

Memorandum

From:	Peter G. Hereña
To:	Professional Staff, Technical Memo File
CC:	
Re:	Advantages of Risk-Based over Prescriptive SIL Selection
Date:	February 12, 2009
Reference:	TM011.pgh

Supplementing Risk-Based Approaches with Prescriptive Templates

One common question that continues to see debate in the safety industry is the decision to use prescriptive-based or performance-based standards. Prescriptive-based standards are mandatory and provide a set of “cookbook” requirements for a given function or hazard. Performance-based standards do not provide a set of specific requirements, instead opting to set a framework for the user to select a level of performance sufficient to satisfy the primary target.

In industry there are examples of successful prescriptive and performance-based process safety standards. API RP 14C: Recommended Practice for Analysis, Design, Installation, and Testing of Basic Surface Safety Systems for Offshore Production Platforms and NFPA-85: Boiler and Combustion Systems Hazards Code each have major prescriptive elements. ANSI/ISA-84.00.01-2004 (IEC 61511 Mod): Functional Safety – Safety Instrumented Systems for the Process Industry Sector is an excellent example of a commonly adopted performance-based standard.

This memo highlights the concerns associated with prescriptive standards and discusses why a performance-based risk program generally provides superior risk reduction at a lower installed cost, along with less maintenance cost. This memo also addresses what can be done to implement a performance-based risk program to yield consistent, but flexible, results. Finally the discussion categorizes the many factors that can affect the selected SIL level of a safety function and provides some anecdotes of how those factors have played roles in various projects.

Prescriptive Standards vs. Performance-based Safety Standards

The positives associated with a prescriptive safety approach are quite clear; it removes the burden of analysis from the end-user, it allows for very similar installation and maintenance costs from one project to another, and it allows for consistent operation and maintenance from one site to another.

However, there are also several drawbacks with the approach. It is not applicable to novel or unique situations, which limits the usefulness of the standard. It can become obsolete over time as technology improves, risk awareness increases and equipment failure modes become better understood. More importantly, it does not allow for local variance and therefore can lead to uneven risk exposure from one location to another.

Furthermore, prescriptive standards are more expensive to install and maintain because they tend to be excessively conservative, which results in an overdesigned safety function to mitigate the hazard risk.

In July 1988 a disaster occurred at an offshore oil platform Piper Alpha, which led to a thorough investigation by the UK HSE and the subsequent release of the so-called "Cullen Report." In the findings, part of Lord Cullen's recommendations was that a "rational, goal-oriented approach" to safety should be implemented. Such an approach would allow for the most effective adoption of technology and methods to achieve the objective. His argument was that as technology and engineering advanced exponentially, the risks associated with those advances can also rise. Therefore, as technology advances the appropriate tools should be deployed to address risk.

One potential drawback for a performance-based risk program is that it can result in uneven implementation of the standard if it is done without sufficient centralized guidance. If many sites are working in a vacuum then some locations could be overestimating (or underestimating) some key risks and mitigation factors. A methodology should be adopted that takes into account that similar hazards exist for commonly used equipment. Yet, there should be enough flexibility to provide for some degree of variation.

Safety Function Templates

One way to achieve a consistent, yet flexible, safety design is to start with a safety function template, which contains a preliminary description of the hazard as it is generally accepted in industry. The template also includes commonly seen initiating events and protection layers. During the SIL selection meeting it is stressed that the results are for the general case, and that it is the team's right and responsibility to modify, add or remove any item that does not fit their situation. It is of utmost importance to protect the voice of the team; failure to do so breeds a cynical attitude that is against the purpose of the exercise.

Several of our clients have requested this kind of service from Kenexis, including a small refiner's multi-site heater study, a large fertilizer producer's multi-site risk study, and a global exploration and refining corporation's refinery templates.

In these cases our clients have reported:

- Greater consistency between sites, which increases acceptance of the results. The team is able to dissect the risk factors that determine why the same function may have different SIL levels at different sites. This becomes important during implementation because the operations and engineers are not as likely to feel that arbitrary SIL levels have been chosen.
- Reduced variance in SIL levels, streamlined documentation and fewer exotic recommendations. This was especially true once the baseline was formed and tuned to the client's risk tolerance and practices (which could be different from the industry risk tolerance and practices).

