

# A Safety Management System for Health Information Technology (HIT)<sup>1</sup>

Nancy Leveson<sup>2</sup>, John Thomas<sup>2</sup>, Stephen Powell<sup>3</sup>, Abigail Williams<sup>3</sup>, Alana Keller<sup>3</sup>

## Executive Summary

Healthcare is one of the most complex and high-stakes systems in our society, yet its safety practices tend to lag behind those of other industries. By adopting proven strategies from communities that have mastered safety management, patient safety can be transformed from a reactive process into a proactive, cost-saving, and highly effective system.

### Three Core Practices

Three proven practices—common in most industries but sometimes overlooked in healthcare—are vital for building an effective safety management system (SMS): (1) hazard analysis, (2) emphasizing preventive vs. reactive safety efforts, and (3) distinguishing safety from reliability. These concepts form the core of any effective SMS.

Hazard analysis, which is the core of safety in most every industry except healthcare, is one of the most effective ways to prevent hazardous conditions and avoid costly response efforts after they occur. Successful prevention leads to enormous cost savings while increasing safety at the same time. Reactive efforts, i.e., investigation after something has gone wrong, will still be needed but to a much lesser extent.

Efforts to improve reliability are not always necessary or helpful if the goal is to improve safety. Safety and reliability are different system properties and, in fact, improving reliability often has little to no impact on improving safety.

An SMS has three, equally important, components: (1) assigned management responsibilities and controls (the safety control structure), (2) the safety information system, and (3) safety culture.

### Management control structure and responsibilities:

Designing the SMS involves (1) defining the expectations (goals), responsibilities, authority, and accountability for safety and (2) assigning each of these to the appropriate parts of the control structure.

The responsibilities for an SMS are commonly specified as “pillars.” Figure A shows the set of SMS pillars that are suggested for a large medical organization like the U.S. Veterans Health Administration (VHA).

The four pillars—establishing culture and policy, controlling hazards, managing operations, and learning from experience—all rest on the foundation of a comprehensive safety information system (SIS). The SIS might be most appropriately managed by the knowledge management function, although all the functions will use and provide inputs to the SIS. The details of what is entailed in each of these responsibilities is included in this report.

Clearly, not all of these activities can be introduced immediately; a plan is needed to determine which should be done first and how changes can be most effectively and efficiently implemented. Some, of course, will already exist when strengthening an existing SMS.

---

<sup>1</sup> This research was partially supported by funding from the VHA.

<sup>2</sup> Massachusetts Institute of Technology

<sup>3</sup> Synensys

A high-level design for an SMS healthcare control structure is provided in this report. The VHA must create their own design details that match their culture and history as well as the lessons learned in all industries from the study of losses in the past.

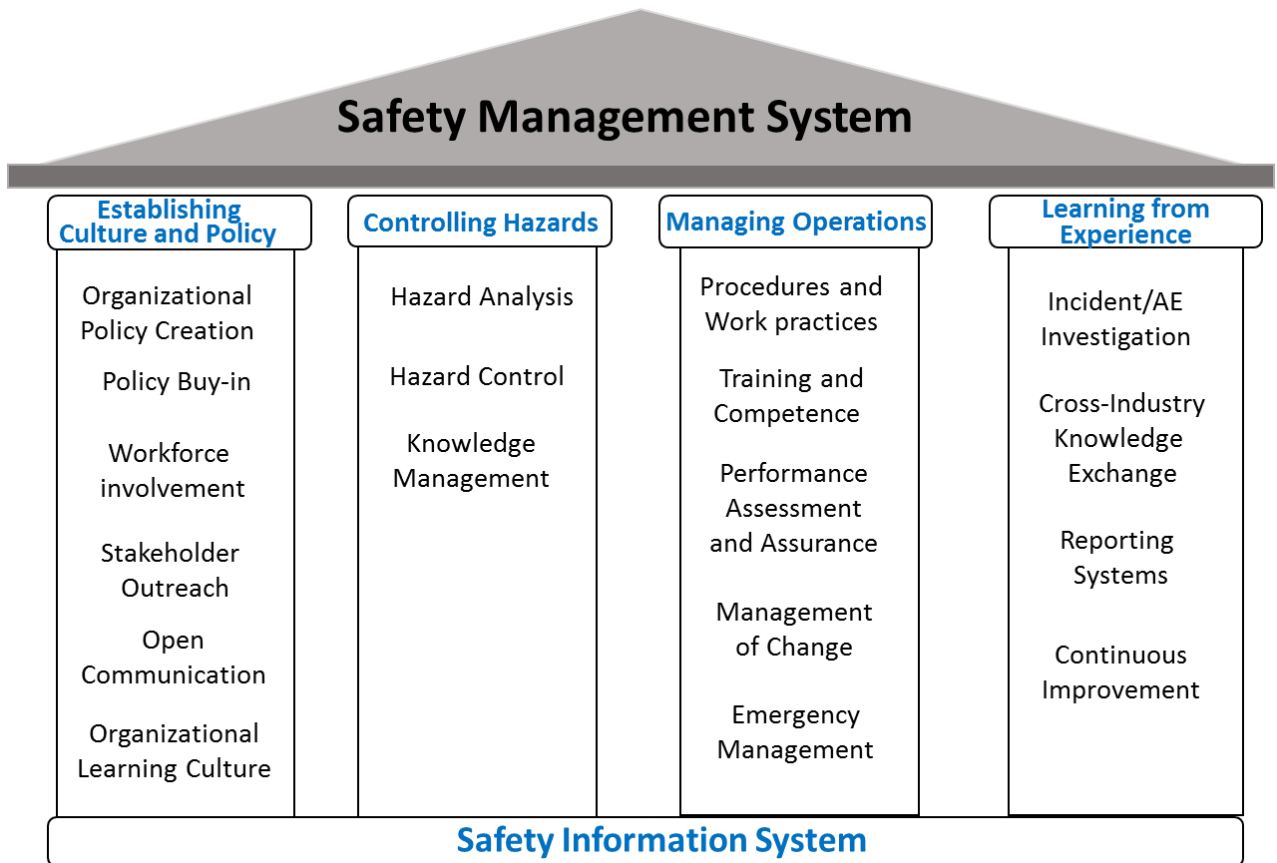


Figure A. The Structure of a Basic SMS for Healthcare

Examining a large number of accidents over most industries shows common systemic factors: (1) ill-defined and diffused responsibility, authority, and accountability; (2) inappropriate placement of system safety responsibility within the organization; and (3) limited communication channels and poor information flow [Leveson, 2023]. These factors should guide the design and improvement of an SMS.

Problems arise especially when responsibility is divided across organizational boundaries: There should be only one person in the organization with overall responsibility for safety. That person should report directly to the person responsible for making decisions about the organization as a whole (i.e., the Deputy Undersecretary for Health).

Some basic guidelines for assigning responsibilities are provided in this report. Of particular concern is the need for communication and coordination: An examination of the organizational factors leading to accidents often turns up communication problems. Not only goals and policies, but the reasons for decisions, procedures, and choices, need to be communicated downward in order to avoid undesirable modifications by lower levels and to allow detection and correction of misunderstandings and misinterpretation. In the opposite direction, the feedback from operational experience and communication of safety issues up the chain of command is crucial for proper decision making. The SIS

provides a useful medium for communication. All feedback channels need to be carefully designed in order to encourage participation. Simply making the channels available is not enough. Communication is also important in coordinating activities and responses to events. People with overlapping responsibilities need communication channels and ways to coordinate their activities to ensure that the 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.

In complex systems, safety must be treated as a control problem rather than a failure or reliability problem. Losses in complex systems do not result only from simple failures or reliability problems. While this report does not specify a detailed design for a Digital Health SMS, System-Theoretic Process Analysis (STPA) could be used to provide information to assist in this process. An example is provided in the accompanying report on the STPA analysis of the Clinical Decision Support System (CDSS) at the VHA. Analyzing the current CDSS management design, we discovered a large number of system flaws that are leading to hazardous behavior. From this information, we were able to identify many improvements that would significantly increase safety.

Implementing these recommendations would go far toward improving the Digital Health safety control structure. The next step might be to extend the STPA analysis beyond CDSS to all of HIT at the VHA and use the results to assist in identifying and designing additional improvements. Part of the effort should be to evaluate how existing tools (such as PIQI, DeX, and Komet) can be used to collect crucial feedback information, help to control hazards, and monitor the operation of the SMS. An STPA analysis could identify the most important ways to use these and other tools to improve the safety and quality of patient care.

### The Safety Information System (SIS)

The SIS is responsible for collecting critical safety knowledge, analyzing it, and sharing actionable insights with the people who need it most. A successful SIS is critical for making informed safety decisions at every level of the organization. Research shows that a well-designed safety information system is, after leadership commitment, the most important factor in reducing losses [Kjellan 1982].

There is important safety information that already exists at the VHA, but it is not accessible to those who need it. One example is the intent and rationale behind the rules in the CDSS, which are identified and recorded during the approval process. This information, however, is difficult to find after the rule is approved and in operation because the information is not linked to the rules. Simply organizing the information that already exists would enable better decisions by clinical and technical teams responsible for fixing and updating the rules over time.

Another key function of the SIS is documenting and tracking known hazards and their resolution. Without this information, it becomes difficult to prioritize issues for overextended technical teams, gauge progress on hazard mitigation, or assess whether teams have the resources needed to address safety risks effectively.

The Digital Health Office at the VHA is very experienced and would be an ideal group to be responsible for developing and managing the SIS. The Knowledge Management group will provide significant inputs to the SIS and might be named as the primary owner of the SIS.

Each of the three SIS processes—collection, analysis, and dissemination—must be carefully designed and controlled, as detailed in this report.

### Safety Culture

A good safety culture is not simply a matter of putting a high priority on safety. Rather, *safety culture is the values and deep cultural assumptions in the industry and organization used to make safety-related*

*decisions*. These values are frequently established and communicated through a written organizational safety philosophy that defines how people in the company are expected to make decisions. Examples are provided in this report, both of a general philosophy for safety as a whole and one for HIT. The safety philosophy lays the foundation for detailed policy, rules, and technology design. Management is responsible for obtaining buy-in and ensuring the philosophy is being followed.

An important factor limiting the improvement of safety in healthcare is a common emphasis on reducing liability over improving safety. The most effective solution to minimize liability in the long term is to strengthen the safety management system so that fewer adverse events occur.

## Table of Contents

Introduction	6
Safety Management Systems	6
SMS Basic Concepts	7
Designing a Cost-Effective HIT Safety Management System	10
Organizational Safety Culture	11
Management Control Structure and Responsibilities	14
Defining Responsibilities, Authority, and Accountability	16
Designing the Control Structure and Assigning Responsibilities	22
Safety Information System (SIS)	29
A Path Forward	30
References	31
Appendix A: The Current VHA Safety Control Structure	34

## Introduction

According to the U.S. Department of Health and Human Services, healthcare information technology (HIT) involves the processing, storage, and exchange of health information in an electronic environment. HIT includes EHRs, laboratory information systems, e-prescribing, imaging systems, health information exchanges, CDSS (Clinical Decision Support Systems), and many other systems that codify, represent, and manage patient health data. While computerized provider order entry (CPOE) HIT has reduced some medication errors (e.g., due to illegible handwritten prescriptions), HIT system hazards have led to numerous losses due to miscommunication, diagnostic errors, under/over treatment, and poor data quality (e.g., semantic loss). Except for pre-market medical devices, most U.S. HIT is not explicitly required to be designed or certified for safety. Current HIT certification instead focuses on system security, privacy, usability, and interoperability. The U.S. healthcare system spends approximately \$300 billion per year for HIT software implementation and maintenance—about 8% of total healthcare costs and rising 5% annually [Davis 2025]. The Department of Veteran Affairs Office of the Inspector General estimates the lifecycle costs of replacing the current VHA VistA EHR with a new EHR will exceed \$50 billion [VA Inspector General 2024].

