

The Physics of System Safety

T.J. Candy

DevSafe Consulting, LLC

Keywords: system safety, MIL-STD-882, software safety



ABSTRACT

This paper summarizes the observations and experiences of a retired Product Safety Principal Engineer with over 40 years of industry experience, including multiple decades in both Design Engineering and System Safety roles. The author holds a Bachelor of Science in Mechanical Engineering (BSME) degree from Michigan State University (MSU) and a System Safety Certificate from the University of Southern California (USC) Viterbi School of Engineering Aviation Safety and Security Program. He is currently a system safety consultant and teaches safety certification courses at USC in Los Angeles to students from around the world.

From this engineering/safety perspective, the author presents a view of the foundational basis for the modern system safety discipline and profession by postulating “The 5 Laws of System Safety”. Also, key observations on the current state of the system safety discipline and profession in industry today which represent departures from one or more of these laws are also presented.

The main thesis of this presentation contends that just as physics guides and bounds the engineering discipline, the foundational principles of the modern approach to system safety, established in the post-WWII expansion in aerospace applications, also guides and binds the practice of the system safety discipline.

This paper is intended to help facilitate self-examination for both individuals and organizations involved in the practice of system safety to compare their perspective with the observations presented in this paper. It builds on a paper (System Safety – A Profession in Crisis) presented by the author at the 2025 ISSS Summit and Training in Austin, Texas. Both the presentation and subsequent YouTube podcast were very well received and received positive feedback from attendees and viewers.

In addition to presenting the author's learnings and perspective, this paper hopes to galvanize both individuals and organizations from very diverse domains to align their core system safety efforts with these “5 Laws” so that the challenges of the present and future can be met with the right safety processes based on concepts which revolutionized product safety performance after WWII.

Contents

1. Background	2
2. Ground Rules & Assumptions	4
3. The 5 Laws of System Safety	5
a. Zeroth Law – Foundation	6
b. 1 st Law - System Safety Goal	6
c. 2 nd Law – Safety Requirements	7
d. 3 rd Law – System Safety Skillsets	9
e. 4 th Law – System Safety Tasking	9
f. 5th Law – Safety Competency	10
4. Observations	12
a. Safety Profession Crisis	12
b. System Safety Disambiguation	13
c. Risk Managing Authority Expansion	15
d. Software Safety Discussion	15
e. The Path Ahead	18
5. Conclusions	19

1. Background

At the start of my professional career as a System Safety Engineer in 2004, many aspects of the system safety process and discipline were brand new to me. I had spent most of the previous two decades as a military gas turbine design engineer, so my initial experience learning the system safety process was ‘on the job’ but working with a product I was quite familiar with.

My interest in joining the safety team was based on prior interactions with flight safety personnel during investigations involving my hardware. I didn’t have a practical understanding of the system safety process and discipline beyond its implicit charter to avoid accidents. But I liked the idea of avoiding accidents.

My initial assignments were supporting fielded products under the watchful eyes of highly experienced safety colleagues. I was able to be productive running risk assessments before I started formal system safety training.

Eventually I was assigned to a new product development program under a MIL-STD-882 based military contract. I was faced with a whole slate of new concepts, taskings and deliverables I had not previously encountered.

It was when I started to receive formal system safety training that I gained a more foundational knowledge of MIL-STD-882 and the safety process in general.

Learning the System Safety Process

Over the course of the following 17 years, through trial and error, I was able to develop a deeper understanding of the system safety process and discipline as well as MIL-STD-882 itself. I was able to work out a safety process and hazard tracking tool “from scratch” which was tailored to our product line. By the time of my retirement in 2024, this process and tool was being used on multiple advanced military product development programs. I helped create a new standard practice guide to be used for all future military development system safety programs that are still in use at the time of this paper.

The hazard tracking tool is compliant to MIL-STD-882 Task 106 requirements and featured many integration points with other safety and reliability tools. The tool predated the Model Based Systems Engineering (MBSE) approach but encompasses many concepts in line with this approach.

It was while receiving my System Safety Certificate from the University of Southern California Viterbi School of Engineering Aviation Safety and Security Program in 2016 that my questions, observations and feedback during some of the classes that the novelty of the process and tool I had created began to be revealed.

It was during one of these classes that I was recommended to come join the USC program as an instructor, which I did in 2021.

Historical Basis




Figure 1 - Engineering For Safety, Stieglitz, 1947

A few years ago, I came across a paper titled “Engineering for Safety” (Stieglitz, 1947). written in 1947 by a Republic Aviation Engineer named William I. Steiglitz. In it, he describes a safety process developed to address the challenges of advancing technology which has many parallels with concepts which made their way into MIL-STD-882.

This paper describes a structured System Safety effort which engages working groups at multiple organizational levels. It also describes the concept of independence of the safety function as well as establishing System Safety as a unique discipline distinct from classical engineering competencies.

I find the content of Mr. Stieglitz’s paper well aligned with MIL-STD-882 as well as my view of the safety discipline and profession.

