# Can we achieve Safety Integrity Level 3 (SIL 3) without analyzing Human Factors?

74th Annual Instrumentation and Automation Symposium for the Process Industries
January 22-24, 2019

**Keith Brumbaugh P.E., CFSE**
Senior Principle Specialist
aeSolutions

## Disclaimer

*The following paper is provided for educational purposes. While the authors have attempted to describe the material contained herein as accurately as possible, it must be understood that variables in any given application or specification can and will affect the choice of the engineering solution for that scenario. All necessary factors must be taken into consideration when designing hazard mitigation for any application. AeSolutions and the authors of this paper make no warranty of any kind and shall not be liable in any event for incidental or consequential damages in connection with the application of this document.*

## Abstract

Many operating units have a common reliability factor which is being overlooked or ignored during the design, engineering, and operation of high integrity Safety Instrumented Functions (SIFs). That is the Human Reliability Factor. In industry, there is an over focus on hardware reliability to the n'th decimal point when evaluating high integrity SIFs (such as SIL 3), all to the detriment of the human factors that could also affect the Independent Protection Layer (IPL). Most major accident hazards arise from human failure, not failure of hardware. If all that were needed to prevent process safety incidents is to improve hardware reliability of IPLs to some threshold, the frequency of near miss and actual incidents should have tailed off long ago - but it hasn't. Evaluating the human impact on a Safety Instrumented Function requires performing a Human Factors Analysis. Human performance does not conform to standard methods of statistical uncertainty, but Human Reliability as a science has established quantitative limits of human performance. How do these limits affect what we can reasonably achieve with our high integrity SIFs? What is the uncertainty impacts introduced to our IPLs if we ignore these realities?

This paper will examine how we can incorporate quantitative Human Factors into a SIL analysis. Representative operating units at various stages of maturity in human factors analysis and the IEC/ ISA 61511 Safety Lifecycle will be examined. The authors will also share a checklist of the human factor considerations that should be taken into account when designing a SIF or writing a Functional Test Plan.

### Introduction to Using Quantitative Human Reliability

It is widely accepted that a SIL calculation is primarily a hardware reliability calculation, which uses probability distributions based on random samples. Qualitative methods (i.e., the safety lifecycle) are used to control systematic failures, such as human impact. This is established practice. However, there are several reasons one may choose to undertake a quantitative evaluation of the human failure component in a Safety Integrity Level (SIL) calculation, or other barrier or Independent Protection Layer (IPL).

1. By considering human reliability it provides a more realistic estimation of the uncertainty associated with the random hardware calculation. Classical confidence limits based on Gaussian style probability distributions are inadequate for treating systematic uncertainty. The issue is potential underlying skewness or Fat Tails of the assumed probability distribution.
2. Related to number 1, from a process safety perspective, by ignoring the human factor in a quantitative risk determination, e.g., Layer of Protection Analysis (LOPA), we are significantly overstating the forecasted mitigated event frequency. This is especially true for LOPA targets of 1e-4 per year or lower. Whenever a model (e.g., LOPA) predicts that a failure will occur with a negligible chance, the probability that the model can fail becomes important.
3. While the Safety lifecycle methods are sufficient to identify and control systematic error associated with a single barrier (IPL, etc.), they are inadequate to predict emergent or non-linear failure modes for multiple barriers in a common threat path. That is, failure modes that arise (emerge) due to the interactions or dependencies between barriers.
4. The qualitative evaluation that comes with a human reliability assessment (for example, a qualitative Task Analysis), is the biggest benefit associated with quantifying human error. Human error rate averages suffer the same problem as hardware reliability averages (small sample sizes, using time-series based on the past history, etc.). However, the qualitative benefit that comes from the opportunity to identify error likely situations or latent failure conditions (used to derive said averages), far outweighs the calculated averages themselves.

To perform quantitative analysis of systematic (e.g., human) impact, a different type of thinking and tools apart from standard hardware reliability concepts are needed. To “think in systems” we must balance our focus from an outcome based statistical data approach (e.g., SIL calculations as typically practiced) to an approach that is subjective and more focused with the causal process of the outcome, rather than the outcome itself (e.g., a Bayesian approach). For outcomes related to process safety, Mother Nature does not tell us how many faces are on the die, or how many slots are on the roulette wheel (this is known as the “gambling with the wrong dice” problem). Moreover, much of the randomness in process safety comes from our lack of certainty in our knowledge about pathways to events (i.e., Bayesian randomness). We calculate the likelihoods of events with no or little thought to the uncertainty in the number. Is a number such as 1e-5 events per year comprehendible (humans as a species have only been on planet Earth for 100,000 years)? Can 1e-5 events per year ever be verified (this is known as the “problem of induction”)?

Even though it may be inconvenient to think about, the practice of Process Safety involves an engineering and hard sciences component, and an equally important sociological/ psychology

based component. Numerical process safety is largely based on classical statistical analysis (e.g., accident statistics, probability of failure on demand, risk targets, etc.). Human thinking and intuition is cause-based (e.g., Hazop decisions, Operational decisions made in real time, etc.). Bayesian methods can be used to bridge the two.

To begin we'll look at two familiar, yet difficult to define concepts: **probability** and **randomness**. Probability is based on the concept of a random event, and statistical inference (i.e., induction) is based on the distribution of random samples.

### Probability from Objective to Subjective Interpretations

What is Probability? This short list is ordered from objective to subjective interpretations.

- A computation (of odds or likelihood) based on frequencies of past events and the associated confidence interval (uncertainty). This is territory we are most familiar with (i.e., SIL calculations, etc.). It also includes Human Error Probability (HEP), defined as the number of errors made divided by the opportunities (based on past occurrences). These are the "known knowns."
- A measure of uncertainty. This is the meta-uncertainty (i.e., the uncertainty of uncertainty). These are the known unknowns (what we know but can't quantify), and the unknown unknowns (what we don't know we don't know)." We rarely even get to this level. The worst of these are the "unknown knowns," i.e., that which we refuse to acknowledge is true. For example, we have poor bypass control of our safety functions but ignore it.
- A belief in an outcome. Here we are starting to look forward in time using a subjective (i.e., Bayesian) belief.
- A belief in the existence of pathways to alternative outcomes. After an event happens, these are known as "counter-factuals." Looking forward, being able to visualize counter-factuals is a uniquely human quality, and is the basis of future AI (artificial intelligence).
- A tool to help think about all of the above.

### Randomness from Classical to Bayesian Interpretations

Randomness is the property that makes statistical calculations come out right (this is a systems definition of randomness). Often, we are fooled by randomness, because there is a causal (or systematic) factor that we have missed. In large aggregates, sufficient randomness exists for the Law of Large Numbers (i.e., averages) to apply (e.g., the number of air molecules in a room, or the number of drivers in the United States). Conversely, many organizations made up of people are considered medium or small numbered systems. Averages can be deranged or skewed in Medium Numbered systems. The foundation of process safety management at a facility rests on people, and is a Medium Numbered system. For this reason, calculating process safety related averages (e.g., SIL calculations, QRA, etc.) often results in an inordinate use of time and money if the focus remains on the number itself. In would be better to invest that additional time and money to identify the potential for skewness or derangement in the number. The tool for this is called Task Analysis.

In the classical interpretation, randomness means "equally likely." This property of randomness has been important since antiquity where randomizers (i.e., lots or dice) were used to make decisions. To ensure fairness and prevent dissention, equally likely outcomes were necessary. The equally likely property of randomness survives today in games of chance (e.g., a balanced die in gambling) for the same reason.

This interpretation of randomness ("equally likely") is not sufficient for process safety. Granted, a *fixed bias* can be added to any equally likely outcome, to make the chance of each outcome unequal (called a non-uniform distribution), and the odds computed. For example, if you wanted to compute the odds of an inferior baseball team winning a 7 game series, a bias can be applied to favor one team over the other. In process safety, the pathways to undesired outcomes are largely unknown and not equiprobable. It is not the same as rolling dice or flipping coins where the likelihood of each outcome is known, and the bias is fixed.