This report was developed within the context of the Veterans Administration (VA) health system, which is the largest integrated healthcare system in the United States. The VA consists of 170 inpatient facilities and over 1300 outpatient care facilities. The VA HIT system is likely the most complex HIT system in the world with 2 EHRs and over 400 interfaces connecting patient diagnostic testing, scheduling, pharmacy, administrative, vaccines, mental health, and many other Veteran care systems. Since the establishment of the Mission Act in 2018, Veterans are receiving more care outside the VA health system (nearly 50%)—requiring significant HIT coordination to maintain a single, complete, and accurate Veteran health record.

The VA has been viewed as a leader in patient safety, HIT, and informatics patient safety for the past 20 years, including the development of the Veterans Health Information Systems and Technology Architecture (VistA) open-source EHR in the 1990s. While the SMS required for the VA will be different than a rural hospital in middle America, this report will help any healthcare system develop a tailored, cost-effective SMS that is able to improve the safety and quality of their HIT systems, care processes, and health outcomes.

## Safety Management Systems

The Safety Management System (SMS) can be defined as the allocation of responsibility, authority, and accountability for safety to groups and individuals working within the system to ensure that accidents/mishaps or adverse events do not occur.

A lot is known about the design of an effective safety management system, but the actual realization of the general principles will differ in different industries. The SMS for aviation, the petrochemical industry, and healthcare must necessarily differ as the hazards are different.

In all industries, a strong SMS has been found to be crucial for improving safety. In addition, a well-designed SMS can lead to significant cost savings, not only by reducing the cost involved in accidents themselves, but by reducing the costs involved in safety controls that are ineffective. Spending limited resources on things that do not improve safety or that have limited impact is not a reasonable substitute for spending fewer resources on activities that significantly reduce adverse events.

While the emphasis of this report is on healthcare information technology, it is not possible to look at only one part of a complex system and ensure safety. Safety is a system property, not a component property. HIT is embedded within the entire healthcare management system, and all the parts interact continually, both directly and indirectly (Figure 1). For example, the EHR interacts directly with clinicians, caregivers, technicians, and patients, and indirectly with standards committees, the CDSS (Clinical

Decision Support System), including both users and developers, drug approval groups, medical societies, developers of EHR systems, regulatory authorities, etc.

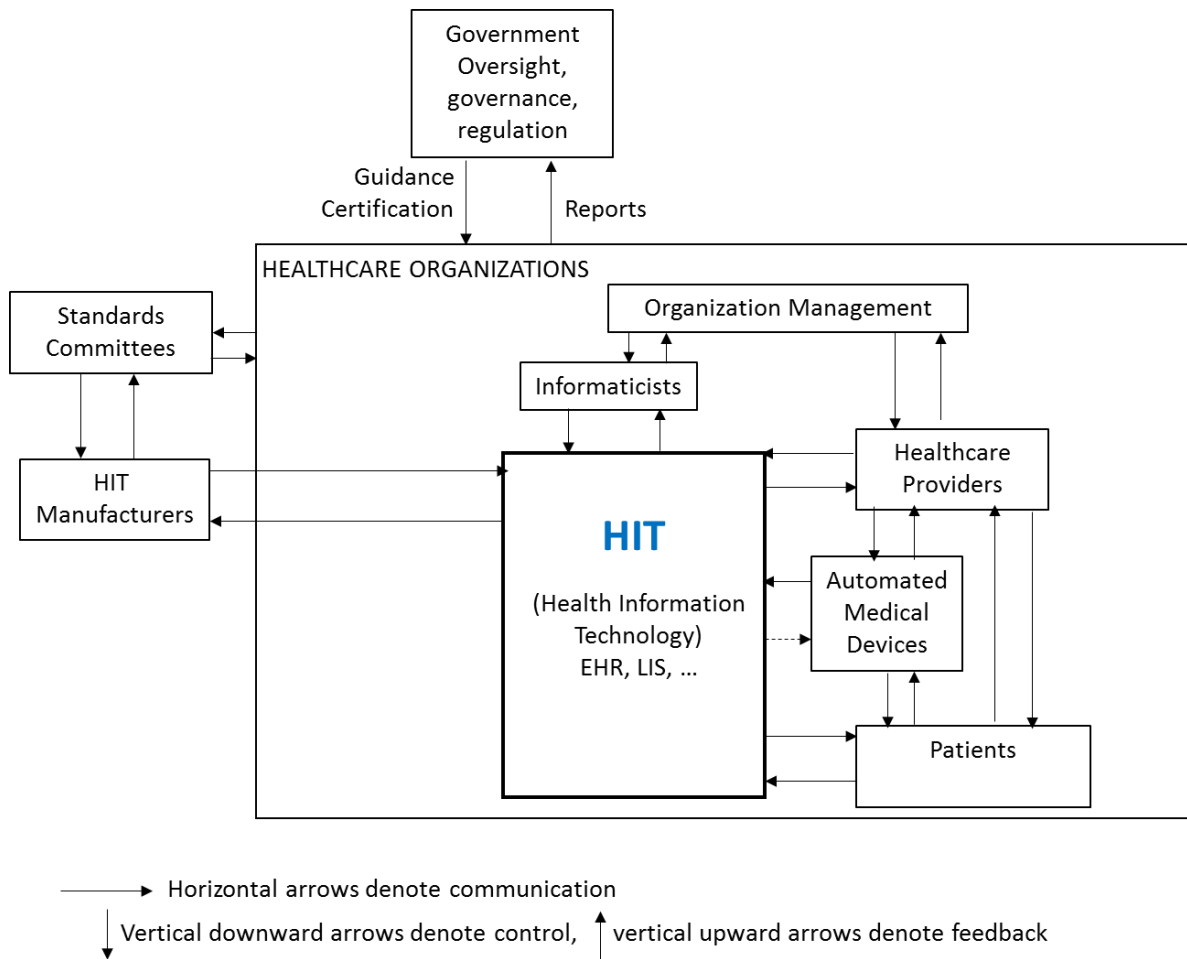


Figure 1: HIT interacts with most parts of the healthcare system

While all these interactions are important and must be considered in any effort to provide safe healthcare, overall emphasis in this report is on management of the safety of HIT and its direct interactions with other components of the general healthcare system.

Before describing how to design a cost-effective SMS for HIT, some general concepts and assumptions need to be explained.

### SMS Basic Concepts

Technology cannot usefully be considered in isolation from the system in which it operates. Computers and software are not by themselves unsafe, that is, they do not catch on fire, explode, or spread infectious organisms. Their operation is only unsafe within a specific context, for example, the information in an EHR leads to an incorrect diagnosis by a healthcare provider that causes appropriate treatment to be delayed. In the example, the information itself is not unsafe; it becomes unsafe only if it is used within a context that can lead to a loss. As a result, although this report will focus on design of controls to ensure HIT safety, the necessary relationship with the total healthcare SMS must be considered.

Some basic concepts underlie the design of any effective SMS no matter the industry, but several of those concepts are not widely taught or used in healthcare: hazards, hazard analysis (proactive prevention vs. reactive efforts), and the difference between safety and reliability.

1. *The most effective safety management systems are built on the basic concept of hazards.*

Virtually every industry uses the concept of “hazard” when dealing with safety except healthcare. It is not enough to simply tell people to concentrate on doing “things right.” Everyone is trying to do the right thing. The solution that is right in one situation may not be right in a different situation, and sometimes the information necessary to determine the “right thing” is not available at the time decisions have to be made.

System safety provides a complementary view to the normal focus on solutions for achieving one’s goals by instead focusing on preventing what could go wrong while working to achieve those goals. These two activities are not equivalent. Together, they provide contrasting views on the required operation of the system.

Separating these two system views has been found to be necessary to significantly improve safety in complex systems. The concept of “hazard” is key to achieving this goal. Hazards, informally, are states of the system that can lead to a loss. Hazards are composed of at least the following information:

- a. [System component\(s\)](#) involved
- b. A potentially [unsafe behavior or condition](#) (later in this report called a *control action*)
- c. [Context](#) in which the condition or behavior is unsafe

As an example: [Physician prescribes a medication](#) when the patient is [allergic to the medication or is taking other contraindicated medications](#).

[System component](#): Physician

[Unsafe behavior](#): Prescribes a medication

[Context](#): When the patient is allergic to the medication the physician prescribes or is taking other contraindicated medications.

Note that prescribing the medication is not itself considered a hazard. The ultimate goal is to eliminate hazards from the system design and operation or, if that is not possible, to minimize their impact. We obviously do not want to eliminate prescribing medication. Only when this action is potentially dangerous within a specific *context* (for example, the patient is allergic to the medication), do we want to eliminate or at least control the action. If a particular action by the caregivers (such as prescribing medications) always resulted in an adverse event, then prescription in general would need to be prohibited. In that case, the behavior is unsafe in all contexts.

In general, every system has some goal(s) along with constraints on how those goals can be achieved. Safety is usually a constraint, although occasionally safety-related goals may be involved (e.g., air traffic control has a goal of eliminating aircraft collisions, although it also has non-safety goals, such as maximizing throughput).

Health information technology usually does not prescribe treatment itself without any input or oversight by a physician or pharmacist, but it can clearly contribute to hazards. IT is often involved in some way in the prescription being filled, for example, the delivery of the prescription to the pharmacist. Again, delivery itself is not the problem; in fact, delivery is the expected and desired goal of the system. A hazard arises only if the delivery leads to a hazardous state, for example, the wrong patient is named or the patient’s EHR notes an allergy, but no warning is issued about that allergy. HIT is also programmed by humans who were following rules provided by someone else, and those rules themselves may be unsafe.

The losses of concern, such as patient harm, and the hazards are defined by the system stakeholders. These stakeholders in healthcare may include patients, clinicians and medical staff, government, hospital management, various medical organizations (such as standards setting groups), i.e., any groups that have responsibility for patient safety within the system being considered). Again, the goal of the SMS is to eliminate or reduce the occurrence of hazards, i.e., the precursors or causes of adverse events. For the HIT SMS, that goal is restricted to the occurrence of IT-related hazards.

## *2. Preventive Safety Efforts are Much More Effective than Reactive and Can Lead to Large Cost Savings*

Preventive safety efforts always occur in every potentially dangerous system. Nobody designs any system without considering the potential losses (e.g., adverse events in healthcare) and how to prevent them. For example, nobody would set up a surgical unit without considering how to ensure proper anesthesia, how to prevent wrong-site surgery, how to provide appropriate and well-trained staff, how to prevent infection, etc. The problem is that this proactive process is usually ad hoc in industries, like healthcare, that do not use hazard analysis. While an ad hoc process may be adequate for relatively simple systems, it is not for the complex systems we are designing and operating today. Using formal processes, called hazard analysis, provides more complete results for complex systems. Virtually every industry today uses formal hazard analysis, except for healthcare.

Preventive measures are much more effective than reactive measures that require waiting until after adverse events have occurred to eliminate or reduce them. While there are obviously costs associated with performing an additional procedure (such as hazard analysis) up front, those costs are swamped by the costs of reacting to each adverse event individually as it occurs as well as the ethical, liability, and other costs of not investing in prevention.

Of course, hazard analysis, like any complex human endeavor, will not always be perfect. Adverse events may still occur. A major goal of root cause analysis (RCA) in those cases is to identify where the hazard analysis went wrong or where unanticipated changes to the system and its environment led to the deficiencies in the hazard analysis. There should be, however, many fewer adverse events to investigate than occur without the use of proactive hazard analysis and prevention efforts. Once improvements are made to the system itself to prevent future similar adverse events, the hazard analysis process used should be examined to determine why that process was not effective in identifying and eliminating the hazard. Appropriate improvements can then be made.