Acquiring the Mission

At a point close to my retirement, all of my experiences and perspectives came together in my mind in what I can only describe as “an epiphany” on the state of the discipline and profession of system safety in industry. That being a subtle, consistent and widespread deviation from what I interpreted as the founding principles of the modern system safety approach.

This deviation is best summarized as an erosion of the independence of the system safety function and subsuming the safety function back to engineering. This was the basis of the paper I presented at the 2025 International System Safety Society (ISSS) Summit and Training in Austin, Texas entitled “System Safety – A Profession in Crisis” (Candy, 2025).

In this paper, I describe what I observed as an erosion from the foundational charter of the system safety discipline and profession. The response to the presentation, and a subsequent podcast on the same topic, has been overwhelmingly supportive, despite the concerning title.

This paper continues that discussion by presenting a foundational perspective which can be used as a gauge to measure this erosion both from a personal as well as organizational perspective.

Describing the Approach

My approach in this paper is to present the system safety discipline from an engineering paradigm. That just as in the Engineering discipline, there are fundamental “physics” that guide the system safety discipline.

The premise is that to the extent these “physics” principles are not followed is the extent to which the safety effort will fail to meet its intended function. Resulting in the return to the pre-WWII “break/fix” approach to safety management.

While the physics behind many of the engineering marvels we see in the marketplace today are more readily mathematically oriented, the “physics” associated with the system safety process follow similar “cause and effect” outcomes.

In other words, violating them brings predictably undesirable results just as violating laws associated with thermodynamics and aerodynamics do the same.

Caveat

As stated in my previous presentations, this paper should not in any way be construed as a critique of the GE Aerospace safety effort. While the bulk of my Engineering and System Safety experience occurred during my employment there, the information presented in this paper comes from a global perspective reinforced by my experience teaching System Safety at USC.

It could be well said that the learnings contained in this paper are a clear example of GE Aerospace’s mission statement to “invent the future of flight”. I was given the freedom to further develop the system safety process as well as write the standard practice guide based on what I’ve learned. I was encouraged to pass along my learnings

to industry by their full support of my serving as an instructor by teaching at USC while still an active GE Aerospace employee.

As a further example of their “walking the talk” on Safety, was when I submitted a series of software patent applications for the HTS tool I had created. All applications were denied by the GE Intellectual Property group citing a long-held corporate conviction of “we don’t compete on safety”.

Any use of the materials presented in this paper used to malign GE Aerospace (or any other entity for that matter) would be evidence of a failure to grasp the core message it is intending to present.

2. Ground Rules & Assumptions

The importance of the GR&A in the conduct of the system safety process, particularly in the product development phase, cannot be overstated.

As with any safety document, a clear declaration of the assumptions made in the presentation of the materials needs to be made. Here are some assumptions associated with this paper.

- The reader has some level of familiarity with the concept, history and/or practice of the system safety discipline. Experience with MIL-STD-882 safety process is helpful but not required.
- The “modern” approach to the system safety discipline is based on lessons learned in the post WWII ICBM development and are encapsulated in following two tenets:
 - The need for independent safety oversight to inform critical product decisions.

- The need for a proactive/predictive methodology to identify and mitigate risk as early in the design process as possible.
- The term “independent oversight” is interpreted as the need for a safety function which is organizationally independent of the engineering or project organizations to perform oversight of safety-related product decisions.
- MIL-STD-882 captures the intent of the modern system safety discipline and provides a suitable guide for execution of the system safety process across a wide variety of product domains throughout the full product lifecycle.

It is acknowledged that:

- MIL-STD-882 provides concepts and not techniques for the definition and implementation of a system safety process.
- MIL-STD-882 must be tailored to the specific product and in alignment with the current product lifecycle.
- While many sources have nuanced versions of the core activities of the System Safety process, for the purposes of this paper they are summarized as:
 - a) Conduct safety assessments as needed to support product decision making process.
 - b) Support identification of risk reduction action(s) necessary to meet established product safety metrics and goals.
 - c) Providing documentation verifying implementation of any necessary hardware and software action(s) to a regulatory function.
 - d) Conducting periodic audits as needed to assure effectiveness and

ongoing compliance to established safety actions.

- The Software Safety section will not address the role of Artificial Intelligence (AI) in the safety process in any substantive way

Any disagreement on one or more of these GR&A items can affect the applicability, utility and/or interpretation of the information contained herein. Direct conversations on any of these items with the author are encouraged.

3. The 5 Laws of System Safety

Just as the laws of physics guide the ongoing application of engineering disciplines (statics, dynamics, kinematics, thermodynamics, etc.), this paper asserts that the discipline of system safety likewise carries with it a set of laws which govern the application of the system safety discipline.

These 5 laws presented here summarize and capture what I believe are the fundamental learnings which were validated by the safety improvements in aerospace in the late 1960's and 70's and may provide guidance to understand more recent negative experiences in safety risk management.

These laws were created by combining my engineering and system safety observations and experiences. The intent in postulating these as laws is to provide a useful gauge for system safety organizations and/or individual professionals to measure their organization and/or tasking against the foundations of the modern system safety process.