Bayesian randomness is based on how much information is known about the process. If no prior information exists, the process is completely random. If the process is well defined (i.e., known), it is considered deterministic. "Systematic" is a degree of knowledge that sits between a completely random process and a well-defined deterministic process.

Process safety as a discipline often sits in this systematic (i.e., systems) region. Our systems are too organized to be considered random, yet too complex to be considered deterministic.

As process safety practitioners, we look into a largely random future. Random is this context means lack of knowledge or understanding about what could happen and how. Often the "how" (the pathways to undesired events) is where the nasty surprises occur. For this reason, looking forward in time, any potential unwanted event is considered random (there is not enough information on how it will occur). After the fact (i.e., after the event), the failure has the potential to be classified as systematic (it is knowable in hindsight). This interpretation of randomness is consistent with Bayesian probability. Whether a failure is considered random or systematic depends on how much is known about the failure. Looking forward in time, very little can be known (all the ways a failure can occur is often less understood than what failures can occur). In hindsight, the dots can be connected and causes determined, so that the failure may be considered systematic. Hindsight bias occurs when we think the future is more knowable than it really is, based on a seemingly deterministic past (which is an illusion, i.e., hindsight is *not* 20/20).

What does this imply for a SIL calculation (or any other reliability calculation)? Include the failures that you know can happen and how often (based on past experience). However, how those failures can happen (i.e., the many pathways) can never be completely known in advance. This represents the largest uncertainty, and usually creeps in due to human, organizational, and other system conditions.

**Bayesian Wars (Bayesian versus Classical Statistics)**

There are two fundamentally different kinds of statistics recognized in the field of mathematics. These are classical statistics and Bayesian statistics (a.k.a., frequentists versus Bayesians). Bayesian thinking is underutilized in Process Safety.

Much of quantitative process safety (e.g., SIL calculations, LOPA, QRA, etc.) rests on classical statistical techniques (i.e., the frequency based approach). However, small sample sizes (i.e., insufficient or un-sampled randomness), non-representative samples, under-reported events, data handling errors, and using time series based on past events (i.e., hindsight bias), all contribute to what is known as the "5-sigma" problem (a.k.a., Fat Tails or skewness). Any outcome that is purported to occur with a likelihood of 1e-4 (frequency or probability) or lower is at risk for more than just being unlucky.

Bayesian techniques are under-utilized in process safety. Bayesian methods can account for the human component impact of the specific site, which is missing in frequentist techniques (i.e., generic data, assumed probability distributions, etc.).

Most engineers are familiar with Bayes' rule as the defining theorem of conditional probability (i.e., P(A │ B), read as "the probability of A given that B is true"). However, this is only a small part of the overall power of Bayes. More importantly, the vertical bar ( │ ) in Bayes' rule (notation introduced in 1931 by Harold Jeffreys) is an empirical claim to faithfully represent the English expression "given that I know." It is an epistemological statement (pertaining to knowledge) rather than a statement of frequencies and proportions. As such, it represents a subjective degree of belief.

Using a Bayesian approach can correct and update the results initially obtained from a frequentist approach (the "prior"). In addition, Bayesian confidence intervals are real-world limits that reflect actual conditions in your plant.

**The Problem of Induction**

Inductive statements are verifiable in some manner (experience, statistics, etc.). Statistical inference (making conclusions about a population based on sampling) is a form of inductive reasoning. Process safety is awash in this type of inductive reasoning. For example, using metrics and indicators to draw overall conclusions about the state of process safety in a facility. The "problem of induction" is that many inductive statements that are made are not verifiable (or even meaningful). Here is a short list of the problem of induction related to process safety:

- How "safety" is defined is not even clear. Is safety the absence of something (i.e., unwanted events)? Is it freedom from harm (again, the absence of something)? We'll see below the problem with trying to "prove a negative" (i.e., prove something doesn't exist, or can't exist). Is safety the presence of something (i.e., success, capacity to do right, culture, resilience, reliability, presence of barriers, meeting tolerable risk, etc.)?

If so, we need metrics and indicators that actually measure these concepts (i.e., are we asking the right questions?).

- The Law of Small Numbers when making a statistical inference. This includes small sample sizes (i.e., insufficient or un-sampled randomness), non-representative samples, under-reported events, data handling errors, and using time series based on past events (to predict the future).
- There is a special case of the Law of Small Numbers that applies when making decisions under uncertainty (i.e., applies to nearly all of Process Safety work). When humans make decisions we do so based on a very limited sample size, mostly our direct personal experience. This has the potential to introduce systematic bias into a system. In other words, the source of all systematic bias is (ultimately) the human mind, which uses heuristics (mental short-cuts) when making decisions under uncertainty. This is known as System 1 (fast) and System 2 (slow) thinking. Examples of heuristic tools such as a PHA-LOPA template, a risk matrix, or a SIL calculation model, all have the potential to introduce systematic bias.
- The perfect bias engine (a.k.a. the human mind).
- The Problem of the Rare Event. This comes under different names such as:
    - "Absence of evidence is not the same as evidence of absence." (i.e., trying to prove that something can't happen, or worse, assuming that something can't happen)
    - The Black Swan (i.e., a rare event, subjectively speaking). If you think Black Swan theory means BS, you are at risk of one happening to you.
    - 5+ sigma events (i.e., events with low probability of occurrence, 1e-4 or less)
    - Fat Tails/ Skewness (i.e., underlying asymmetrical probability distributions)
    - The Turkey Trap (i.e., Using the past to predict the future, you could suffer the same fate as a Thanksgiving turkey - (i.e., things were going just great, until they went really bad)
    - Fooled by Randomness (i.e., "My LOPA calc shows I meet 1e-6, so I'm safe.")

The problem of the rare event can be demonstrated with a simple thought experiment. Suppose I have an urn containing a large number of both red and black balls. The more information I have (gained by sampling the urn, 'n'), the more confident I can be about the outcome, i.e., the proportion of red to black balls (notice Bayes showing up here again). As I sample the urn, common statistical methodology says my confidence in the relative proportion of red to black balls will increase as the square root of 'n' times increase in the sample size. That is, for example, if I double the number of samples, my confidence level increases by 1.41.

Now suppose the urn contains a skewed distribution of balls: very few red balls (representing a rare event) and a large number of black balls. As I perform sampling, my knowledge about the absence of red balls will increase very slowly. Much more slowly than the expected square root of 'n' rate assumed above. If a red ball is found (drawn from the urn), we've confirmed the presence of a red ball, but it's too late (the rare event has occurred!).

**Using Human Reliability Assessments to Support Quantitative Analysis**

A Human Reliability Assessment (HRA) can help bridge the gap with any and all of the issues discussed above (i.e., uncertainty, systems issues, Bayesian probability, and The Problem of Induction).

Traditionally, a Human Reliability Assessment is a tool used to evaluate human performance in the context of:

- Human error as the initiating cause for hazard event scenarios.
- Human response to abnormal events as a safeguard or independent protection layer (IPL).
- Failure to restore safety critical equipment following ITPM (inspection, test, and preventive maintenance) activities, or other error made as part of maintenance.

HRA may be used to "sharpen the pencil" in these cases to provide more realistic estimates of unrecovered human error probability. Moreover, the qualitative capability of HRA is greater than the quantitative. The methodology can handle the same operator making and recovering from his or her own error, as well as handling human redundancy. The method is particularly well suited to account for dependencies within and between individuals.

There are a myriad of HRA methodologies, some public, some proprietary, and extending from 1st generation to 3rd generation methods.