Until relatively recently, the existing hazard analysis techniques did not apply to human-intensive, adaptive systems like healthcare. But new, more powerful hazard analysis techniques have been developed and are being extensively used in other industries. Experimental demonstrations in a variety of healthcare applications have proven their efficacy in healthcare. These applications include radiology/radiotherapy, anesthesiology, laboratory data exchanges, clinical decision support systems, medication administration, operating room procedures, pandemic response, medical devices, and more [Baker 2022; Baker 2025; Cilvis-Cividjian et al. 2020; Chen et al. 2021; Ghorbani et al. 2024; Martinazzo et al. 2021; Pawlicki et al. 2016, Leveson et al. 2023; Raman et al. 2016; Samedy et al. 2022; Samost-Williams and Nanji 2020; Samost-Williams et al. 2025; Thomas et al. 2025; Wong and Pawlicki 2022].

## *3. Safety and reliability are different system properties.*

To create an effective safety management system, it is critical to understand the difference between reliability and safety. For example, system interactions may reliably prevent allergy information from being communicated to the pharmacist. That might happen if a poorly designed HIT automatically (and reliably) hides certain fields (perhaps to prevent overwhelming users with too much information) or an automated procedure (reliably and purposely) omits a step to enter allergy information into a certain field. This system can be reliable in the sense that it always delivers the same information to the pharmacist, but it certainly would not be safe if it doesn't include the information necessary to prevent

adverse events. Most major accidents in safety-critical industries occur when the system is operating with high reliability. Repeatedly and reliably performing the same potentially unsafe behavior does not lead to a safe system.

Confusion can be reduced by carefully defining terms, as is true in any scientific approach to solving problems. Here are the definitions widely used outside of healthcare.

*Safety*: Elimination (non-occurrence) of losses.

Safety is defined in terms of *hazards* or states of the system that can lead to a loss.

*Reliability*: (in engineering) The probability that a piece of equipment or a system will perform its intended/specified function for a specified time and under specified environmental conditions.

Reliability is usually defined in terms of *failure*, that is, not performing the specified function.<sup>4</sup>

*Repeatability*: Provides the same values under the same conditions.

Repeatability is mostly applicable to assembly lines where the goal is to produce products that lie within an acceptable standard of deviation from each other. Outside of assembly lines or systems where correctness is defined in terms of tolerances (upper and lower limits), repeatability is normally not considered an important system quality. Repeating low quality or mediocre processes is not normally a desirable system goal.

*Correctness*: Freedom from error, where *error* is defined by some accepted or specified standard.

One way to understand the difference between these properties is that *context* is a critical factor in safety, but reliability and repeatability can be measured without considering context. A clinical decision support system implemented in an electronic health record (EHR) may reliably implement a decision rule without failure, but that rule may be unsafe in some contexts.

The confusion between reliability and safety arose about 50 years ago. Before the introduction of information technology and computers into systems, most accidents were the result of hardware failures or human errors.

In the case of hardware, increasing reliability did make it safer. But the problem changed as technology changed, and computers were invented and introduced into most systems. As systems became more complex, accidents started to occur where each component operated “correctly” (according to specification), but hazardous states were created by the interactions among operating (non-failed) components.

With the introduction of computers into most systems, the difference between safety and reliability became even more obvious. Software does not “fail” in the way that hardware does. Software instructions are executed unless there is an underlying computer hardware failure, which today is quite rare. Given the same inputs, the software will provide the same outputs.

Safety problems arise in the design of the software and, most often, in flawed requirements. For example, the software design does not check for or issue warnings about allergies under certain conditions (in a particular context) because nobody included this behavior in the software requirements, or the requirements were not implemented correctly. In both cases, the software would be considered to be incorrect, but not unreliable. It would be unreliable only if, given the same inputs, sometimes it checked for allergies and sometimes it did not.

Related to these quality distinctions is the use of redundancy to improve safety. Again, in the much simpler systems of the past, when accidents were primarily caused by simple hardware failures and human errors—and thus the unreliability of hardware and humans—redundancy was a useful tool to

---

<sup>4</sup> The definitions used outside of engineering are too vague to be useful, i.e., trustworthy (which is defined as able to be relied on and is therefore just a synonym for reliability) or performing consistently well (“well” is undefined).

improve safety. Today, however, redundancy is much less useful, unless the problem occurred because of a random hardware failure.

James Reason made an important distinction between *slips* and *mistakes* [Reason 2000]. For simple slips, such as hitting button A when meaning to hit button B, redundancy can potentially be useful. For example, humans are often required to enter data redundantly to avoid simple typing errors. However, from a usability standpoint, redundantly entering data twice is disliked, and humans often find ways around these requirements. Even more important, when considering more complex *mistakes*, e.g., the human was supposed to hit button A but thought they were supposed to hit button B or they entered potentially dangerous data that they did not know was dangerous, redundancy does not help. Most hazards in healthcare today involve mistakes rather than slips.

Safety may not even involve the failure of the system components. Two perfectly reliable components—they satisfy their requirements and do not fail—may interact in ways that create hazards. For example, the IT or human performs a behavior reliably in a context in which that behavior may be dangerous. A physician may reliably (and repeatedly) prescribe a specific medication for a given condition. The prescription may be reliably transmitted to the pharmacist. The pharmacist may reliably fill that prescription. The patient may reliably take the medication as prescribed. But if the patient has an allergy to that medication, those actions—although reliable—are not safe. The reliability of specific system components can be determined without reference to the context in which that behavior occurs. As stated earlier, safety of a behavior (or a system component) can only be defined with respect to the context in which the behavior occurs.

Of course, reliability is an important system quality for many reasons. But assuming that reliability will assure a different system quality, such as safety, is a sure path to having preventable adverse events. Systems can be perfectly reliable and be unsafe. Systems can be unreliable and still be safe. Sometimes reliability and safety conflict, that is, enhancing one can actually decrease the other. High reliability is equivalent to safety only if the context of the behavior involved is irrelevant. That condition is almost never true. Reliability can be determined without reference to context. Safety, i.e., hazards and hazardous behavior, can only be defined with respect to context.

Safety engineering starts with identifying what is needed to eliminate or mitigate hazards and then focuses resources and effort on those aspects of the system design and behavior. It does not start by making everything and everyone reliable or behaviors repeatable.

## **Designing a Cost-Effective HIT Safety Management System**

This report describes how to design a safety management system for HIT. However, as explained above, such an SMS must be part of the overall healthcare SMS. An important design problem for a HIT SMS is determining how the parts of the entire organizational SMS must interact and communicate to avoid hazards.

There is no one correct SMS system. Safety is only one aspect of any system. Different industries and different organizations need an SMS that works within their own culture and their other organizational or industry goals. Therefore, instead of providing a specific design for a HIT SMS, this report provides design guidance for such an SMS without overspecifying only one way to achieve the goals.

There are three basic components of an SMS: (1) safety culture, (2) assigned management responsibilities and controls (the safety control structure), and (3) the safety information system. All are equally important.

## **Organizational Safety Culture**

Social dynamics and organizational culture greatly influence individual decision making. Edgar Shein, considered to be the “father of organizational culture,” defined safety culture as:

*The values and assumptions in the industry and/or organization used to make safety-related decisions* [Shein 1985].

A common misunderstanding is that safety culture is reflected by the concern about safety by those in the industry or organization. In fact, everyone can be highly concerned about safety, but the organization may have a very poor safety culture. While healthcare involves people with great concern for patient safety, they may unintentionally make decisions and act in a way that can cause harm. The goal of the SMS is to eliminate that behavior.

Figure 1 shows the relationship between organizational culture and other components of the organization's operations.

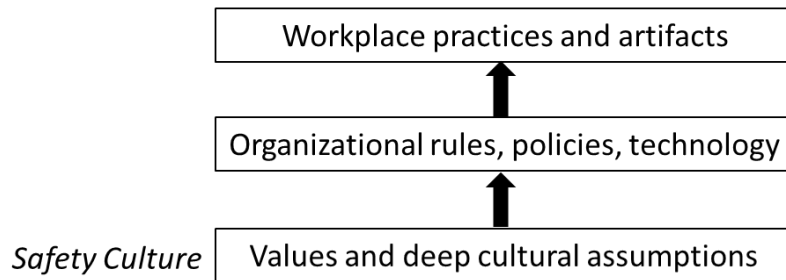


Figure 1. Safety culture provides the foundation for the operation of the organization.

The *organizational culture*, in general, is the values and deep cultural assumptions that form the foundation underlying the creation of the organizational rules, policies, and technology (including the information technology). The rules, policies and technologies, in turn, are used to design workplace practices and artifacts. The *safety culture* is the subset of the organizational or industry culture that reflects safety values. Note that the design, effectiveness, and safety of HIT (technology in the second level of Figure 1) is impacted by the safety culture.

Top management establishes the value system under which decisions are made in an organization. The first step, therefore, in improving the safety culture is for management to establish and communicate what is expected in the way of safety-related decision making and behavior.

A sincere commitment by management to safety is often cited as the most important factor in achieving it [Kjellan 1982]. Employees need to feel that they will be supported if they exhibit a reasonable concern for safety in their work and if they give priority to safety in short-term conflicts with other goals, such as schedule and cost. In the long term, conflicts between safety and other goals are less likely to occur. For example, short-term conflicts can arise between safety and productivity. But in the long term, safety and productivity (and profitability) usually go hand in hand. Accidents, of course, are expensive—but more than that, poorly designed procedures introduced to prevent losses may be very costly, while not improving safety significantly. Large studies and individual company results have shown that increasing safety goes hand in hand with increasing productivity and profits [Leveson, 2012, 2023].

While most everyone in healthcare truly cares about safety, that is not enough. Values about tradeoffs with other properties and assumptions about what is needed to reduce adverse events require more than simply recognizing the importance of safety. Great effort, particularly in the past 30 years, has been exerted to reduce the large number of adverse events in healthcare, resulting in only scattered and limited success. Some of the difficulty stems from the fact that, despite a sincere desire, assumptions about how to achieve the goals are flawed.

Management establishes the value system under which decisions are made in an organization. The first step, therefore, in improving the safety culture (and system safety) is for management to establish and communicate what is expected in the way of safety-related decision making.

### *Creating the Desired Culture*

The first step in creating an effective safety culture is to create a safety philosophy, i.e., identify the principles under which the organization wants its members to make decisions. The goal is to identify the values that are to be used in decision making, especially when there are conflicts between goals. There are lots of possible values, so it is not possible or desirable to identify the philosophy that should be used in all organizations, even within an industry. In addition, a shorter list will be more effective than a long one—between 5 and 7 statements is best. More detailed policy statements can be created in addition to the overarching philosophy statement.

Some examples of what might be included in patient safety policy statements:

- Healthcare safety can only be evaluated within the specific context within which the healthcare is provided.
- Safety and reliability are different system properties. Achieving one does not require achieving the other, and, in fact, improving one may actually decrease the other.
- Increasing quality and safety leads to decreasing cost and schedule (efficiency) and, in the long run, increased profits (efficiency and better use of resources). Preventing adverse events is good business.
- Patients (including those with disabilities) and their health advocates should be able to access and participate in their care decisions.
- All adverse events are preventable. The goal of investigation is to identify the systemic factors leading to unsafe behavior and not just the people involved.
  - Adverse events are an important window into systems that are not operating safely and should trigger comprehensive causal analysis and improvement actions instead of blame and punishment (i.e., promoting a Just Culture).
- Effective communication and the sharing of information is essential to preventing losses.
- Safety commitment, openness, and honesty is valued and rewarded in the organization. Each employee will be evaluated on his or her performance and contribution to our safety efforts.

The organizational safety philosophy is an important input to the development of the organizational rules, policies, and information technology. Safety policy related to HIT will probably appear in specific supplemental policy documents derived from the overall policy statement. Examples include:

- The safety of specific HIT can only be evaluated within the specific contexts in which it is to be used.
- Because significant safety problems have not occurred in other HIT installations does not mean that they will not occur in other contexts. Safety must be ensured for each context in which the HIT is installed.
- HIT should prioritize safety and usability over ease of programming and software development. Emphasis should be on optimizing the workflows for the healthcare providers, not on easing the job of IT groups and software specialists. Assuming users can easily adjust to the requirements of HIT and can adopt standard workflows without degrading the quality and safety of patient care will lead to unnecessary adverse events.
- Proactive hazard analysis applied to HIT is required to identify the scenarios leading to adverse events. Testing and simulation are not adequate for assuring the safety of HIT.
- Reacting after adverse events is too late (although necessary). Hazardous behavior must be designed out of the system before adverse events occur.