- Better capability to analyze and address concerns that affect multiple sites. This analysis commonly started with a more detailed review of the hazard, possibly including a consequence modeling study, a scenario fault tree or other advanced statistical model. If the scenario was recognized as a real hazard that affected multiple sites, the need for broad mitigating action was better justified.
- Less time required for the meetings because every project was not a complete re-write of what is commonly accepted information.

One case in particular demonstrates the utility of a safety function template. A client who is a major fertilizer manufacturer requested a cluster of studies meant to develop SIL targets for their ammonia production plants. An ammonia unit contains numerous safety and equipment protective functions, including ten common functions:

- Steam reformer low steam-to-gas ratio
- Steam reformer main fuel gas low pressure
- Steam reformer main fuel gas high pressure
- Steam reformer firebox high pressure
- Secondary reformer high temperature/high air-to-gas ratio
- Methanator high temperature
- Synthesis Gas Compressor KO Drum high level
- Synthesis Gas Compressor auxiliary (lube, seal oil) functions
- CO2 Absorber low level
- Ammonia compressor KO Drum high level

This study included ammonia plants operating at various locations within the United States that used two different process unit designs. With respect to safety system design, the various plants did not historically have a high degree of coordination between sites. The results of the Safety Integrity Level (SIL) selection meetings at the first location formed the base for discussion at subsequent locations. Then Kenexis performed a review to compare the SIL targets selected at the plants for the above 10 common functions.

There are two ways to analyze the results; on an aggregate bases and on a per-function basis. Analyzing the per-function basis involves looking at the “spread” of selected SILs for a given function. That is, the differential between the highest SIL selected and the lowest SIL selected for each function. Table 1 shows the SIL spread for the ten critical functions:

SIL Spread (Difference Between High and Low SIL Selected)	Number of the above 10 Safety Functions with the SIL Spread
0	4
1	5
2	1

Table 1: SIL Spreads for 10 Common Functions in Ammonia Plants

A “SIL Spread” of 2 happened for 1 of the 10 functions. For this case, 4 of the 5 plants had chosen SIL-2 for the function and one had chosen SIL 0 (also known as SIL “A,” or a safety function whose performance target is less than SIL 1). The results were reviewed for the site that had chosen SIL-“A” and that team confirmed that there were critical differences between their site and the others that justified the selection of their target.

The Cost of Prescriptive Design

An aggregate analysis can be performed by determining how frequently the selected SIL target agreed with the most frequent SIL target for that function. First, for each function a “SIL Mode” is determined. The “SIL Mode” is the statistical mode (most frequently observed number) of the selected SILs for a given function. As an example, if the Ammonia compressor KO Drum high level safety function had selected SIL targets of 1,1,1,2,2 then the SIL Mode would be 1 for that function. Another safety function could have a different SIL mode, depending on the severity of the hazard and other factors.

Of interest is the frequency with which each site agrees with the SIL Mode of the function, the frequency with which the safety function is “overdesigned” (requires a LOWER SIL target relative to the SIL Mode), and lastly the frequency with which the safety function is “underdesigned” (requires a HIGHER SIL target relative to the SIL Mode). Table 2 below shows the results. Note that 2 of the critical functions have not yet been analyzed at one site, which results in the total number of analyzed functions of 48.

SIL target selected	Frequency
Matches the SIL Mode	39 (81%)
Less than SIL Mode (potentially overdesigned)	7 (15%)
Greater than SIL Mode (potentially underdesigned)	2 (4%)

Table 2: Frequency of SIL selected relative to SIL Mode

If a prescriptive standard were utilized that applied the SIL Mode to all functions, the data imply that a considerable number of them could be overdesigned. That could result in requiring one or more additional sensors; at the installed cost of \$8,000-10,000 per sensor the resulting increased cost could be significant. If any functions required additional final elements, particularly valves, the cost increase would only escalate. Perhaps worse than the additional expenditure is that a small portion of the functions could end up underdesigned, and thus lead to elevated risk exposure.