The Health and Safety Executive (HSE) in the UK considered that it would be useful to be up to date with developments in the field of quantitative HRA methods and to have knowledge of the capability of the tools and an understanding of their strengths and weaknesses, to improve consistency, and determine acceptability of their use. To this end, in 2009 the HSE published document **RR679** "Review of Human Reliability Assessment Methods," to summarize review of 72 HRA methods that were "potentially relevant to HSE major hazard directorates." Of the 72 potential HRA tools, 17 were considered useful for major hazard directorates. For example, THERP (Technique for Human Error Rate Prediction) is one of the 17 methods considered to be useful by the HSE.

The general steps to perform an HRA are:

1. Identify the task to be studied. This can come from a PHA/ LOPA, SIL analysis, audit, alarm study, etc.
2. Determine if you will be solving for a probability (e.g., as part of an independent protection layer) or frequency (as an initiating cause). This will affect how you set up the human factors worksheets.
3. Gather process safety information (applicable procedure, P&IDs, etc.).
4. Perform the task analysis focusing on Human Factors (see **Table 1**):

- Talk-through of task with Operations and Engineering
- Walk-through with Operations demonstrating task (live or simulated)

5. Create a task analysis table
6. Quantify the task analysis table
7. Recommend error reduction techniques
8. Report results to Team and incorporate back into PHA/ LOPA, etc.

Analysis of Human Factors is an important feature of task analysis. Every task analysis will include at least one of the scientific areas found in **Table 1**.

**Table 1**. Human Factors is a Science

| | |
|---|---|
| Alarm Management | QRA (Quantitative Risk Analysis) |
| Automation | Risk Perception |
| Crew Resource Management (non-technical Team work) | Safety Culture |
| Design & Installation | Situation Awareness (What?, So what?, What now?) |
| Fatigue, effects of | Stress, effects of |
| Human-machine interaction | Codes and Standards |
| Operating and maintenance Procedures | Left Blank |

**Sources of Error and recovery factors**

Below will identify some potential negative human interactions with a Safety Instrumented Function (SIF) design through the entire SIF Lifecycle. This list is not intended to be exhaustive, but is provided as an example of sources of error and to encourage an end user to think of their own unique sources of error and to address recovery factors that could be utilized to overcome these errors. Recovery factors are an integral part of Human Reliability Analysis (HRA). Humans make and recover from their own errors every day. In addition, human redundancy acts as a recovery factor. Any realistic estimate of human error should include recovery factors.

**Analysis Errors - LOPA Team**

Layers of Protection Analysis (LOPA) Team errors could be introduced from an inexperienced team. The team might not be clear on the severity of the hazard, or could call for a SIF design that isn't practical or has independence issues. The team may also poorly understand LOPA Templates and could credit a SIF to close a LOPA gap, but in actuality there is still a gap from either oversight or overconfidence.

Some LOPA Team recovery factors could include utilizing an experienced LOPA facilitator, the plant personnel could call for a more senior team member when a junior team is struggling.

Bad SIF design discovered can be recovered by a competent SIS engineering team. Gaps left in a LOPA could be discovered by competent Functional Safety Assessment (FSA) assessor.

NOTE: For all listed recovery factors, the effectiveness is entirely dependent on the competency of the individual players and can vary greatly. A competency assessment of each person mentioned as a recovery factor could help determine the recovery available.

**Analysis Errors - LOPA Facilitator**

LOPA Facilitator errors could include an inexperienced facilitator or poorly trained facilitator. In a LOPA, the inexperienced facilitator could miss the real hazard and lead the team to pick a lower severity than is the real potential. The facilitator could also apply IPLs with independence issues between the SIF. This could result in a SIF that should actually have a higher design target or be re-designed.

The facilitator could also apply Conditional Modifiers that aren't valid. For example, the CCPS book "Guidelines for Risk Based Process Safety (CCPS 2007) Appendix B – Peak Risk Concepts" warns that this concept has potential to be misapplied and credit given where it should not have been.

Some LOPA Facilitator recovery factors include Plant management review of LOPA results. A competent SIF design team could find IPL issues. A competent independent review of the report could find misapplied conditional modifiers or other poor credits.

**Analysis Error Recovery Overall**

Overall, there will not be many recovery factors available for a misidentified hazard (severity called out lower than actual), but competent reviews after a LOPA should be able to detect most errors in IPL independence and misapplied conditional modifiers. Additionally, a Revalidation PHA or Project PHA may serve as a recovery factor; however these can occur many years later.

**Design Errors - Process Engineer**

Process engineering errors could manifest as bad process data, which could include

- Incorrect working temperature/ pressures, could lead to sensors reading the wrong value
- Incorrect maximum design temperature and pressure, could lead to an instrument failing prior to the time of a SIF demand
- Incorrect shutoff differential pressure identified which could lead a SIF valve to not have enough force to close when commanded
- Overly optimistic process safety time or incorrect hazard set point beyond the point of no return (i.e. in the case of a runaway reaction). E.g. SIF is set to a set point that will be too late to prevent the hazard and will never respond in time to a real hazard.

Some Process engineer recovery factors include competent back checks of all data by other process engineers. The SIS design team can also back check process safety times provided if a number does not look right. In addition, instrumentation is generally over designed such that if the operating limits are called out too low, the instrument could still remain functional.

Overall, recovery will be low of process engineering errors. Other disciplines would need to be on their toes to catch base process data errors. A competent independent back check has the

most chance of success. Potentially some errors, such as incorrect shutoff differential pressure, can be recovered during commissioning and start-up.

**Design Errors - Piping Engineer/ Designer**

Examples of piping engineer/ designer errors include bad P&IDs. Perhaps the wrong line class or material requirements are listed for critical components, leading to undersized equipment or incompatible materials which could cause the equipment to fail. P&IDs are considered governing documents in most cases. Potentially if a line has been marked as class 150# but in actuality the design pressure ought to have the line class at 300# or 600#; a 150# valve might fail to prevent the hazard when demanded at the actual hazard pressures and temperatures.

Some piping engineer/ designer back checks include competent squad checks, process engineering back checks, and instrumentation engineering back checks. These are all generally performed over the course of most projects. Additionally, PHA teams scrutinize P&IDs closely. There are also P&ID reviews by the project team before a PHA.

Overall, recovery from piping errors should work themselves out based on reviews by other phases.

**Design Errors - Safety System Engineer**

Examples of safety system engineer errors include numerous documentation and specification errors. At best these errors might underspecify a system, at worst there is potential for erroneous requirements and designs. Some documents of concern include Safety Requirements Specifications (SRS). An incomplete SRS might misidentify potential problems, independence issues could be missed, potentially the engineer could neglect to call out various requirements that should have been analyzed (any parts of IEC/ISA 61511 Clause 10). Another document at risk is SIL calculation/ verification reports. These could be done incorrectly and might incorrectly conclude a SIF is far better than if it had been properly modeled. An inexperienced engineer might not know to challenge failure rate data that is overly optimistic (outside of reality). The engineer could botch the calculation, such as assuming redundancy of components where there is none (e.g. modeling a 2oo2 system as a 1oo2 system). Common cause factors could be misapplied lower than justified, or not even applied at all. The inexperienced engineer could poorly document assumptions which the client needs to know, such as diagnostic credits/ deviation alarming credited in the SIL calculation; the client might not know they must program the diagnostic.

Some safety system engineer back checks could include document reviews by competent and senior practitioners. The senior practitioner could identify missed items, however if the errors are extensive it might be hard to find them all. An FSA could discover if credits were applied to meet a design target (such as diagnostics in a SIL calculation), but were not documented in the SRS. Items like this could still potentially slip through the cracks. The client team could catch some errors, however if the client sought out an outside firm in the first place to design the SIF based on lacking the expertise internally, the client might not provide a sufficient back check.

Overall, if an inexperienced safety system engineering team is utilized, the probability of a good back check is poor at best. When an inexperienced firm is utilized, especially in high SIL applications, high levels of independence on the back checks and FSA's should be utilized. If a firm's competence is unknown by the client, the client should highly consider utilizing as much

independence as possible on the FSA and other assessments of any high SIL documents. This includes up to seeking the opinion of an independent company considered an expert in the industry. Note that this is in line with IEC/ ISA 61511 (2018 and 2004) clause 5.2.6.1.2, and especially guidance from part 2 “the requirement for an independent assessor may have to be met using an external organization.”