Let's begin. Just as the laws of thermodynamics has a Zeroth law which is foundational to the study of how energy

flows in relationship to the thermal environment, so system safety has a zeroth law which is foundational to the management of safety risk.

a. Zeroth Law – Foundation

The consequence of violating any of the System Safety laws is a default to a reactive System Safety Management approach commensurate to the degree of deviation.

There are three fundamental approaches to safety risk management:

1. Predictive - The identification of safety hazards through analysis, their cause(s) and risk reduction actions prior to a loss event occurring.
2. Proactive - The examination of product usage information to identify risk reduction actions prior to a loss event occurring.
3. Reactive - The examination of a loss event and help identify action(s) necessary to preclude a recurrence (default safety management approach).

Prior to the discovery of the modern approach to System Safety in the post WWII era, the predominant safety management approach was reactive (“after the fact” also known as “fly, fix, fly”, “break/fix” or “fly around and find out (a new FAFO)”).

In this approach, the safety function operated primarily as emergency room (ER) responding to crises rather than a doctor’s office addressing issues before they reach a crisis stage.

Entities such as the FAA, ASME, DOT and others were created after tragedies had occurred to provide some level of assurance that the same accident could be avoided in the future. While this could accurately be seen as a form of proactive (“before the

mishap”) safety management, the ever-escalating demands of new technologies and processes meant a predictive (“during development”) approach was needed.

The MIL-STD-882 200 Series tasking provides a conceptual framework for using different forms of analysis to discover vulnerabilities which could be assessed for safety impact and addressed as early as the conceptual stage of development. Allowing for a predictive safety management approach.

b. 1st Law - System Safety Goal

The fundamental goal of the System Safety discipline is to provide independent oversight of critical product and/or process decisions as early in the lifecycle as possible in order to enable a proactive or predictive Safety Management approach.

It is the assumption of this paper that fidelity to these two principles; independent oversight and early lifecycle engagement, is what led to the vast improvement in safety performance in the Aerospace industry of the 1970’s.

Independent oversight proscribes a separation between the safety function and the engineering/project functions to assure safety input is as unbiased by pressures specific to the engineering/project functions as possible. It is the contention of this paper that the Safety function is intended to inform critical product decisions, not make them. Which is in the purview of the other functions.

Erosion of this independence since the inception of the modern system safety process has been evidenced as early as the 1980’s with the Challenger disaster and continues to this day, which is discussed in more detail in my earlier paper.

One consistent observation from my students is the number of occasions the safety team is first engaged well into the product lifecycle.

MIL-STD-882 Task Section 200 (Analysis) lays a progression of safety analysis concepts which allow specific forms of safety analysis to be conducted as early as the conceptual stage.

It is intuitive to understand that as a product matures, the cost and complexity required to implement changes can increase significantly. This is illustrated graphically for the military procurement cycle in Figure 3.




Figure 3 - Safety in Procurement Cycle

The farther a product design matures from the concept stage, the economics associated with design changes can drive program and engineering leaders to consider less complex solutions. This is captured in the MIL-STD-882 Order of Precedence (Figure 2).

Mitigation Action Order of Precedence

- Eliminate hazards through design selection
- Reduce risk through design alteration
- Incorporate engineered features or devices
- Provide warning devices
- Incorporate signage, procedures, training, and PPE

Figure 2 - Order of Precedence - MIL-STD-882 § 4.3.4

The Order of Precedence identifies mitigation approaches in order of

decreasing effectiveness in reducing risk exposure if design change is not viable for any reason. This list establishes the goal of implementing changes at the highest level in the Order of Precedence in order to achieve the most effective risk management option given the product constraints.

From this, the motivation to engage the System Safety discipline as early as the concept stage can be easily seen as providing the best opportunity to most effectively and economically manage safety risks.

c. 2nd Law – Safety Requirements

The fundamental output of the System Safety discipline is the identification and verification of risk mitigation actions (requirements) necessary to meet established system safety risk acceptance criteria.

While there are many perspectives involved in any discussion of System Safety, the fundamental goal of the safety process is the generation of safety requirements.

MIL-STD-882 provides a fundamental overview of the System Safety process (see Figure 4). While there are many nuanced variations to this, they all follow a very similar approach to this.




Figure 4 - MIL-STD-882 Figure 1

The Physics of System Safety

An ordered stepwise progression of planning, hazard identification, analysis and assessment with the goal of determining necessary risk reductions action(s) to effectively manage safety risk.

These identified actions being the basis of what are generally referred to as “derived safety requirements”. This reference is used as their basis to trace back to high-level safety requirements instated at the product (customer) level on which the product safety effort itself is based.

Therefore, it could be accurately stated that the goal of the system safety process is the determination of derived safety requirements, which represent actions needed to comply with all established product safety requirements. This means that the verification of the successful implementation of these derived safety requirements forms the basis of compliance with the product safety requirements.