SIL Selection: One Size Fits All?

Although a template can provide more consistent results with better team satisfaction and better time efficiency, some variation of results inter-site and even intra-site should be expected. When a safety function is implemented, four factors typically determine the required safety performance (SIL target):

- 1) The Consequence associated with the hazard
- 2) The Initiating Events that can potentially lead to the hazard
- 3) Safeguards or Protection Layers that can prevent or mitigate the hazard
- 4) Hardware Details and Testing Frequency of the safety function

For similar protective functions, a number of elements impact the above factors, which can in turn affect the required safety function performance. This paper shall touch on just some of these elements, and will give some anecdotal examples of how unforeseen variation contributed to an atypical SIL target.

Consequence:

Physical Layout – this has a profound effect on hazard severity. An event such as a pump or compressor seal fire that is located far from routine pedestrian and vehicle traffic, such as in a tank farm, probably has a significantly reduced potential for injury or loss-of-life. An excellent example of this comes recently from one of our refinery clients. They have several hydrotreating and hydrocracking units at one site; it is common for hydroprocessing units to have several pumps (charge, wash water, amine) where the pump differential is so high (1900 psi) that a Safety Instrumented Function (SIF) protects against backflow. In the event of pump failure the flow could run backwards through the pump, where loss of containment of flammable material could occur if the SIF does not close backflow prevention valves at the pump discharge.

Normally this hazard is serious enough, but when the plot plan was examined it was realized that these pumps were located next to a satellite operation building, and also in very close proximity to a pathway that was constantly in use by operation and maintenance personnel. In most other units personnel might be near the pump only a few times per day; therefore the risk was much higher than normal for this case. In this case, underestimating that risk could have led to an unacceptably high risk exposure to personnel in the unit area. The solution was not necessarily more instrumentation, but to reduce the frequency of personnel in the area and thus reduce the risk of injury in the case of an undesirable incident.

Past Exposure to a Hazard (Sensitization) – As humans, we are hardwired to recall more sensational events in our lives, allowing the more common events to fade into the background of our memory. Institutional memory can work the same way, setting higher sensitivity for certain events that have previously happened onsite. This aversion to risk exposure from recurring events is sometimes justified because in the event of recurrence the corresponding punitive measures could be more severe than if it were a first-time incident.

A large US refiner operates a complex that has a Fluidized Catalytic Cracker (FCC) unit. An FCC unit will commonly use a protective function that prevents against “reversals.” An FCC unit has two main vessels, a Reactor and a Regenerator. The catalyst passed through the Reactor where it contacts with the heavy oil to be cracked. As part of the reaction coke is deposited on the catalyst. The coke is oxidized in the Regenerator, which regenerates the catalyst and allows the cycle to continue. However, if a process upset occurs that allows hydrocarbon into the regenerator (a “reversal”) the result could be a rapid uncontrolled combustion of heavy oil in the oxygen-rich regenerator. This event is not expected to result in an explosion or external fire, although it could result in residual hydrocarbon, SO_x, NO_x and catalyst to spew from the regenerator to the atmosphere in a large cloud. The primary concern is not safety, but rather the negative environmental and public relations impact associated with a “reversal.”

At one refinery location a community park and major interstate highway were located downwind of the FCC unit. "Reversal" events had happened previously that resulted in acute visibility reduction on the highway, although to that point it had not resulted in a highway accident. The "reversal" events had caused some temporary damage to the community park and a negative public relations event. This required the integrity levels of the protective functions to be much higher than normal to reduce the chance of the event recurring.

Initiating Events:

Frequency of Utility Disruptions – A major event, such as a sudden loss of power or a rapid change to the fuel quality (pressure, heating value) that feeds a fired heater, is a site-specific issue. The frequency of these events can drive the overall risk reduction a SIF must provide.

A recent example is from a major chemicals manufacturer who operated several fired heaters at one site. One problem that was site-specific was serious fluctuations in the fuel gas pressure, which could be severe enough to result in burner flame-out. If the fuel gas pressure returned to normal after the loss of flame, the accumulation of raw combustibles could create an explosion hazard. The event frequency at that site was several times per year, an order of magnitude more frequent than is normally considered. In this case the end-user realized there was a significant degree of residual risk, so they considered providing a higher SIL target for the function. At a location where an end-user utilizes multiple independent fuel gas sources, or where the lone fuel gas source is extremely reliable, the SIL target could be reduced.