**Hardware/ Software Errors - SIS PLC application programmer**

Some examples of Safety Instrumented System (SIS) Programmable Logic Controller (PLC) application programmer and vendor errors are as follows. There are ample opportunities for an application programmer to misread the requirements for programming the SIS PLC from descriptions in the SRS and Software Requirements Specifications (SWRS). It is very easy for numbers (e.g. transmitter scaled ranges and trip points) to be "fat fingered" and entered incorrectly, potentially being an order of magnitude off. There can be problems with interpreting the design documents, or the potential for programming logic in a way that was never specified. These errors are even more probable when an outside company other than the team/ firm which specified the SIS and SIF design documents (SIL Verifications and SRS documents etc.) is used to program the SIS PLC. At worse these two companies and teams will never interface to address issues and questions. After a safety system engineering firm has delivered its documents from the SIL verification and SRS design phases, often times the documents are “hands off” from that point forward and the original team will have no idea if what they specified has been followed. Likewise, sometimes clients can look to save money on the SIS PLC hardware by pulling spare PLCs or even spare DCS modules "off the shelf" with no consideration if these devices are actually SIL rated and SIL capable or not.

Recovery factors for SIS PLC errors is primarily through a through Factory Acceptance Test (FAT) and Site Acceptance Test (SAT). The most important check for logic errors is during the FAT. Even so, recovery is not guaranteed, especially if the SIS PLC application programming is done by a separate firm from the safety system engineering firm and the engineering firm is not consulted to check the FAT plan. It is possible the person that wrote the FAT did not understand the safety system engineer’s specifications. Often the person writing the FAT plan is the application programmer himself. If the application programmer didn't understand the logic in the first place, how confident could an owner operator be that the FAT Plan will detect all errors? Note, the authors have witnessed this occurrence many times when they have been able to back check a FAT plan. Quite often the FAT Plan does not line up with what was specified, up to the system not even doing what was intended logically. Review of the FAT documents propagated the misunderstanding and, if left to run their course would have shown the system passed FAT even though there were fundamental flaws in the application program compared to the requirements.

To help with application programmer errors, the safety system engineering firm which developed the SRS and SWRS documents should be involved in the subsequent FAT Plan review at a minimum, and can detect many errors in programmed logic. For hardware error recovery factors, a hardware audit/ audit of the design and engineering phase of the project should be able to find specification errors prior to the cost of correcting the errors becoming a non-starter. You would not want to find out your engineering team programmed and FAT'ed/ SAT’ed a non SIL capable DCS logic solver after the system is installed and operating in the plant.

Overall, a lot of errors in the SIS PLC hardware and software phase can be identified and recovered. The earlier the detection of a hardware specification error the better. A competent FSA stage 2 should be sufficient. In addition, software errors are avoidable if the safety system engineering firm which developed the SRS and SWRS is involved with this phase, at least during the FAT.

**Equipment Engineering Errors**

Some examples of equipment (e.g. instrumentation, mechanical, etc.) engineering errors include:

- Ignoring the safety system engineer's requirements in the SRS documents. For example, if the SRS called out a special high integrity SIL certified device per the SIL calculation model but the instrumentation engineer opts to use off the shelf instruments that are on hand, this could lead to a SIF not having as high of an integrity as needed. Some back checks by other disciplines can be assumed to catch these errors; however the level of client sophistication will determine how effective this would be.
- Poor corporate standards, which could call out incompatible materials for equipment and process. In the author’s own experience, there was a case where a plant kept having process leaks from flanges with potential exposure near misses. The process material was Ethylene Oxide (EO), a highly toxic and flammable material. The root cause was eventually discovered to be an incompatible gasket material called out for EO service. The plant has since commenced the expense of replacing every single gasket in EO service with the correct type.
- Mistakes could be made in the instrument specification with numerous opportunities. This is easy to occur when dealing with multiple similar equipment but with differences which may or may not be realized, and could be major sources of potential failure down the line.

Recovery Factors from equipment engineering errors include FSAs. This could find the errors in specification, where devices ordered do not match the requirements. Corporate standard errors can be discovered with periodic reviews. Maintenance records could be audited to determine if there are issued with the corporate standards calling out incompatible materials. A back check of all instrument engineering deliverables by a senior competent engineer should be able to discover any errors in the interpreting corporate standards, SRS, or duplication errors; however the recovery factor could be low if the back checker is also not familiar with the source documents. This is especially true of an outside EPC new to the plant's standards or unfamiliar with safety system documentation.

**Purchasing Errors**

Purchasing phase errors can occur after the devices are specified and have been sent out for purchase. When devices return from quotation from multiple sources it's possible that the devices quoted by vendors are not technically acceptable. There are many owner operators where price is the overriding factor. It happens often that the technical aspects of a quote are completely overlooked and a bid is won solely on price alone. The problem comes when the device is not technically acceptable.

Recovery factors include having the original equipment engineer involved in the purchasing phase with a technical bid review. This should be performed prior to the purchasing department disqualifying the bid on a price basis alone. This could serve to correct the mistakes of carrying forward an incompatible device quoted from the vendor.

Overall, recovery from purchasing errors depends on a corporation's purchasing culture. If the company is fully price driven it's highly possible that non-technically acceptable devices have made their way into service and could be a problem waiting to happen.

Note that errors in the purchasing phase are not limited to equipment alone but to the entire spectrum of a project. This includes engineering firms chosen to do any phase of the work - analysis, design, specification, commissioning, auditing, operating, and maintenance, etc.

**Commissioning Errors**

Examples of commissioning errors could include a rush job where workers are under pressure to deliver a job on time and under budget. If a schedule is too aggressive it could lead to workers damaging equipment during the install, valves or other equipment installed backward etc. There could be a sloppy job done wiring up the system, possibly due to reuse of messy junction boxes or messy marshalling cabinets.

Some recovery factors occur during SAT, where most devices and loops will be tested from “end to end” (i.e. drive a signal from the transmitter and witness the valve move). SAT can work out a good portion of errors (such as wiring problems). Routine inspections and proof testing can also uncover errors.

Overall a good portion of errors can be recovered with testing, however some errors still might not be detected even with testing. For example, most SATs don't test a device (such as a valve) up to the actual working pressure. If there was an error with the valve install (e.g. an underrated spring is used on a spring return actuator, or the valve was installed backwards), the error might not be discovered until a real demand occurs at which point the valve will not act appropriately on the actual demand pressure. In testing the valve could still have opened or closed correctly leading to a false send of security.

If equipment is damaged on install, some of the damages may result in a delayed failure of the device even after testing has occurred and shows the device is working. There is an example of a remote sealed differential pressure based level transmitter's diaphragm seal being ruptured on the installation. The device was still able to function but eventually the ruptured cell could have leaked out the capillary fill fluid and introduced error. This problem was only detected due to a random inspection - the proof test never called out unbolting the capillary flange and inspecting the diaphragm.

**Operations Errors**

Example of errors from operations are many. One of the most troublesome would be uncontrolled bypassing of any or all SIF components, for whatever reason, for any length of time. SIF trips are sometimes bypassed indefinitely due to operations "not liking" a SIF for reasons such as poorly designed SIF with a high amount of spurious trips being viewed as a nuisance. The authors have been in more than one Hazop where they have heard of interlock set points raised/ lowered (by board operators) so that the trip never comes in for reasons such as avoiding spurious trips.

Recovery factors include communicating expectations to Operations about what/ why safety functions are in place and periodic auditing of bypassing (don't let it be done on the "honor system"). Procedural PHA's related to start-up and other non-steady state operation represent an opportunity to discuss troublesome SIFs and unaccounted for bypassing of said SIF, as well as Operational expectations related to the SIF.