An overview chart of an example product safety development which graphically illustrates how the System Safety effort progresses from customer requirements to derived safety requirements is shown in Figure 6. The green boxes represent the sources of compliance to a typical set of customer safety requirements;




Figure 6 – Sample Safety Process Overview Chart

- Compliance with quantitative (numerical) safety metrics. Usually via Fault Tree Analysis

- Verification of implementation of derived safety requirements to meet established risk acceptance criteria (see Risk Acceptance Matrix)

The burden of the creation of derived Safety Requirements in the System Safety process is levied on Safety Working Groups which are described in MIL-STD-882 Task 105. These groups focus on processing data captured in the Hazard Tracking System (HTS) described in MIL-STD-882 Task 106 and shown on Figure 6.

In this process each working group (usually divided by product team boundaries) processes data collected in the HTS to determine what corrective actions (if any) are required for causal scenarios identified for their specific group.

The communication of “acceptable” safety risk is typically accomplished by some form of the Risk Assessment Matrix from MIL-STD-882 Table III (see Figure 5).

RISK ASSESSMENT MATRIX				
SEVERITY	Catastrophic (1)	Critical (2)	Marginal (3)	Negligible (4)
PROBABILITY				
Frequent (A)	High	High	Serious	Medium
Probable (B)	High	High	Serious	Medium
Occasional (C)	High	Serious	Medium	Low
Remote (D)	Serious	Medium	Medium	Low
Improbable (E)	Medium	Medium	Medium	Low
Eliminated	Eliminated			

Figure 5 - Risk Assessment Matrix

This method uses severity of consequence and probability of occurrence (quantitative or qualitative) to rank gradations of risk levels from high to low.

It is generally interpreted that High and Serious are considered “unacceptable” risk requiring formulation of additional risk reduction actions or a specialized risk assessment to be provided to the regulator to determine final acceptability.

Scenarios falling into any of the other categories on Figure 5 are considered “acceptable” residual risk and no further action is required from a safety standpoint.

Upon completion of the safety effort, verifications of all derived safety requirements identified (and any specialized risk assessments from the High or Serious category) are presented to the regulatory agency in suitable format to serve as the basis for regulatory approval.

d. 3rd Law – System Safety Skillsets

The fundamental skillset of a System Safety expert is the ability to plan and conduct a product system safety effort. This involves analysis of products/processes in order to identify hazards and risk causal scenarios which allow safety risk to be assessed, and action(s) identified to reduce risk to within established risk acceptance criteria.

“Safety is a specialized subject just as are aerodynamics and structures. Every engineer cannot be expected to be thoroughly familiar with all developments in the field of safety any more than he can be expected to be an expert aerodynamicist.”

William Steiglitz
“Engineering for Safety”
September 1947

Along with the lessons learned from the post WWII era ICBM development, the emergence of Systems Theory in the same era forms the basis of what eventually was encapsulated in MIL-STD-882, and by association, all safety standards and practices which followed.

The early System Safety Engineers were engineers first. Even well before MIL-STD-882, engineers such as William Stieglitz had published a paper in 1947 titled

“Engineering for Safety”. In it, he captured his learnings from that time which presented many parallels to what eventually was encapsulated in MIL-STD-882.

Among them;

- The role of working groups meeting regularly in the review and oversight process
- The need for a Safety Committee independence from engineering in the safety function
- The recognition of a specialized skillset required for the safety function

This paper goes further to establish a higher-level Safety Board which operated at the executive level to assure safety focus above the product level and into senior leadership.

Any modern approach to system safety which does not acknowledge the specialized skillset required for a System Safety practitioner or expert runs the risk of returning to the pre-ICBM approach to safety risk management.

e. 4th Law – System Safety Tasking

The fundamental tasking of the System Safety expert is to create, initiate and oversee the conduct of a tailored system safety process which ultimately meets product/process regulatory safety requirements.

The high-level tasking for a typical product safety effort includes the following fundamental elements:

- Creation of a Safety Plan which defines the safety effort for the specific product which is aligned with the overall product development plan
- The definition, execution and management of a safety process which is tailored to the specific product intended to satisfy the

customer safety requirements and engages the full organization.

- Presenting documentation validating compliance to all safety requirements to customers and/or regulatory agencies.

The System Safety Practitioner must be knowledgeable and skilled in all appropriate safety analysis techniques commensurate with the application, intended operation and lifecycle phase of the product. This is best accomplished by engaging trained, experienced and certified System Safety Professionals who operate organizationally independent of the product engineering function. The practitioner should maintain professional currency through ongoing training and regular interaction with fellow safety professionals.

One of the primary tasks associated with the System Safety discipline is the identification of safety hazards associated with the product and its use.

Figure 7 shows a Hazard Identification Matix which divides the hazard identification burden in terms of whether the existence of hazards are known (based on past experiences) or unknown, and the discoverability of the hazards, whether are knowable (discoverable) or unknowable relative to some form of analysis or action.

Existence	Discoverability	Hazard Identification Method
Known	Knowable	Corporate Memory (Lessons Learned, Safety Centers, etc.), Standards, Policies, Best Practices
Known	Unknowable	Undocumented Lessons Learned, unshared safety information, uncaptured expert knowledge
Unknown	Knowable	Safety Analysis (200 series tasks, others), SME engagement, open reporting
Unknown	Unknowable	Opportunistic Inspection (Fleet Leader), Analytics, Trial & Error (break/fix)

Figure 7 - Hazard Identification Matrix

The **Known/Knowable** hazards are generally those discovered in the past. Lessons learned and best practices and/or standards compliance generally cover these hazards.

