Demonstrating ALARP/SFAIRP

DISCLAIMER: The contents of this document should not be construed as legal advice. It should be used as a general guide only.

What cost is proportionate for further risk reduction?

This paper describes a way of demonstrating whether risk has been reduced to as low as reasonably practicable (or so far as is reasonably practicable), as required by law in some jurisdictions.

This method is based on:

- estimating the current level of risk associated with a potentially hazardous situation, and
- showing that the additional cost for further risk reduction would be disproportionately high compared to the expected benefit.

The method depends on deciding the range of cost that a 'reasonable person' would consider to be proportionate when compared to the societal cost of a fatality.

We might naturally shy away from putting a value on human life, but fatalities have a real cost that is borne by society. The range of proportionate costs varies between different societies.

We can usually agree on some range of values $*Vmin* to $*Vmax* as being proportionate. We do not need precise values. $*Vmax* should be chosen so that a reasonable person would agree that costs 2 or 3 times greater than $*Vmax* would be grossly disproportionate.

The current level of estimated risk may be expressed in different ways, depending on the hazardous scenario being considered.

For example, the risk might be expressed as an estimated rate of fatalities '*f* '. It could also be expressed as a rate of serious injury. Fatality rate *f* could be expressed as a rate per annum, or per participation hour, or per exposure hour. It could be the total expected number of fatalities over the projected life of a facility. It could also be expressed as probability of fatality per event.

The cost of a risk control can then be described as proportionate if its estimated cost is in the range *$Vmin x f* to *$Vmax x f*. The cost is obviously evaluated over the same basis as the risk.

A cost 2 or 3 times greater than $*Vmax* x *f* would be grossly disproportionate.

The risk is reduced progressively by selecting risk controls one at a time. Risk controls that are expected to be more dependable and more effective are usually preferred.

The more effective controls are prioritised within an industry-accepted effectiveness model, or hierarchy of controls. The lifecycle costs of each potential risk control are evaluated to decide whether the costs are proportionate.

f is a variable; it represents the current level of residual risk. The residual risk is as low as reasonably practicable only when there are no more risk controls available at a cost that is not grossly disproportionate.

A simple test can be applied at each stage:

Risk controls need to be considered if their expected cost per time or per event is less than *$Vmax* x *f, regardless* of how low *f* might already be.

CC BY SA

Risk matrix scaling

The selection of colours on a corporate risk matrix is usually implicitly related to proportionate cost of further risk treatment. The scaling and colouring of risk matrices varies, depending on the context.

Many organisations are reluctant to discuss how they determine risk matrix scaling. They will need to demonstrate the basis to their regulators if either SFAIRP or calculus of negligence applies.

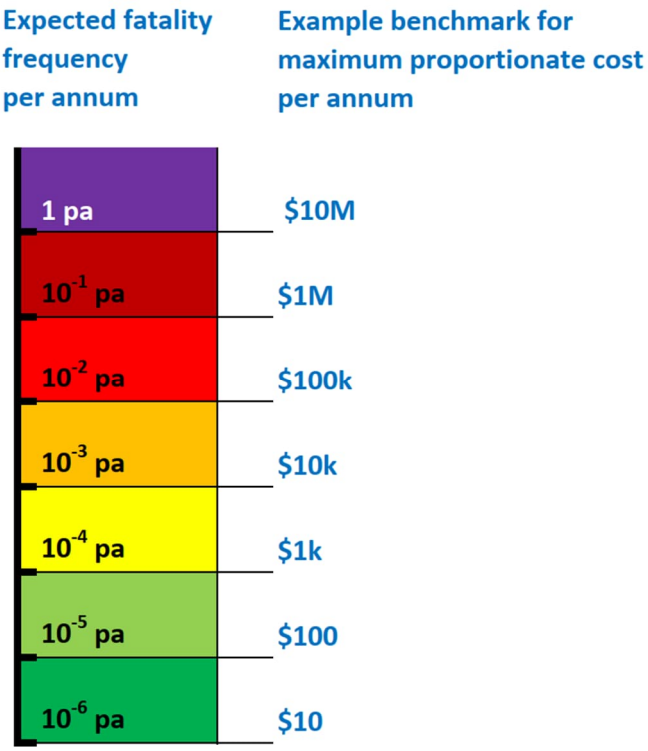

This diagram shows how the range of proportionate annualised cost increases with the expected risk of fatality per annum. The same principle applies for risks expressed per event or per hour of exposure, or on any other basis.

'Acceptable' risk may require further risk reduction

Boxes that are coloured green might be interpreted as representing a level of risk that is broadly acceptable. However, even at this level, further risk reduction is still required according to the SFAIRP principle (and calculus of negligence) if the cost is not grossly disproportionate.

For example, the risk of an unauthorised person climbing an electrical power pylon might be classed as low enough to be 'acceptable'. Fewer than 1 in a million people are killed this way each year in Australia, too few to be counted in a meaningful way.

Danger signs are usually required by the regulations, but barbed wire might not be explicitly required. The annualised cost of signage and barbed wire for each pylon is likely to be in the order of only $10 to $100 per pylon.

It might be difficult to argue after a fatality, that the cost of barbed wire was disproportionately high, even though it might not have been explicitly required.

CC BY SA

High-integrity controls

Risk control methods or 'protection layers' are generally designed to achieve 1 or 2 orders of magnitude in risk reduction [10]. The risk reduction factor (RRF) is modelled as a factor of 10 or 100.

Sequential selection and application of risk controls reduce the residual risk progressively by 1 or 2 orders of magnitudes at each step. Proportionate costs for further risk reduction can then be evaluated.

A risk control reduces the likelihood from *f* to *f* / RRF. The maximum benefit in risk is *f* (1-1/RRF). The maximum proportionate value in adding any risk control is $*Vmax* x *f* x (1 – 1/RRF).

