New System Safety Process for New Business Opportunities

by Mark A. Vernacchia
Milford, Michigan

This paper outlines an approach to developing and implementing a systems safety process for new business applications in industries without existing system safety processes. It describes the desired characteristics of such a process and extends the discussion to include the basic needs and content of any systems safety process in general.

New business applications are being created every day, and safety must be considered in these endeavors. This paper explores how to determine if such applications warrant a system safety process so that potential risk is identified and appropriate prevention, detection, and mitigation content and capabilities are designed, developed, implemented, tested and verified prior to the application deliverable being placed into customer hands.

The role of management in determining what is unacceptable risk and how risk is allocated between participating business and customer entities is also discussed. In this paper, "system" is used as a generic term for the product the company plans to deliver.

Introduction

Businesses are being challenged to develop new products to satisfy ever-changing market needs and desires. New and established companies are branching out into new business opportunities where system safety may need to be addressed. The challenge is, "Does a business need a system safety evaluation as part of its new opportunity design and development?" The essence of this question comes down to the risk presented to the company, customers or society if this new business opportunity is brought to market and the likelihood and severity associated with this risk. Basically, it's all about risk.

Risk may appear in many forms. Each business needs to decide what risk means to them. Automotive companies associate risk with harm to humans — either to drivers or occupants within a vehicle, or to pedestrians surrounding a vehicle. Aerospace and defense companies associate risk not only with harm, but with things like failure to complete a set of mission objectives. Companies can also experience risk in the form of credibility loss, where the public loses faith in them, resulting in revenue losses or lost business opportunities. One way to think about this is to consider when "bad" things would happen and — if they do — what is the risk? So how does a business decide what risk means to them?

Companies need a process that identifies potential risks and the seriousness of those risks, and determines the effort required to prevent or manage those risks. Included in this process must be a structure that acknowledges that management is responsible for risk management, for defining the level of risk that is acceptable, and for approving that level of risk for the system.

New Business Opportunity Characteristics and System Safety Evaluation

New business opportunities may not easily fall under existing automotive, military and aviation-based standards such as ISO-26262, MIL-STD-882E and ARP-4761, respectively. This may depend on marketplace maturity, expected users, planned life cycle, planned volume of units and the allocation of risk mitigation responsibility between the company and the users.

New business opportunities may have unwanted implications for performing a system safety evaluation. These might include:

- Suppliers deciding not to provide a response to a technical request for quote (RFQ) because they are not able, or are unwilling, to support technical necessities driven by existing system safety standards such as those mentioned earlier
- The use of "off-the-shelf" items whose suppliers adopt a "take it or leave it" response

to questions regarding internal technical details of their product
- Perception that because a product is being sold to the public, it is by definition "safe"
- Judgment as to what entity, company or customer will be responsible for managing any potential risk situations
- Most important, a hesitancy to acknowledge that system safety is a key part of a balance of parameters, just as important as cost, time and resources

All of these things influence how a system safety evaluation would be done.

Attributes of a Useful and Effective System Safety Process

A useful and effective system safety process should provide requirements that prevent or manage potential risk to a level that is acceptable to the company, customers and society at large. Attributes of a good system safety process would include, but are not limited to:

- Capability to understand needs and expectations and accommodate constraints
- Methods to identify potential risks and assess the seriousness of those risks
- The ability to support balance and optimization decisions
- The ability to provide the *right requirements* at the *right time* in the *right level of detail* to support engineering process needs that prevent or manage risks, even when various levels of system detail are available
- The capacity to verify system performance and behavior to requirements and to validate system capabilities against stakeholder needs and expectations
- Methods for risk allocation
- Methods for company management to decide and approve whether risk is acceptable

Such activities allow risk to be identified and assessed in a straightforward manner. Once this is done, requirements may follow to prevent or manage risk situations, along with methods to verify that those requirements result in an acceptable risk, per management approval.

System Stakeholders' Needs and Expectations

Stakeholders are individuals or entities who have a vested interest in how the system will work and what it will deliver to the marketplace. Stakeholders are usually comprised of management, marketing and engineering representatives who have decided, or were instructed, that the company should develop a particular new system based on business opportunities.

