Simplified safety function design methods

Safety functions can be designed and analysed using simple methods without complicated calculations.

Simple methods can be used to design a safety function to achieve any given safety integrity level. The safety integrity level achieved depends primarily on whether the equipment is suitable for the intended application, and on whether it is readily accessible for inspection, testing and maintenance. The safety integrity level can be improved by applying fault tolerant architecture.

Simple methods can also be used to estimate the safety integrity level that has been achieved in the past by any existing function. The estimate is based on the number of failures that have been recorded in operation, and on how often the functions have been inspected and tested.

Worked examples are provided to show how to apply simplified methods. The results are compared with fully detailed calculations based on IEC 61508-6. The comparison shows that complicated calculations are not necessary.

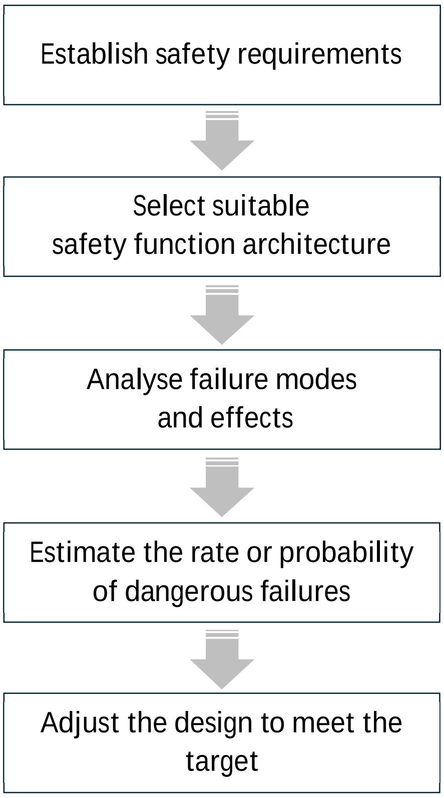

Establish safety requirements

Select suitable safety function architecture

Analyse failure modes and effects

Estimate the rate or probability of dangerous failures

Adjust the design to meet the target

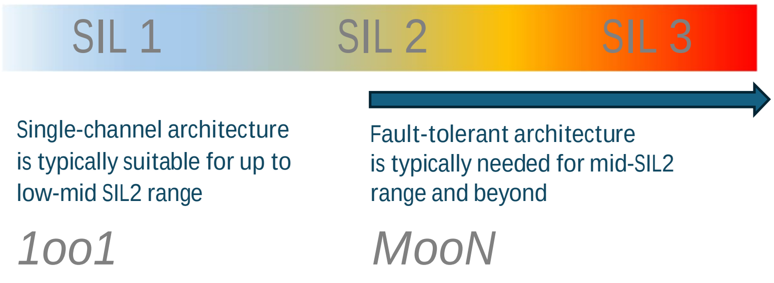

Single-channel architecture is typically suitable for up to low-mid SIL2 range

1oo1

Fault-tolerant architecture is typically needed for mid-SIL2 range and beyond

MooN

Safety function performance can be improved to meet any target by:

1. Reducing failure rates by selecting appropriate equipment
2. Reducing failure rates through planned preventive maintenance
3. Adding continuous diagnostics to detect dangerous failures
4. More frequent inspection and testing to reveal undetected failures
5. Applying fault tolerant architecture
6. Reducing common cause failures

Background

The functional safety standards IEC 61508 and ANSI/ISA 84 describe complicated calculations that may be applied to estimate the probability of failure for safety functions.

ISO/TR 12489 provides a more detailed analysis of several different approaches to estimating failure probability.

The failure probability models all assume that device failure rates remain constant over the useful operating life of the devices. The second edition of IEC 61511 introduced a new requirement that forced a reassessment of that basic assumption:

> '*reliability data uncertainties shall be assessed and taken into account when calculating the failure measure*'.

Reliability data are always uncertain because equipment *failure rates are never constant in the real world*. Equipment failure rates depend not only on the suitability of equipment for the specific application, but also on how effectively the condition of the equipment is maintained.

Uncertainty intervals in reliability data are always at least an order of magnitude wide. That means that failure rates should be expected to *vary by more than a factor of 10* between different applications. Failure rates should also be expected to vary over time in any individual application – unless deliberate efforts are made to prevent or correct degradation.

The level of precision implied in the IEC 61508-6 and ANSI/ISA 84 calculations cannot really be justified because of the wide range of uncertainty in reliability data.

Simple approximations of failure probability are just as accurate as the detailed models.

Low demand mode

1) The performance of low demand mode safety functions depends on the rate of undetected dangerous failures λ_{DU}, the test interval T, and the hardware fault tolerance

2) The rate of undetected dangerous failures can be estimated from past performance, but performance can be improved by reliability centred maintenance or by applying diagnostic functions

3) Any given SIL target can be achieved by improving the failure rate, testing more often, or by applying fault tolerant architecture.

Estimating the SIL achieved by past performance

Failures are classed as dangerous if they prevent successful action of a safety function. The total number of dangerous failures includes dangerous failures of all devices necessary for successful action by one channel of the safety function.

The historical mean time between undetected dangerous failures is abbreviated as $MTBF_{DU}$. It can be estimated by taking the total aggregated time in service τ and dividing by the total number of dangerous failures n_{DU} that were not detected by continuous diagnostic functions.[1]

$$MTBF_{DU} \approx \tau \ / \ n_{DU}$$

Dangerous failures may be revealed during periodic inspection and testing, or they may be revealed when a safety function does not perform correctly on demand.

The $MTBF_{DU}$ is the reciprocal of the average rate of undetected dangerous failures λ_{DU}.

For example, $MTBF_{DU} \approx 30$ y if around 3 dangerous undetected failures were counted over one year in a population of 100 similar devices (100 / 3 ≈ 30). We could expect the total number of failures to vary at least in the range from 2 to 4 per year, so the $MTBF_{DU}$ would be in the range from 25 y to 50 y. The number of failures could vary more widely if environmental factors and maintenance practices were to vary.

The uncertainty in failure rates is always at least +/- 30% because all failure rates vary. Failure rates vary from place to place depending on the environment and application. Failure rates vary over time because of changes in maintenance personnel, practices and resources.

The risk reduction factor (*RRF*) that has been achieved by a low-demand safety function can be estimated from the number and types of failures that have been recorded in operation, and the weighted average test interval T that has been achieved[2].

For a single channel architecture: $RRF \approx 2 . MTBF_{DU}/T$

For any fault tolerant MooN architecture: $RRF \approx 3/2 . MTBF_{DU}/(\beta . T)$

[1] Systematic failures may be excluded if the cause has been identified and eliminated for all similar devices.

[2] The weighted average test interval takes proof test coverage into account. Refer to page 6 for further explanation.