The **Known/Unknowable** hazards deal with losing access to past experience due to failure to adequately document past experience or allowing a safety expert depart without somehow capturing their knowledge (brain drain”).

The **Known/Unknowable** is the focus of the modern system safety process. These are “discoverable” hazards which are exposed via some form of safety analysis. Failure to fully complete the safety analysis effort results in residual risk from unknown hazards that would have been exposed had the analysis had been fully completed.

The **Unknown/Unknowable** hazards capture the reality is that no safety analysis will find 100% of the safety vulnerabilities of a given product or its use. Therefore the “Unknown Unknowable” will always be present, especially when new technologies are introduced by the product.

The system safety expert must assure the safety process being used is sufficient for the product needs as well as oversee and assure this process is dutifully followed and nothing is allowed to fall “through the cracks”.

f. 5th Law – Safety Competency

The definition of a System Safety expert is a System Safety practitioner with suitable maturation typically requiring a minimum of 10 years of system safety-specific training and product/process-specific system safety experience to achieve.

This paper proposes a terminology that defines a “Safety Expert” as a trained (preferably certified) and experienced individual with 10 or more years operating in a System Safety organization. A “Safety Practitioner” is defined as a minimally trained and/or experienced individual

operating within a System Safety organization.

The identification of the 10-year guideline relative to system safety competency is largely empirical based on the author's experience and observations as a safety certification instructor.

This guideline is consistent with something called the Dunning-Krueger Effect (see Figure 8).




Figure 8 - Dunning Krueger Effect

This chart reflects the inherent cognitive bias involved with self-assessment of competency. Basically, without training and experience (exposure), you don't know what you don't know.

Empirical observation has found that safety interactions between individuals with 10 or more years of experience were more efficient and productive, even across organization and business boundaries.

A challenge to the industry is that at the time of this paper, no centralized, global industry-recognized credentialing exists to provide bona fides in the System Safety discipline. The International System Safety Society (ISSS) has begun a Credentialing Initiative which is underway (with the author's participation) which is seeking to address this.

In the United States, one of the most notable sources of safety credentialing is the Board

of Certified System Safety Professionals (BCSP) which provides levels of certifications for safety professionals in multiple industry segments in the United States. In addition, several universities (such as Embry-Riddle) confer academic degrees in System Safety while other universities and organizations (such as the University of Southern California) offer certificates upon completion of qualifying non-academic courses.

Outside the United States, the concept of Suitably Qualified and Experienced Professionals (SQEP) to provide some level of assertion of system safety competency, however, there is no centralized guidance discovered to date.

One of the major challenges facing any entity trying to establish a centralized competency within the System Safety Discipline.

A major challenge to an credentialing effort is the wide range of domains which utilize the fundamental system safety process.



Figure 9 - System Safety Domain Chart

Figure 9 gives an example of the range of domains utilizing the System Safety process and illustrates the different factors which influence how an individual interacts with the system safety process.

Each "line of assignment" (ex: Aviation – military – piloted – sustainment – sub-

contractor – safety team member) carries with it acronyms, analysis methods, standards, tools, etc. which are specific to the domain they are in. Not to mention country divisions and regulatory agencies. Meaning there are potentially hundreds of permutations of individual assignments that are significant to the determination of competency as a System Safety expert. Expertise in one permutation can have differing levels of Dunning-Krueger placement in another permutation.

While this paper does not offer solutions, the intent is to understand the complexity of the issue in assessing competency associated with any credentialing effort.

4. Observations

The following section is a series of observations by the author based on experience and student interactions which use specific System Safety issues for which compliance to the 5 Laws of System Safety would provide the groundwork to address.

These observations are;

- A restatement of the crisis within the System Safety Profession
- A disambiguation of the term “Safety”
- The need for formalization of the role of the safety risk manager
- The need for a clearer definition of “Software Safety”
- Thoughts on the path ahead

a. Safety Profession Crisis

In a paper presented at last years’ ISSS Summit and Training entitled “System Safety – A Profession in Crisis”, a case was made to

support the assertion that System Safety as a profession in a crisis mode.

In summary, the paper asserts that an erosion of independence within the safety function and lack of sufficiently trained safety experts are the prime drivers for the crisis. The net result of this erosion is the safety function returning back to the engineering function, where it was at the time of the founding of the modern safety approach in the post WWII ICBM era. The effect of this is the industry observing a movement from predictive to reactive safety management approach.

The response to this paper from a wide range of attendees, as well feedback from a subsequent podcast on this same topic, was overwhelmingly positive and supportive.

This assertion is not meant to cast aspersions on engineers or project managers trying to “take over” the safety function as it is to acknowledge the realities of human nature and marketplace drivers.