The needs and expectations of the proposed system behavior and deliverables are then transformed into high-level system requirements. These requirements should be summarized in a document that describes, in a straightforward manner, how the system is expected to operate under all expected conditions throughout its life cycle. This document becomes a narrative describing the system behaviors when operators interact with it, how it will execute expected activities and what form the results (deliverables) of its behaviors will assume. The content outlined here can be summarized in a concept of operations document, usually referred to as a "ConOps." The ConOps will become the validation criteria to assess if the resulting system satisfies stakeholder needs and expectations.

Understand and Accommodate System Constraints

System constraints define the boundaries of the system scope and the limitations the design team will need to accommodate during system design and development. These constraints usually involve cost, timing, regulatory, resources (both people and equipment), customer needs, etc.

Safety constraints will be added to this list as the system begins the concept phase of the design. This list will be utilized in engineering and management trade-off study evaluations that determine the "optimized" system balance that will best fulfill the ConOps requirements within the constraints imposed on the system.

Risk Identification and Assessment

A system safety process needs to identify and assess potential risks and their seriousness. An approach to accomplish this is to evaluate the system elements' intended functions, noting what inputs or feedback they need to perform each function and what commands or requests the function elicits.

Each input/feedback and command/request item can be evaluated against a set of guidewords that concern themselves with inappropriate behaviors such as "Not Provided When Needed," "Provided When Not Needed," "Provided Too Soon/Too Late/Out of Order," "Stopped Too Soon/Maintained Too Long" or "Provided But Too Much/Too Little." Typically, evaluation methods such as hazard operability (HAZOP) and/or system-theoretic process analysis (STPA from MIT) are useful in this effort. Many links are available on the web for HAZOP. The following link is useful for learning about STPA: http://psas.scripts.mit.edu/home/

The outcomes of the guideword evaluation lead to a determination of whether the results pose a safety hazard, where under the right conditions, something "bad" could happen. The *bad* here is related to the company's definitions of and decisions about what they do not want to lose, or what losses they do not want to inflict on users or society. *Losses* in a system safety context denote *bad* things. Let's look at an illustrative example.

A new company wants to develop and market a small mobile delivery device that an operator can ride on to transport heavy packages within a factory. The stakeholders have provided the following information:

- A hand-operated control allows the operator to control the speed by measuring how hard the operator squeezes a lever, similar to a bicycle hand brake, that will send an electrical signal to the wheel motors.
- A hand-operated control allows the operator to slow the vehicle by providing hydraulic pressure to a disk-caliper mechanism on the drive axle.
- The operator steers the device by shifting their weight side to side on a small platform they stand on that will send an electric signal to a motorized rack-and-pinion steering mechanism.
- The device is powered by a plug-in rechargeable 48V lithium-ion battery.
- Wheel motors need to be purchased as off-the-shelf components.
- Device must be capable of carrying 1,000 lbs. of cargo.
- Device must cost no more than $2,500.

Reading the stakeholder needs and constraints, some potential risks may be easy to identify. This device is a moving transport, so it could hit people or objects. It also has a lithium-ion battery, which have been linked to fires. So, two *bad* things can happen. The first centers on harm to people and the second encompasses not only harm to people, but also property damage due to a fire or loss of the device itself. Let's examine these in more detail by performing a guideword-based evaluation that looks at a few example functions.

The major elements (and their primary function) of the device can be described as a speed control system (convey operator speed intent), a steering control system (convey operator steering intent), an electric motor to propel the device (make device move), a brake device to slow the device (make device stop) and a motorized steering mechanism (steer the device).

Figure 1 shows an example evaluation using five functions and two guidewords that results in finding unwanted situations where a hazard may be present. The

Function	Not Provided When Needed	Provided When Not Needed	Potential Hazard	Possible Mishaps
Receive Propulsion Command		Device received propulsion request when operator did not request it	Unwanted Device Motion	Pedestrian Injury Operator Injury Collision with Fixed Object
Provide Steering Command		Operator movement on device causes unwanted steering input while turning	Loss of Directional Control	Pedestrian Injury Operator Injury Collision with Fixed Object
Brake Activation	Brakes not applied when requested by operator		Loss of Braking	Pedestrian Injury Operator Injury Collision with Fixed Object
48V Provided				Burns Electrocution
Motor Torque		Motors activate while operator using hand brake	Motors Overheating	Burns from Hot Surfaces
		Motor torque used as holding function on incline	Motors Overheating	Exposure to Fire

Figure 1 — Functions, Guidewords, Hazards and Mishaps.

functions are listed first, followed by guideword assessment, potential hazards and possible mishaps.