- In complex, adaptive systems such as healthcare, the safety of HIT must be continually evaluated through feedback, identification of potential changes that could impact safety, and the use of *leading indicators* [Leveson 2015] of increasing risk.

#### *Implementing the Safety Philosophy (Management Responsibilities)*

A written statement of the organizational 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 administrators and managers. To be successful, there needs to be real commitment by those at the top, not just sloganeering and making statements that are belied by the behavior of the organization's leaders: Management needs to model and reward the desired behavior.

Once the philosophy, policies, etc. are established, executive management is responsible for communicating them to everyone in the organization, getting “buy in,” and ensuring that the philosophy and policies are followed. Feedback channels need to be established and monitored for feedback about the policies, reporting violations of the policies, and providing for general feedback about safety in the healthcare organization. Feedback should be used to determine:

- whether the philosophy and policies are acceptable to those in the organization, or changes are needed as the organization and healthcare itself change over time,
- if those in the organization are following the safety philosophy (that is, decisions are based on the assumptions and values specified in the safety philosophy), and
- suggested improvements in both the philosophy and policies themselves and how they are being implemented.

The organizational values (culture) is communicated in two basic ways. The first is by modeling the behavior desired, and the second is communicating the expected behavior.

#### *Resolving Conflicts Between Safety and Liability*

An important factor inhibiting the improvement of safety in healthcare is a common emphasis on reducing liability over improving safety.

The role of lawyers in an organization is to reduce legal liability. The role of system safety is to eliminate or decrease adverse events. Sometimes these roles conflict, not just in healthcare but in any open society. Excessive “lawyering”—leading to fear of sharing data or investigating adverse events or limiting the scope of an investigation of incidents because the results could lead to legal liability—usually leads, ironically, to more such events and increased liability.

The most effective solution to minimize liability in the long term is to strengthen the safety management system so that fewer adverse events occur. When adverse events are reduced through proactive hazard analysis, robust reporting systems, and a culture of transparency, there are simply fewer incidents that could lead to litigation or regulatory action. Organizations that focus on safety improvements not only protect patients but also build stronger legal cases by generating clear documentation of due diligence, continuous improvement efforts, and adherence to SMS best practices (summarized in this report). Over time, this focus greatly reduces both the number of errors as well as the human and financial costs that arise after an error is made.

## **Management Control Structure and Responsibilities:**

When using systems thinking and systems theory<sup>5</sup>, safety is treated as a control problem rather than a failure or reliability problem. The goal is to eliminate or reduce losses through controlling the occurrence of hazards (hazardous behavior). The ultimate goal is to eliminate hazards, but complete elimination is not practical in many or even most cases, so the goal of elimination must be extended to include reducing the occurrence of hazards and mitigating their impact if they do occur.

Control is not used here in the dictatorial sense by ensuring that everyone behaves in an identical manner, but in the sense of controlling only the behaviors that will lead to adverse events. Control over safety must allow individuals to optimize other goals and behaviors as well as allow useful and necessary adaptation of the system as the requirements and environment change.

Because of conflicts between goals in any real, complex system, mechanisms should exist to allow informed decision making about such conflicts. For example, in many situations, a risk of doing something also involves a different risk in not doing it. A very low risk of an adverse reaction to a treatment may, in many cases, involve a tradeoff with a higher likelihood of curing the patient. The goal should be to enhance individual decision making about safety by providing better information to the decision maker about the potential risks and benefits involved.

In this report, the design principles for creating an effective HIT SMS are described using a hierarchical control structure, which is basically a standard management structure model with a few components added. Figure 2 shows the basic components used in the definition of the SMS structure. A manager has defined responsibilities. Along with the responsibilities, the manager must have appropriate authority (called control actions) to carry out his or her responsibilities. Decisions about what control actions to use are made using a decision-making process of some type along with a mental model of the current state of the system. A common reason for flawed decisions is that the decision maker does not have an accurate model of the controlled system. For example, a manager thinks those under his or her supervision are following specified policies when they are not. Mental models are kept up to date through various types of feedback. Feedback channels also form the basis for assigning accountability.

---

<sup>5</sup> *Systems thinking* is a way of making sense of the complexity of the world by looking at it in terms of wholes and relationships rather than by decomposing the world into its parts and considering them separately. Systems thinking draws on and contributes to *Systems Theory*. About 100 years ago, scientists started to develop what they considered to be a theory of complexity, which is now commonly referred to as Systems Theory, in order to better explain natural and human-designed systems. Ludwig von Bertalanffy (biology) is usually credited with its origination [von Bertalanffy 1934, 1968]. Others involved include James Gibson (psychology), Margaret Mead and Gregory Bateson (anthropology), Russell Ackoff and John Sterman (management), Ervin Laszlo (physics and philosophy), Kenneth Boulding (economics), Norbert Wiener (Mathematics), and Talcott Parsons (sociology).

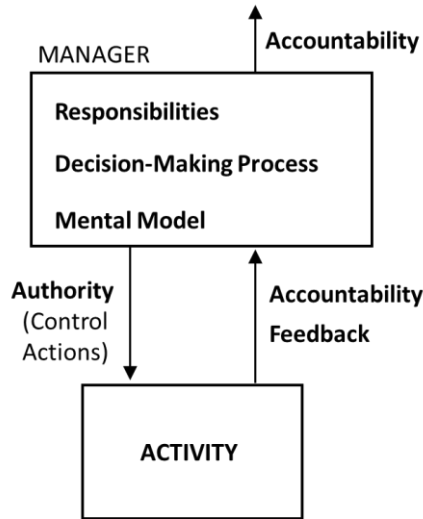


Figure 2. The control structure is made up of “feedback control” components. A template is shown here.

The organizational structure must ensure hazards are eliminated or, if that is not possible, controlled and mitigated. It also must support the positive safety culture values as communicated in the documented safety philosophy. For a government agency, like the Department of Veterans Affairs (VA), the overall control structure for healthcare safety might look like Figure 3:

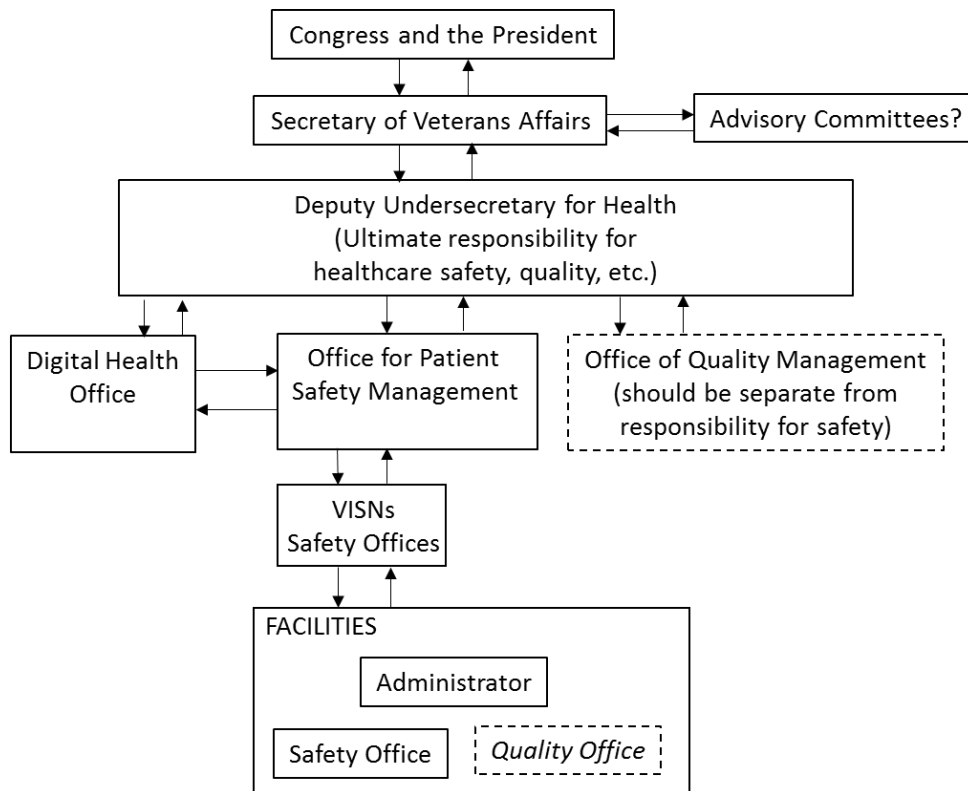


Figure 3. A high-level view of part of a potential VHA safety control structure.

The overall goal when designing an SMS is to design a control structure that eliminates or reduces losses. Satisfying this goal requires a clear definition of the expectations, responsibilities, authority, and accountability for safety-related tasks at all levels of the control structure. In addition, to operate effectively, the structure requires appropriate feedback and coordination among entities. There should be 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 behavior through physical design, processes and procedures, and social interactions and culture.

Examining a large number of accidents over most industries shows common systemic factors: (1) ill-defined and diffused responsibility, authority, and accountability; (2) inappropriate placement of system safety responsibility within the organization; and (3) limited communication channels and poor information flow [Leveson, 2023]. These factors should guide design and improvement of an SMS.

There are two important aspects of designing the SMS: (1) defining the expectations (goals), responsibilities, authority, and accountability for safety and (2) assigning each of these to the appropriate parts of the control structure. The next section describes the responsibilities that need to be assigned. The following section provides guidance on where to put them in the control/management structure, considering only the HIT safety controls.

### **Defining Responsibilities, Authority, and Accountability**

As with any effective management system, there must be responsibility, authority, and accountability assignments. A belief that “everyone is responsible for their own and others’ safety” leads to excessive accidents. If everyone is responsible for safety, then nobody is. Of course, everyone should act in a way that keeps themselves and others safe. However, that does not lead to a coherent and effective safety program for the organization as a whole. Leadership and management are important. The overall requirements for achieving safety goals must be identified and assignments of responsibility, authority, and accountability made to ensure that the goals are achieved.

In addition, as with any effective management system, the appropriate assignment of responsibility, authority, and accountability is critical. Responsibility lies at every level of the organizational structure, although appropriate responsibilities will differ at each level.

The SMS responsibilities are often depicted as “pillars” under a common roof. There is no set number or type of pillars; they differ among industries. For example, expectations and responsibilities may differ for a plant that manufactures dangerous chemicals, an aviation transportation system, and organizations providing healthcare.

Figure 4 shows what we believe are the appropriate pillars for healthcare. HIT is not separated out because HIT is a part of all of responsibilities; dealing with it separately can result in missing the interconnections between the HIT and the other aspects of controlling safety and thus in missing opportunities for decreasing adverse events. That does not mean, however, that a group focusing on HIT safety and having HIT expertise is not necessary, only that it will report to a higher level having a wider scope and will require interactions with other components of the SMS. More about this in the next section.

There are four pillars shown. That does not mean to imply anything about where the responsibilities are assigned. The *Office for Patient Safety Management* should be coordinating all these activities throughout the organization, but the Digital Health Office may (especially at first) need to create these for HIT alone and report up to the higher level that exists to manage and coordinate all the efforts.

The four pillars are (1) Establishing Culture and Policy, (2) Controlling Hazards, (3) Managing Operations, and (4) Learning from Experience. Underlying all of the pillars is a comprehensive Safety Information System, which is used by everyone but managed by Pillar 2: Knowledge Management.

The basic activities in the pillars are first described, and then specific ways to put these into practice at the VHA, at least for HIT, is described. Clearly, not all of these activities can be introduced

immediately; a plan is needed to determine which should be done first and how changes can be most effectively and efficiently implemented.