For risk controls with RRF ≥ 10, the maximum proportionate value can be approximated more simply as ≈ $*Vmax* x *f*. The maximum proportionate value 0.99 x $*Vmax* x *f* of a risk control with RRF of 100 is only 10% higher than 0.9 x $*Vmax* x *f* for RRF of 10. The difference is negligible.

Low-integrity controls

The method can be extended to low-integrity controls. Probability of failure of risk controls can be estimated to within half an order of magnitude at best (i.e. a factor of 3). [13, 14, 15, 16, 17]

A low integrity control has a probability of failure > 0.1, typically in the range 0.3 to 0.5. The RRF is in the range of 2 to 3. The maximum benefit of a risk control with RRF in the range 2 to 3 is ≈ 0.5 to 0.7 x $*Vmax* x *f*.

A proportionate value for the cost of low integrity controls can be taken to be 0.5 x $*Vmax* x *f*. Low integrity controls include process control functions that are not subject to strict control.

Restricted occupancy would be a low-integrity control if it were unmanaged and unmonitored. For instance, a barbed wire fence is not a dependable control on its own. Higher integrity access control might be achieved with locked gates, restricted profile keys, security clearances, and entry permit systems.

'Enabling conditions and conditional modifiers' [11] can be either high or low-integrity controls.

Precise estimates of probability and RRF are not plausible. The integrity of risk controls depends on many systematic factors. Integrity and reliability can vary widely, just as risk varies (see Variability in risk, below).

Reverse application to evaluate removing a risk control

The $*Vmax* x *f* evaluation method may be applied in reverse to assess whether the cost saving is defensible.

A risk is currently assessed as *f* fatalities per unit time or per event. An existing risk control is estimated to reduce the risk by a factor of RRF. The overall average cost of the risk control is $*C* per unit time or per event.

The likelihood of fatality can be expected to increase by RRF x *f* if the risk control is removed.

The likelihood increases from *f* to RRF x *f*. The maximum benefit of the control is $*Vmax* x *f* x (RRF -1).

The cost of the existing risk control is proportionate if $*C* < $*Vmax* x *f* x (RRF -1). This can be approximated to $*C* < $*Vmax* x *f* x RRF because of the lack of precision in RRF.

The cost saving in removing a control might be defensible in law if the cost of the control is clearly grossly disproportionate: $*C* >> $*Vmax* x *f* x RRF.

CC BY SA

Variability in risk

Hazardous events are usually caused by failure of equipment, action by humans or animals, by environmental factors, or by a combination of these.

The first step is to identify potentially hazardous situations and to assess their likelihood.

Likelihood can be thought of as the expected probability of hazardous events in a future period, i.e. the expected rate of future events. Estimates of likelihood can be informed by the frequencies of similar events measured in the past. Some people are drawn into precise modelling based on historical data, but likelihood cannot be predicted and estimated accurately. Assessing likelihood of hazards may seem to be difficult because likelihood depends on many contributing factors, all of which can vary.

However, assessing likelihood is not really a question of '*what are the expected event frequencies?*'. The question is '*what future event frequencies are feasible in this situation, and how can we achieve them?*'

A likelihood that is assumed in demonstrating SFAIRP effectively sets targets for future performance in operation and maintenance. The targets selected by the evaluation team need to be feasible for the specific context.

Many books and papers have published to describe accident events, and to discuss the causes and contributing factors that led to the events. Some are listed below in the list of references.

It should be clear from those many references that there are two main factors that need to be controlled at the managerial level:

- Safety culture
- Quality culture.

Effective application of each has been demonstrated to reduce likelihood of hazards by at least a factor of 10.

A lack of due diligence in safety and quality has been demonstrated consistently to increase risk by at least a factor of 100. [6]

This is really where SFAIRP needs to start, if only because the cost of effective safety culture and quality culture is never disproportionately high. On the contrary, both can and should result in lower overall lifecycle cost.

Unfortunately, many people are naturally drawn into trying to save money by cutting corners and taking short cuts. '*We haven't got time to do it properly, but we'll find time to repeat it 3 or 4 times over, until we finally get it right*'. Performance is optimised by providing simple and effective safety (e.g. ISO 31000 or ISO 12100) and quality frameworks (e.g. !SO 9000 and ISO 55000) to ensure:

- Clear expectations, objectives and requirements are set for activities
- Requirements are fully understood and defined before activities start
- The level of effectiveness and independence in review, checking, inspecting and testing the outputs produced is appropriate, given the complexity, novelty and risk involved
- Equipment is safe to use and fit for service (suitable for the intended application and environment)
- Effective environmental protection is provided and maintained to avoid excessive degradation
- Equipment condition is monitored and maintained at a level of effectiveness that is appropriate, given the costs resulting from breakdowns
- Careless, reckless or negligent behaviour is understood to be unacceptable.

Examples

We might decide that the societal cost of a single fatality in a workplace is somewhere in the range between $1M and $10M. The range does not need to be defined precisely. It is not a fixed value that is right or wrong. It should be acceptable to any reasonable person in our society.

Example 1: Road safety

The rate of fatalities caused by motor vehicle accidents in Australia is typically around 5×10^{-5} pa.

The rate in some south-east Asian countries is typically 3 or more times higher, at around 1.5 to 2×10^{-4} pa. The rate in Hong Kong is 3 times lower, at around 1.5×10^{-5} pa.

Australian statistics show that the fatality rate varies over more than an order of magnitude within Australia. It depends on age and gender of a person. It depends on the types of vehicles involved. Motorcyclists face higher risk. Bus passengers face lower risk.

Fatality rates vary widely across the country. People in remote regional areas are at least 10 times more likely to die in traffic accidents than people in major cities.

Male drivers between 17 and 25 years of age and drivers older than 75 experience a fatality rate of around 8×10^{-5} pa, almost twice the average rate. They are more likely to kill other road users and pedestrians as well as themselves.

Female drivers aged in their 30s and passengers younger than 16 experience significantly lower fatality rates, typically below 1×10^{-5} pa.