If any of the functions misbehave as described in the guideword columns, they can create situations where people or property damage could occur. Risks associated with harm to people or property damage were envisioned before the guideword evaluation, but now traceability is developing from the function, to the unwanted behavior, to the hazard and finally, to the mishap. There could be one high-level goal that says, "No people shall be injured," which is really the mishap. The lesson here is not to focus on the mishap, but to prevent or manage the hazardous situations that could lead to the mishap, as these are things that can be controlled. High-level requirements to prevent or manage hazards can be developed at this point.

Review of a guideword-based evaluation can further define the risks (harm, property damage, loss of function, etc.) and the associated hazards, allowing the company to understand the potential implications of a specific device design on the marketplace.

Safety Balanced Against System Opportunity Imperatives

Risk identification and the associated high-level requirements to prevent or manage hazardous scenarios are useful for technical trade-off and balance optimization discussions. Early concept development may result in a number of workable designs, say for the factory transport device above, that emphasize cost and timing but compromise performance, or hazard prevention or management. Other designs provide what some feel is the required "safety" for risk management, but drive high costs or delay delivery.

This interactive process between stakeholders, design team engineers, system safety engineers and program management is greatly aided if the safety implications associated with each design option are known. The key point is that risk needs to be acknowledged and understood by the company's management structure, as they will have to accept the risk associated with the selected design.

Right Requirements, Right Time, Right Detail

The impact of system safety-driven requirements on the design should be evaluated from a top-down approach. A common observation with engineers is that they want to get into the details quickly. This has disadvantages of getting into "functional silos" too quickly, or not understanding the interactions between major parts of the system being designed. This compromises the ability to perform high-level trade-off studies and optimization discussions, and causes downstream design changes to occur later in the engineering process that are more costly and time consuming than they might be if they were found earlier. Typically, the worst situation is when issues are found after a product has been introduced into the marketplace.

A way to help with this situation is to adopt an approach that focuses on developing the right requirements when they are needed by the engineering process and with the right level of detail to be useful in the engineer-

Design Phase	Right Requirement	Right Time	Right Detail
Concept	Human-Machine Interactions Functional System Interactions	Concept Reviews Architecture Alternatives Supplier Identification	Safety Goals to Prevent Bad Things
Requirements	Provides Boundaries to Sub-Systems Do Not Constrain Sub-System Design	Architecture Selection Function Allocation to Major Sub-Systems Supplier Selection	High-Level Requirements to Constrain Sub-Systems ("Do What You Need to Do in Your Area But Meet These High-Level Requirements")
Design and Development	Requirements To Drive Test Cases Requirements Trace Back to High-Level Requirements	Prototype Sourcing Production Sourcing Prototype Testing Development Testing	Sufficient to Design Sub-Systems Sufficient to Design Specific Test Cases
Verification	Test Cases Test Results	Production Testing	Sufficient to Execute Test Cases Sufficient to Document Test Results

Figure 2 — Right Requirements, Right Time, Right Level of Detail.

ing process. Keep things as simple as needed. Too much detail too quickly causes wrong requirements to be developed or, worse, to not be developed at all.

The concept phase includes the guideword-based risk identification effort. Knowing the risks, a risk assessment determination is appropriate.

Risk Assessment

Risks have been identified by the guideword process, but does the company have to, or choose to, address them? Two main activities may help: determining how serious these risks are and determining what type of requirements the risk level will need to provide a picture of acceptable risk.

Existing system safety standards typically assess risk using two parameters — severity and likelihood. In ISO-26262, likelihood is expanded into two additional parameters: the likelihood of being in different operating conditions when the hazard may occur and the ability of humans to influence the outcome of the hazardous situation, known as controllability. Figure 3 illustrates a simple rendition of a risk matrix showing a "Low" to "High" transition from lower left to upper right.