The insights gained from my multiple decade exposure to the Engineering and System Safety disciplines is helpful in recognizing inherent traits that can unwittingly contribute erosion factors.

- Product advocacy – A prudent engineer typically takes complete ownership of their product including safety performance, which is commendable. This can lead to an adversarial posture to any oversight function. What I like to call the “don’t call my baby ugly” mode.
- Self-sufficiency – A prudent engineer is always looking to expand their knowledge and capabilities, which is also commendable. In fact, suitably trained engineers can conduct the safety

functions. But they cannot provide effective oversight of their own work.

- Chafing at oversight – It is human nature to not enjoy people looking over one's shoulder, as safety can sometimes appear. It takes an enlightened perspective to recognize the need for, value of and welcome the oversight function.
- Cost is King – If safety oversight is seen as redundant to engineering then it becomes an attractive target for cost reduction. Under-funding, under-staffing and under-training is the most common issue facing safety teams globally.
- Checkbox Mentality – If safety is seen as a perfunctory exercise or “blind compliance” to some safety guidance, then it is vulnerable to misapplication of the process and/or the actions proscribed. The fact is that System Safety is more art than science which defies a “cookbook” approach that many outside of safety see in safety specification and standards.

Another key contributing factor to erosion is the use of when, for example, an underqualified and/or insufficiently experienced safety person is sent to provide safety input during a conceptual study. This lack of competency can easily be seen as impediments to progress when attempting to learn the product and the safety process at the same time.

This can lead to the concept team concluding that they are better off handling the safety aspects themselves, resulting in excluding the safety team from the early design process, which is the best time to address safety issues.

Also, past negative interactions between engineering (or project) and the safety team

can also result in exclusion from program activities. This can happen if the safety team misinterprets their charter to take a more authoritarian or controlling role in the design process, which is inconsistent with the original oversight charter for the safety discipline.

Another is the failure of safety team members to understand when past safety guidelines need to evolve as they can create unnecessary or even useless product or process constraints.

Experiences like these help explain the erosion leading to the current crisis. Organizations which don't recognize this erosion exists and/or fail to actively manage the balance needed for proper safety engagement usually defaults to more of the System Safety function being relegated to the Engineering function.

It is the contention of this paper that recognition and concurrence with the foundational truths of the system safety discipline and profession will provide the logical path to arrest this erosion.

Without changes, my assessment is that the discipline and profession of System Safety will continue to erode.

b. System Safety Disambiguation

One of the most confounding challenges facing any attempt to address issues associated with System Safety as a discipline or profession is the wide range of specialized sub-disciplines and interpretations contained in the word “safety”.

Safety is a word with great power. At its mention, one can get people's attention, bring sweeping change, or even stop things in their tracks. Wielding the "s" word can be like swinging a "sword" (word play intended).

Bringing great power to the bearer. And with great power comes great responsibility. The truth is that safety is more like a Swiss Army Knife than Excalibur.



Figure 10 - Safety Army Knife

In the hands of a true safety expert, this word can cut through bias, posturing and even inexperience to support a more focused and informed decision-making process. In the wrong hands, it can be abused for nefarious purposes, create inefficiencies or even mislead the decision-making process.

The problem is that like the army knife, there are many "blades". For the purposes of discussion, here are definitions for a few key "blades":

- System Safety - Structured approach to hazard identification, risk assessment and corrective action verification
- Functional Safety - System safety concepts applied to sub-systems or subfunctions (e.g., ADAS, Gas Turbines)
- Software Safety - Functional Safety in the context of software-controlled systems
- Flight Safety - System safety concepts applied to the initiation, sustainment or termination of flight
- EHS or Industrial Safety - System safety concepts applied to environmental or personnel risks

- Operational Safety - Compliance with established safety guidance

Some disciplines which are distinct but have ties to System Safety;

- Human Factors
- Reliability
- Cybersecurity
- Survivability

When some safety conversations are engaged, it is easy to see how the participants can become confounded and even confrontational when the word safety is used without awareness or respect for the different disciplines and skillsets engaged in each of the "blades".

The reality is that each of these "blades" have similar but distinct processes, standards, training and acronyms. For this reason, there is a need to be keenly aware of these differences in the context of any safety discussion, even among safety experts.

For example, I have heard organizations including both System Safety and Environmental Safety (EHS) without sufficient disambiguation in high level safety discussions leading people to bring System Safety issues to EHS and vice-versa.

This ambiguity can create harm should even safety experts who aren't sufficiently trained/experienced take on cross-tasking in other safety domains. Anyone who has taken home electrical or plumbing jobs instead of calling in a pro can attest to the quality and confidence difference, especially on challenging jobs.

Understanding the many facets and domains of safety can also help understand the risks when hearing from safety experts (including myself!) who are speaking outside their own domain. Especially if their input clashes with other known safety experts.

The truth is that whenever safety experts clash, individuals outside the safety discipline are forced to decide whom to believe. Clashing itself isn't bad, however, when it is fueled by a lack of understanding or "lumping together" of the different safety domains that exist, it is counterproductive.

c. Risk Managing Authority Proposal

The reality is that in any system safety effort there are people who take leadership roles in formulating, overseeing and/or managing the system safety effort whether it is formally identified or acknowledged.