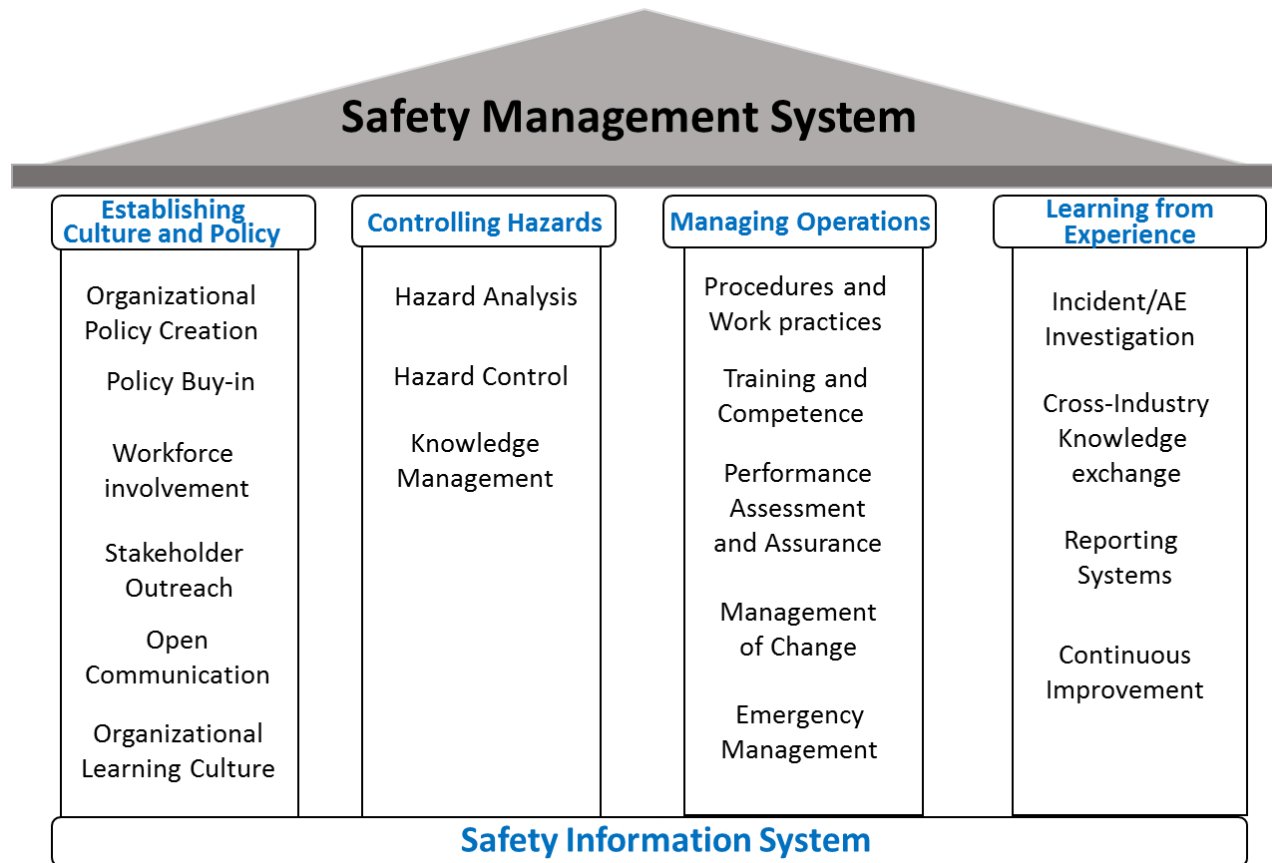


Figure 4: The Structure of a Basic SMS for Healthcare

**Pillar 1: Establishing Culture and Policy**

*Organizational Safety Policy Creation:* Responsible for (1) identifying values and assumptions to be used in decision making and (2) creating a safety policy that establishes those values.

*Safety Policy Buy-in:* Responsible for ensuring safety policy is disseminated throughout the organization, modeled by the leaders, and followed throughout the organization.

*Workforce Involvement:* Responsible for establishing a culture where the workforce fully participates in the safety program.

*Stakeholder Outreach:* Responsible for ensuring that stakeholders’ inputs are solicited and influence the safety program. Stakeholders include employees, local communities, regulators, industry groups, and shareholders.

*Communication:* Responsible for ensuring that communication channels are open and operating correctly.

*Organizational Learning Culture:* Responsible for the ensuring that the organization establishes a culture of learning and continual improvement with respect to patient safety; more specifically, creating a culture that ensures that the organization is continually learning from adverse events and close calls and also continuously improving the SMS using what has been learned. If losses are to be reduced over time

and healthcare organizations companies are not going to engage in constant firefighting, a process for continual improvement needs to be implemented. Identified flaws must not only be fixed, i.e., symptom removal, but the larger operational and development safety management systems must be improved, as well as the process that allowed the flaws to be introduced in the first place. The overall goal is to change the culture from a *fixing orientation*—identifying and eliminating deviations that are symptoms of deeper problems—to a *learning orientation* where systemic causes are included in the search for the source of safety problems [Leveson, 2023].

To accomplish this goal, a feedback control loop is needed to track and assess the effectiveness of the safety management system and its controls. Were hazards overlooked or incorrectly assessed as unlikely or not serious? Were some potential failures or design errors not included in the hazard analysis? Were identified hazards inappropriately accepted rather than being fixed? Were the designed controls ineffective? If so, why? When numerical risk assessment techniques are used, operational experience can provide insight into the accuracy of the models and probabilities used.

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.

## Pillar 2: Controlling Hazards

*Hazard Identification and Analysis:* Responsible for ensuring that hazard analysis is being used effectively. Hazard analysis is the process of identifying potential hazards—such as incorrectly dismissing an alarm, wrong site surgery, or incomplete or missing patient information—as well as generating the causal scenarios that can lead to the occurrence of these hazards. Responsibilities here include ensuring that the most appropriate and effective HA methods are being used, that those performing them are trained and competent, and that the HA process is being used and is used correctly.

Some healthcare organizations occasionally use what is called *Healthcare Failure Modes and Effects Analysis* (HFMEA).<sup>6</sup> FMEA was created in the late 1950's by reliability engineers to identify the probability of events that could result from failures of electronic components in military systems. While still occasionally used for hardware reliability failure assessments, it is too simplistic even for today's hardware and is not appropriate for complex systems that involve software and humans. The bottom line is that not enough is learned to justify the cost. More sophisticated tools were created for hardware systems in the 1960s (Fault Tree Analysis and HAZOP), but again they do not handle software and the human factors in modern systems [Leveson 2023].

A new hazard analysis technique, called STPA (System-Theoretic Process Analysis), has been developed and demonstrated to be highly effective for complex, adaptive systems like healthcare [Leveson 2012, Leveson and Thomas 2018]. A separate report has been generated that demonstrates its use on the VA Clinical Decision Support System (CDSS) [Thomas et al. 2025]. We have also demonstrated it on the hazards associated with healthcare laboratory data exchanges [Leveson et al. 2023]. Other recent uses in healthcare, by ourselves and others, include analyzing safety of operating room procedures, pandemic response, medication administration, radiology/radiotherapy procedures, medical devices, etc. It is widely used in most other industries.

For example, in a HIT CDSS, new rules may be created to recommend prescriptions if certain patient conditions are met. A hazard analysis would proactively identify how the new rules could inadvertently introduce new hazards, and the analysis would produce mitigations and constraints to prevent those hazards. More specifically, the hazard analysis may raise concerns about a proposed rule that triggers a popup alert with lengthy text and a single "OK" button. If the alert is often irrelevant, that is, the rule is too sensitive, then clinicians may quickly develop the habit of dismissing it rather than critically

---

<sup>6</sup> Sometimes this technique is more generally labeled Functional Failure Modes and Effects Analysis (FFMEA).

evaluating it. With no way to track “accept” vs. “reject” responses, HIT teams would remain unaware that the rule is ineffective, leaving it unfixed.

*Hazard Control:* Responsible to ensure that the identified causes of hazards (the potential scenarios leading to hazards identified by hazard analysis) are eliminated or mitigated and that the system is monitored to ensure that the implemented controls are effective.

In the HIT CDSS example above, mitigations could include designing the rule with more specific trigger conditions to prevent false positives and options to “Accept recommendation” and “Reject recommendation,” which can be tracked to understand the effectiveness of the rule over time.

*Knowledge Management:* Responsible for ensuring that the knowledge about hazards, their causes, and their prevention is maintained and available over time. For example, this knowledge may lead to identifying the rules to be used in the CDSS, etc. One critical requirement to allow safe operations is the recording of assumptions about (1) operations in the hazard analysis, (2) design rationale, and (3) links between the hazards and the specific design features used to control them. Without such documentation, the cost of reanalysis when contemplating changes during operations can be enormous and, in some cases, impractical.

For example, if a change to a complex system is desired or required, the maintainer will need to know if the change will impact a design feature purposely used to eliminate or mitigate an identified hazard. Most of the required information should be identified and documented in the original hazard analysis and design rationale, but it needs to be provided in a form that makes it easy to locate and use during operations.

Poor decision making often results from decision makers having an incorrect assumption of the risk being assumed. Decision makers therefore must understand how their decisions will impact the designed controls in the safety control structure.

One common problem is that informal risk assessment may change as time passes without a loss. The actual risk has probably not decreased, but our perception of it does. That leads to a change in priorities and in the resolution of conflicting goals.

One way to combat this erosion of risk perception is to provide ways to maintain accurate risk assessments in the minds of the system controllers. The better information controllers have, the more accurate will be their risk assessments and therefore the better their decisions. Accident analysis and various kinds of audits should involve a careful investigation of the accuracy of risk perception in the mental models of those controlling high risk projects. Unearthing the reason for this migration during an accident or incident causal analysis when unrealistic risk perception was involved can help to identify ways to design the safety control structure to prevent it or detect it when it occurs.

The Knowledge Management group will provide significant inputs to the SIS and might be named as the primary owner of the SIS.

### Pillar 3: Managing Patient Safety

*Procedures and Work Practices:* Responsible for the creation, auditing (for compliance), and updating of safe procedures and work practices. For a healthcare organization as diverse as the VHA, requiring standard procedures and practices may not be desirable in terms of the quality of healthcare provided. One concern is that too little flexibility leads to shortcuts and nonconformance. As shown in the accompanying Workflow Hazard Analysis paper [Leveson et al. 2025], it is possible to specify and verify the safety of abstract workflows from which multiple safe workflows can be generated.

When healthcare professionals do not follow standard written procedures, the response should not be forced compliance but instead an attempt to understand why they felt the need to deviate from the procedures and why they thought it was warranted. Healthcare workers may need to adapt procedures to deal with a changing system and to satisfy multiple goals, such as safety, efficiency and productivity.

Instead of forcing compliance with specified procedures, the best approach is to (1) monitor and identify the gap that exists between specified procedures and actual practice, (2) understand why the gap exists, and (3) update and rewrite the procedures accordingly.

*Training and Competence in System Safety:* Responsible for documentation, training, and auditing the effectiveness of both the procedures used and the training in them. Everyone in the safety management system, 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 safety control structure—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 following 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. Finally, decision makers must know how their decisions will impact the designed controls in the safety management system.

Training should include “why” as well as “what.” Understanding the rationale behind the safety rules they are asked to follow will help reduce (1) complacency, (2) what appears to be reckless behavior but to the person made perfect sense, and (3) unintended changes or shortcuts leading to hazards. The rationale includes understanding why previous adverse events occurred and what changes were made to try to prevent a reoccurrence.

People who interact with complex systems today need to learn more than just the procedures to follow: they must also have an understanding of the logic used in the automation they are using or supervising and in the safety controls. Learning about recent events and trends should be part of this training. Assessing the effectiveness of the training, perhaps through regular audits and performance assessments, can be useful in implementing an effective training improvement and learning process. RCA results are an important source of information about training effectiveness.

*Performance Assessment and Assurance:* Responsible for creating assessment and assurance activities, ensuring they are being performed, and identifying updates and improvements needed. Assessments can be used to provide important information about the effectiveness of training. The operational information system must be monitored to ensure that models and assumptions used during initial decision making and design are correct and that the system is constructed, used, and maintained in the manner assumed by the designers.