We could infer that fatality rates are related to drivers' attitude and experience, or capability. Younger drivers are at increased risk due to lack of experience. Older drivers may be at increased risk due to delayed sensory input processing and reduced multi-sensory integration capabilities. Some drivers may develop a sense of invulnerability if they have not previously had any accidents or infringements.

Compulsory driver re-training can be imposed on drivers that are convicted of serious traffic infringements. Traffic Offender Intervention Programs (TOIP) typically cost about $200. Intensive driver safety training courses can cost up to $500. These costs are proportionate because drivers that are guilty of serious infringement notices are likely to have fatality rates higher than 10^{-4} pa. They also put other road users at risk.

Australian law requires the driver of a motor vehicle to ensure that seat belts are correctly worn by the driver and all passengers travelling in the vehicle. Approved restraints are required to be used for children that are too small to use seatbelts. The cost of ensuring that restraints are correctly installed and used is negligible compared to the risk, even if the risk is below 10^{-5} pa.

Example 2: Recreational aircraft maintenance costs and pilot review costs

Aviation legislation in some countries requires owners of recreational aircraft to carry out regular preventive maintenance as a condition of aircraft registration. Aircraft maintenance typically costs about $10 to $30 per flight hour.

Pilots are required to undergo biennial flight reviews with a flight instructor. The averaged cost of review flights would typically be less than $2 per flight hour.

Is it reasonable to expect owners to bear this cost?

Aviation Occurrence Statistics published by the Australian Safety Bureau fatality reported a fatality rate of 36.5 fatalities per million hours (4×10^{-5} per participation hour) flown in recreational aviation in the

10-year period 2010 to 2019. [18] The statistics suggest that engine failure or malfunction is one of the most common causes of accidents in recreational aviation.

The *Flight Safety Australia* magazine regularly publishes analysis of aviation-related accidents and incidents. The magazine summarises international reports as well as Australian reports.

Fatal accidents inevitably result from a combination of many contributing factors. Three factors seem to be prominent in the safety reports:

- Poor judgement and decision making by pilots and ground crew
- Ineffective maintenance practices
- Flights continuing into adverse weather conditions.

We can infer from the statistics that the fatality rate would be at least 3 times higher if pilots were not regularly trained in how to respond to deteriorating conditions or to equipment failure. The fatality rate is also likely to be at least 3 times higher if preventive maintenance is not carried out effectively as required.

The expected fatality rate would be at least 1×10^{-4} fatalities per participation hour (i.e. per hour spent flying in an aircraft). The fatalities may involve passengers and bystanders as well as the pilot.

All fatalities and serious injuries result in significant loss and cost. The cost is borne by society, not just by the pilots and their families.

In this example, the cost of biennial flight reviews and the cost of mandatory maintenance are clearly proportionate to the cost of the fatalities that might be prevented.

Example 3: Gas regulator leakage

A hypothetical power generation utility has carried out a periodic risk review on an 80 MW gas turbine powered generator. The turbine and generator are fully enclosed in an acoustic enclosure on a skid-mounted package. It has been in service for more than 20 years.

Fuel gas is supplied to the skid package at 300 kPa. A gas pressure regulator is mounted on the skid package in a sealed chamber.

The review team have identified that there is a risk of gas accumulating in the chamber. Maintenance technicians have reported smelling gas leaks in the chamber during a recent inspection.

The review team estimate that the frequency of a significant gas leak could be at least 10^{-2} pa (i.e. 1 in 100 chance of a substantial leak within a one-year period, 0.01 pa). It might be as high as 10^{-1} pa (0.1 pa), depending on the condition of the seals in the regulator. They cannot agree on a precise number.

The fuel is natural gas. The team assume that the probability of ignition of a substantial leak would be about 0.1. They agree that it would probably cause a contained fire rather than an uncontained explosion with catastrophic consequences. They review statistics reported by the manufacturer and studies published in journals. They decide that the chance of catastrophic explosion following ignition would be less than about 0.1.

Their estimate of the expected frequency of catastrophic explosion is about 10^{-4} pa or 10^{-3} pa. This is estimated as 10^{-2} pa or 10^{-1} pa x 0.1 (probability of ignition) x 0.1 (probability of catastrophic failure).

They decide that catastrophic explosion could be expected to have at least a 10% chance of causing 1 or 2 fatalities, because the turbine package is within 5 m of the maintenance team's crib-room.

The risk of fatality is therefore estimated as being in the range 10^{-5} pa to 10^{-4} pa.

The company's guidelines suggest that annualised costs in the range of $100 to $1k pa would be proportionate if the risk really is as high as 10^{-4} pa.

They consider installing a fixed gas detector into the chamber to trigger an alarm in the plant-wide fire and gas system. The annualised cost is expected to be in the order of $3k pa. That cost includes design, procurement, installation and maintenance. The cost seems high, but perhaps not grossly disproportionate.

They then consider modifying the enclosure to improve natural ventilation and to duct any fugitive gas emissions to vent in a safe area. The annualised cost is expected to be less than $1k. They agree that the risk of fatality would then be in the order of 10^{-5} pa. They decide to improve ventilation instead of installing a fixed gas detector.

They also decide to reduce the planned overhaul interval for the gas pressure regulator from 8 years to 6 years. It is expected to reduce the likelihood of a gas leak to below 10^{-2} pa. The annualised cost of the overhaul would increase from about $400 to about $600. The additional cost is expected to be less than $200 pa. That would not be grossly disproportionate even if the risk were lower than 10^{-5} pa. It seems reasonable irrespective of the actual level of risk, given that the level of risk is uncertain.

Finally, they also decide to add leak testing into the annual inspection. A test using soapy water should add less than 10 minutes to the inspection. The additional cost should be less than $20 pa. It might not reduce the risk much, but the cost is so low that it is reasonable to do this anyway.