RRF achieved with *T* ≈ 1 year

$MTBF_{DU}$ achieved	Single channel 1oo1	Dual channel 1oo2 with β = 0.1
$MTBF_{DU}$ ≈ 10 y	*RRF* ≈ 20 (SIL 1)	*RRF* ≈ 150 (SIL 2)
$MTBF_{DU}$ ≈ 30 y	*RRF* ≈ 60 (SIL 1)	*RRF* ≈ 500 (SIL 2)
$MTBF_{DU}$ ≈ 100 y	*RRF* ≈ 200 (SIL 2)	*RRF* ≈ 1,500 (SIL 3)
$MTBF_{DU}$ ≈ 300 y	*RRF* ≈ 600 (SIL 2)	*RRF* ≈ 5,000 (SIL 3)

The *RRF* is inversely proportional to the weighted average test interval *T*. For example, increasing the test interval from 1 y to 3 y reduces the *RRF* by a factor of 3:

RRF achieved with *T* ≈ 3 y

$MTBF_{DU}$ achieved	Single channel 1oo1	Dual channel 1oo2 with β = 0.1
$MTBF_{DU}$ ≈ 10 y	*RRF* ≈ 7 (non-SIL)	*RRF* ≈ 50 (SIL 1)
$MTBF_{DU}$ ≈ 30 y	*RRF* ≈ 20 (SIL 1)	*RRF* ≈ 150 (SIL 2)
$MTBF_{DU}$ ≈ 100 y	*RRF* ≈ 70 (SIL 1)	*RRF* ≈ 500 (SIL 2)
$MTBF_{DU}$ ≈ 300 y	*RRF* ≈ 200 (SIL 2)	*RRF* ≈ 1,500 (SIL 3)

Designing for future performance

The target set for the *RRF* of a safety function effectively leads to targets for the average test interval T and for the $MTBF_{DU}$ of each channel in the function.

The overall $MTBF_{DU}$ of a safety function channel can be improved by:

- Selecting devices that are appropriate for the application
(failure modes and failure rates depend on the type of device)
- Adding diagnostic functions to detect failures during normal operation
- Reliability centred maintenance

$MTBF_{DU}$ of each channel and T to meet SIL targets

Target SIL	Single channel 1oo1	Dual channel 1oo2 with $\beta = 0.1$
SIL 1 with *RRF* > 10	$MTBF_{DU} > 5.T$	
SIL 1 with *RRF* > 30	$MTBF_{DU} > 15.T$	
SIL 2 with *RRF* > 100	$MTBF_{DU} > 50.T$	
SIL 2 with *RRF* > 300	$MTBF_{DU} > 150.T$	$MTBF_{DU} > 20.T$
SIL 3 with *RRF* > 1,000	Not Recommended NOTE 1	$MTBF_{DU} > 70.T$
SIL 3 with *RRF* > 3,000	Not Recommended NOTE 1	$MTBF_{DU} > 200.T$

NOTE 1 Single channel architecture is not recommended for SIL 3 because it is difficult to achieve sufficiently high $MTBF_{DU}$.

Example

A single channel safety function relies on a pneumatically actuated valve as a final element. About 90% of the dangerous failures of the safety function would result from the valve sticking or jamming. This type of safety function will typically achieve an overall $MTBF_{DU}$ in the range of about 40 y to 80 y.

A single channel with $MTBF_{DU}$ of about 50 y would need to be tested at least once a year to achieve *RRF* =100. (2 x 50 /1 =100) That would be the bare minimum performance target for SIL 2. There would be no margin for uncertainties in the performance of the equipment or the maintenance team.

Setting a target of *RRF* =300 would provide a reasonable margin for uncertainty in achieving SIL 2. *RRF* =300 would need $MTBF_{DU} > 50.T$ The higher *RRF* target could be achieved by more frequent testing, (reduce T) more effort in maintenance (increase $MTBF_{DU}$) or by applying a fault tolerant architecture.

- *RRF* =300 could be achieved by testing 3 times a year with $MTBF_{DU} = 50$ y, because 150/3 = 50 y.
- *RRF* =300 could be achieved with annual testing by improving the $MTBF_{DU}$ to about 150 y. The failure rate of pneumatically actuated valves can be improved by servicing the equipment at shorter intervals. For example, the valve assembly could be overhauled at intervals of 8 y instead of 15 y. Service intervals could be based on the measured condition of the valves.
- *RRF* =300 could be achieved with a fault tolerant dual channel architecture with $MTBF_{DU} > 20.T$ (assuming a common cause failure fraction β of 0.1) For example, testing every 2 y would be sufficient if $MTBF_{DU} = 40$ y.

Explanation

The MooN SIF model and a simplified MooN model for safety functions explain in detail how the probability of failure can be estimated for single channel and multiple channel architectures.

Probability equations are expressed in terms of *RRF* and $MTBF_{DU}$, the reciprocals of PFD_{AVG} and λ_{DU}.

Single channel and non-redundant architectures

The probability of failure for single channel (1 out of 1 voting) architecture can be estimated as:

$$PFD_{AVG} \approx \quad \lambda_{DU}.T/2 \qquad \text{or} \qquad RRF \approx \quad 2.MTBF_{DU}/\,T$$

λ_{DU} is the average rate at which undetected dangerous failures are expected to occur. These are the failures which are not detected by continuous, automatic on-line diagnostic functions.

All N channels need to work correctly in N out of N voting architecture. The probability of failure can be estimated as:

$$PFD_{AVG} \approx \quad N.\lambda_{DU}.T/2 \qquad \text{or} \qquad RRF \approx \quad 2.MTBF_{DU}/\,(T.N)$$

Performance targets for $MTBF_{DU}$ can be set to achieve any *RRF,* depending on *T*:

$$MTBF_{DU} \approx RRF.T.N/2$$

The model assumes that failures detected by diagnostic functions will be restored to service within the target mean time to restoration (*MTTR*), typically 2 or 3 days. The diagnostic coverage (*DC*) is assumed to be ≤ 95% for these approximations. The *DC* is the fraction of dangerous failures which can be detected by continuous, automatic on-line diagnostic functions and restored within the target *MTTR*.

The probability of failure due to detected failures needs to be added to the estimate if *DC* > 95%, but only if equipment operation is expected to continue as normal after a failure has been detected. The alternative would be to put the equipment into a safe state automatically when dangerous failures are detected.

$$PFD_{AVG} \approx \quad \lambda_{DU}.T/2 + \lambda_{DD}.MTTR$$

Weighted-average test interval *T*

The weighted average test interval *T* depends on the proof test coverage *PTC*:

$$T = PTC.T_1 + (1 - PTC).T_2$$

PTC is the fraction of undetected failures that can be revealed by routine inspection and testing at average time intervals T_1.

The remaining failures might be revealed by inspection and test with complete coverage. Otherwise, the failures might remain undetected until they are revealed when the safety function fails to act successfully on demand in response to a developing hazard.

T_2 is the average time interval needed to reveal the remaining undetected failures. T_2 could be the interval between full inspection and test, or it could be the time between demands on the function (the reciprocal of the safety function demand rate).

Fault tolerant architectures

Fault tolerant MooN architectures have a total of N channels. At least M of the N channels need to work correctly for the safety function to act successfully. The performance is improved by a factor of approximately $3/(4.\beta)$ in comparison with a single channel architecture, where β is the common cause failure fraction. This is the fraction of failures that can be expected to affect at least M of the N channels in a similar way. The fraction depends on the MooN architecture. It increases with M/N because the architecture becomes more susceptible to failure of multiple channels. It decreases with the level of fault tolerance, N-M, because more coincident faults can be tolerated.