System safety discussions should occur where the risk level of the hazards is assessed relative to each other and relative to what is observed in the marketplace. Empirical data should be used to determine breakpoints along each risk axis whenever it is available. This was the approach used in standards such as ISO-26262, MIL-STD-882E and ARP-4761A, and these standards provide three to four breakpoints along each axis.

It becomes more difficult to determine risk breakpoints if the new system is developed for a new marketplace that does not have such historical data. Here, system safety discussions should accommodate perspectives such as potential customer input (how bad does a customer think the hazards are), the breadth of potential impact if a hazard occurs and the ability of humans to manage the hazardous situation.

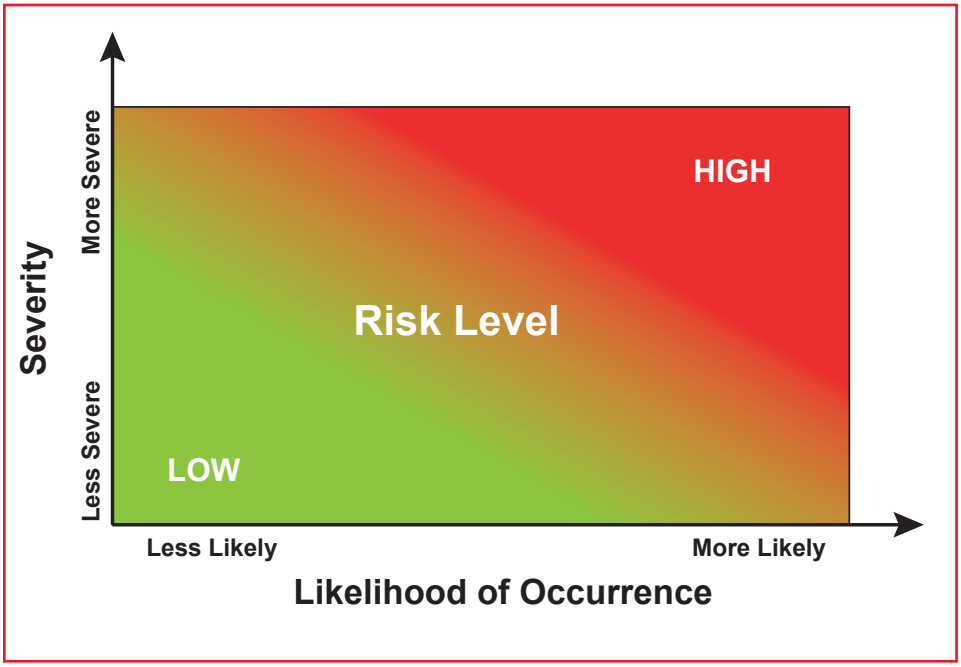

Figure 3 — Risk Based on Severity and Likelihood.