Option		Cost	Benefit	Decision
1	Fixed gas detector installation	≈ $3k pa	Improved detection and alarm response; potentially reduces risk of fatality to below 10^{-5} pa	Considered high cost and ultimately not selected Not required if any of the other options are implemented.
2	Improve ventilation, vent to safe area	< $1k pa	Reduces fatality risk to below 10^{-5} pa	Selected as cost is proportionate and seems reasonable
3	Increase regulator maintenance	< $200 pa (incremental increase from ~$400 pa to ~$600 pa)	Reduces gas leak likelihood to $<10^{-2}$ pa; Reduces fatality risk to below 10^{-5} pa	Selected as considered reasonable despite uncertainty, the cost is proportionate
4	Add leak testing to annual inspection	< $20 pa	Risk reduction by a factor of perhaps 3, uncertain	Negligible cost, should be implemented anyway regardless of risk

In this example, the proportional benefit of leak tests would be about 0.5 x $10M x 10^{-5} pa ≈ $50 pa.

Clearly, a cost of $20 pa is proportionate and reasonable.

Example 4: Reactor start-up

The $*Vmax x f* evaluation method produces the same conclusions, regardless of the chosen rate basis. The method may also be applied in reverse to assess whether existing risk controls is proportionate.

The start-up process of a reactor can be considered as an example.

A chemical reactor has been in operation for 27 years. The reactor runs for about 10 weeks between shutdowns for maintenance and catalyst renewal. The reactor is usually restarted 4 times each year.

Hazardous events are more likely to occur during the startup sequence than during normal steady-state operation. Significant releases of flammable material were recorded on 2 separate occasions, 25 years ago and 22 years ago. The operating company prepared detailed procedures in response to those events. The procedures include independent inspection of the equipment before it is restarted.

The independent inspection costs $3k per startup. There have been no similar incidents in more than 85 startups since the incident 22 years ago.

The company applied Williams' *Human Error Assessment and Reduction Technique* (*HEART*) for the development of the detailed procedures. The procedures were designed to reduce the risk of release by at least a factor of 30.

Statistics have been compiled from other users around the world. The statistics suggest that significant loss of containment incidents occur in between 1 in 30 and 1 in 300 startups with this type of reactor. The chance that a loss of containment of this type would escalate to cause fatalities seems to be about 1 in 100.

The statistics suggest that frequency of events and the likelihood of escalation both depend on a combination of safety culture and quality culture.

The company is now under new management. The new management team engaged an external management consultant to carry out benchmarking studies. The consultant suggested that maintenance costs are high in comparison with other operators.

The production manager has been directed to reduce maintenance costs by at least 20%.

A new simplified procedure has been proposed, in which only 30% of the work would be inspected independently. *HEART* analysis suggests that the likelihood of error can be expected to increase by a factor of about 3. The cost saving will be $2k per startup.

Evaluation per startup

The international benchmark for likelihood of a significant event seems to be in the range between 1 in 30 and 1 in 300 per startup. The overall likelihood of fatality is about 0.01 per event. The benchmark risk appears to be in the range 3×10^{-5} and 3×10^{-4} fatalities per startup.

Simplifying the existing procedure is expected to save $2k per startup and increase the risk by at least a factor 3.

The company assumes from the benchmarking that the likelihood of fatality is currently 1×10^{-4} fatalities per startup. They assume it will increase by a factor of 3 to around 3×10^{-4} fatalities per startup if the procedure is simplified. This would be consistent with the upper end of the benchmark range.

The cost of an existing risk control is proportionate if $\$C < \$Vmax \times f \times (RRF - 1)$. A proportionate value for a risk control with RRF = 3 would be around $\$8M \times 10^{-4} \times (3 - 1) = \$1.6k$.

$2k is not grossly disproportionate because it is only about 20% higher than $1.6k.

CC BY SA

It is difficult to justify simplifying the procedure unless a more cost-effective way can be found to achieve a similar level of risk reduction.

Evaluation per annum

The result of the calculation is the same on an annualised basis. There typically 4 startups pa.

The detailed inspection is estimated to avert 8×10^{-4} fatalities pa at a cost of $8k pa.

The implied cost of averting a fatality is about $10M (estimated as $8k / 8×10^{-4}).

Evaluation over projected facility life

Similarly, evaluating costs over the whole projected life leads to the same conclusion. The company expects to continue operation for another 25 years at least and another 100 operational cycles. The detailed inspection is estimated to avert 2×10^{-2} fatalities at an overall lifetime cost of $200k. The cost is not grossly disproportionate.

Example 5: Applying multiple risk controls with sensitivity analysis

A petrochemical company is installing a new fractionation column in an existing plant. The level controller in the reboiler controls the flow of bottom product through a control valve. Failure of the level control function may cause flooding of the column. The inflow to the column is controlled by flow control valve so that the total mass flowing into the column is balanced against the total mass out flow of bottom product, distillate product and overhead vapour. The inflow controller should limit the rate at which liquid could build up in the column.

The risk workshop team agree that flooding of the column could result in overloading of the overhead vapour scrubber unit. That could be expected to result in a significant release of vapour leading to at least a 10% chance of 3 to 4 fatalities.

Causal event

The level controller (LIC) can be expected to fail at a rate of at least 0.1 pa, but the rate should be less than 0.3 pa.

Process control action

The inflow controller (FIC) can be expected to reduce the risk of flooding if the level controller fails, but it has a similar rate of failure. The likelihood of common cause failure affecting both controllers at a similar time is taken to be at least 0.3 because the controllers are not fully independent.

Estimated likelihood without protection layers

The estimated fatality rate with process control but without other protection layers would be in the range 1×10^{-2} pa to 3×10^{-2} pa.

The estimate is based on the causal event (failure of LIC) having a frequency in the range 0.1 pa - 0.3 pa, x 0.3 (probability of common cause failure affecting the FIC) x 0.1 (probability of escalation) x 3 (fatalities), i.e. 0.1 pa x 0.3 x 0.1 x 3 $\approx 10^{-2}$ pa, and 0.3 pa x 0.3 x 0.1 x 3 $\approx 3 \times 10^{-2}$ pa.

The company considers $10M as the maximum value that would be proportionate to prevent or avert a fatality. Risk controls costing less than $300k pa would have a proportionate cost of averting fatality ($10M x 3 x 10^{-2} pa).