The probability of failure for MooN architectures can be estimated using this simple approximation:

$$PFD_{AVG} \approx {}^{2}/_{3}.\beta.\lambda_{DU}.T \qquad \text{or} \qquad RRF \approx {}^{3}/_{2}.MTBF_{DU}/(\beta.T)$$

A target can be set for the *MTBF*$_{DU}$ needed to achieve any given target for *RRF,* depending on *T*:

$$MTBF_{DU} \approx {}^{2}/_{3}.\beta.RRF.T$$

IEC 61508-6 Annex D describes a model for estimating β. Values of β are selected from the range 0.01, 0.02, 0.05 and 0.1 for sensors and final elements with a 1oo2 architecture. The model will typically result in an estimate of $\boldsymbol{\beta} \approx 0.1$ for 1oo2 architecture.

IEC 61508-6 Annex D provides scaling factors to be applied for other MooN combinations.

SINTEF published a study in 2015 titled '*Common Cause Failures in Safety Instrumented Systems; Beta-factors and equipment specific checklists based on operational experience*.' (SINTEF reference A26922). The study reviewed the common cause failure fraction that is typically achieved in operation. The study concluded that β is typically > 0.1 for 1oo2 architecture, usually between 0.12 and 0.2.

Simplified safety function design methods

The following chart can be used as a quick guide to selecting voting architecture and test intervals to achieve any given target for risk reduction. It shows a single channel architecture tested every 6 months delivers a similar risk reduction to a dual-channel fault tolerant architecture tested every 4 years. Note extending test intervals may not be appropriate for devices requiring frequent operation to prevent certain failure modes e.g. valves can jam if there is a long duration without use.

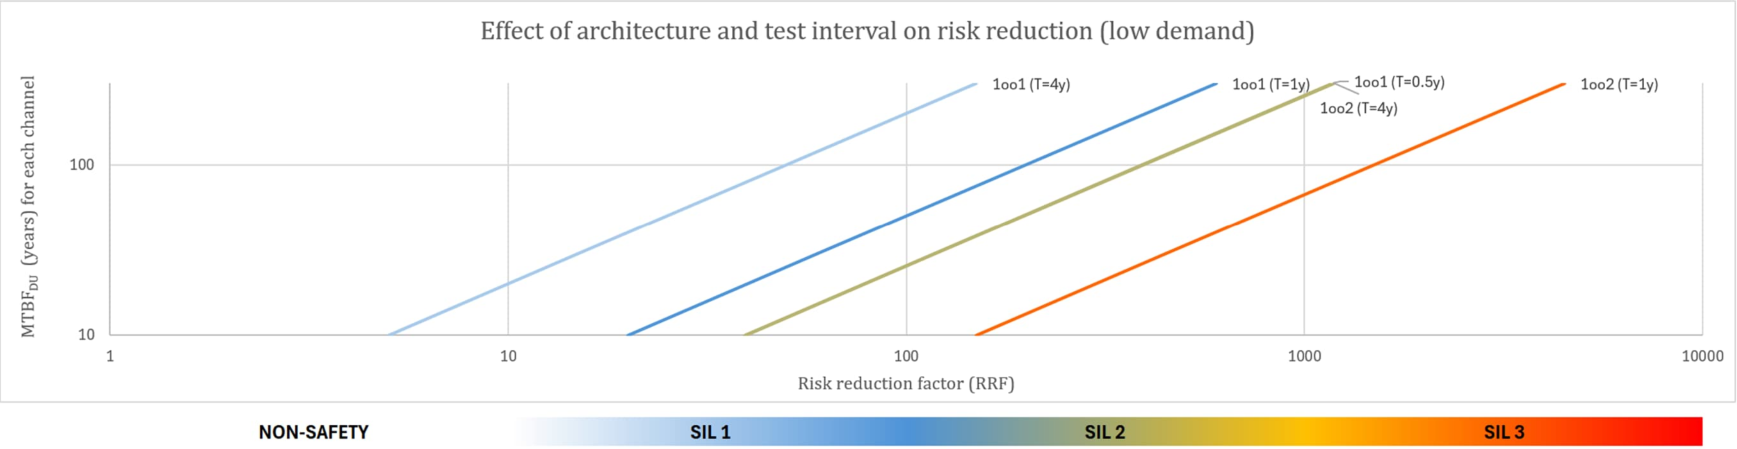

Typical values of *$MTBF_{DU}$* of each channel in process sector applications

Device types	*$MTBF_{DU}$*
Sensor subsystems	300 y to 1,000 y
Logic subsystems	1,000 y to 10,000 y
Actuated valves, spring return	30 y to 100 y
Electrical contactors or relays	100 y to 300 y

The lower values of *$MTBF_{DU}$* will be achieved if equipment is allowed to degrade until it fails to meet its performance targets. The higher values of *$MTBF_{DU}$* can be achieved with regular preventive maintenance and/or continuous diagnostics.

A single channel architecture using actuated valves as final elements can typically achieve *$MTBF_{DU}$* > 50 y with regular maintenance.

High demand or continuous mode

1) The performance of high demand or continuous mode safety functions depends on the rate of undetected dangerous failures λ_{DU} and on the hardware fault tolerance

2) The rate of undetected dangerous failures can be estimated from past performance, but it can be improved by reliability centred maintenance or by applying diagnostic functions

3) Any given SIL target can be achieved by improving the failure rate or by applying fault tolerant architecture.

Estimating the SIL achieved by past performance

The performance achieved by high-demand or continuous mode safety functions can be estimated from the rate of undetected dangerous failures λ_{DU} of a single channel or its reciprocal, $MTBF_{DU}$.

$$MTBF_{DU} \approx \tau \ / \ n_{DU}$$

As for low demand, the $MTBF_{DU}$ is estimated from the total aggregated time in service τ divided by n_{DU}, the total number of undetected dangerous failures. The failure rate is the reciprocal value $\lambda_{DU} \approx n_{DU} \ / \ \tau$.

Failures are classed as detected if they revealed by a diagnostic test and a prompt fault reaction is initiated in response to keep the equipment in a safe state. The test can be automatic or manual but would normally be carried out at least daily to be classed as a diagnostic.

Failures are classed as undetected if they revealed by a periodic test much less frequent than the demand rate, or if they are revealed when the safety function fails in a dangerous way during normal operation.

All dangerous failures in a single channel architecture are treated as undetected if fault reactions are not completed within the process safety time.

The *DC* is the fraction of failures that were detected.

Some *DC* can usually be achieved in high demand and continuous mode safety functions if necessary. The elements used in these functions can usually be fully exercised at least once a day.

Probability calculations are not necessary for high demand and continuous modes SIFs. SIL is characterised in terms of a *dangerous failure rate*, in contrast to low demand mode where SIL is characterised by average probability of failure on demand.

A dual channel fault tolerant architecture improves the performance by the factor $1/\beta$, where β is the common cause failure fraction, the fraction of failures that can be expected to affect both channels in a similar way.