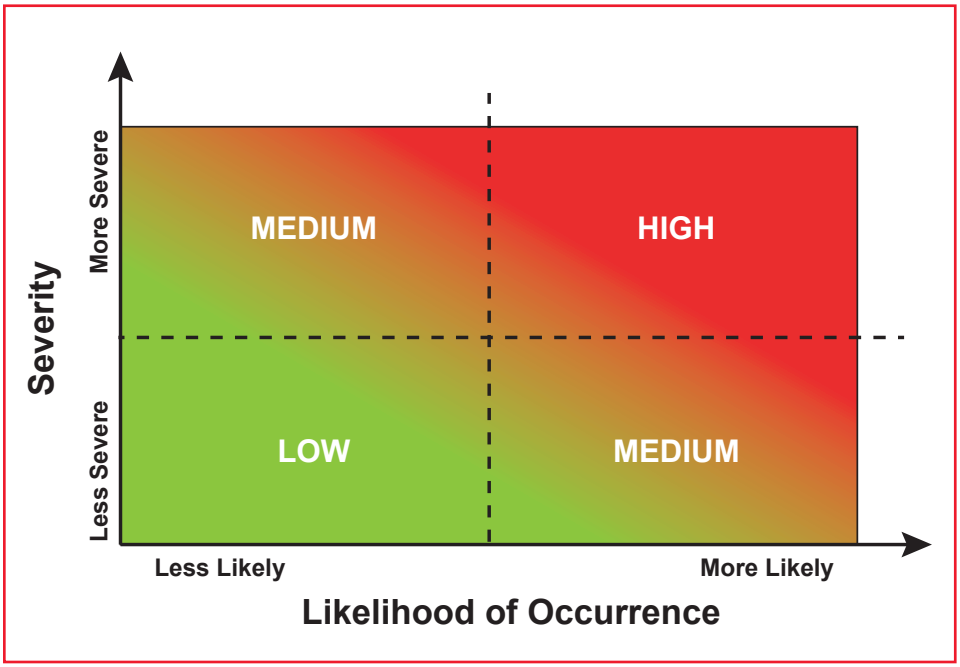

Figure 4 — High, Medium, Low Risk.

Having risk levels decided and allocated is an initial step. The system safety process also should consider how these different risk levels will be used relative to requirement development. A real benefit of risk assessment is that it can help answer the question: "How good is good enough?" Engineers continually struggle with this question as they want clear, simple answers to what requirements they need to satisfy. This desire is true for safety-driven requirements too. Therefore, risk levels can be a helpful way to define requirements that provide the required system content, performance and behavior.

Higher risk levels should drive more hardware and software diagnostics and mitigation capability than lower risk levels. Addressing single-point faults may be appropriate for all risk levels, but higher levels may want to include dual point faults also. This illustrates the need for experienced systems safety engineers to lead these discussions so that clear, concise, consistent and executable assessment methods are developed.

The hazards identified for the example system described previously could be classified using three levels — LOW, MEDIUM and HIGH. These are illustrated in Figure 4.

Figure 5 shows the results of a system safety assessment regarding assigned risk levels and the resulting requirement impact for the hazards identified in our example factory delivery device.

The information shown in Figure 5 would be the result of a system safety process where HIGH risk hazards require dual-point fault protection, MEDIUM risk hazards require single-point fault protection and LOW risk hazards require only warning labels or user manual instructions.

The risk assessment levels could also drive the types of evaluations necessary to provide objective data for risk acceptance decisions. HIGH risk levels may require critical software assessments, while MEDIUM and LOW risks would not. MEDIUM and LOW risks may require FMEA analysis, but HIGH risks would require FMEA and FTA analyses. HIGH and MEDIUM risks would require system-level testing, but LOW risks would require only sub-system testing.

Verification to Requirements and Validation to Stakeholder Needs and Expectations

Test cases and test results will verify that requirements were satisfied. The risk assessment assigns risk levels that impact requirements. Requirements are developed at different points in the program timeline and should have a level of detail no more than what is required to complete tasks at that time (keep things as simple as possible). Test results should be used as objective data to demonstrate that risk is managed to an acceptable level.

Another aspect would be to validate that the system fulfills stakeholder needs and expectations. The design may satisfy all specified requirements, but may not provide the planned utility, desired benefits and expected behaviors necessary to provide value to both the company and the intended marketplace.

Risk Allocation

It's all about risk, and now, who is responsible for it. The choice basically comes down to deciding if the company feels the system they are providing should manage the risk to an acceptable level, accommodating the cost, timing and content necessary to achieve this, or whether the customer should manage the risk by having operators or users manage the risk in a level they can accept. This seems like an extreme approach. The real answer is more likely shared responsibility. A system safety process that provides an approach to balance between these options is desirable.

Organizations will typically struggle with this activity. Perspectives and attitudes such as the following might be present:

Potential Hazard	Possible Mishaps	Risk Assessment	Requirement Impact
Unwanted Device Motion	Pedestrian/Operator Injury Collision with Fixed Object	HIGH	Dual-Point Fault Protection
Loss of Directional Control	Collision with Fixed Object Pedestrian/Operator Injury	HIGH	Dual-Point Fault Protection
Loss of Braking	Collision with Fixed Object Pedestrian/Operator Injury	HIGH	Dual-Point Fault Protection
Electrical Arcing	Burns Electrocution	MEDIUM	Single-Point Fault Protection
Motors Overheating	Burns from Touching Hot Surfaces	LOW	Warning Labels
	Exposure to Fire	HIGH	Dual-Point Fault Protection

Figure 5 — Risk Assessment and Requirement Impact.