*Management of Change:* Most losses occur after some type of change. Adaptation and change are an inherent part of any system and are required for an organization to thrive. Because changes are necessary and inevitable, processes must be created to ensure that safety is not degraded by the changes. 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 efforts to prevent them. The goal is to allow change as long as it does not violate the safety constraints.

Common types of IT changes that must be controlled are:

1. Physical and electronic changes: new equipment is introduced, or the equipment may degrade or not be maintained properly. Software needs to be continually changed as knowledge increases and healthcare practices change. Examples are the (i.e., updating of CDSS rules, generic drug name changes, new diagnostic test parameters, etc.).
2. Human changes: Human behavior and priorities often change over time or humans may create workarounds.

3. Organizational changes: Changes in the safety controls and safety management system.
4. Physical or social environment changes in which the IT system is used or with which it interacts. These changes may be permanent or of limited duration.

Different procedures are needed depending on whether the change is planned or not.

*Planned Changes:* Responsible for the creation of a VHA Management of Change (MoC) policy and ensuring people know how to use it, it is being used (followed), and it is effective. Planned changes require at least a partial hazard analysis to determine if any new hazards are being introduced or previous controls become less effective due to the change.

*Unplanned Changes:* Safety depends on the accuracy of the assumptions and models underlying the design and hazard analysis processes. The information system must be monitored to ensure that these models and assumptions are not violated by unplanned changes in the IT system environment. Leading indicators can be designed to identify potentially hazardous changes. Hazard analysis results can be used to create effective leading indicators.

*Emergency (Contingency) Management:* Responsible for identifying potential emergencies, proactively preparing for them and designing contingency plans, ensuring that training has been done, and doing post mortems. An important part of contingency planning in healthcare today is the prevention of health information security breaches and proactive activities to reduce the potential impacts.

#### Pillar 4: Learning from Experience

The final pillar involves responsibility for improvement over time.

*Close Call and Adverse Event Investigation:* Responsible for learning from adverse events and close calls. No matter how good the hazard analysis efforts are, they will not be perfect, and investigation and response to prevent reoccurrence will be necessary. However, placing major emphasis on reacting to adverse events and close calls is not an effective way to manage safety in a complex, human-oriented, adaptive system like healthcare, even if the reactive efforts are very good. There are too many causes of adverse events to wait for each one to occur before efforts are initiated to prevent them.

Given a strong proactive hazard analysis program, there should be few adverse events. For example, in the aviation industry there are few accidents, but those that do occur along with potentially serious incidents (close calls) are investigated thoroughly, and changes are implemented to react to them. Learning from experience will also be necessary to identify new causes of losses as healthcare changes over time, for example, “fixes” that are no longer effective (or perhaps never were) or new types of events.

One important use of RCA is to improve proactive hazard analysis efforts. The results of adverse event investigation should be used to determine why the original hazard analyses did not identify effective controls (e.g, limits in the methods used, inadequate fixes, changes in the system and in the environment, complacency, etc.) or why the controls became ineffective over time.

Because RCA processes are not universal, we examined the process used at the VHA [VHA NCPS, 2025]. The process has some good features, but it needs to be strengthened. For example, it is unlikely to find causes related to HIT flaws. In our study of CDSS, we found that nobody could remember any CDSS flaw that had been identified by an RCA. There had been many flaws identified, but the reports came from users who voluntarily reported problems they had run up against, not from the formal RCA process. To be able to identify HIT-related causes of adverse events and close calls, RCAs must look beyond the simple cause–effect diagrams and single chain-of-events methods specified in the VHA RCA step-by-step guide [Leveson 2023, Leveson 2024].

Investigation requires more than just writing a report. There must be assignment of responsibility to ensure that appropriate measures are taken to strengthen any aspects of the safety management

system 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 the safety management process until another similar incident or adverse event occurs.

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 past incidents and accidents, 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 organizational processes 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.

Trained and professional investigation teams should be considered instead of inexperienced practitioners performing investigation of adverse events and close calls, as is currently the prevailing practice. For example, a number of leading U.S. healthcare centers have recently formed dedicated groups for adverse event investigations with appropriate training on how to do so effectively.

*Cross-Industry Knowledge Exchange:* Responsible for knowledge exchange not only cross-healthcare but cross-industries and national borders. Create communication channels to share knowledge.

*Reporting Systems:* Responsible for establishing and monitoring a reporting system that is easy to use, provides the necessary information to fix problems or at least identifies existing problems, and encourages reporting. At a minimum, the reporter should be told that their report has been received and that it will be evaluated and later they should get information about the result of that evaluation. If reporters feel that their reports have fallen into a black hole, they will stop reporting.

*Continuous Improvement:* Responsible for creating an effective process to ensure that the organization is continually learning from safety incidents and improving the SMS and the safety of the healthcare provided. SMS designs are rarely perfect from the beginning and the world changes over time. Responsible for establishing feedback channels, ensuring they are used and usable, and getting important feedback information to those that can respond to it. Special checks should be established to determine whether implemented improvements were actually effective.

### **Designing the Control Structure and Assigning Responsibilities**

The previous section described safety-related responsibilities that need to be assigned to appropriate places and individuals in the SMS. The overall goal of the safety management system is to assure informed decision making about safety. While groups are useful in generating alternative solutions to problems and communicating decisions to others (see *working groups* below), ultimate responsibility for decisions must be assigned to individuals.

Figure 3 shows a high-level view of a suggested SMS design. A next step is to create a more detailed design that suits the specific organization. For comparison, the current VHA SMS organizational structure can be found in Appendix A. Once again, there is no “correct” design of safety management systems, and the VHA should create their own design that matches their culture and history.

There are, however, some principles that should guide the design. These principles derive from the study of losses in the past, in healthcare as well as other industries [Leveson 2023]. One of these lessons learned is (as stated earlier) that accidents (losses) are very often associated with a fractured, organizationally dispersed safety staff and with ill-defined responsibility and authority for safety. Problems arise especially when responsibility is divided across organizational boundaries: There should be only one person in the organization with overall responsibility for safety. That person, who would

head the Office for Patient Safety Management, should report directly to the person responsible for making decisions about the organization as a whole (i.e., the Deputy Undersecretary of Health).

Basic principles that should underlie the general SMS design include:

- 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 filtered by groups with potentially conflicting interests. Some decision makers may also need fast access to information.
- Safety is involved in almost all organizational activities. Therefore, direct communication channels to most parts of the organization are required.
- 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.
- An appropriate part of the government (such as Congress) provides oversight. Oversight and input to the VHA Deputy Secretary may also be provided by VA advisory committees, which should have assigned responsibility for oversight of patient safety and the Office of the Inspector General of the VA.

Deputy Undersecretary for Health: The Deputy Undersecretary for Health (DUSH), is responsible and accountable for the achievement of the organizational goals. There are always multiple goals for any organization, and tradeoffs may need to be made when the multiple goals cannot all be fully achieved. The DUSH needs a way to get the information necessary to make informed decisions and the power or authority to implement those decisions. The DUSH also needs a way to ensure that those lower in the organization understand the principles (safety policy) to use to achieve the organizational safety goals.

Office for Patient Safety Management (OPSM): The DUSH needs to have the most pertinent and up-to-date information with which to make decisions. At the enterprise level, there should be one person responsible for patient safety who heads the OPSM and who reports directly to the DUSH to avoid time delays and communication channel blockages. The primary responsibility of this individual is patient safety in the VHA as a whole, that is, the implementation of the SMS pillars. Therefore, this person's primary responsibility, beyond providing safety-related information to higher levels, is to ensure that all the pillars are operating effectively (i.e., that all the required activities are taking place and that they are effective). This person must also be able to communicate directly with top management and provide input to all types of management decision making. These responsibilities imply that the person in this position must report directly to someone with influence (in this case, the DUSH) and be seen within the organization as having the support of senior management.

Some details:

- The head of the OPSM provides leadership and coordination to a group assigned to support him or her. 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 and that approach must be identified at the top of the organization and communicated throughout the entire organization. The OPSM integrates the information received from lower levels throughout the VHA and make sure it is available to all decision makers.
- The head of the OPSM should be supported by a group of people who have been assigned responsibilities for the activities in the SMS pillars. The Office for Patient Safety Management provides both the overall perspective of how the VHA as a whole controls safety but also the long-term vs. short-term view.

- Management of safety needs to be separated from management of other desired organizational properties. Many accidents have resulted from one group being responsible for multiple desired organizational qualities, such as safety and quality. The DUSH has ultimate responsibility for achieving all the desired system properties and making the decisions for short- and long-term tradeoffs. In the short term, safety may appear to conflict with efficiency, productivity, and other important system goals. But in the long term, safety improvement goes hand in hand with improvement of other goals, such as efficiency, productivity, quality of care, etc.
- Some have argued that lower levels of the management structure have greater information about the specific safety issues and that decision making about safety should be pushed down to the lower organizational levels. This philosophy, however, has been an important causal factor in many major accidents. At the lower levels of the control or management structure, short-term or narrow considerations may be more cogent, and the bigger picture is not available. In addition, the lower levels lack perspective: Although they have detailed information about their specific parts of the system, they do not have the equivalent visibility and information about other parts of the system. They also cannot anticipate or control the behavior of other components in the larger system. Responsibility for system safety decision making necessarily requires more visibility about the larger system state than is possible at lower levels of the management hierarchy. A decision that seems perfectly safe at one level can be seen to be dangerous for the organization or system as a whole at a higher level.
- There should be safety-related activities at all levels of the organization, coordinated by the OPSM. The high levels should have broad responsibilities for implementing the SMS pillar activities, with each successive level having more focused responsibilities appropriate to the level at which they operate. Each level must provide oversight of the level below by ensuring they are using appropriate procedures, that they are carrying them out correctly, and that the efforts are effective. There also usually needs to be a management focal point with responsibility for ensuring that the safety management system is designed and working properly at each level of the management system.

Creating such a comprehensive SMS for the entire VHA might prove to be too great a goal in one step. An alternative might be for the DHO first to create their own subset of a VHA-wide SMS and to use the experience obtained to learn how it might be extended to the larger VHA as a whole.

Communication and Coordination: An examination of the organizational factors leading to accidents often turns up communication problems. Not only goals and policies, but the reasons for decisions, procedures, and choices, need to be communicated downward in order to avoid undesirable modifications by lower levels and to allow detection and correction of misunderstanding and misinterpretation. In the feedback channel, the feedback from operational experience and communication of technical uncertainties and safety issues up the chain of command is crucial for proper decision making.

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 the safety control structure. Feedback is critical in managing any activity. Having access to too much information can be as bad as having too little when it overwhelms the ability to identify the most important information for effective decision making. 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 in this report.

Communication is also important in coordinating activities and responses to events. People with overlapping responsibilities need communication channels and ways to coordinate their activities to ensure that the 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.

One very effective means for communication and coordination devised by the U.S. Department of Defense (DoD) for development projects is *working groups*. DoD projects can span many years, 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 DoD, but there is a Prime Contractor that provides the system engineering and coordination among the subcontractors. For such a structure to work, communication is critical. Working groups have been successful in such coordination and communication efforts and can be adapted for less complex projects. Working groups can be equally effective in non-development environments such as healthcare.

A safety working group provides an interface between two hierarchical components of the safety control structure or between two or more components at the same level. Members of these groups are responsible for coordinating safety 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.

An enterprise working group may be composed of the safety managers for different divisions or programs (including the Digital Health Office) while, at the lower levels, working groups may be composed of representatives from safety groups at different VISNs, VAMCs, or CBOCs.

Collaborative environments can provide frameworks for stakeholders to come together, exchange ideas, and collectively address shared challenges. But the most effective path to future improvement involves leadership and the ability of one person or a small group to make the decisions necessary to propel improvements. Depending on large groups to make decisions and effectuate change is almost always less effective.

Assigning responsibility, and especially accountability, to a group (as opposed to an individual) is an ineffective way to achieve goals. Sometimes a natural leader arises within a group and assumes responsibility to achieve the group's assigned goals, but in other cases such leadership does not arise. Groups may be formed to advise the person with the ultimate responsibility to make decisions, but responsibilities and leadership need to be defined when forming the group.