SIL achieved in continuous mode and high-demand mode, based on the overall rate of dangerous failures per hour measured in the past few years

Channel $MTBF_{DU}$	Single channel λ_{DU} (per hour)	Dual channel (1oo2) with $\beta = 0.1$ overall λ_{DU} (per hour)
$MTBF_{DU} \approx 10$ y	> 1×10^{-5} (non SIL)	> 1×10^{-6} (SIL 1)
$MTBF_{DU} > 12$ y	< 1×10^{-5} (SIL 1)	< 1×10^{-6} (SIL 2)
$MTBF_{DU} \approx 30$ y	4×10^{-6} (SIL 1)	4×10^{-7} (SIL 2)
$MTBF_{DU} \approx 100$ y	1×10^{-6} (SIL 1)	1×10^{-7} (SIL 2)
$MTBF_{DU} \approx 300$ y	4×10^{-7} (SIL 2) NOTE 1	4×10^{-8} (SIL 3)
$MTBF_{DU} \approx 1{,}000$ y	1×10^{-7} (SIL 2) NOTE 1	1×10^{-8} (SIL 3)
$MTBF_{DU} \approx 3{,}000$ y	4×10^{-8} (SIL 3) NOTE 1	4×10^{-9} (SIL 4) NOTE 2
$MTBF_{DU} \approx 10{,}000$ y	1×10^{-8} (SIL 3) NOTE 1	1×10^{-9} (SIL 4) NOTE 2

NOTE 1 Single channel architecture is not recommended for SIL 2 or SIL 3 because of the uncertainty in failure rate data.
NOTE 2 Fault tolerance of at least 2 is recommend for SIL 4 because of the uncertainty in failure rate data.

$MTBF_{DU} \approx 10$ y is equivalent to $\lambda_{DU} \approx 1.14 \times 10^{-5}$ per hour. That should be rounded to one significant figure because the uncertainty in failure rates is always at least +/- 30%.

$MTBF_{DU} > 12$ y is sufficient to claim marginal SIL 1 performance, but it leaves no margin for uncertainty.
$MTBF_{DU} \approx 30$ y achieves SIL 1 performance with a reasonable margin.
$MTBF_{DU} > 120$ y would be the minimum needed to claim SIL 2 with a single channel.

Designing for future performance

The simplified design method sets a target for the $MTBF_{DU}$ of each safety function channel, based on the maximum allowable overall rate of dangerous failures per hour needed for the whole function to achieve the required SIL.

As for low demand, the overall $MTBF_{DU}$ of a safety function channel can be improved by:

- Selecting devices that are appropriate for the application
- Adding diagnostic functions
- Reliability centred maintenance.

A target for the overall *DC* in a channel can be set to achieve a target $MTBF_{DU}$, based on the estimated $MTBF_D$ that is expected in the safety function channel.

The *DC* for a new function can be estimated by conducting a FMEA. The FMEA identifies and estimates the expected failure rate for each mode of failure of each element in the safety function subsystem. Diagnostic functions are designed specifically to detect the identified failure modes. The reliability of each diagnostic function needs to be considered. A diagnostic function that is claimed to detect 99% of failures would need to be subject to an appropriate level of quality procedures and techniques.

Example

$MTBF_D \approx 10$ y is equivalent to a failure rate of 0.1 pa, or about 1.1 x10^{-5} per hour. That would be too high for SIL 1.

If > 12% of all dangerous failures can be detected, then the rate of undetected dangerous failures is reduced to < 10^{-5} per hour, good enough for SIL 1.

DC required for a single channel architecture to achieve target SIL in continuous mode or high-demand mode

Target SIL	Target λ_{DU} (per hour)	$MTBF_D \approx 10$ y	$MTBF_D \approx 30$ y	$MTBF_D \approx 100$ y
SIL 1	< 1x10^{-5}	*DC* > 12%	*DC* not needed	*DC* not needed
SIL 2	< 1x10^{-6}	*DC* > 92%	*DC* > 74%	*DC* > 12%
SIL 3	< 1x10^{-7}	*DC* > 99% NOTE 1	*DC* > 97% NOTE 1	*DC* > 92% NOTE 1

NOTE 1 Single channel architecture is not recommended for SIL 3 with channel $MTBF_D$ < 100 y because it is difficult to achieve sufficiently high levels of reliability in *DC*. Fault tolerance is recommendation due to uncertainties in failure data.

The performance targets for each channel in a fault tolerant dual channel safety function can be lower by a factor of about $1/\beta$.

The target for *DC* in a fault tolerant architecture relates to faults that may affect multiple channels in a similar way within the planned periodic test interval.

The main concern is with failures that can affect multiple channels in the same way at the same time.

Failures that affect only one of the channels will not significantly affect the integrity level, provided that the failures are corrected promptly after they are revealed. Single channel failures could be revealed either by continuous diagnostics or by periodic inspection and testing at least once a year.

DC required for a fault tolerant dual channel architecture to achieve target SIL in continuous mode or high-demand mode, assuming $\boldsymbol{\beta} = 0.1$

Target SIL	Target $\boldsymbol{\lambda}_{DU}$ (per hour)	$MTBF_D \approx 10$ y (each channel)	$MTBF_D \approx 30$ y (each channel)	$MTBF_D \approx 100$ y (each channel)
SIL 1	$< 1\text{x}10^{-5}$	*DC* not needed	*DC* not needed	*DC* not needed
SIL 2	$< 1\text{x}10^{-6}$	*DC* > 12%	*DC* not needed	*DC* not needed
SIL 3	$< 1\text{x}10^{-7}$	*DC* > 92%	*DC* > 74%	*DC* > 12%

Refer to ISO 13849-1:2023 §6.1.8 for a similar simplified approach.

Explanation

The MooN SIF model and a simplified MooN model for safety functions explain in detail how the overall dangerous failure rate (*DFR*) is determined for any MooN architecture.

The term ‘dangerous failure rate’ is used here instead of the IEC 61508 parameter ‘probability of failure per hour’, *PFH*. One reason for using *DFR* is that the units of time used for measuring failure rates can vary (we could use failures per year). Another reason is that probability is a dimensionless number between 0 and 1. Probability does not have units of measurement.

Single channel architecture

For single channel (1 out of 1 voting) architecture the dangerous failure rate of the function is the same as the undetected dangerous failure rate of the channel:

$$DFR \approx \quad \lambda_{DU}$$

This is on the basis that the safety function includes fault reactions that will put the equipment into a safe state in response to detected failures. Otherwise, the overall dangerous failure rate of a single channel architecture is simply the rate of *all* dangerous failures:

$$DFR \approx \quad \lambda_D$$

Example

SIL 1 requires the overall rate of dangerous failures to be < 10^{-5} per hour.

$MTBF_D \approx 10$ y corresponds to a failure rate of $\lambda_D \approx 1.1\text{x}10^{-5}$ per hour.

DC > 12% would be sufficient to achieve SIL 1, though with no margin for uncertainty.

DC > 92% would achieve $\lambda_{DU} \approx 1\text{x}10^{-6}$ per hour, sufficient to achieve borderline SIL 2 performance.

Estimates of *DC* should not be expected to be precise. *DC* > 10% and *DC* > 90% would be close enough, given the uncertainty in the estimates.

Fault tolerant architectures

The overall dangerous failure rate of a safety function with MooN architecture is effectively the rate of dangerous failures that affect at least M of the N channels in a similar way. It depends on the common cause failure fraction β and the rate of dangerous failures in a single channel:

$DFR \approx \beta.\lambda_{DU}$ with diagnostics and fault reaction, or else $DFR \approx \beta.\lambda_D$

A fault tolerant dual channel architecture with $MTBF_D \approx 10$ y, $\beta = 0.1$ and $DC = 90\%$ would achieve $\beta.\lambda_{DU} \approx 1\text{x}10^{-7}$ per hour, at the lower border of the SIL 3 range.

