be disseminated to the appropriate people in a useful form.
- Find ways to collect data that minimize distortion of the data (filtering, suppression, and unreliability).
- Keep detailed information on actual safety-related incidents and ensure that information gets to the people who need it.
- Create ways to improve the comprehension and reliability of data collection.
- Try to include statistical significance on numeric data if possible.
- Identify what data needs to be collected and provide guidance on the importance of the events that are occurring.
- Collect data to validate the hazard causal analysis results and to identify factors that were thought to be eliminated or mitigated.
- Ensure that data is analyzed and not just collected.
- Present data in a way that people can learn from it and apply it to their daily jobs.
- Keep data up to date.
- Document design rationale and intent and provide traceability to assist in the change process.
- Tailor the information provided and the presentation format to the needs of those receiving it.
- Providing too much data, particularly raw data, can be as dangerous as providing too little.
- Adapt the presentation of information to the cognitive styles of the users and integrate it into the environment in which safety-related decisions are made.
- Use the safety management structure design process to determine what information is needed, when it is needed, and how it will be used.

Summary

In general, effective safety management requires:

- Commitment and leadership at all levels
- A strong corporate safety culture
- A clearly articulated safety vision, values and procedures, shared among stakeholders
- An SMS with appropriate assignment of responsibility, authority, and accountability.
- Feedback channels that provide an accurate view of the state of safety at all levels of the safety management structure

- Integration of safety into development and line operations (not just a separate and independent group or a separate subculture)
- Individuals with appropriate knowledge, skills, and ability
- Stakeholders with partnership roles and responsibilities
- A designated process for resolving tensions between safety priorities and other priorities
- Risk awareness and communication channels for disseminating safety information
- Controls on system migration toward higher risk
- An effective and usable safety information system
- Continual improvement and learning
- Education, training, and capability development

For more information, the reader is referred to Chapter 13 of *Engineering a Safer World* and to Chapters 11 and 12 in *Safeware*. The Appendix of this paper also contains a list of the responsibilities that need to be included in an effective SMS.

Appendix D: Responsibilities to be Included in the Safety Management Structure

The list below (from *Engineering a Safer World*, Chapter 13) can be used in creating an effective SMS, as a checklist for those analyzing their existing SMS or initiating an activity to improve it, and in identifying inadequate controls and control structures when performing incident and accident analysis. It is not meant to be exhaustive and will need to be supplemented for specific industries and safety programs.

This list only contains general responsibilities and does not indicate how they should be assigned. Appropriate assignment of the responsibilities to specific people and places in the organization will depend on the management structure of each organization. Each general responsibility may be separated into multiple individual responsibilities and assigned throughout the safety management structure, with one group actually implementing the responsibilities and others above them supervising, leading or directing, or overseeing the activity. Of course, each responsibility assumes the need for associated authority and accountability, as well as the controls, feedback, and communication channels necessary to implement the responsibility.

General Management

- Provide leadership, oversight, and management of safety at all levels of the organization.
- Create a corporate or organizational safety philosophy and more detailed policy to implement it. Establish criteria for evaluating safety-critical decisions and implementing safety controls. Establish distribution channels for the philosophy and policy. Establish feedback channels to determine whether employees understand it, are following it, and whether it is effective. Update it as needed.
- Establish corporate or organizational safety standards and then implement, update, and enforce them. Set minimum requirements for safety engineering in development and operations and oversee the implementation of those requirements, including any contractor activities. Set minimum physical and operational standards for hazardous operations.
- Establish incident and accident investigation standards and ensure recommendations are implemented and effective. Use feedback to improve the standards.
- Establish management of change requirements for evaluating all changes for their impact on safety, including changes in the safety management structure. Audit the safety management structure for unplanned changes and migration toward states of higher risk.
- Create and monitor the organizational safety management structure. Assign responsibility, authority, and accountability for safety.
- Establish working groups.
- Establish robust and reliable communication channels to ensure accurate management risk awareness of the development system design and the state of the operating process. These channels should include contractor activities.
- Provide physical and personnel resources for safety-related activities. Ensure that those performing safety-critical activities have the appropriate skills, knowledge, and physical resources.
- Create an easy-to-use problem reporting system and then monitor it for needed changes and improvements.
- Establish safety education and training for all employees and establish feedback channels to determine whether it is effective along with processes for continual improvement. The education should include reminders of past accidents and causes and input from lessons learned and trouble

reports. Assessment of effectiveness may include information obtained from knowledge assessments during audits.

- Establish organizational and management structures to ensure that safety-related technical decision making is independent from programmatic considerations, including cost and schedule.
- Establish defined, transparent, and explicit resolution procedures for conflicts between safety-related technical decisions and programmatic considerations. Ensure that the conflict resolution procedures are being used and are effective.
- Ensure that managers who are making safety-related decisions are fully informed and skilled. Establish mechanisms to allow and encourage all employees (including front-line operators) and contractors to contribute to safety-related decision making.
- Establish an assessment and improvement process for safety-related decision making.
- Create and update the organizational safety information system.
- Create and update safety management plans.
- Establish communication channels, resolution processes, and adjudication procedures for employees and contractors to surface complaints and concerns about the safety of the system or parts of the SMS that are not functioning appropriately. Evaluate the need for anonymity in reporting concerns.