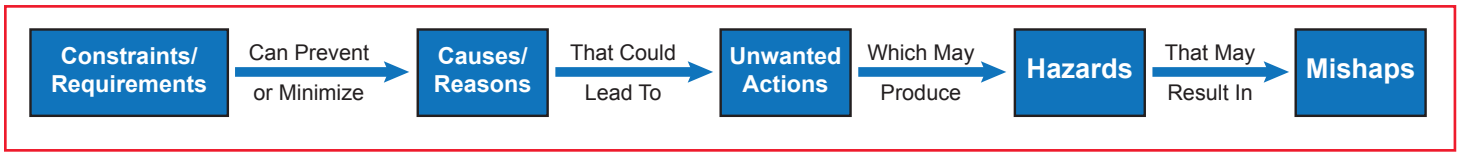

Figure 6 — Constraints/Requirements Linkage to Mishaps.

- Groups who think system safety is an impediment to progress and just there to make things hard. They will not understand why there is concern since "we are using off-the-shelf components and they are obviously safe as they are available on the web."
- Groups who think the system provided needs to manage any and all potential hazard risk and that the program or marketing teams just need to understand this and stop complaining about having to implement system safety requirements.
- Groups who have no, or limited, human factors or human behavior experiences and think that operators of the system will be fine dealing with any misbehaviors.
- Groups who think customers should know what they are buying as risks are self-evident.

These issues of understanding need to be addressed through continuous joint work between all participants involved in the system program. One method that has proven useful to educate people is to explain that the reason requirements are needed is to prevent or minimize causes or reasons that could lead to unwanted action that may produce hazards where risk is present and may result in mishaps. Figure 6 illustrates this logic flow.

In the vast majority of these discussions, individuals will understand the reason for having system safety requirements, as they are a means to manage potential risks. Typically, engineers do not want the system they design and deliver to cause harm or to not attain the desired goals.

Understanding why system safety needs to be part of the program content and execution risk proposition, the remaining big topic is how to decide who is responsible for managing this risk. The spectrum can range from the company managing all the risk to the customer managing all the risk. As usual, the "truth" is somewhere between the extremes. Some helpful suggestions may assist in finding this compromise using our factory delivery device example system.

- The steering is controlled by the operator standing on some sort of pivot plate where they shift their weight side to side to control direction. This was a stakeholder need to provide some "whiz-bang" technology to impress customers. Using this control mechanism, the guideword evaluation identified that unwanted steering could occur if the operator's weight shifted to the outside of a turn, causing the device to begin steering out of the turn prematurely. The company could decide to invent ways to control this leaning motion by putting dampers on the pivot plate, ignoring a sudden application of counter-steering and only allowing slow steering inputs, or maybe adding some extra driver action to indicate to the system that they actually want to counter-steer. Seeing how hard and costly this invention would be (company managing risk), or how effective the operator would be to interact with dampened pivot plate or extra control input (customer managing risk), the best course of action may be to change the steering design to a conventional steering wheel or joystick where the customer would control the steering and, therefore, would be responsible for managing the risk.

- Brakes are applied by a hand control that allows the operator to slow the vehicle by providing hydraulic pressure to a disk-caliper mechanism on the drive axle. The hazard associated with this function is Loss of Braking and has been classified as a HIGH risk, meaning dual-point fault protection is required. A number of options for the company to manage the risk might be:
  - o Add a second independent hydraulic system driven by the single hand level
  - o Add a dual-level system with each controlling a redundant system
  - o Have the single hydraulic system split the hydraulics to each side of the device

  All these options add complexity and cost. Maybe a better solution would be to divide the responsibility between the company and the customer by having the company specify a highly reliable brake system and inform the operator that they need to test the brake system manually by squeezing the hand control and feeling pressure resistance.

- Burns from touching hot surfaces were classified as LOW risk, driving appropriate warning labels to be placed by such surfaces. Here, the risk is managed by the customer, as they will have to heed the warning and not touch the surface. It should be mentioned that there is some responsibility for the company to place the warning labels appropriately and to have effective content on the label that conveys the risk. In addition, the company

should design the system in a way that eliminates or minimizes the likelihood of inadvertent contact.