The following chart can be used as a quick guide to selecting voting architecture and diagnostic coverage to achieve any given target for risk reduction. It shows adding 90% diagnostics has the same effect as adding a second channel.

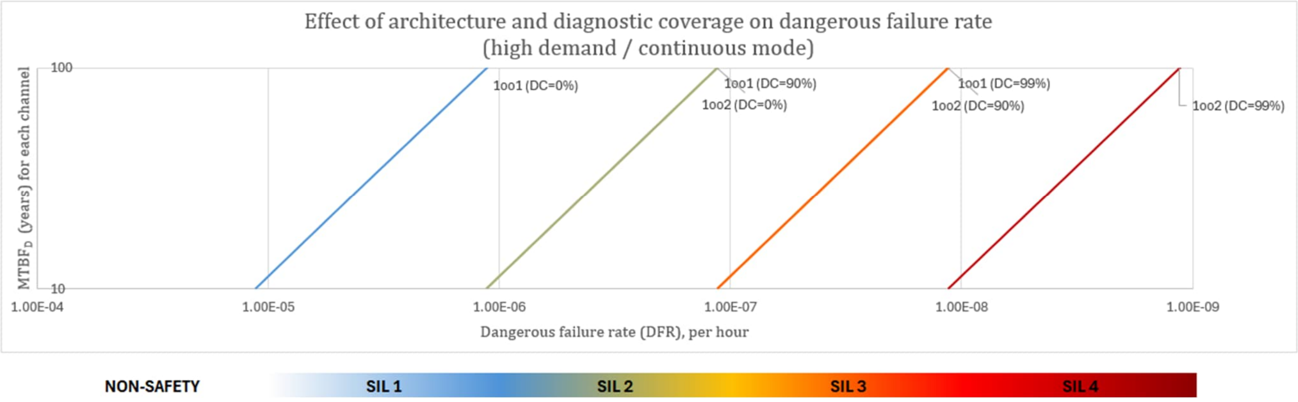

An overall $MTBF_D$ of about 10 y can typically be achieved for a single channel safety function in either process sector or machinery applications. The safety function would need to use well-tried components and well-tried principles, or components that have been demonstrated as being suitable for service in a similar environment and similar application.

An overall $MTBF_D$ of about 100 y could be achieved with regular and effective condition monitoring and reliability centred maintenance techniques.

A volume of operating experience of about 2 or 3 times the target $MTBF_D$ in device-years is enough to demonstrate suitability. There should have been no more than about 3 failures in that period. For example, 4 failures counted in 100 devices over 3 years would suggest that $MTBF_D < 80$ y. The experience should be over at least 2 or 3 calendar years to allow for seasonal variations.

Simple rules for choosing between low-demand, high-demand and continuous modes

Sometimes the mode of operation for a safety function is a matter of debate. Different people will have different opinions. Comparison of the different modes shows that the results are similar in any borderline case, no matter how the function is analysed.

The basic rules are:

- Continuous mode functions always have some sort of continuous action (or inaction, such as maintaining isolation) to keep equipment in a safe state.
- Demand mode functions act only on demand to put equipment into a safe state in response to a separate hazardous event.
- A dangerous failure in a continuous mode safety function causes a hazardous event within a short time (minutes or hours rather than months or years).
- A dangerous failure in a demand mode function has no impact until some other unrelated event causes a hazard. The failure of a demand mode safety function could remain unrevealed for many years before a hazard occurs.
- A safety function operates in a high-demand mode if the demands on the function are more frequent than the routine periodic inspection and testing of the function. Dangerous faults in the function are more likely to be revealed by a failure on demand rather by inspection and testing.
- High-demand mode safety functions are evaluated in the same way as continuous mode safety functions.
- Safety functions that are designed to meet high-demand or continuous mode requirements will always meet low-demand requirements.

The performance of high-demand functions often depends on high frequency diagnostics or tests rather than periodic annual tests. Diagnostic intervals are typically measured in seconds, minutes or hours.

Periodic test intervals are typically measured in months or years. Periodic inspection and testing is still valuable for high-demand functions because it allows progressive degradation to be detected and corrected before failure results.

The performance of low-demand functions can take advantage of periodic inspection and testing *instead of* continuous diagnostics. The inspection and testing would usually need to be at annual intervals if diagnostic functions are not practicable. Inspection and test intervals could be extended to periods of more than 5 y if some level of continuous diagnostic coverage can be achieved.

Example

A risk study has identified a hazardous scenario. The chance of a single fatality is estimated to be 1 in 100. The team cannot agree on whether the scenario might occur once a year or once in 10 years. Some team members suggest that it could occur as often as 3 or 4 times per year.

The maximum tolerable frequency is set at 10^{-4} pa in that area of the facility. The residual risk needs to be well below the maximum tolerable frequency.

A low-demand mode safety function could be considered if the demand rate is less than once per year. A mid-range SIL 2 function with a target of *RRF* =300 would reduce the risk of fatality to about 3×10^{-5} pa if the demand rate is 1 pa.

The residual risk can be estimated as the demand rate multiplied by the chance of fatality and divided by the *RRF*: $1 \text{ pa} \times 10^{-2} / 300 \approx 3 \times 10^{-5}$ pa

RRF of 300 could be achieved using a 1oo2 dual channel architecture with $MTBF_{DU} > 20.T$

With $T = 1$ y we would need $MTBF_{DU} > 20$ y for each channel.

The more pessimistic team members remind us that the demand rate could be as high as 4 pa.

The alternative is to design a high-demand mode safety function to achieve a similar level of residual risk. The advantage of high-demand mode is that the residual risk is completely independent of the demand rate. The dangerous failures of the function are now the only cause of the hazardous event.

The target failure rate for the high-demand mode safety function can be determined by dividing the residual risk target by the probability of failure of the other layers that reduce the chance of fatality.

The target failure rate for a high-demand mode safety function in this example could be set at 3×10^{-5} pa divided by $10^{-2} \approx 3 \times 10^{-3}$ pa $\approx 3 \times 10^{-7}$ per hour.

That target could be achieved using a 1oo2 architecture with $MTBF_{DU} > 30$ y for each channel.

The $MTBF_{DU}$ targets need to be slightly higher for high-demand mode compared with low-demand mode. The difference in cost is relatively small in borderline cases (i.e. when the demand rate is close to the rate of inspection and testing, $1/T$). The higher $MTBF_{DU}$ could be achieved by adding some level of diagnostic coverage or by servicing the equipment at 8 y intervals instead of 10 y intervals.

Annual inspection and testing would be appropriate in either mode.

The low-demand mode design is always more cost effective if the demand rate is < 0.1 pa.

A SIL 1 function with *RRF* = 30 would be sufficient in this example if the demand rate is less than 1 in 10 y: $0.1 \text{ pa} \times 10^{-2} / 30 \approx 3 \times 10^{-5}$ pa

RRF = 30 could be achieved using a single channel with $MTBF_{DU} > 15$ y and $T = 1$ y.

Explanation

The frequency of hazardous events H in low-demand mode scenarios is the demand rate D multiplied by the probability that all of the protection layers fail.