MIL-STD-882 alludes to several authorities participating in the safety process, some which are identified within the construct provided by DoDi 5000 acquisition policies.

This paper proposes the identification of a formal Risk Managing Authority (RMA) role and Safety Risk Manager (SRM) role to be identified as individuals for each product or process being managed for safety risk.

These individuals act as authorities responsible for the oversight of safety efforts acknowledging the integration aspects needed for product engaging multiple layers of suppliers.




Figure 11 - Sample RMA/SRM Functional Interaction

In this context, the RMA is the highest managing authority primarily responsible for the safety of the integrated system and

serves as the focal point for all subsystem safety efforts and primary interface with the regulatory agency.

The RMA creates product level safety and risk acceptance requirements as well as reviews all safety documentation, interfaces and coordinates with all associated SRM(s) and coordinates with a Regulator (see Figure 11), which functions in a similar role to the RMA as the final authority on risk acceptance.

The SRM is responsible for the safety of their respective sub-system and any contributing suppliers similar to the RMA.

As indicated previously, all safety efforts function in such a structure by default. The problem comes when the RMA and SRM roles are not embodied in individuals who acknowledge their roles, or worse yet, lack the training and/or experience to sufficiently function in these roles. Another issue is the personnel turnover that can happen on large-scale product development efforts.

Having rotating or multiple individuals in RMA or SRM roles can create significant inefficiencies in the safety process.

Having these roles clearly delineated and individually identified within the safety planning documents would greatly aid the efficiency and function of the safety efforts, particularly for large complex product development programs.

d. Software Safety Perspective

Similar to the disambiguation needed for the word safety, there is a need for software safety to be broken down into sub-disciplines for the same reason.

Addressing software safety in general is one of the most vexing technical challenges facing the system safety profession today.

More fundamental than the wide variety of disciplines within software safety today is the integration of software into the system safety discipline.

Systems developed in the early post WWII era typically had very little in the way of software controls, so a higher focus was naturally kept on hardware related vulnerabilities. Indeed, many early safety practitioners downplayed the role that software played in safety management. “There’s no such thing as a software hazard” was heard regularly not too many years ago.

With the onset of the digital processing era and more complex functionality carried out by software came the early views of software safety. MIL-STD-882 Rev B (1984) introduced the Software Control Category (SCC) and Rev C (1993) introduced the Software Criticality Index (SwCI) as a means of addressing software risk by assigning a level of rigor (LoR) in order to provide guidance to software developers in terms of levels of analysis and testing rigor to subject software controlling elements as a function of their importance to safety. The more important the function, the higher the analysis and testing rigor that should be applied.

From my perspective, software safety discipline can be divided into two clear categories;

- Within the historical System Safety tasking
- Outside the historical System Safety tasking

Within historical System Safety Tasking

From a generic perspective, the System Safety process is agnostic to hardware and software.

Many early safety practitioners asserted that software was deterministic in nature and not probabilistic and so any failures were driven by hardware failures rather than software.

But as software complexity increased software clearly played a larger role in safety management.

The key engagements between system safety and software safety are:

- Safety Critical Functions (SCF) are identified as a product of the Functional Hazard Analysis (FHA) process as functions are assessed independent of whether they are implemented in hardware or software features. The SCF’s are communicated to the software development team.
- Level of Rigor guidance via SwCI for validating software analysis and testing. This is usually captured in a System Safety Program Plan (SSPP) or similar.
- Corrective actions identified during the system safety effort which are implemented in software form the basis of derived software safety requirements which are flowed down to the software development teams as HLR’s (High Level Requirements). These HLR’s are typically captured in the Hazard Tracking System (HTS) and are the basis of traceability evidences presented at certification.
- Additional software requirements generated as HLR’s are decomposed as a part of the software development process are communicated to the system safety team and incorporated in the HTS.
- Incorporation of software hazards and corrective actions into the System Safety process that were identified by a Software System Safety Process described in the following section are entered into the HTS

The key coordination is software related data tracked in the HTS.

Outside historical System Safety Tasking

Specialized skillsets, which usually reside within the software development team function, are required to conduct a system safety process specific to software. This process is captured within the Software Safety Program Plan (SwSPP) and assures coordination with the SSPP.

The typical functions within this role are:

- Coordination of HTS software HLR's into the software development process
- Oversight of software analysis and testing commensurate with SwCI LoR guidelines.
- Ongoing coordination with hardware SRM relative to HTS updates affecting software as well as pertinent results from software analysis and testing efforts.

Whether this process is carried out by individuals within or outside an organization's system safety function is dependent on how the organizational structure is implemented.

Safety organizations which do not have staff members with sufficient software development experience usually "deputize" a member of that team to conduct or oversee the conduct of specialized software analyses designed to identify causal scenarios which can potentially compromise one or more of the SCFs.

These individuals are essentially software SRM's (see previous section) and would greatly benefit from formal safety training and certification.