Overall, the recovery factors will depend on the sophistication of the plant safety lifecycle management, as well as policies and procedures which are backed up by routine audits.

**Maintenance Errors**

An example of maintenance error could be anytime that hardware (or software) of a safety function is touched, there is a chance to leave it in a failed state (or cause some other failure). This could be from an instrument technician to the DCS engineer who also maintains the SIS. One of the authors, based on their experience performing Safety Lifecycle audits, was shocked to learn the actual rate that DCS engineers access SIS safety code to make minor "improvements" (each an opportunity for an error to be made). Also, often times set points and timers are changed during testing (including proof tests). There is always the chance that the set points and timers are not returned to their "As Found" state.

Recovery factors include back checks and sign offs any time one group touches the safety system. For example, the field technician radios back to the console operator after placing an instrument back in service to verify it is reading correctly. In addition, periodic audits can be performed to look for forgotten bypasses or forced logic that may not be visible from an operator graphic.

Overall, the frequency of the opportunity for error depends on the testing interval (the more frequent will actually *increase* the opportunity for error). It is important to note there is a trade-off between more frequent testing to achieve PFD, and the increased opportunity to make an error (again, either on the safety function in question or an "adjacent" safety function). Plant sophistication for back checks from competent staff will provide an increased probability of recovery.

**The "How" – Human reliability examples factored into hardware safety integrity.**

In this section, some conceptual example calculations will be provided which will examine a few of the previously listed sources of error due to the human component in SIF design, engineering, and operation; which will be factored into a traditional SIL calculation. The results will give an idea of the "real" (i.e., less uncertain/ same order of magnitude) safety integrity of a Safety Instrumented Function (SIF).

Examining IEC/ISA 61511 for further guidance, the following abridged definitions in relation to this concept are of note:

- System - a set of elements which interact according to design - and may include hardware, software, and human interaction.
- Safety Integrity - Average probability of a safety instrumented system satisfactorily performing the required safety instrumented functions under all the stated conditions within a stated period of time.

    - Note, in determining safety integrity, *all causes of failures* (both random hardware failures and systematic failures) which lead to an unsafe state should be included.
- Systematic Failure - Failure related to a certain cause which can only be eliminated by a modification of design, manufacturing process, operation procedures, documentation, or other relevant factors. [Note all of these are related to human error in some way].
- Hardware safety integrity - Safety Integrity of the Safety Instrument Function (SIF) relating to random hardware failures in a dangerous mode of failure [i.e. SIL Calculations]
- Systematic safety integrity - cannot usually be quantified (as distinct from hardware safety integrity). [Note an HRA could be performed as one method to quantify the human component].

An important take away from the above definitions, Safety Integrity should factor in both hardware and systematic failures. SIL calculations cover the hardware failures, however quantifying the systematic failures and then adequately factoring them into the achieved safety integrity is a difficult prospect and not well defined. Human errors are generally thought to be "systematic" in nature, whilst SIL calculations generally only look at the hardware safety integrity. The analysis provided in this paper is not typically performed due to the systematic failures component of a system are not typically quantifiable without thorough analysis (and as this paper has pointed out, this might not even "cover the bases"). Systematic failures are typically assessed qualitatively, however systematic failures are generally only understood in hindsight (i.e., after the fact) and cannot be adequately factored into the performance of a safety system as a whole before the fact.

The following calculations will demonstrate an approach which may serve as means of estimating a baseline for the human error safety integrity contribution (prior data), and then gathering data to perform a Bayesian analysis on the safety integrity as a whole.

**The Calculation Setup**

The three calculations will look at a simple SIL 3 SIF with various amounts of systematic error via human interaction factored in. In addition, a base case is modeled to show the SIF results given no addition of human factors into the SIL calculations (a pure hardware safety integrity analysis). The approach used in the analysis of the inclusion of systematic failures into the hardware model is a typical fault tree for dormant failures with the addition of a various extra sources of human failure represented as probabilities with recovery factors and dependencies factored in as per a typical Human Reliability Analysis.

In order to simplify the analysis, typical human error probability has been assumed across the board. A basic human error probability of failure of 0.003 for general human error of commission or omission has been used per appendix G of the HRA Handbook (Swain and Guttmann, 1983).

Recovery factors are provided to correct a human error; basic numbers are again utilized from the HRA handbook. Factors are Zero Dependence, Low, Medium, High, and Complete Dependence. Values used for recovery factor are failure probabilities of .003, .05, .15, .50, and 1 respectively.

Note that only a single human error contributor is modeled for each leg of the base calculation. In reality there would be many more potential human interaction points to model for

an exhaustive and thorough analysis. This is beyond the scope of this paper as the intent is to demonstrate the concept of systematic human contribution to a SIF. Note only modeling a single interaction as has been done is optimistic and overlooks many points where humans can introduce error.

## Base Calculation

The base calculation is a simple SIL 3 SIF using generic components.

Dual temperature transmitters with Thermocouples voted 1oo2 with a 10% beta factor and a 1 year test interval. A high SIL 3 rated logic solver with a 1 year test interval. Triple redundant valves with actuator and solenoids, with a 5% Beta factor and a 1 year test interval voted 1oo3.

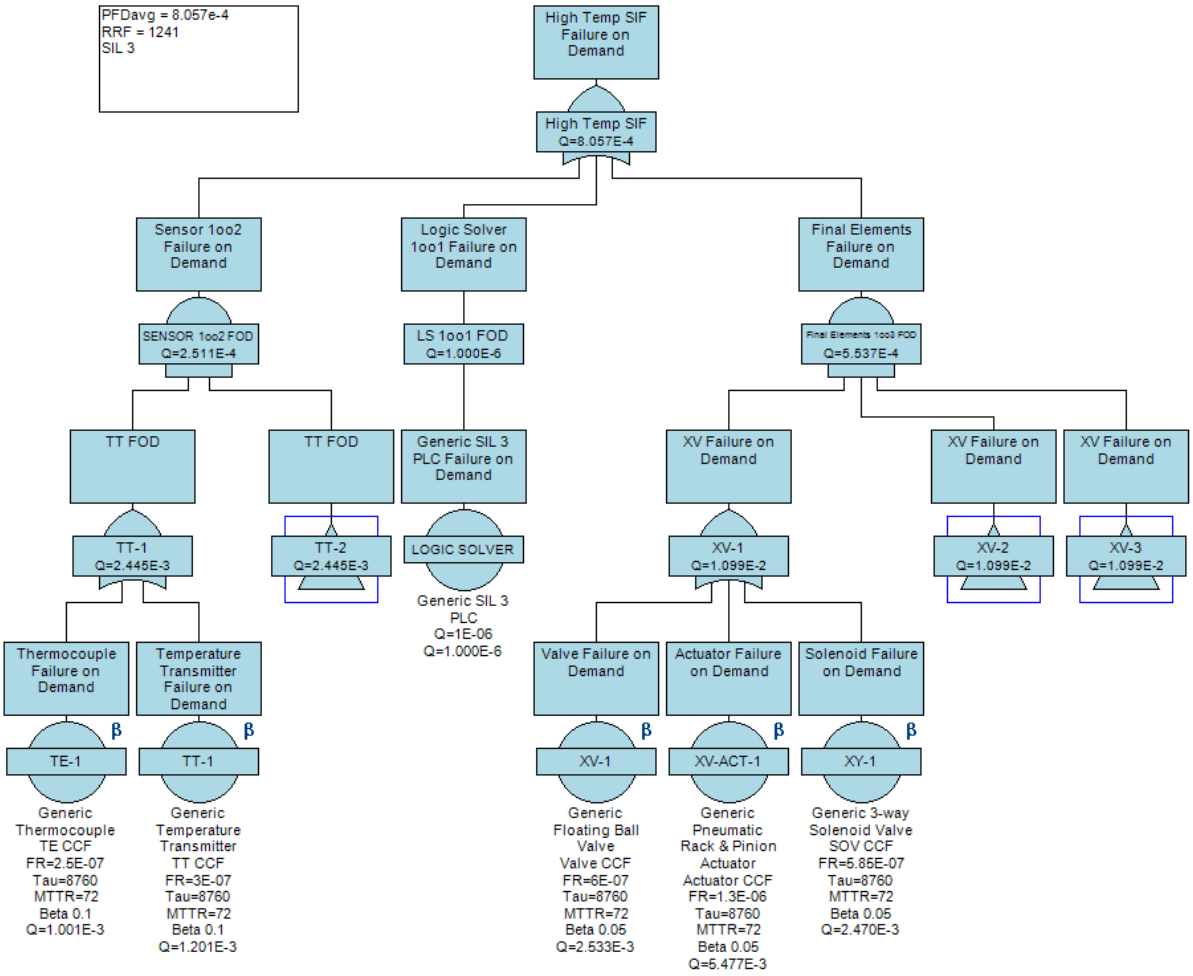


**Figure 1**: Base SIL Calculation

This SIF easily achieves SIL 3 at an RRF of 1200 (SIL 3 is at least an RRF of >1000).