For example, with a single channel low-demand mode safety function as the only protection layer:

$PFD \approx \lambda_{DU}$ x $T/2$

$H \approx \lambda_{DU}$ x $T/2$ x D

The frequency of hazardous events H in high-demand or continuous mode scenarios is simply the rate of undetected dangerous failure:

$H \approx \lambda_{DU}$

Both high-demand and low-demand functions would achieve the same residual risk of hazardous consequences when $D \approx 2/T$.

$H \approx \lambda_{DU}$ x $T/2$ x $2/T \approx \lambda_{DU}$

However, the demands are twice as frequent as the periodic inspection and tests at this borderline point. The approximation $PFD \approx \lambda_{DU}$ x $T/2$ is no longer valid when $D > 1/T$.

Failure in the safety function is more likely to be revealed on demand when the demands are more frequent than the tests. The probability of failure on demand for the safety function is then $PFD \approx \lambda_{DU}$ x t, where t is the time elapsed since the last successful operation of the function. That could have been either on demand or in the most recent inspection and test.

$H \approx \lambda_{DU}$ x t x D

The average time between demands is $1/D$, so $H \approx \lambda_{DU}$ x $1/D$ x D at the time of the demand, and obviously $1/D$ x $D = 1$, so $H \approx \lambda_{DU}$ whenever the demands are more frequent than the tests.

The demands could be applied 100 times per year, but the hazard only occurs if an undetected dangerous failure has occurred in the safety function. The rate of the hazardous events is limited to the rate of undetected dangerous failures in the function.

Worked example 1 – low-demand mode

Imagine a process plant that has a high-pressure trip function. The objective is to isolate the main gas supply if the pressure in a vessel exceeds 10 MPa.

The process studies show that there is sufficient capacity in the system to withstand some leakage through the isolation valves. The valves need to limit the flow to below 1% of full open capacity for successful safety action.

The Company's preferred equipment list includes the items summarised in the following tables. The company has more than 10 calendar years of operational history with this equipment.

All failures of this equipment were recorded on corrective work orders. Root cause analysis was carried out for each failure and each failure was classified as safe or dangerous.

EZ-Pi pressure transmitter, 0.5 MPa to 50 MPa range (company standard hook-up)	
Safe range for measurement deviations	+/- 100 kPa
Safe range for response time constant	< 1 second
Total number of dangerous failures recorded 2 failures involved calibration drift due to temperature 3 failures involved degraded response times due to scaling	5
Total recorded operational service time	2,500 device-years

TriButter triple offset butterfly valve, 200 mm (non-TSO) Quarter turn spring-return pneumatic actuator (air to open) Acme 3590 3-way solenoid valve	
Safe range for closing time	< 8 seconds
Total number of dangerous failures recorded 1 failure involved excessive leakage due to seat damage (excessive leakage was classed as >0.5% of maximum design flow rate) 5 failures involved degraded closing time	6
Total recorded operational service time	240 device-years

SafeBall trunnion mounted ball valve, 200 mm (non-TSO) Quarter turn spring-return pneumatic actuator (air to open) Acme 3590 3-way solenoid valve	
Safe range for closing time	< 8 seconds
Total number of dangerous failures recorded 5 failures involved degraded closing time	5
Total recorded operational service time	250 device-years

The valves are fully closed and opened at least once every year. The stroking time of each valve is automatically recorded. The records are analysed to find deterioration in stroking times. The change in output flow or pressure is also recorded for most valves.

The Company's current maintenance policy is to inspect and test the valves once every year and to overhaul the valves at 16-year intervals. The maximum time permitted for continued operation with a safety function out of service is 3 days.

Consider the following questions:

1. What safety integrity level could be achieved by a single channel high pressure trip function that uses a ball valve as the final element?
2. What safety integrity level could be achieved for the same function with a dual channel 1oo2 architecture, and two ball valves as the final elements?
3. What sort of architecture would we need to achieve SIL 3 with a risk reduction factor exceeding 1,000 by a reasonable margin?

Assume that the *PFD* of the logic solver subsystem is $< 10^{-5}$.

Step 1 – Establish safety requirements

The functional requirements and the performance requirements for the safety function need to be defined objectively, with clear acceptance criteria.

The requirements need to be defined with enough detail so that equipment degradation, faults and failures can be categorised as either safe, dangerous, or of no effect.

The requirements for diagnostic functions need to be defined in enough detail so that the diagnostic coverage can be estimated for each failure mode. Diagnostic requirements would normally be based on FMEA or RCM studies.

Step 2 – Establish suitability for service

The next step is to establish that selected equipment is suitable for the intended service.

We need to have a dossier of information that meets the objectives of a safety manual. In this example we would expect to see a dossier or a report that meets the requirements of IEC 61511-1 §11.5.3, requirements for selection of devices based on prior use.

This information is essential before any estimate of failure probability can be made.

The dossier needs to include, for each device:

- Detailed specifications
- Instructions and constraints on use
- Requirements for maintenance
- Details of recorded failures, failure effects and remedial actions
- Total time in service over which the failures were recorded.

Step 3 – Failure mode and effects analysis

Carry out a failure mode and effects analysis for the entire safety function, including the interfaces to the process and including the cables and junction boxes.

For this example we will assume that:

- The trip setting is 10 MPa and an error of +/- 0.5 MPa is acceptable.
- The flow needs to be shut off within 15 seconds.
- Leakage of 0.5% x full open valve capacity is acceptable (i.e. Class II shutoff)

We need to identify which failure modes are dangerous and estimate the rate at which dangerous failures can be expected.[3]

Sensor sub-system failure rate

Assume that our FMEA concludes that the dangerous failure rate for the pressure transmitters can be expected to be similar to the rates previously recorded on site: 5 dangerous failures in 2,500 device years of service. The uncertainty is at least +/- 1 failure per 2,500 years.

The rate is approximately $\lambda_D \approx 0.002$ pa or 220 FITS, corresponding to $MTBF_D \approx 500$ y.

The range of uncertainty is at least 0.0016 pa to 0.0024 pa, or 180 to 280 FITS. The corresponding $MTBF_D$ is in the range of 420 to 620 y.

That range of dangerous failure rate is consistent with the ranges reported on silsafedata.com.

Diagnostic functions could be applied to reduce the rate of undetected dangerous failures. The types of failures that were reported failures may be detectable. Both internal diagnostics and external sensor comparison functions can be considered. The FMEA should include reliability targets for the diagnostics, typically either 0%, 90% or 99% for each individual failure mode.

Valve sub-system failure rate

Assume that our FMEA concludes that the dangerous failure rate for the ball valve can be expected to be similar to the rates previously recorded on site for that type of valve: 5 dangerous failures in 250 device years of service. The uncertainty is at least +/- 1 failure per 250 years.

The rate is approximately $\lambda_D \approx 0.02$ pa or 2200 FITS, corresponding to $MTBF_D \approx 50$ y. The range of uncertainty is at least 0.016 pa to 0.024 pa, or 1800 to 2800 FITS. The $MTBF_D$ is in the range of 40 to 60 y. That rate is consistent with the ranges reported on silsafedata.com.

We do not expect to achieve any diagnostic coverage for the valves because they are usually operated only once or twice each year.