Equipment failure rate – This is a factor that will not only be dependent upon site, but is also dependent on the exact model and type of equipment. Equipment failure rate is also linked to localized issues such as climate, instrument air quality, the propensity for fouling, maintenance practices and equipment age. It is not uncommon for a pump or compressor to have failure rates ranging from multiple times per year to once every few years at the same site. This can also apply to non-rotating equipment if a history of specific valve or instrument trouble exists.

Additional (or Fewer) Initiating Events – Naturally, as additional credible initiating events are added or removed when analyzing the scenario, the level of risk reduction required could be affected. This is especially true when the frequency of the initiating event being added or removed is the most frequent (and thus dominating) initiating event.

An example of this is relayed to us from a client who had several hydrocrackers operating in numerous refineries nationwide. All hydrocrackers employ a high reactor temperature SIF that depressurizes the reactor in the event of severe temperature excursions so that the temperature does not climb high enough to cause loss of containment.

One of the events that could start a temperature “runaway” is an increase in olefinic material in the feed, the hydrogenation of which is a highly exothermic reaction. At one refinery studied, the feed to the unit had large olefin concentration fluctuations several times per year. Most other initiating events that lead to temperature excursions are related to control loop failures, which generally have a frequency of once in ten years. The high-frequency event, which may not apply to all sites, could lead to either excessively conservative safety functions or unmitigated risk exposure.

Safeguards or Protection Layers

Operator intervention – The capability to claim operator intervention as a mitigating credit is typically based upon several factors, including the existence of an independent sensor with alarm, the likelihood of the operator to acknowledge, diagnose and initiate a response to the alarm, as well as sufficient time for the operator to take action that prevents the hazard. The assumptions that permits the credit may be plausible at one site and implausible at another, even though the safety functions are in essence the same.

As an example, consider a client at a chemical complex located in the United States. Some of their processes have concerns related to gas blow-by. This occurs when a gas-liquid separation vessel loses the liquid level and allows high-pressure gas to rapidly pressurize downstream vessels and equipment. In such cases if the pressure is in excess of the downstream design pressure rating the result could be a severe loss of containment. In many blow-by scenarios the capability to take operator intervention is dependent upon how quickly the level could be lost relative to the operator response time.

The client considered implementing low-level SIFs across the site (to shut in the liquid outlet before the level is lost) when the service did not allow pressure relief devices to provide sufficient risk reduction. Upon further review the team realized that for a large proportion of these scenarios the vessel was physically close enough to the operator control room that SIFs for these scenarios were not deemed necessary. SIFs were added in situations where some of the equipment was more remote and the chance of intervention was low. Only a careful review of the equipment data sheet, plot plan and operator reaction time could have allowed this cost saving.

Equipment Auxiliary Systems –Auxiliaries on pumps and compressors are subsystems that allow the piece of equipment to properly operate. Compressor auxiliary systems include lube and seal oil pumps, alarm systems, seals and advanced control systems. There are numerous cases where, for example, a dual seal has provided a protection layer that reduced a SIF performance requirement (SIL rating) or entirely obviated the need for the SIF. There are other cases, especially in the case of older equipment, where certain critical alarms were either not installed, had been disabled or did not provide an audible or visual indicator to the operator. A prescriptive standard does not

easily allow for such a variance in equipment; in the event that it did the result could be a prohibitive expense to get every piece of equipment upgraded to the same standard.

Hardware and Testing Frequency

Sensor Type – A sensor's failure rate is one of the critical components that determine whether or not a certain level or risk reduction has been achieved. Sensor technology, including sensing method, setup and diagnostics can have widely varying failure rates. Even given the same measurement technology, different manufacturers and models can have widely varying failure rates. For example, there are at least three commonly used types of level sensors: radar, differential pressure and float. Between published sources, maintenance records and manufacturer-supplied information the dangerous undetected failure rate (λ_{DU}) between these three can be between $4E-8 \text{ hr}^{-1}$ and $6E-6 \text{ hr}^{-1}$ for common brands!