Alarm and operator action

A separate and fully independent high-level alarm has already been included in the design. The alarm is categorised as safety critical. The operators will be trained in how to respond effectively to the alarm.

The team assume a likelihood of 0.1 for the operator failing to respond to a high-level alarm, given that it will be managed as a high integrity alarm. The likelihood of failure could be higher than 0.3 if alarm integrity and operator response are not managed effectively.

Safety function action

The proposed design includes a separate high-level trip safety function to shut down the inlet to the fractionation unit.

The risk workshop team needs to evaluate an appropriate safety integrity level (SIL) target for the safety function. To begin with the likelihood of fatality is estimated without taking credit for a safety function.

Applying and evaluating protection layers sequentially

The team follows this hierarchy of effectiveness in considering the risk controls:

1. Provide statutory protections as required by applicable regulations (pressure vessel design, pressure relief valves, vent scrubber or flare system)
2. Minimise the expected frequency of causal events by improving reliability of process equipment
3. Improve the inherent strength or robustness of the process equipment to reduce the likelihood of failure escalation and the expected severity of consequences
4. Provide safety instrumented functions in an independent safety system
5. Provide clear indication and alarms to the operator and train operators in how to respond effectively
6. Improve the effectiveness and reliability of the basic process control system in responding to failures and disturbances so that the process can be kept within a safe operating envelope
7. Provide external engineering controls to reduce the risk of vapour ignition, such as explosion protection for electrical equipment (may be required by acts of law or regulations)
8. Provide mitigative barriers such as blast protection or water deluge systems
9. Implement procedural controls such as restricted occupancy and permits-to-work.

A flare system is not currently available at this plant unit. It is possible to install a new header to the flare system, but the annualised cost is expected to be at least $100k over the life of the unit. The team dismisses this option because the cost seems to be disproportionately high. It is not a statutory requirement in this (hypothetical) case, and additional environmental constraints were not considered at that time.

Process control optimisation

The first option to consider is to improve the reliability and effectiveness of the level control function. The failure frequency of the level controller could be reduced to about 0.1 pa rather than 0.3 pa through applying more rigorous and effective procedures in design, operation and maintenance.

The reliability of the flow controller could also be improved to around 0.1 pa, but the reliability of 2 separate controllers in combination is limited by common cause failures and systematic factors. The failure frequency of the combination will typically be in the order of 0.03 pa.

The failure frequency of the combination could arguably be reduced further to around 0.01 pa if the controllers were sufficiently independent and separate. This would usually require the application of a standard such as IEC 61511 for the design, implementation, operation and maintenance of the controllers. The controllers would typically be designated as ‘non-SIL (or SIL a) safety critical process

control' and segregated from non-safety controllers. It is difficult to achieve that level of risk reduction in a basic process control system, so it is not usually a preferred option.

The standard maintenance plan for the level controller is based on a 4-year interval for periodic inspection and testing. The controller loop could be categorised as a safety critical process control loop. The periodic inspection and testing interval would be shortened to 1 year. The annual inspection would include analysis of setpoint deviation time and magnitude. The loop would be subject to modification control. The operator would not be able to override the setpoint or put the loop into manual mode. The tuning parameters would be locked. The additional cost for the upgrading the control loop would be around $1k to $3k pa.

The cost of upgrading the flow controller is similar, around $1k to $3k pa. The flow controller could be designed to improve independence from the level controller. A failure mode and effects analysis (FMEA) is necessary to identify and reduce the likelihood of common cause failures affecting both controllers.

Upgrading the two control loops and ensuring independence between them should reduce the failure frequency of the combination to below 0.03 pa.

The improvement in process control is effectively a low-integrity risk control, reducing the risk by a factor of 3 (half an order of magnitude).

The residual risk would be in the order of 0.03 pa (LIC and FIC) x 0.1 (probability of escalation) x 3 (fatalities): 0.03 pa x 0.1 x 3 ≈ 1 x 10^{-2} pa.

A proportionate maximum value for a low-integrity control, reducing risk from 3 x 10^{-2} pa to 1 x 10^{-2} fatalities pa, would be around 0.5 x $\$Vmax \times f$ = 0.5 x $10M x 3 x 10^{-2} pa ≈ $150k pa.

A cost of $2k to $6k pa for improved process control would clearly be proportionate and reasonable

The process control engineer suggests that upgrading both control loops should also improve the yield from the column by around 5%. The commercial benefit is expected to be worth around $200k pa, at least 30 times greater than the cost.

Alarm management

The risk with an optimised process control layer would be about 1 x 10^{-2} pa.

Further risk reduction is required if it can be achieved for less than about $30k pa ($10M x 3 x 10^{-3} pa).

It is easy to justify the cost of further risk reduction using a safety critical alarm and a formal alarm management process.

Alarms cannot simply be taken for granted. The probability of the operator failing to respond successfully to an alarm is usually > 0.1.

A risk reduction factor of 10 can only be claimed if:

- the reliability of the alarm is managed and maintained effectively, and
- the operators are all trained and practised in how to respond effectively to the alarm, following a clearly defined response plan.

The annualised cost of a safety critical alarm can be expected to be in the range of about $3k to $5k.

The risk with an optimised process control layer together with a safety-critical alarm would be about 1 x 10^{-3} pa.

The cost of further risk reduction in addition to the alarm is reasonably proportionate if it can be achieved with cost below about $10k pa ($10M x 10^{-3} pa).

Safety function SIL determination

The company policy and SIL determination guidelines set a target of 10^{-5} fatalities pa. The team applies the guideline to set a target risk reduction factor (RRF) of > 100 for a SIL 2 safety function:
10^{-3} pa/ 10^{-5} pa = 100

A single channel safety function achieves RRF ≈ 2 x $MTBF_D$ / T, where $MTBF_D$ is the overall mean time between dangerous failures in the channel, and T is average time taken to detect or reveal dangerous failures.