Step 4 – Estimate the performance achieved

1. What SIL could be achieved by a single channel architecture?

The overall undetected dangerous failure rate of a single channel is ≈ 0.02 pa + 0.002 pa ≈ 0.022 pa, $MTBF_D \approx 45$ y.

With annual inspection and testing, our look-up table tells us that will easily achieve SIL 1 ($MTBF_{DU} > 15$) but is not quite enough for SIL 2 (requires $MTBF_{DU} > 50$).

[3] Example FMEA can be found in the MooN Safety Function Calculation Tool.

We can calculate *PFD* ≈ 0.022 pa x 1 y / 2 ≈ 0.011. The contribution to *PFD* from the logic solver subsystem is < 10^{-5} and can be neglected. The risk reduction can be estimated as *RRF* ≈ 2 x 45 y/1 y ≈ 90. This confirms *RRF* is towards the upper end of the SIL 1 range, not enough for SIL 2.

The full IEC 61508-6 equations give a similar result:

Sensor *PFD*	≈ 1×10^{-3}
Logic solver *PFD*	≈ 1×10^{-5}
Valve *PFD*	≈ 1×10^{-2}
Total *PFD*	≈ 1.1×10^{-2} corresponding to *RRF* ≈ 90.

Calculations based on Markov models will also give similar results.

2. What SIL could be achieved fault tolerant dual channel architecture?

With annual inspection and testing, our look-up table for $MTBF_{DU}$ tells us fault tolerant dual channel architecture will easily achieve SIL 2: RRF 300 ($MTBF_{DU}$ > 20) but is not quite enough for SIL 3 (requires $MTBF_{DU}$ > 70).

The *PFD* achieved by a fault tolerant dual channel architecture (1 out of 2 voting, or '1oo2') would be approximately

$$\begin{aligned} PFD &\approx 2/3 \times \beta \times \lambda_{DU} \times T \\ &\approx 2/3 \times 0.1 \times 0.022 \times 1 \quad \text{(assuming that } \beta = 0.1\text{)} \\ &\approx 0.0015 \end{aligned}$$

The risk reduction would be about

$$\begin{aligned} RRF &\approx 3/2 \times MTBF\ /\ (\beta.T) \\ &\approx 3/2 \times 45/0.1 \\ &\approx 700 \end{aligned}$$

The full 1oo2 equations given in IEC 61508-6 produce a slightly lower estimate for *PFD*:

Sensor *PFD*	≈ 1.01×10^{-4}
Logic solver *PFD*	≈ 10^{-5}
Valve *PFD*	≈ 1.13×10^{-3}
Total *PFD*	≈ 1.24×10^{-3}, corresponding to *RRF* ≈ 800.

This is towards the upper end of the SIL 2 range, not enough for SIL 3.

An estimate from the simplified model is typically about 10% to 20% more conservative than a detailed calculation.

Detailed calculations should not be expected to be more accurate because future failure rates cannot be predicted with accuracy. The future $MTBF_{DU}$ of the valves could be anywhere in a range of at least 40 y to 60 y. It depends on how well the condition and performance of the valves is maintained. The corresponding uncertainty range for *RRF* would span from about 600 to 900.

The uncertainty results from variability in human factors, environmental factors and systematic factors. These factors can be controlled to some extent, but they cannot be modelled or predicted with precision.

Step 5 – Propose an architecture to achieve SIL 3

The simplified model clearly shows that *RRF* depends primarily on $MTBF_{DU}$, *T*, and β:

$$RRF \approx 3/2.MTBF_{DU}/(\beta.T)$$

The performance of a dual channel 1oo2 architecture can be improved to SIL 3 by any of these strategies:

- Reducing the average test interval
- Reducing the failure rate through reliability centred maintenance
- Reducing the rate of undetected failures through diagnostic functions
- Reducing the likelihood of common cause failures.

Other architectures such as 1oo3 and 2oo3 could also be considered.

1. Shorter inspection and test intervals

Inspecting and testing the valves fully at 6-month intervals would double the *RRF* of the 1oo2 architecture to approximately 1400.

If full tests are not practicable then we could consider implementing inspection and partial stroke tests during the year. The intervals can be staggered. One valve could be tested after 4 months and the second after 8 months. A proof test coverage of about 60% is typically achievable. The inspection and test of each single valve should reveal most failures with a common cause that might affect the other valve. The weighted average test interval *T* could be estimated as:

$$T = PTC.T_1 + (1 - PTC).T_2$$

$$T \approx 0.6 \times 0.33 + 0.4 \times 1 \quad \approx 0.6$$

The risk reduction would be about *RRF* ≈ 3/2 x 45/0.06 ≈ 1100

It would be better to round the *RRF* to 1 significant figure of precision: *RRF* ≈ 1000.

There is too much uncertainty in the failure rate estimates to justify 2 significant figures of precision.

2. Shorter overhaul or renewal intervals

The simple design rule $MTBF_{DU} > 70.T$ for *RRF* > 1000 shows that $MTBF_{DU}$ of 70 y would be just enough for SIL 3 performance if *T* = 1 year.

Preventive maintenance could be expected to improve *MTBF* of the valves by at least a factor 3, from 50 y to 150 y. The risk reduction would be about *RRF* ≈ 3/2 x 150/0.1 ≈ 2000

The test and inspection interval for the valves could be kept at once a year. The quality of inspection and test could be improved through independence, using two technicians instead of one.

The planned service interval could be shortened from 16 y to 8 y. The service might include changing stem packing and seals, lubrication and cleaning. The service activities could be based on reliability-centred-maintenance studies.

Full records would need to be kept of the condition of each valve assembly component found at each inspection, test and service.

The level of detail in the inspection and testing, and the level of independence in verification of results need to be appropriate for SIL 3. As a minimum we would expect detailed check sheets and independent (in-situ) review by suitably experienced technician. It should not be simply a 'desktop review' by the maintenance team leader.

Consider staggering the service intervals. For instance, service the first valve after 8 y and the second valve 4 y later (subject to condition). Bring service forward if required.

Note that this strategy might leave us with the 'Resnikoff conundrum': the failure rates might be too low to measure. For example, there might be only 12 valves in SIL 3 service at this facility. A target failure rate of 1 failure in 150 device-years of service corresponds to only 1 failure in 12 calendar years. Some safety functions need to be designed so that failures can be prevented rather than counted.

Clearly, the *RRF* cannot be estimated with any precision or accuracy. The point of the analysis is only to provide a rational basis for the maintenance strategy and to set objective performance targets for the maintenance team.

3. Improve diagnostic coverage

Diagnostic coverage could be considered for the pressure transmitters, but it would not make any significant difference. The sensors contribute less than 10% of the total safety function *PFD*.

Diagnostic coverage would only be practicable on the valves if they could be operated at least several times each week.

4. Reduce common cause failures

The *RRF* could be improved by a factor of about 2 or 3 by reducing the likelihood of common cause failures.

A combination of the following strategies could be considered:

- Review the root cause analysis of previous failures to identify failures that had similar causes
- Conduct a detailed FMEA to identify further potential causes of common cause failure
- Prepare detailed maintenance, inspection and test procedures based on FMEA or RCM
- Use diverse equipment design and selection for the safety function channels:
    - Use a ball valve for one channel and a butterfly valve for the other channel
    - Use a pneumatic actuator for one channel and a hydraulic actuator for the other channel
    - Use different types of solenoid valves
- Use different people to inspect, test and maintain the equipment
- Stagger the inspection and test intervals
- Stagger the service intervals

5. 1oo3 architecture

The benefit of 1oo3 voting architecture is limited by common cause failures, unless the third channel is completely separate and diverse.