## Model 1 - Human Error Introduced to Base Calculation with Poor Recovery

The next calculation examines the same base calculation which previously achieved SIL 3 but now factors in a few different human error components. This model assumes a basic human error at some point during the engineering, operations, or maintenance phases, which if left uncorrected will defeat the SIF once a real demand is experienced. The model assumes also that the recovery factors available to detect the initial human error are either poor or non-existent. This could be thought to represent a plant which is not very sophisticated in their implementation of the IEC/ISA 61511 standard, perhaps a relatively new adopter of the standard that has not fully developed their lifecycle processes and therefore will have a hard time detecting errors. Note skilled personnel are still assumed to have made the original error (0.003 human error probability).

As discussed above, not all sources of human error have been modeled as this is beyond the scope of the paper. Mainly the intent is to show the impact human error could have on achieving a SIL 3 target when human interaction is factored into a hardware safety integrity calculation to estimate an overall safety integrity. In addition, not all recovery factors have been exhaustively modeled. Mainly the most obvious recovery factors have been chosen.

The points of error with recovery factors are as follows:

- Sensors – Specification error in either the engineering, design, or purchasing of the sensor equipment; which *will* lead to a covert SIF failure of the device.
    - Recovery Factor "RF1A" (see figure 2) is an independent back check of the original engineers work. Here it is assumed no back checks are performed due to the lack of sophistication – therefore there is complete dependence.
    - Recovery Factor "RF1B" (see figure 2) is detecting the error via testing or other audit activity. It is assumed that at least a Pre Startup Safety Review is performed, though detecting this error would be more a chance occurrence – therefore there is high dependence/ low recovery.
- Logic Solver – Maintenance error in Logic Solver, inadvertent logic modification; which *will* lead to a covert SIF failure of the device.
    - Recovery Factor "RF2A" (see figure 2) is detecting the error via testing such as FAT or SAT or other audit activity. It is assumed that this step is not performed due to the lack of sophistication – therefore there is complete dependence/ no recovery.
    - Recovery Factor "RF2B" (see figure 2) is an independent back check of the original engineers work. Here it is assumed no back checks are performed due to the lack of sophistication – therefore there is complete dependence.
- Final element – Operations error in bypass or failure to properly return to service; which *will* lead to a covert SIF failure of the device.
    - Recovery Factor "RF3A" (see figure 2) is an independent back check of the original engineers work. Here it is assumed no back checks are performed due to the lack of sophistication – therefore there is complete dependence.
    - Recovery Factor "RF3b" (see figure 2) is detecting the error via testing for return to service or other audit activities. It is assumed that this step (auditing) is not performed due to the lack of sophistication – therefore there is complete dependence/ no recovery.

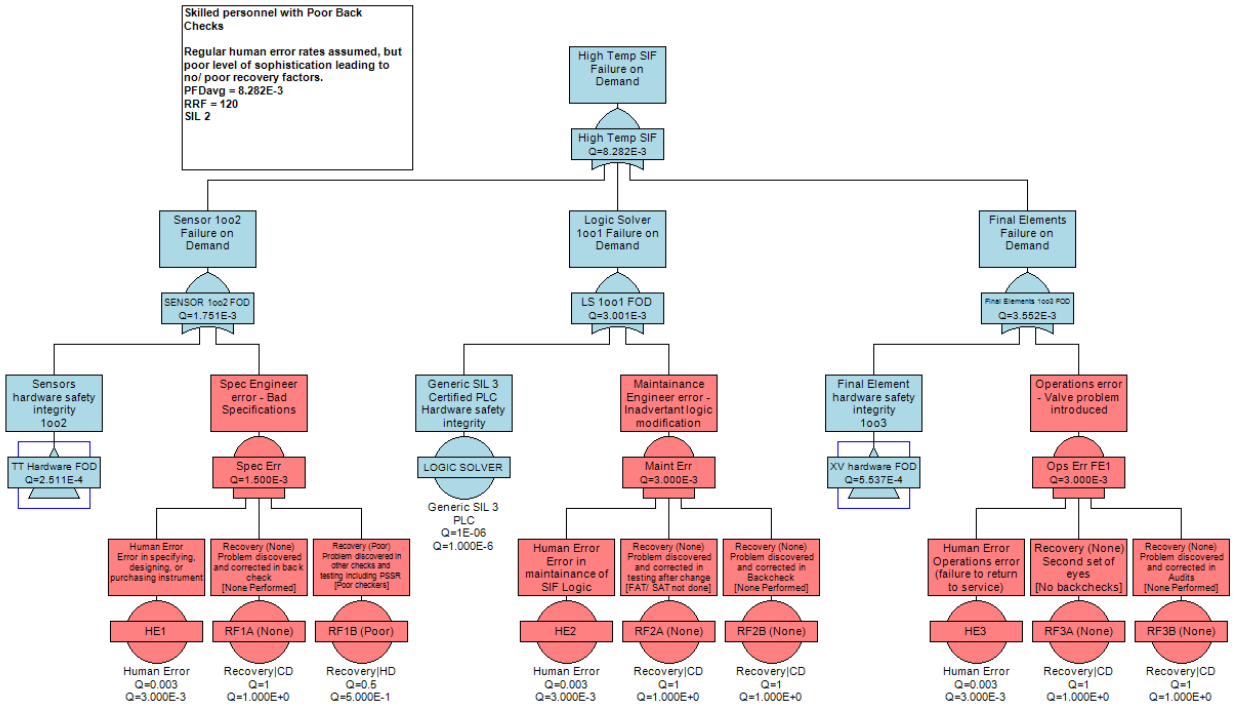


**Figure 2**: Human Error Introduced to Base Calculation with Poor Recovery

As seen in the diagram, the exact same SIL 3 SIF as previously modeled now fails to meet SIL 3. SIL 2 is barely even achieved (RRF of 120). The human component is a major factor in keeping the SIF from achieving SIL 3, and if even more human interactions are factored in, it's possible the best the plant can manage is SIL 1.

## Model 2 - Human Error Introduced to Base Calculation with Decent Recovery

The next calculation examines the same calculation previously modeled, but now with better recovery factors. The recovery factors available to detect the initial human error are better than model 1 but still not the best they could be. This could be thought to represent a plant which is further along in implementing the IEC/ISA 61511 standard – perhaps with a decent amount of experience managing the safety lifecycle with better defined processes and review cycles, but still has not fully developed their processes and procedures with thorough back checking requirements.

The points of error with recovery factors are as follows:

- Sensors – Specification error in either the engineering, design, or purchasing of the sensor equipment; which *will* lead to a covert SIF failure of the device.
    - Recovery Factor "RF1A" (see figure 3) is an independent back check of the original engineers work. Here it is assumed a back check is completed, but it's not thorough and relies on the back checker's experience. The checker might also be on the same team – therefore high dependence is assumed.