This divergence might not have a sizeable impact when evaluating their performance for SIL-1 SIFs because that target is relatively easy to meet for most configurations. However, if higher performance (such as SIL-2) is desired, these differences can become very important. Consider a safety function where level transmitters are used in 2oo3 voting configuration, through a SIL-3 rated PLC and then to 1oo2 shutoff valves. Using generic data and assuming a 4-year test interval, a sensor with a λ_{DU} of $5.7E-7/\text{hr}$ could allow the function to meet SIL-2 requirements, whereas a sensor with a λ_{DU} of $1.3E-6/\text{hr}$ might not.

The case for existing units becomes much more complex as there is no good answer to the question of "is what I have good enough?" Most new safety systems will use a SIL-rated PLC with transmitters that have a standardized configuration. But for existing units these items are often not clear. If a regular off-the-shelf PLC is utilized as part of a SIF it will likely not have SIL-2 performance. If a relay panel is used it might meet SIL-2 targets, but the arrangement of the relays and other panel hardware become critical in determining performance. If the function is energize to trip (ETT) the failure modes are generally very different than for standard de-energize to trip (DTT) functions, with different performance characteristics. If these are not addressed properly the result could either be expensive over-instrumentation or insufficient risk coverage provided by the safety function.

Function Architecture – "Architecture" for safety functions refers to how the sensors and outputs are voted together (1oo1, 2oo3, etc...). Different architectures provide for varying safety performance (SIL levels) and nuisance trip protection. Most functions are fairly straightforward but when multiple sensors or final elements are used in complex configurations, the safety performance can be difficult to determine.

A client recently was using a safety function to prevent uncontrolled exothermic reaction (temperature "runaway") in two of their process units. The two units originally utilized a common design 30 years ago, but since then each had gone through a

number of modifications. For unit #1, in order for the safety function to be successful a single valve needed to close. For unit #2 it was determined that, in order for the function to succeed, valve A needed to open AND either valve B had to open (opposite of the valve failure orientation) or valve C had to close. In other words the outputs had an architecture of A and (B or C). To make things more complex, the valves were a “mixed bag” of motor operated valves, air-operated shutoff valves and air-operated control valves. Does either unit have sufficient risk reduction? A performance-based standard can answer the question with some authority. A prescriptive standard is likely to lead to a complete redesign of one of the units; an expensive undertaking that might not have any justification from the point of view of risk reduction.

Utilizing non-SIS Equipment – The ISA-84 guidelines generally do not permit the use of process control equipment as part of a SIF if that same equipment could place a demand on the SIF. However, there are some cases where process equipment can be used, and when this is permitted it could result in lower installation and maintenance cost without degrading safety performance.

An example of such a situation is with a hydroprocessing unit pump backflow SIF previously discussed. The two typical initiating events for this function are Pump Failure and possibly Power Failure. Failure in the control loop at the discharge of the charge pump is not typically considered because if it has failed close to stop forward flow, it should also stop reverse flow through the pump. In such cases the discharge control valve could be modified to include an interposing solenoid between the actuator and positioner that is activated by the SIS. The SIS will de-energize the solenoid, release the air from the (presumably fail-closed) valve and thus close the valve regardless of the normal control signal. This practice can provide a cost benefit, although to safely implement it the scenario must be individually reviewed to ensure that the equipment is sufficiently separated from any initiating events.

Conclusion

The debate between prescriptive-based and performance-based standards will undoubtedly continue. Each has its own advantages and disadvantages, but a performance-based safety standard is flexible enough to provide the optimal result in cases where numerous factors affect the required safety function performance. The use of a safety function template, undertaken with centralized support and administered by a consistent facilitator, can greatly reduce SIL variance between disparate sites and it can account for the many factors that come into play when selecting a SIL target. A safety function template provides a common base for all units, which can assist an administrator at the corporate level by providing insight as to when broad mitigating action may be required. The performance-based safety standard is a tested and well-proven method that can be applied to any situation and yield the most cost-effective solution.