Level of Rigor - The Missing Link

One long term observation relative to coordination of System Safety and Software Safety deals with aligning SwCI with

software assurance and integrity levels. Both are manifestations of forms of LoR's, but have subtle differences which often lead hardware and software safety practitioners to disconnect when trying to resolve shortfalls in compliance.

From my interpretation, SwCI is a function of both criticality and an additional function captured in the SCC, while software LoR's expressed as Design Assurance Level (DAL), as in RTS DO-178 and DO-254 electronic devices primarily incorporate criticality thresholds to define the different LoR's.

Reconciling this difference requires coordination between the classical System Safety organization and Software Development organization.

In the absence of this coordination, many organizations resort to implementing the highest level of rigor for all software, which is a valid conservative assumption but can drive higher development costs, or maintaining a "two level" (defining minimum safety testing) approach and putting software into one of the two categories.

While this is a functional approach, as software gets more complex, the need to have a multi-level hierarchical level of rigor definition will be necessary for both safety and economic reasons.

AI considerations

With the rise in focus on AI, it is becoming apparent that there may be an expanded set of skills to required which could require even further coordination with the legacy system safety function.

While this paper in no way takes on directly addressing potential safety vulnerabilities exposed by the use of AI techniques, an observation is that there are two distinct

focal areas relative to AI which fall within the purview of system safety:

- Use of AI software in either full or partial control of product or process control such as in autonomous vehicles.
- Use of AI software to assist in the safety process itself. Either by conducting analysis or supporting risk management decision making.

The first bullet affects risk management within operation of the product or process; the second bullet affects the operation of the system safety process itself.

At the current time, methodologies such as the overall safety process, identified in ISO 26262 and mitigation approaches such as Safety Of The Intended Function (SOFIT – ISO 21448) have been used to address risk from AI techniques used in products.

The advent of new software tools which utilize AI techniques to assist the safety process itself have presented challenges for the as evolution of the system safety discipline is clearly needed.

In this transition time, a higher reliance on expert system safety oversight is needed as these definitions are refined.

This may also recognition of a unique title for System Safety professionals who carry the specialized training, skillsets and role responsibilities associated with software development in general and emerging AI capabilities in particular.

e. The Path Ahead

Resolving the crisis in system safety not only requires concurrence from the System Safety community at large but also must include addressing the environment that Safety professionals around the world operate in.

While my previous paper addressed the system safety crisis, one of the major reasons this paper was created was to find common ground amongst system safety professionals on the fundamentals of the discipline and profession. It is my intention to create and submit a paper on addressing the environmental aspects.

Improving the environment that the system safety profession and discipline is a task levied first on the leadership of the safety organization itself. However, it is my contention that the largest influence in terms of affecting the effectiveness of a system safety organization is higher in the organizational chart.

When safety professionals meet and discuss improving various safety techniques, standards and processes, what is not often addressed is the role of senior leadership (CEO's, VP's, etc) in their oversight of the safety organizations within their business.

While "Safety is Job 1" or some variation is very loudly and consistently declared from the leadership of most organizations, their role beyond the Safety Management Systems (SMS) Promotion and how that sets the environment that Safety Teams operate in is not clearly delineated.

It is my observation that the efforts of even the most experienced and capable system safety expert can be tempered or even muted by any intentional or unintentional shortcomings in the implementation of the fundamental aspects of the safety discipline.

The types of CEO tasking include;

- Assuring compliance with established organizational safety policy/policies
- Assuring Safety Team vitality

- Safety Team is staffed adequately to support all current product safety efforts with the necessary level of expertise (training and experience)
- Assuring “pipeline” of system safety expertise is maintained to reduce vulnerability to personnel changeover, job poaching and retirements

As a final point on the path ahead, the accident/incident investigation process is the most highly visible engagement of the safety process as well as the most evolved.

That said, after a major accident, it takes conviction and focus to move beyond the

single “smoking gun” root cause analysis to delve into a wider form of analysis, such as with STAMP/STPA/CAST methodologies, to uncover contributing factors beyond the immediate to determine corrective action which can have farther reaching impact even beyond the specific product involved.

This is particularly true if identification of where the system safety process *should* have found the source of the incident/event, it can lead to process corrective actions which have the potential to have positive impact on all future products.

5. Conclusions

The observations presented represent the culmination of someone who has operated for multiple decades in both the Engineering and Safety realm, as well as from the perspective of a system safety instructor who has taught and interacted with students from around the world.

What I have attempted to document in this paper are the foundational truths I’ve observed within the system safety discipline and profession globally and across many domains and product applications.

It is my hope that the information contained in this paper forms the basis for future conversation and discussion which leads to some level of consensus on a path to return the system safety discipline and profession to its fundamental charter and function. I believe this is absolutely essential to effectively address the challenges of new products and technologies and experience the same progress and successes evidenced in the safety progress of the 1960’s and 70’s.

It is my highest hope that both individuals and organizations alike take in the information in this paper and if in agreement with its central premise, conduct self-assessment and determine if the way they “do safety” today needs to better align with what I believe are the most fundamental aspects of the discipline and practice for system safety