Development

- Implement special training for developers and development managers in safety-guided design and other necessary skills. Update this training as events occur and more is learned from experience. Create feedback, assessment, and improvement processes for the training.
- Create and maintain the hazard log. Establish and maintain documentation and tracking of hazards and their status.
- Establish working groups where appropriate.
- Design safety into the system using system hazards and safety constraints. Iterate and refine the design and the safety constraints as the design process proceeds. Ensure the system design includes consideration of how to eliminate or reduce contextual factors that cause or contribute to unsafe operator behavior that, in turn, contributes to system hazards. Distraction, fatigue, etc. are risk factors resulting from system design flaws, not human operator flaws.
- Document operational assumptions, safety constraints, safety-related design features, operating assumptions, safety-related operational limitations, training and operating instructions, audits and performance assessment requirements, operational procedures, and safety verification and analysis results. Document both what and why, including tracing between safety constraints and the design features to enforce them.
- Perform high-quality and comprehensive hazard analyses to be available and usable when safety-related decisions need to be made, starting with early decision making and continuing through the system's life. Ensure that the hazard analysis results are communicated in a timely manner to those who need them. Establish a communication structure that allows communication downward, upward, and sideways (i.e., among those building subsystems). Ensure that hazard analyses are updated as the design evolves and test experience is acquired.
- Train engineers and managers to use the results of hazard analyses in their decision making.
- Maintain and use hazard logs and hazard analyses as experience is acquired. Ensure communication of safety-related requirements and constraints to everyone involved in development.
- Gather lessons learned in operations (including accident and incident reports) and use them to improve the development processes. Use operating experience to identify flaws in the development safety controls and implement improvements.

Operations

- Create an operations safety management plan
- Develop special training for operators and operations management to create needed skills and update this training as events occur and more is learned from experience. Create feedback, assessment, and improvement processes for this training. Train employees to perform their jobs safely, understand proper use of safety equipment, and respond appropriately in an emergency.
- Establish working groups.
- Maintain and use hazard logs and hazard analyses during operations as experience is acquired.
- Ensure all emergency equipment and safety devices are operable at all times during hazardous operations. Before safety-critical, non-routine, potentially hazardous operations are started, inspect all safety equipment to ensure it is operational, including the testing of alarms.
- Perform an in-depth investigation of any operational anomalies, including hazardous conditions (such as water in a tank that will contain chemicals that react to water) or events. Determine why they occurred before any potentially dangerous operations are started or restarted. Provide the training necessary to do this type of investigation and proper feedback channels to management.
- Create management of change procedures and ensure they are being followed. These procedures should include hazard analyses on all proposed changes and approval of all changes related to safety-critical operations. Create and enforce policies about disabling safety-critical equipment.
- Perform safety audits, performance assessments, and inspections using the hazard analysis results as the preconditions for operations and maintenance. Collect data to ensure safety policies and procedures are being followed and that education and training about safety is effective. Establish feedback channels for leading indicators of increasing risk.
- Use the hazard analysis and documentation created during development and passed to operations to identify leading indicators of migration toward states of higher risk. Establish feedback channels to detect the leading indicators and respond appropriately.
- Establish communication channels from operations to development to pass back information about operational experience.
- Perform in-depth incident and accident investigations, including all systemic factors. Assign responsibility for implementing all recommendations. Follow up to determine whether recommendations were fully implemented and effective.
- Perform independent checks of safety-critical activities to ensure they have been done properly.
- Prioritize maintenance for identified safety-critical items. Enforce maintenance schedules.
- Create and enforce policies about disabling safety-critical equipment and making changes to the physical system.
- Create and execute special procedures for the startup of operations in a previously shutdown unit or after maintenance activities.
- Investigate and reduce the frequency of spurious alarms.
- Clearly mark malfunctioning alarms and gauges. In general, establish procedures for communicating information about all current malfunctioning equipment to operators and ensure they are being followed. Eliminate all barriers to reporting malfunctioning equipment.
- Define and communicate safe operating limits for all safety-critical equipment and alarm procedures. Ensure that operators are aware of these limits. Assure that operators are rewarded for following the limits and emergency procedures, even when it turns out no emergency existed. Provide for tuning the operating limits and alarm procedures over time as required.
- Ensure that spare safety-critical items are in stock or can be acquired quickly.

- Establish communication channels to plant management about all events and activities that are safety-related. Ensure management has the information and risk awareness they need to make safe decisions about operations.
- Ensure emergency equipment and response is available and operable to treat injured workers.
- Establish communication channels to the community to provide information about hazards and necessary contingency actions and emergency response requirements.