With different groups making decisions influencing safety at all levels of the organization, some person or group is required to integrate the information and make sure it is available to all decision makers. Within each VISN, for example, there may be a person responsible for safety for all the facilities that are part of the VISN that reports up to the OPSM. A working group, composed of the safety managers for each VISN, might be useful to provide communication and coordination among the VISNs.

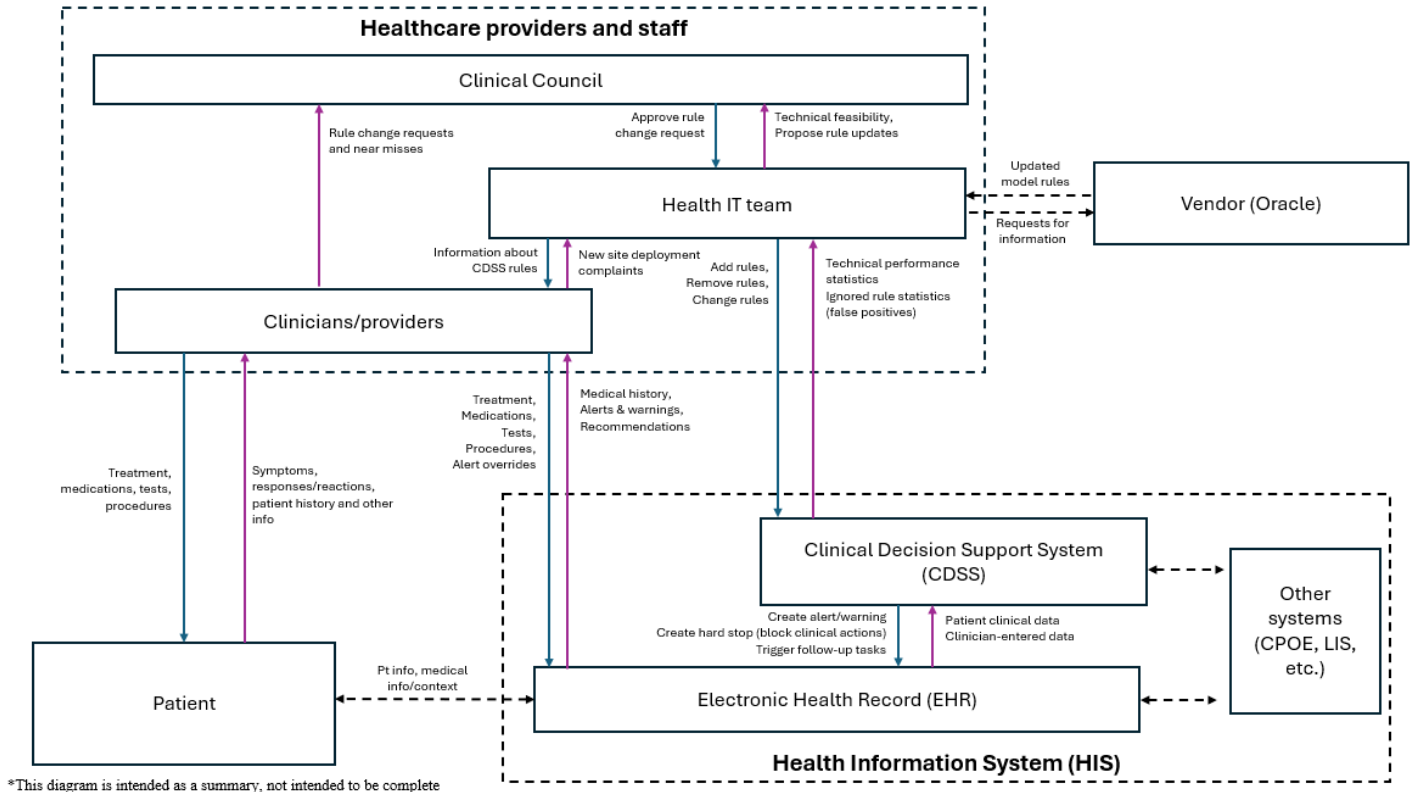
#### The Digital Health Office Safety Control Structure Design:

What about the SMS for the Digital Health Office in particular? As noted earlier, HIT safety cannot and should not be isolated from overall patient safety. However, a responsibility of Digital Health should be to provide information to the Office for Patient Safety Management and to use the information provided by the office to improve the safety of HIT.

The authors of this report have not specified a detailed design for Digital Health. STPA could be used, however, to provide information to assist in this process. For example, we performed an STPA analysis on CDSS, which identified flaws in the current CDSS management structure that could lead to adverse events [Thomas et al. 2025]. We were surprised to find that 93% of Chief Medical Information Offices (CMIOs) report experiencing at least one CDSS malfunction in the last year. Most of these were overlooked by the testing and were only caught by voluntary reports by those who experienced them.

This result is not surprising as testing is known to not be very effective in finding information system errors.

In applying STPA to this problem, we first created the current safety control structure for CDSS, shown in Figure 5.



We then identified unsafe control actions by those in the CDSS control structure and the scenarios that could lead to these unsafe control actions. The details are in the report, but from the scenarios we identified specific flaws and general systemic safety control factors that could lead to adverse events related to CDSS. All of these were reviewed by VHA CDSS experts to determine that they were plausible.

The systemic flaws include:

**Delayed/Missing Feedback about the Model Rules:** A recurring communication gap exists between the EHR vendor and the health IT team. The vendor does not consistently provide timely or reliable notifications regarding the release of new model rules or updated versions of existing rules. As a result, critical updates are not always implemented, creating the potential for outdated or incomplete model rules to remain active within the clinical environment. In addition, near misses have been observed in cases where changes to event sets, order names, or code sets introduced incompatibilities that disrupted the function of implemented model rules.

- **Lack of Customization of the Model Rules:** Vendor-provided model rules often do not conform to the VHA's established standardization framework. In practice, the health IT team is frequently required to "fix" model rules by developing supplemental rules to close functional gaps. This then increases the overall number of rules and introduces additional complexity into the rule catalog. When subsequent vendor updates are released, these downstream modifications can themselves break or conflict with the updated model rules.

- Reactive Processes rather than Proactive Ones: There is no systematic or proactive monitoring in place to identify rule degradation, logic failures, or gaps in coverage prior to the occurrence of a clinical impact. As a result, issues are often discovered only after a clinician has encountered an alert failure or inconsistency during patient care. Moreover, the current audit report is focused on technical measures, not clinical effectiveness, so there is no long-term analysis of whether rules are firing appropriately for the clinical environment. Finally, effective hazard analysis is not done by the vendor, and there is no government requirement for it to be done. The same is true for sophisticated human factors analysis.
- Unused and Missing Feedback: The responsibility for evaluating rule performance data, which contains important clinically relevant insights, is outside the health IT team’s current scope and resource capacity. Therefore, feedback about the performance of rules is not available. The absence of governance and sustained data stewardship creates a systemic blind spot where ineffective rules may persist undetected, and opportunities to improve rule fidelity and reduce clinician burden are not pursued. Audit tables and rule logs are temporary—disappearing after only 30 days—so, diagnosing rule issues becomes difficult after 30 days.
- Resource Constraints: The current 4000 rules are being managed by only three to four people. High turnover and attrition, along with organizational changes through reorganization and policy changes results in potential loss of institutional knowledge. Manpower limitations are exacerbated by technical infrastructure limitations.
- Competing Priorities: Although rules have to be standardized to a certain extent across the VA and DoD, different stakeholders have disparate and competing interests.
- Lack of Clinical Rule Ownership: There is no tracking of changes, such as clinical guidelines, that affect rules; no continuous monitoring of clinical relevance and only perhaps a periodic reassessment; no checks of rule effectiveness, such as false-positive and false-negative rates; and no checks for unintended consequences after implementation. There is a handful of specialties with “superusers,” but this ownership is fragmented, purely informal and/or voluntary, and not clearly documented.
- Missing/Inaccessible Feedback about Rule Intent and Rationale: Intent and rationale of CDSS rules are carefully documented, however, this information cannot be recovered or referenced after deployment.
- Inadequate Learning from Past Events and Root Cause Analyses (RCA): The health IT team does not receive structured feedback or communication from root cause analyses (RCAs), even in cases where adverse events are directly attributable to CDSS rules.

From this information, we were able to identify improvements that would significantly increase safety. The full set is in the accompanying report, but examples of recommended improvements include:

- Strengthen vendor communication about model rule changes.
  - Improve vendor processes to notify VA when model rules are updated.
  - Investigate the reasons for the communication gaps and establish a simple and easy-to-use way for the vendor to notify of any model rule change.
- Improve proactive monitoring processes.
  - Move beyond purely reactive change requests; develop systems that actively monitor for external changes (e.g., guideline updates, new drug codes, position changes).
  - Consider automated monitoring tools that link rules to source guidelines or vocabularies, allowing for proactive identification of affected rules.
- Strengthen coordination between sites, councils, and deployments.

- Prevent rule divergence across sites by improving coordination between councils, deployment teams, and DoD partners.
- Avoid redundant or duplicate rules across VA sites.
- Create ways to monitor rule effectiveness (e.g., override rates, false positives).
  - Analyze how often rules are fired and overridden to identify ineffective or burdensome rules.
  - Consider ways to collect feedback from clinicians about rule usefulness (e.g., “this rule was useful” button).
  - Use this feedback to guide rule refinement and reduce alert fatigue.
- Build tools to extract and analyze rule metadata: This approach has already shown success for pharmacy teams and should be expanded.
  - Continue expanding reporting tools that allow solution teams to proactively analyze rules (e.g., searching for drug codes, facilities, positions affected by rules).
- Ensure that the health IT team is appropriately staffed and able to not only react to tickets and build rules based on tickets, but also to be able to monitor what’s already in the system. Create a team for cross-collaboration.
- Establish formal ownership for each rule.
  - Assign responsibility to specific clinical teams or solution groups for monitoring and maintaining rules after implementation.
  - Ensure each rule has a clearly documented owner responsible for monitoring relevant clinical guideline changes.
- Improve tracking of organizational responsibility over time.
  - Account for organizational changes that make rule ownership ambiguous over time.
  - Create a way to track ownership that survives reorganizations, staff turnover, and changing structures.
- Improve RCA integration with CDSS review.
  - Ensure RCA investigations examine whether CDSS rules contributed to adverse events or could prevent similar events in the future.

Implementing these recommendations would go far toward improving the DHO safety control structure. The next step might be to extend the STPA analysis beyond CDSS to all of HIT at the VHA and use the results to assist in making additional improvements in the safety control structure. Part of the effort should be to evaluate how existing tools (such as PIQI, DeX, and Komet) can be used to collect crucial feedback information and to monitor the operation of the SMS. An STPA analysis could identify the most important uses for these tools.

### The SMS and COTS

Managing Digital Health safety in a COTS (commercial off-the-shelf) environment is difficult. While COTS software provides many advantages for governmental organizations, giving up control to a commercial organization can lead to serious problems. While COTS software is common in healthcare, responsibility for safety cannot be delegated to a commercial vendor as it has important conflicts with their most basic goal of profitability. This fact implies that some type of rigorous oversight and control is needed. The VA is ultimately responsible for patient safety in the VHA system.

This problem of inadequate control over COTS safety is not new. In radiation therapy, as an example, professional societies have stepped up to provide the control over safety that the FDA did not or could not exercise. These professional societies created standards that healthcare organizations can use in selecting the equipment to be purchased. Customer pressure has had a significant impact on the companies providing this equipment. The American Medical Informatics Association (AMIA) might be an appropriate professional group to step in for general HIT systems.