Risk acceptance by the customer may also help in the balance and optimization of the system content. An example might be in a military application, where a highly trained soldier operating a system accepts more risk of a potential system misbehavior. In this case, the system would include less risk-management content than it would if the user was an average individual.

The end result of risk allocation is a clear understanding of the related risk(s) and a clear understanding of who is responsible for managing risk. This information is necessary to move to the next portion of the system safety process — risk acceptance and approval.

Management Risk Acceptance and Approval

A final aspect of a system safety process is to have a risk management structure that supports the safety culture in place, and a risk approval process that decides and documents that the system has achieved a level of acceptable risk.

A risk management structure should enable a safety culture that supports the people responsible for determining system risks and defining the requirements to prevent those risks or to manage those risks to an acceptable level. The safety culture should support the use of systems engineering and systems thinking to facilitate an environment that questions conclusions and assumptions, minimizes complacency, commits to excellence and fosters the company taking responsibility for safety matters. As reference, a more detailed description of "safety culture" may be found in ISO-26262 Part 2 – Management of Functional Safety – Annex B, from which the initial content of this paragraph was paraphrased.

A risk management structure should have a risk approval process that contains the following key aspects and/or characteristics:

- A method to understand risk based on its severity and likelihood of occurrence
- Clear understanding of which entity (i.e., company, customers, users, society, etc.) will be responsible for which aspects of the risk
- A method to approve the risk allotment to the company, customers and/or society by a management team with appropriate authority to make such decisions
- Documentation of the risk approval and supporting rationale

A final aspect of any risk approval process is that of "independence." Independence aims to have an unbiased point of view based on objective data and to be free from conflicts of interest. Independence in this context relates to organizational independence and is especially true for three main groups — those authoring the safety analyses and assessing risk, developers of the product itself and, finally, project managers. Working together, and with proper independence, these three groups form a platform upon which system safety assessment and risk approval may stand.

Conclusion

It's all about risk. The attributes of a good system safety process were detailed at the beginning of this paper. These steps enable risk to be associated with "bad things" that may happen during the system life cycle. System safety requirements are written to prevent or manage the severity and likelihood of occurrence of these bad things. Verification should demonstrate the system safety requirements that detect and mitigate bad things. Any remaining risk needs to be understood and approved by the proper management level and the target customers, and eventually be accepted by the marketplace at large.

A new interpretation of Figure 6 is helpful to summarize the importance of a system safety process.

Figure 7 illustrates how system safety requirements, developed using a robust system safety process, may identify risks that may lead to unwanted results such as mission loss, loss of intended utility and/or human harm. A management structure with proper organizational independence can decide the level of acceptable risk allocated to the company, customers or society at large.

About the Author

Mark A. Vernacchia is a GM Technical Fellow and is the principal system safety engineer for all GM propulsion systems worldwide. Mark earned a BS in mechanical engineering from Purdue University and an MS in engineering sciences from Rensselaer Polytechnic Institute (RPI). Mark is a professional engineer in the State of Michigan and is recognized as an Expert Systems Engineering Professional (ESEP) by the International Council of Systems Engineers (INCOSE). ◉

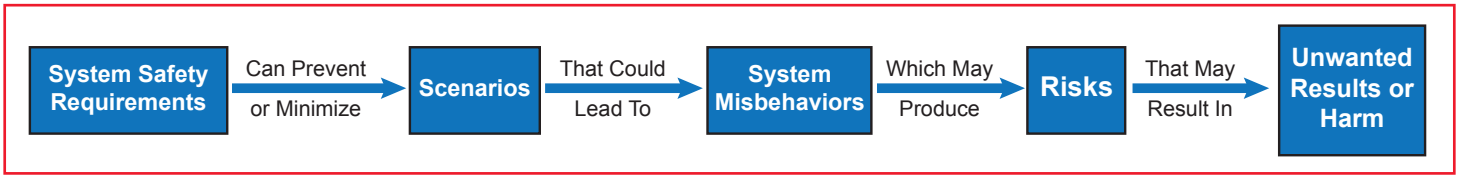

Figure 7 — System Safety Requirements Preventing or Managing Risk to Avoid Unwanted Results or Harm.