    - Recovery Factor “RF1B” (see figure 3) is detecting the error via testing or other audit activity. It is assumed that at least a Pre Startup Safety Review is performed, though detecting this error would be more a chance occurrence. This example the plant is assumed to also perform Functional Safety Assessments but the standard is only loosely followed, maybe no checklist is used for the assessment – therefore there is medium dependence/ recovery.
- Logic Solver – Maintenance error in Logic Solver, inadvertent logic modification; which *will* lead to a covert SIF failure of the device.
    - Recovery Factor “RF2A” (see figure 3) is detecting the error via testing such as FAT or SAT or other audit activity. It is assumed that this step is performed, but perhaps the procedures have not been written thoroughly and the programmer who made the errors also wrote the FAT plan and some logic errors might not be detected – therefore there is medium dependence/ recovery.
    - Recovery Factor “RF2b” (see figure 3) is an independent back check of the original engineers work. Here it is assumed a back check occurs of the engineer’s work, however the back checker is from the same team as the engineer doing the work and “looking over the shoulder” – therefore there is medium dependence.
- Final element – Operations error in bypass or failure to properly return to service; which *will* lead to a covert SIF failure of the device.
    - Recovery Factor “RF3A” (see figure 3) is an independent back check of the original engineers work. Here it is assumed back checks are performed but the back check procedure is not very thorough and does not require a manager sign off – therefore there is high dependence.
    - Recovery Factor “RF3B” (see figure 3) is detecting the error via testing for return to service or other audit activities. It is assumed that audits and testing are performed, but the procedures may not be very through – therefore there is medium dependence/ recovery.

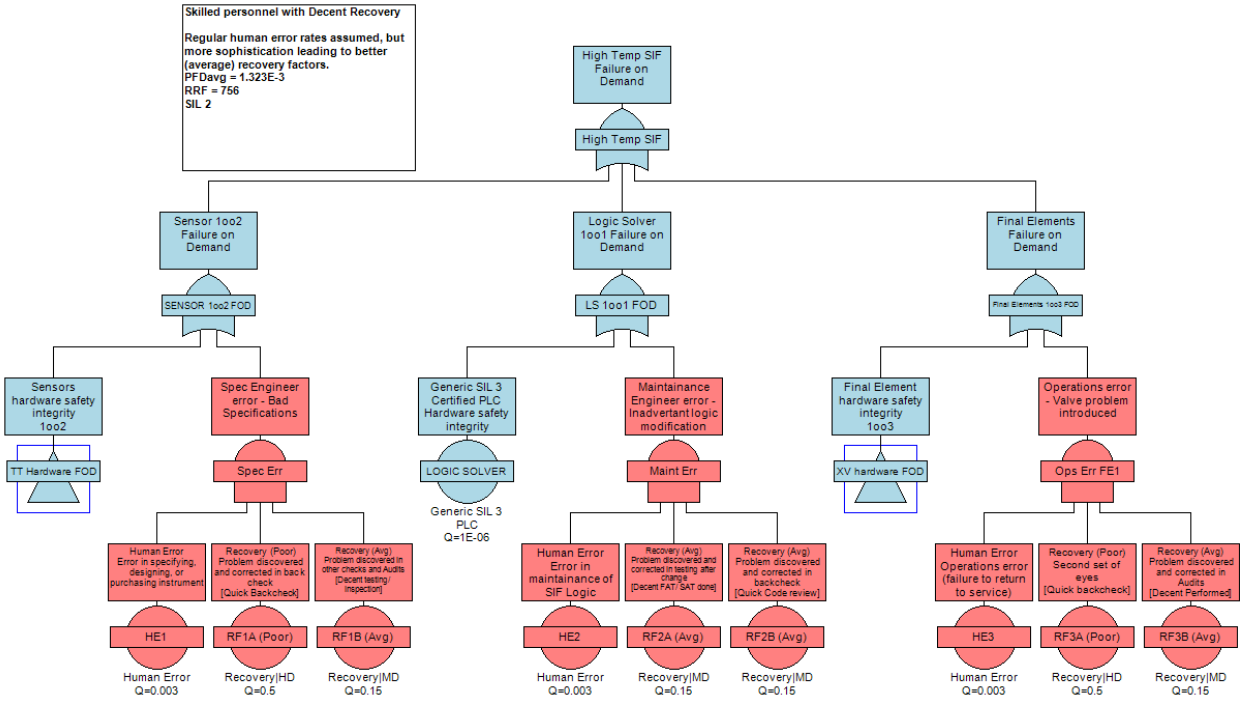


**Figure 3**: Human Error Introduced to Base Calculation with Decent Recovery

As seen in the diagram, the exact same SIL 3 SIF as previously modeled again fails to meet SIL 3. SIL 2 is more readily achieved though (RRF of 756). The human component is still a factor in keeping the SIF from achieving SIL 3, yet it is less of a factor than in the previous example.

### Model 3 - Human Error Introduced to Base Calculation with Good Recovery

The final calculation examines the same calculation previously modeled, but now with even better recovery factors. The recovery factors available to detect the initial human error are considered good, and far better than model 1 or 2. This could be thought to represent a plant which is highly sophisticated in implementing the IEC 61511 standard – with ample experience managing the safety lifecycle with well-defined processes and review cycles.

The points of error with recovery factors are as follows

- Sensors – Specification error in either the engineering, design, or purchasing of the sensor equipment; which *will* lead to a covert SIF failure of the device.
    - Recovery Factor “RF1A” (see figure 4) is an independent back check of the original engineers work. Here it is assumed a back check is completed and the checker is a master practitioner in their field. Furthermore the checker is from an independent team than the engineer who made the error. However some errors require high amounts of scrutiny to detect – therefore medium dependence is assumed.
    - Recovery Factor “RF1B” (see figure 4) is detecting the error via testing or other audit activity. It is assumed that Pre Startup Safety Review is performed, though detecting this error would be more a chance occurrence. This example the plant is

assumed to also perform Functional Safety Assessments and the standard is closely followed with well-defined checklists that have been developed over time to catch many errors that have been discovered in the past. Some errors could theoretically still slip through though – therefore there is low dependence/ high recovery.

- Logic Solver – Maintenance error in Logic Solver, inadvertent logic modification; which *will* lead to a covert SIF failure of the device.
    - Recovery Factor "RF2A" (see figure 4) is detecting the error via testing such as FAT or SAT or other audit activity. It is assumed that this step is performed, and although the testing may be more through, it is still difficult to detect some changes in logic if they were not anticipated – therefore there is low dependence/ high recovery.
    - Recovery Factor "RF2B" (see figure 4) is an independent back check of the original engineers work. Here it is assumed a back check occurs on the engineer who made the error, and the checker is a separate engineer not associated with the unit under control. However, some logic changes can still be difficult to conceptualize – Therefore there is low dependence.
- Final element – Operations error in bypass or failure to properly return to service; which *will* lead to a covert SIF failure of the device.
    - Recovery Factor "RF3A" (see figure 4) is an independent back check of the original engineers work. Here it is assumed back checks are performed and the back check procedures are very thorough and require a manager sign off, however the other operator is likely not an independent person and would be from the same unit and might be subject to the same line of thinking as the original operator – therefore there is low dependence.
    - Recovery Factor "RF3B" (see figure 4) is detecting the error via testing for return to service or other audit activities. It is assumed that audits and testing are performed, and while the procedures may be more thorough some errors could still slip through if testing is not done up to demand pressures and temperatures which is unlikely – therefore there is low dependence/ high recovery.