A typical channel in a safety function can usually achieve $MTBF_D$ > 30 y. If the function is tested at 1-year intervals it can achieve RRF = 60, in the upper half of the SIL 1 range.

SIL 2 performance can be achieved by improving reliability of a single channel safety function.

Borderline SIL 2 performance with RRF of 100 is feasible with $MTBF_D$ > 90 y and inspection and testing at 2-year intervals. RRF of about 200 is feasible with $MTBF_D$ > 90 y and annual testing.

The annualised cost of a single channel SIL 2 function achieving RRF = 200 over a 30-year operating life is expected to be $8k to $13k. It is a reasonable cost compared with the target of $10k pa for proportionate cost.

Residual risk

The team needs to show that the residual risk of 5 x 10^{-6} fatalities pa with a SIL 2 safety function (RRF 200) is ALARP or SFAIRP.

The risk should be reduced further if it can be achieved for less than about $50 pa ($10M x 5 x 10^{-6} pa). It is difficult to find further dependable risk reduction for a cost of less than $50 pa, so this level of risk can be claimed to be ALARP / SFAIRP.

SIL 3 in place of alarm and SIL 2

Some team members suggest that the risk reduction provided by a SIL 3 safety function would be more dependable than the combination of a SIL 2 safety function with a safety critical alarm. The overall risk reduction would be similar.

The overall cost of design, implementation and maintenance for SIL 3 function with RRF of 2,000 is likely to be around $15k to $30k.

The combination of an alarm (RRF 10) and SIL 2 function (RRF 200) would cost about $11k to $18k.

The additional cost for SIL 3 seems to be reasonable given the target of $30k pa ($10M x 3 x 10^{-3} pa).

The one disadvantage of having a SIL 3 safety function alone is in the potential for loss of production. An operator responding to an alarm may be able to restore safe operation without shutting down the process. The SIL 3 function will shut the process down more dependably, but production would be interrupted.

The decision depends on the balance between the cost of lost production and the potential cost of a hazardous event. In this case the team decides to retain the alarm and SIL 2.

Reducing RRF to reduce cost

It could be argued that the policy of setting the target at 10^{-5} fatalities pa is too conservative.

About $2k pa could be saved by extending the interval between full test and inspection from 1 y to 4 y. The RRF would be reduced by a factor of 4, from 200 to about 50 or 60, in the SIL 1 range.

The residual risk would be increased from 5 x 10^{-6} fatalities pa to about 2 x 10^{-5} fatalities pa.

At that level, a cost of $200 pa is proportionate ($10M x 2 x 10^{-5} pa). The additional cost of $2k pa for full testing at 1 y intervals seems to be disproportionately high. Reducing the target RRF to 50 by extending the test interval to 4 y, and classifying the function as SIL 1 might be justified in this case.

Could we justify applying all risk controls together for maximum effectiveness?

It could also be argued that the risk should be reduced beyond the target if the cost is reasonable.

The fatality rate with process control but without other protection layers was estimated to be in the range 1 x 10^{-2} pa to 3 x 10^{-2} pa. At that level, risk controls costing less than $300k pa would have a proportionate cost of averting fatality.

The total cost of applying all 3 controls including a SIL 3 safety function is less than only $40k pa, well below $300k pa, so isn't that a reasonable and proportionate cost?

Risk Controls	Risk Reduction Factor	Expected Range of Cost
Process control optimisation	3	$2k to $6k pa
Safety critical alarm and response by trained operator	10	$3k to $5k pa
SIL 3 safety function	2,000	$15k to $30k pa
Overall total	60,000	$20k to $40k pa

The residual risk with all 3 controls would be less than 5 x 10^{-7} fatalities pa (3 x 10^{-2} pa / 60,000), about 20 times lower than the target of 10^{-5} fatalities pa.

The counter argument is that a SIL 1 function with RRF = 50 in combination with an alarm would be sufficient to meet the target of 10^{-5} fatalities pa.

Risk Controls	Risk Reduction Factor	Expected Range of Cost
Process control optimisation	3	$2k to $6k pa
Safety critical alarm and response by trained operator	10	$3k to $5k pa
SIL 1 safety function	50	$6k to $10k pa
Overall total	1,500	$10k to $20k pa

The cost would be about $10k to $20k pa lower, but the overall total RRF would be only 5,000 rather than 200,000. The residual risk would be about 2 x 10^{-5} fatalities pa, (3 x 10^{-2} pa / 1,500).

The extra $10k to $20k for the SIL 3 function would avert less than 2 x 10^{-5} fatalities pa.

The implied cost to avert one fatality with the additional cost for a SIL 3 function is more than $500M, which is grossly disproportionate.

Proportionality of an existing risk control can also be assessed in reverse using $*Vmax* x *f* x RRF. The additional RRF for the SIL 3 option in this case is about 40 (= 2,000/50). It would achieve *f* < 5 x 10^{-7} fatalities pa.

A proportionate value for the additional cost of SIL 3 would be \$10M x 5 x 10^{-7} x 40 ≈ \$200 pa.

An overall RRF of about 1,000 and a residual risk of 10^{-5} fatalities pa is clearly sufficient in this example. It could be achieved either with the combination of safety critical alarm with a SIL 1 function with RRF 50, or with a SIL 2 function alone, without the alarm.

The conclusion is that risk controls should be applied sequentially until the cost for further risk reduction is grossly disproportionate.

Workplace health and safety legislation

Workplace health and safety legislation in the UK, Australia and Aotearoa / New Zealand requires that the risk of harm to people in (or around) a workplace is reduced to '*as low as reasonably practicable'* (ALARP), or *'so far as is reasonably practicable'* (SFAIRP).

The term '*reasonably practicable*' is not used in the United States. A similar principle '*as low as reasonably achievable'* (ALARA) is applied in the United States for safety relating to ionizing radiation.