### Quantitative vs. qualitative Information in the control of safety

Quantitative data, particularly aggregated data, is the easiest to collect and track. The VHA collects a lot of aggregated quantitative information. The important question is whether this data provides the information needed to adequately control patient safety. We can get so focused on numbers that we lose the original goal of the oversight, which is to identify specific hazards and control them. Qualitative analysis, such as the specific causal factors leading to adverse events and close calls, is the best way to determine what changes will be most effective in changing the aggregated numbers. The numbers themselves do not provide the information needed.

### **Safety Information System (SIS):**

Creating a comprehensive and useful safety information system is a large investment, but it has been found to be second in importance only to upper management commitment in reducing losses [Kjellan 1982]. The most practical approach is to start with the most useful information and add slowly so that the long-term cost is reasonable. A lot of important information already exists, such as the intent and rationale behind the rules in the CDSS. This information, however, is difficult to find when needed, so reorganization to ensure easy retrieval is all that is required. Another important place to start might be the documentation and tracking of hazards and their resolution. The very experienced DHO at the VHA would be an ideal group to be responsible for developing and managing the SIS.

Each pillar contributes information to the SIS and retrieves the information needed to carry out their responsibilities. HIT provides both the storage of and access to general VHA healthcare safety information, but information about the operation of the HIT itself is also important to include in the SIS.

In general, SIS information must be collected, analyzed, usable, and accessible to those who need it when they need it. All hazards need to be recorded, not just the most critical; otherwise there is no record that a particular condition has already been evaluated. A complete hazard log and audit trail will show what has been considered and why decisions were made. Because hazard information may become very large, ways to prioritize and summarize the most significant information are usually necessary and useful for reviews and management oversight.

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 over time. After major losses, it is often found that the information to prevent the loss existed but was not used or was not available to those involved. Often lots of information is collected only because it is required for government reports or is easy to collect and not necessarily because it is important for the operation of an effective SMS.

Information may be collected by organizations or by industries; usually both occur. The sharing of information across an industry and national borders can increase knowledge about hazards and about effective and ineffective control measures. Even within a single organization, an SIS can provide valuable information and feedback about the hazard analysis process and about the need for additional controls or modifications.

The SIS information system should include information not only about losses but also close calls, which can warn of and provide the information necessary to prevent losses in the future. Civil aviation and nuclear power are industries where “near misses” are widely reported, and the information is used to improve safety.

No matter how the information is collected, understanding its limitations is important.

*Collection:* To be most useful, safety information must be accurate and timely. Data may be distorted by the way it is collected because of systematic filtering, suppression, and unreliability. Checklists tend to focus attention on limited categories of conditions and factors. The most useful data does not just focus on proximal events and actors but instead focuses on the systemic factors involved, such as management problems or organizational deficiencies. 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 as a serious causal factor.

*Analysis:* Simply collecting the information, of course, is not enough: it needs to be analyzed and summarized. Problems may arise from the difficulty of systematizing and consolidating a large mass of data into a form usable for learning. Raw, quantitative data can be misleading.

One common problem arises from the ability today to use automation to collect large amounts of data. Airlines today are suffering from this type of data overload due to large-scale automated aircraft and flight information collection. As a result, the safety information system may contain only summary statistical data that can be easily processed by a computer and not the information about specific hazards, trends, and changes over time that is needed to learn from events before major losses occur.

Hazard analysis tools can help not only to identify what types of data need to be collected but to provide guidance on the importance of the events that are occurring. Data to evaluate the hazard analysis results and to identify causal factors that were thought to be eliminated or controlled should be part of the information collection and analysis process.

*Dissemination and Usability:* Disseminating information in a useful form may be the most difficult part of an SIS. Data is not the same as information; the data needs to be processed and presented in a form that people can learn from and apply to their daily activities. Hazard analysis results or information from practice comprising hundreds or even thousands of pages, perhaps in a tabular or graphical form, is not going to be very helpful in providing answers to individual questions.

To assist in dissemination, information should be integrated into the environment in which safety-related decisions are made. For example, during an RCA investigation, it should be possible to determine if other, similar events have occurred.

## **A Path Forward:**

In summary, effective safety management requires:

- Commitment and leadership at all levels;
- A strong organizational safety culture with a clearly articulated safety vision, values, and procedures;
- Stakeholders with partnership roles and responsibilities;
- A safety control structure with appropriate assignment of responsibility, authority, and accountability at all levels of the organization;
- Feedback channels that provide an accurate view of the state of safety at all levels of the safety control structure;
- Integration of safety into the provision of healthcare and not just an independent group or separate subculture;
- Individuals with appropriate knowledge and skills;
- A designated process for resolving conflicts between safety and other priorities;
- Risk awareness and communication channels for disseminating safety information;
- Controls on system migration toward states of higher risk, particularly due to unplanned changes;
- An effective and usable safety information system;
- Continual improvement and learning; and

- Education, training, and capability development.

Most of these activities are not costly. They simply involve switching VHA resources from current activities that are not directly reducing adverse events. There are a few startup costs, mostly those involved in education, but they will allow a reduction in total costs over time. The most costly activity, but probably the potentially most important, is creating a comprehensive safety information system containing the information that is needed to make significant improvements in safety. However, the creation of an effective SIS can begin slowly with the most important information added first and then the SIS augmented over time. We identified what that information is for CDSS in our accompanying report.

Improving the management structure itself can start with the information provided by our STPA CDSS analysis. Extending this same modeling and analysis process to other parts of the HIT safety management system could be another early step toward creating a comprehensive HIT SMS.

A final goal should be to improve the safety culture. Simply considering safety to be important is not enough. First, a safety policy, at least for HIT, needs to be created. Second, the differences between safety and other qualities, such as reliability and repeatability, need to be thoroughly understood or precious resources will be spent on things that do not improve safety. Third, a proactive (vs. reactive) culture needs to be created by adopting the use of powerful hazard analysis tools. A final critical cultural change involves the understanding that a focus on limiting liability is very likely to increase the number of adverse events, and ways to make tradeoffs between these goals need to be established.

## References

Baker, Elizabeth White (2022), *Safety in Hospital Medication Administration Applying STAMP Processes*, S.M. Thesis, System Design and Management, MIT.

Baker, Elizabeth White (2025), Using STPA to Analyze Specimen Handling and Tracking in the OR, in preparation.

Bertalanffy, Ludwig von, (1934), Untersuchungen über die Gesetzlichkeit des Wachstums. I. Allgemeine Grundlagen der Theorie; mathematische und physiologische Gesetzlichkeiten des Wachstums bei Wassertieren. *Arch. Entwicklungsmech.*, 131:613-652.

Bertalanffy, Ludwig von (1968), *General System Theory: Foundations, Development, Applications*, New York: George Braziller, revised edition 1976: ISBN 0-8076-0453-4

Chen, S., Khastgir, S., and Jennings, P. (2021), Analyzing National Responses to COVID-19 Pandemic Using STPA, *Safety Science*, vol. 138, June, <https://doi.org/10.1016/j.ssci.2021.105195>

Davis, Mickey (2025), Unlocking the value of IT in healthcare, downloaded from <https://www.nordicglobal.com/blog/unlocking-the-value-of-it-in-healthcare#:~:text=Amid%20ongoing%20pressure%20on%20resources,only%20contribute%20to%20this%20trend.>

Ghorbani, J., Marquez, M., and Samedy, P. (2024), *Application of CAST in Site Identification Safety in Interventional Radiology (IR) at Memorial Sloan Kettering Cancer Center*, International STAMP Workshop, Sept. 23, 2024  
Downloadable from <http://psas.scripts.mit.edu/home/mit-stamp-workshop-presentations/>

Kjellman, Urban (1982), An evaluation of safety management systems at six medium-sized and large firms, *Journal of Occupational Accidents*, 3:273–288.

Leveson, Nancy G. (2012), *Engineering a Safer World*, Cambridge, MA: MIT Press.

Leveson, Nancy G. (2015), A systems approach to risk management through leading safety indicators, *Reliability Engineering and System Safety*, 136:17-34, April, <http://dx.doi.org/10.1016/j.ress.2014.10.008>

Leveson, Nancy G. (2023), *An Introduction to System Safety Engineering*, Cambridge, MA: MIT Press.

Leveson, Nancy G., Baker, E., Harrington, P., Kim, E., Powell S., Thomas J. (2025), Workflow Hazard Analysis, submitted for publication.

Leveson, Nancy G., Thomas, J., Rose, R., Harrington, P., Powell, S., and Keller, A. (2023) System Safety within Laboratory Data Exchanges, <https://psas.scripts.mit.edu/home/wp-content/uploads/2023/10/System-Safety-within-Laboratory-Data-Exchanges-Report.pdf>

Leveson and Thomas (2018), *STPA Handbook*, downloadable from <http://psas.scripts.mit.edu/home/books-and-handbooks/>

Leveson, Nancy (2024), CAST for Healthcare Handbook, downloadable from <http://psas.scripts.mit.edu/home/books-and-handbooks/>

Martinazzo, A., Martins, L.E. and Cunha Sebastião Vagner Aredes, T. (2021), Safety Analysis of a Low-Cost Insulin Infusion Pump Using STPA: A Case Study with Brazilian Company, International STAMP Workshop. Downloadable from <http://psas.scripts.mit.edu/home/mit-stamp-workshop-presentations/>

Pawlicki, T., Samost, A., Brown, D., Manger, R., Kim, G-Y, and Leveson, N. (2016), Application of systems and control theory-based hazard analysis to Radiation oncology, *Journal of Medical Physics*, 43(3):1514-30, Mar. doi: 10.1118/1.4942384

Leveson, N., Thomas, J., Harrington, P., Rose, R., Powell, S., and Keller, A. (2023), System Safety within Laboratory Data Exchanges, Report for the U.S. FDA, September. Downloadable from <http://psas.scripts.mit.edu/home/wp-content/uploads/2023/10/System-Safety-within-Laboratory-Data-Exchanges-Report.pdf>

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

Raman, J., Samost, A., Leveson, N., Dobrilovic, N., Oldham, M., Dekker, S., and Finkelstein, S. (2016), When a checklist is not enough: How to improve them and what else is needed, *Journal of Thoracic and Cardiovascular Surgery*, 152(2):585-592, doi: 10.1016/j.jtcvs.2016.01.022

Samedy, P., Chu, B., Bellamy, M., Ghorbani, J., O'Keefe, D., and Marquez, M. (2022), Introducing STPA to Interventional Radiology Within a Large Hospital (Memorial Sloan Kettering Cancer Center), International STAMP Workshop, June 7. Downloadable from <http://psas.scripts.mit.edu/home/mit-stamp-workshop-presentations/>

Samost-Williams, A. and Nanji, K.C. (2020), A systems theoretic process analysis of the medication use process in the operating room, *Anesthesiology*, 133; 332-3410.

Samost-Williams, A., Sinyard, R.D., Tabayoyong, L.L., and Nanji, K.C. (2025), Application of a systems theory-based accident analysis technique to perioperative safety reports from the COVID-19 pandemic, *Journal of Patient Safety*, June, DOI:[10.1097/PTS.0000000000001372](https://doi.org/10.1097/PTS.0000000000001372)

Shein, Edgar (1985), *Organizational Culture and Leadership*, San Francisco, CA: Jossey-Bass Publishers.

Silvis-Cividjian, N., Verbakel, W., and Admiraal, M. (2020), Using a systems-theoretic approach to analyze safety in radiation therapy—first steps and lessons learned, *Safety Science*, 122:105-419, <https://doi.org/10.1016/j.ssci.2019.104519>

Thomas, J., Kim, E., Harrington, P., and Leveson, N. (2025), STPA Analysis of an Automated Clinician Decision Support System (CDSS), in preparation, July.

VA Inspector General (2024), VA Needs to Strengthen Controls to Address Electronic Health Record System Major Performance Incidents. VA OIG, 22-03591-231, downloadable from <https://www.vaoig.gov/sites/default/files/reports/2024-09/vaoig-22-03591-231.pdf>

Wong, L. and Pawlicki, T. (2022), A system-based operational assessment of external beam radiotherapy, *Medical Physics*, 49(7), July. <https://doi.org/10.1002/mp.15704>

## **Appendix A: The Current VHA Safety Control Structure**