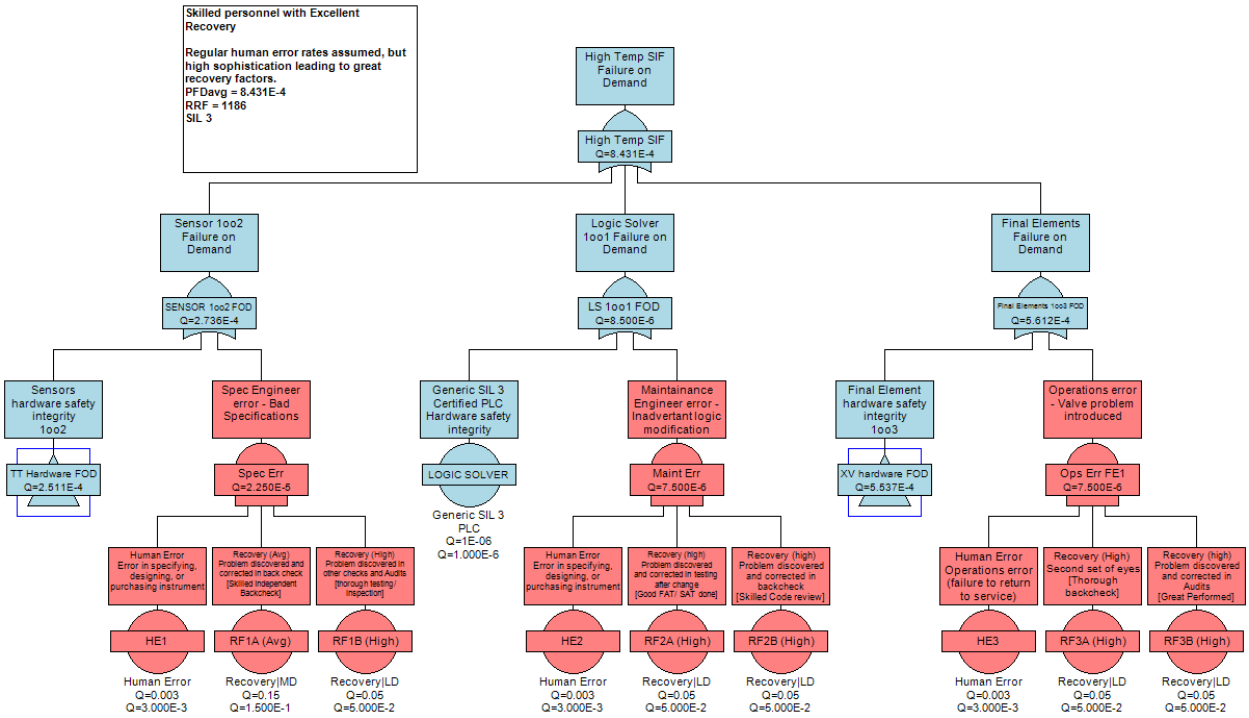


**Figure 4**: Human Error Introduced to Base Calculation with Good Recovery

As seen in the diagram, the SIL 3 SIF previously modeled now once again meets SIL 3 (RRF of 1100 achieved). The human component has been minimalized due to more experience and sophisticated procedures and back checks, and a greater adoption of the ISA/IEC 61511 safety lifecycle standards. The human component still factors in, but is a marginal contributor to the overall probability of failure.

## Conclusion

As was seen in the examples, human error can be a significant detractor to the overall safety integrity achieved by a SIF. This is especially true for SIL 3 targets where the potential contribution from human error begins to dominate the calculated random hardware integrity. In addition, if an owner/ operator is not very sophisticated in their management of a nominal SIL 3 SIF, there is a high likelihood that the best the SIF will actually operate at is SIL 2 or worse. Conversely, a plant that is sophisticated in their management of SIFs likely does not need to factor in the human component in a SIF PFD calculation, because it would be a marginal contributor, however they will have demonstrated and documented this in other ways (audits, FSA, etc.).

As for modeling human error components, this paper proposed the idea of a Bayesian approach. Using a Bayesian approach can correct and update the results initially obtained from a frequentist approach. In addition, Bayesian confidence intervals are real-world limits that reflect actual conditions in an operating plant. In order to start out though, a baseline of human error contribution to overall safety integrity must be established (i.e. estimated). The approach of this paper for modeling the baseline human error contribution to an overall safety integrity was via a theoretical HRA. Note the approach in this paper was not a thorough analysis; the intent was

simply to provide a rough idea of how an analysis *could* be performed, and some factors to be aware of. As it should be obvious, conducting a full HRA for every step of the SIF design could be *very* time consuming, however most models can be performed once and then applied to all SIFs as a factor. A drawback of a thorough HRA is this could lead to *over confidence*, because as noted from the ISA/IEC 61511 standard, human error contribution can be poorly understood. That said, there does need to be a starting point to assess an operating plant's systematic safety integrity, and the HRA is the best tool available to discover human errors and recovery factors. It is best to implement a Bayesian approach after the baseline is developed. The beauty of a Bayesian approach applied after the baseline estimate is completed, are that errors in the initial assumptions and HRA should reveal themselves.

Some other ideas for factoring in the human error component could be through confidence intervals as proposed in the IEC 61508 standard for various applications. Some SIL calculation tools add other ways to factor in human interaction and potential for errors. One such tool proposes a concept labeled "Maintenance Capability," though a basis for the number used would need to be established. A caution with this approach is if the estimate is *only* for maintenance, it would not factor in all the other points of human interaction outside of maintenance and will lead to over confidence. As this paper points out, there are many more sources of human error than maintenance alone.

Some ideas to reduce human error contribution are firstly to fully implement the ISA/IEC 61511 safety lifecycle which calls out for assessments and audits throughout the entire lifecycle. Outside of the standard, more thorough back checks should be performed. Checklists could be utilized. As a plant gains experience, the previous failure points can be added into the checklist. This should be considered an evergreen process or the risk of not improving and never being able to achieve SIL 3 will be a given. Some various ideas on back checks and check lists could include:

- FAT test plan written by the person who understands how the system should work (i.e., not by the logic programmer).
- SAT must test each device individually to detect errors in wiring
- A full physical inspection should be undertaken prior to process introduction to ensure valves facing right way, equipment not damaged where it can't be seen (broken thermowell/ DP cell punctured)
- Back check of design basis process data
- Back check of instrument specifications
- Periodic review of plant specifications to make sure no incompatible materials are specified as experience is gained.
- Consult maintenance for problems that may be occurring, but are not being documented.
- Audit Operation's interaction with a SIF (e.g., bypassing practices etc.)
- Audit for unauthorized, hidden or forgotten bypasses (e.g., forces, or jumpers left on).
- Audit the Control Engineer's interaction with a SIS (e.g., how often changes are required, which increases the potential for human error).

**Bibliography**


1. Bennett, Deborah J., 1998, *Randomness*, Harvard University Press.
2. HSE (Health and Safety Executive), 2009. *Review of Human Reliability Assessment Methods*. Research Report RR679. United Kingdom.
3. Kahneman, Daniel., 2011, Thinking *Fast and Slow*, Farrar, Straus & Giroux, New York.
4. Leveson, N., 2012, *Engineering a Safer World: Systems Thinking Applied to Safety*, MIT Press.
5. McLeod, R.W., 2015. *Designing for Human Reliability – Human Factors Engineering in the Oil, Gas, and Process Industries*, Gulf Professional Publishing, Massachusetts.
6. Mlodinow, Leonard, 2008, *The Drunkard's Walk – How Randomness Rules our Lives*, Pantheon Books, New York.
7. OESI (Ocean Energy Safety Institute), 2016. *Human Factors and Ergonomics in Offshore Drilling and Production: The Implications for Drilling Safety*. College Station, TX.
8. Pearl, Judea, 2018, *The Book of Why – The New Science of Cause and Effect*, Basic Books, New York.
9. Swain and Guttmann. NUREG/ CR – 1278. *Handbook of Human Reliability Analysis*. 1983.
10. Taleb, N.N., 2004. *Fooled by Randomness – The Hidden Role of Chance in Life and in the Markets*. Random House, New York.
11. Taleb, N.N., 2010. *The Black Swan*, 2nd Ed., Random House, New York.
12. Weinberg, G., 2001, *An Introduction to General Systems Thinking*, Dorset House, Silver Edition.