United States case law includes consideration of *proportionate cost* in relation to duty of care in taking precautions against loss. The term *proportionate cost* was established in 1947 by the case *United States v. Carroll Towing Co.* The case set a precedent for determining whether a legal duty of care has been breached by considering the *calculus of negligence (Hand formula).* The formula compares the cost of taking precautions against the expected value and likelihood of the loss that might be incurred by other parties.

The UK Health and Safety at Work etc. Act 1974 requires

> '*it shall be the duty of every employer to ensure, so far as is reasonably practicable*, *the health, safety and welfare at work of all his employees*.
>
> [...]
>
> '*In any proceedings for an offence under any of the relevant statutory provisions consisting of a failure to comply with a duty or requirement to do something so far as is practicable or so far as is reasonably practicable, or to use the best practicable means to do something, it shall be for the accused to prove (as the case may be) that it was not practicable or not reasonably practicable to do more than was in fact done to satisfy the duty or requirement, or that there was no better practicable means than was in fact used to satisfy the duty or requirement*.'

The UK Act does not specifically relate the practicability of risk treatments to their expected cost.

The term 'reasonably practicable' has been used in UK case law since the case of *Edwards v. National Coal Board* in 1949. The ruling was that the risk must be significant in relation to the sacrifice (in terms of money, time or trouble) required to avert it: risks must be averted unless there is a gross disproportion between the costs and benefits of doing so.

The HSE UK published guidance on this topic in 2001. The guide '*Reducing risks, protecting people - HSE's decision-making process*' sets out a framework which:

> *'... makes clear that both the level of individual risks and the societal concerns engendered by the activity or process must be taken into account when deciding whether a risk is unacceptable, tolerable or broadly acceptable ;*
>
> [...]
>
> *'– as control measures are introduced, the residual risks may fall so low that additional measures to reduce them further are likely to be grossly disproportionate to the risk reduction achieved, though the control measures should still be monitored in case the risks change over time;'*

In European law the '*Protection of the safety and health of workers*' – is covered by Directive 89/391/EEC – Article 5(1) – Employer's duty to ensure the safety and health of workers in every aspect related to the work – Employer's liability).

The Directive does not consider whether risk reduction measures are reasonably practicable, nor whether costs are disproportionate. The preamble to the Directive includes:

'… the improvement of workers' safety, hygiene and health at work is an objective which should not be subordinated to purely economic considerations'.

The European Court of Justice ruled in 2007 that in applying the SFAIRP principle in UK law, the UK has not failed *'to fulfil obligations to fulfil its obligations under Article 5(1) and (4) of Directive 89/391'.*

Safe Work Australia developed model work health and safety (WHS) laws in 2011. The model laws have subsequently been implemented in most Australian jurisdictions. The law states that:

'A duty imposed on a person to ensure health and safety requires the person:

(a) to eliminate risks to health and safety, so far as is reasonably practicable; and

(b) if it is not reasonably practicable to eliminate risks to health and safety, to minimise those risks so far as is reasonably practicable'

The Australian law imposes a duty of care on any '*person conducting a business or undertaking*'. The law also covers self-employed persons:

'A self-employed person must ensure, so far as is reasonably practicable, his or her own health and safety while at work.'

The term *reasonably practicable* is defined specifically in relation to cost:

'*reasonably practicable*, *in relation to a duty to ensure health and safety, means […] taking into account […] the cost associated with available ways of eliminating or minimising the risk, including whether the cost is grossly disproportionate to the risk.*

The model laws explicitly include the requirement to consider whether:

'*the cost of further risk reduction would be grossly disproportionate to the risk.'*

The Australian law does not consider whether risk might be considered as being 'acceptable' or 'tolerable'. Further risk reduction is always required if the associated cost is not disproportionate.

The cost to society of avoidable fatalities

This method is based on the concept of '*implied cost to avert a fatality*'. That relates to averting avoidable fatalities that are caused by injury or illness to people who are in the middle of their working life. It becomes more difficult to discuss the value of averting death for people approaching the end of their natural term of life.

It is natural to feel uncomfortable about putting a value on human life because it involves a subjective judgement. There is no standard or broadly accepted way of evaluating averted fatalities. Different methods are used in different societies.

The Wikipedia article on 'Value of a statistical life' describes the different ways that this issue has been addressed in different societies. It summarises the criticisms and concerns that have been expressed regarding valuation of life.

The article explains that '*the value also includes the quality of life, the expected lifetime remaining, as well as the earning potential of a given person especially for an after-the-fact payment in a wrongful death claim lawsuit.*'

The Australian statute law explicitly requires that the cost of risk reduction must be considered in managing workplace risk. Consideration of cost is unavoidable. There must be some defensible method for evaluating the cost of a fatality.

The Australian Government publishes a guidance note on the Value of statistical life (VSL) addressing this issue. The note is intended to provide:

> '*…guidance on how* [Government] *officers preparing the cost-benefit analysis in regulatory impact statements or impact analysis should measure and articulate the benefit of reducing the risk of fatality or physical harm*'

The revision published in January 2026 suggested the value of a statistical life was in the order of $6M (2025 dollars).

The Australian VSL is specifically for Government use. For example, the VSL is used in the evaluation of improvements proposed for highways and roads. It is applied in healthcare and epidemiology. It does not apply directly to the general public, or to private enterprises. It does however provide a useful benchmark to establish a range of proportionate costs.

The UK Government publishes 'The Green Book', which is

> '*…intended for anyone involved in developing, reviewing or approving proposals that use public money. It is mandatory for UK government departments and arm's length bodies to follow the Green Book and its accompanying business case guidance when developing proposals that involve significant public spending.*'

It includes guidance related to costs and benefits to society, including consideration of 'Life and heath':

> *'A proposal might involve a risk to human life. This should be estimated using a generic price known as the Value of a Prevented Fatality (VPF). This measures the social value of small changes in fatality risks where levels of human safety vary between options. VPF is not the value of a life. It is the value of a small change in the risk or probability of losing a statistical life. Not to value this in appraisal would effectively treat human safety as having zero value. VPF allows alternative levels of fatality risk to be taken into account in option design. The latest estimate of VPF can be found in the TAG data book, published by the Department for Transport.'*

Guidance from the UK Health and Safety Executive 'Appraisal values or 'unit costs'' suggests that that total societal cost is in the order or £2M per fatality (2024 value). This is equivalent to about $4M in Australian dollars.

On that basis one might argue that a cost of less than $10M per fatality prevented is proportionate for workplace safety both in Australia and in the UK.

Self-employed persons may argue that a lower range of cost is proportionate if they cover their own medical treatment, rehabilitation and insurance costs.

Similarly, people may choose to accept higher risk outside working hours if they are prepared to bear the cost. Some people choose to ride motorcycles on public roads.

Many people are safer at their workplace than they are at home or while travelling to work.

References

1. Kletz, T. '*What Went Wrong?: Case Studies of Process Plant Disasters*', Gulf Publishing. 1998
2. Hopkins, A. '*Lessons From Longford: The Esso Gas Plant Explosion'*, Wolters Kluwer. 2000
3. Hopkins, A. '*Failure to Learn: the BP Texas City refinery disaster*', Wolters Kluwer. 2008
4. Hopkins, A. '*Disastrous Decisions: The Human and Organisational Causes of the Gulf of Mexico Blowout*', Wolters Kluwer. 2012
5. Hopkins, A. '*Organising for Safety: How structure creates culture*', Wolters Kluwer. 2019
6. Miller, K. '*Quantifying Risk and How It All Goes Wrong*', IChemE. 2018
7. IMechE Safety & Reliability Group - Working Group 2 '*ALARP for Engineers: A Technical Safety Guide*', IMechE. 2024
8. Lauridsen, K., Kozine, I., Markert, F., Amendola, A., Christou, M. and Fiori, M. '*Assessment of Uncertainties in Risk Analysis of Chemical Establishments'*, Risø National Laboratory, Roskilde, Denmark, Risø-R-1344(EN), 2002
9. Center for Chemical Process Safety *'Guidelines for Developing Quantitative Safety Risk Criteria'*, Wiley. 2009
10. Center for Chemical Process Safety *'Guidelines for Initiating Events and Independent Protection Layers in Layer of Protection Analysis',* Wiley. 2015
11. Center for Chemical Process Safety '*Guidelines for Enabling Conditions and Conditional Modifiers in Layer of Protection Analysis',* Wiley. 2013
12. Health and Safety Executive (UK) '*Out of control: Why control systems go wrong and how to prevent failure*', UK HSE publication. 2003
13. Smith, D. J. *'Reliability, Maintainability and Risk'*, 10th Ed. Butterworth Heinemann. 2021
14. Nowlan, F.S. and Heap, H. *'Reliability-centred Maintenance'*, Springfield Virginia. National Technical Information Service, US Department of Commerce. 1978
15. Moubray, J. '*Reliability-centred Maintenance*', 2nd Ed. Industrial Press Inc. 1997
16. Okoh, P., Haugen S., *A study of maintenance-related major accident cases in the 21st century*' Process Safety and Environmental Protection Volume 92, Issue 4. July 2014
17. Bukowski, J.V. and Stewart, L. '*Quantifying the Impacts of Human Factors on Functional Safety*' *exida*. 2016
18. Williams, J.C., '*HEART – A proposed method for achieving high reliability in process operation by means of human factors engineering technology*' in Proceedings of a Symposium on the Achievement of Reliability in Operating Plant, Safety and Reliability Society (SaRS). NEC, Birmingham. 1985
19. Australian Transport Safety Board. '*Aviation Data and Statistics Report AR-2020-047*' ATSB. 2020
20. *Flight Safety Australia* magazine, https://www.flightsafetyaustralia.com/
21. Bureau of Air Safety Investigation. *'Human Factors in Fatal Aircraft Accidents, Department of Transport and Regional Development'*.1996
22. Lally, M. '*A cost–benefit analysis of COVID-19 lockdowns in Australia*'. Monash Bioethics Review. 2022
23. Pak, A., Adegboye, O.A., McBryde, E.S. *'Are We Better-Off? The Benefits and Costs of Australian COVID-19 Lockdown*' Frontiers in Public Health. 2021
24. Australian Bureau of Statistics. '*Effects of COVID-19 strains on the Australian economy*', ABS 2022
25. Costantino, V., Grafton, Q., Kompas, T., Chu, L., Honeyman, D., Notaras, A., MacIntyre, C. R. '*The public health and economic burden of long COVID in Australia, 2022–24: a modelling study*', Medical Journal of Australia. 2024
26. 'Long COVID cost Australian economy about $9.6 billion in 2022', University of Melbourne . 2024
27. UK legislation Health and Safety at Work etc. Act 1974

28. Edwards v. National Coal Board. Archived 10 April 2019 at the Wayback Machine (1949) All ER 743 (CA)
29. Health and Safety Executive (UK) '*Reducing Risks, Protecting People*' UK HSE publication. 2001
30. HM Treasury 'The Green Book – UK Government Guidance on Appraisal', HM Treasury. 2026
31. Harmonised Australian legislation Model Work Health and Safety Act, adopted in NSW, Qld, SA, Tas, NT, ACT and WA
32. Victorian legislation Occupational Health and Safety Act 2004
33. Department of the Prime Minister and Cabinet – Office of Impact Analysis 'Value of statistical life', Australian Government. 2026
34. Aotearoa / New Zealand Health and Safety at Work Act 2015

Creative Commons Licence

The document is published by I&E Systems Pty Ltd. It was prepared by Mirek Generowicz and reviewed by Adrian Hertel and Tom McCarthy.

The authors acknowledge comments and contributions from Harvey Dearden.

It was released by I&E Systems Pty Ltd for public use under a Creative Commons BY-SA Licence in March 2026

https://creativecommons.org/licenses/by-sa/4.0/legalcode

Attribution — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

ShareAlike — If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.

You are free to:

Share — copy and redistribute the material in any medium or format.

Adapt — remix, transform, and build upon the material for any purpose, even commercially.