The SINTEF PDS Method Handbook and IEC 61508-6 Annex D both suggest that β_{1oo3} = 0.5 x β_{1oo2}.

Adding a third channel will therefore only improve the *RRF* by a factor of approximately 2.

6. 2oo3 architecture

2oo3 architecture could be considered for the sensors. It is not commonly used for valve assemblies though it is possible.

The *PFD* achieved by 2oo3 architecture is a factor of approximately 1.5 x to 2 x worse than 1oo2 architectures if the test intervals are unchanged.

The SINTEF PDS Method Handbook and IEC 61508-6 Annex D present slightly different scaling factors to account for 2oo3 voting architectures in their common cause failure β models.

The SINTEF PDS Method Handbook suggests that β_{2oo3} = 2 x β_{1oo2} and IEC 61508-6 Annex D Handbook suggests that β_{2oo3} = 1.5 x β_{1oo2}.

Staggering the inspection and test intervals for the 3 separate channels effectively reduces the average test interval by a factor of 3. Failures or deterioration identified on any one channel should lead to investigation of the remaining 2 channels at the same time. The *PFD* achieved by a 2oo3 architecture would then be similar to 1oo2, because the likelihood of common cause failures is reduced.

The overall *RRF* of the safety function would not be significantly affected by applying 2oo3 voting for the sensors. The sensors contribute less than 10% of the total safety function *PFD.*

The advantages of 2oo3 voting include:

- Inspection and testing can be carried out without bypassing the trip function
- The spurious trip rate is reduced
- Sensor comparison diagnostics may be applied to improve diagnostic coverage

Step 6 – Update safety requirements

The requirements and design process is iterative. Update the safety requirements specification or referenced documents with the architecture, diagnostic and test interval outcomes.

Worked example 2 – high-demand mode

Consider the same example in high-demand mode:

1. What safety integrity level could be achieved by a single channel high pressure trip function that uses a ball valve as the final element?
2. What safety integrity level could be achieved for the same function with a dual channel 1oo2 architecture, and two ball valves as the final elements?
3. What sort of architecture would we need to achieve SIL 3 with a reasonable margin?

Steps 1 to 3 – As above

1. Establish safety requirements

2. Establish suitability for service
3. Failure mode and effects analysis

Step 4 – Estimate the performance achieved

1. What SIL could be achieved by a single channel architecture?

The overall undetected dangerous failure rate of a single channel is ≈ 0.02 pa + 0.002 pa ≈ 0.022 pa, $MTBF_D \approx 45$ y

Our look-up table shows us SIL 1 can be achieved. *DC* is not needed for SIL 1.

SIL 2 would require *DC* > 70% and SIL 3 would require *DC* > 97%.

2. What SIL could be achieved fault tolerant dual channel architecture?

The overall dangerous failure rate of a dual channel architecture with $\beta \approx 0.1$ is
≈ 0.1 x 0.022 pa
≈ 0.0022 pa, corresponding to $MTBF_D \approx 450$ y

Our look-up table shows us SIL 2 can be achieved. *DC* is not needed for SIL 2. SIL 3 would require *DC* > 70%.

Step 5 – Propose an architecture to achieve SIL 3

The target for SIL 3 is $\lambda_{DU} < 1\text{x}10^{-7}$ per hour, or $MTBF_{DU} > 1{,}200$ y.

The dangerous failure rate *(DFR)* would need to be improved by a factor of about 3 to achieve SIL 3 in a dual channel architecture.

SIL 3 would require *DC* > 70%, or else a reduction in $MTBF_D$.

The 2 factors that directly determine *DFR* are β and λ_{DU} (or its reciprocal $MTBF_{DU}$).

$$DFR \approx \quad \beta.\lambda_{DU}$$

The performance can be improved to SIL 3 by any of these strategies:

- Reducing the failure rate through reliability centred maintenance (RCM)
- Reducing the rate of undetected failures through diagnostic functions
- Reducing the likelihood of common cause failures.

1. Shorter inspection and test intervals

It might seem that inspection and testing are not relevant because the inspection and test interval does not appear in the equation for estimating dangerous failure rate.

Regular and methodical inspection and testing is essential in managing the dangerous failure rate of any equipment. The main objective of inspection and testing is to reveal degradation before it results in failure.

Annual inspection and testing is usually sufficiently frequent for SIL 1 and SIL 2. Daily or weekly inspections and/or tests could be considered for SIL 3.

The frequency of inspection and testing can be based on RCM or FMEA studies. It depends on the type of deterioration that can be expected and on how quickly the deterioration might lead to failure.

The likelihood of common cause failures can be reduced by staggering inspection and test intervals for each of the N channels in MooN architectures.

2. Shorter overhaul or renewal intervals

The simple design rule suggests that $MTBF_{DU}$ > 120 y is just enough for SIL 3 performance with $\beta \approx 0.1$.

Preventive maintenance could be expected to improve *MTBF* of the valves by at least a factor 3, from 50 y to 150 y. That would be sufficient to achieve SIL 3 but without any margin for uncertainty.

3. Improve diagnostic coverage

SIL 3 would require the overall safety function to have *DC* > 70%.

DC > 70% should be feasible for the pressure transmitters, but the sensors contribute less than 10% of the total safety function λ_{DU}. Diagnostic coverage would also be need on the valves. That may be feasible in high demand mode applications if the valves are operated fully at least several times each week.

Diagnostic coverage would usually be practicable on valves in continuous mode applications.

4. Reduce common cause failures

As for low demand mode, the likelihood of common cause failures can be reduced by at least a factor of 2 or 3. The same strategies can be used, as described in the section 'Reduce common cause failures' above.

Worked example 3 – low-demand mode with generic data

For this example we will analyse the same safety function but with failure rates similar to those found in industry databases and commercial software packages.

Typical sensor system failure data

These failure rates are typical of sensors such as pressure sensors and include failures in cabling and logic solver input channels.

Dangerous detected failure rate (FITS)	1,000
Dangerous undetected failure rate (FITS)	200
Dangerous not-detected failure rate (FITS)	10
Safe detected failure rate (FITS)	500
Safe undetected failure rate (FITS)	70

MTTR (days)	3
Periodic test interval T_1 (years)	1
Full coverage test interval T_2 (years)	6
Proof test coverage	0.95
Common cause failure fraction	0.1

Typical final element system failure data

These failure rates are typical for a pneumatically actuated shutdown valve, including solenoid and actuator, and failures in cabling and logic solver output channels.

Dangerous detected failure rate (FITS)	300
Dangerous undetected failure rate (FITS)	2800
Dangerous not-detected failure rate (FITS)	100
Safe detected failure rate (FITS)	300
Safe undetected failure rate (FITS)	300

MTTR (days)	3
Periodic test interval T_1 (years)	1
Full coverage test interval T_2 (years)	8
Proof test coverage	0.96
Common cause failure fraction	0.1

Again, assume that the *PFD* of the logic solver subsystem is < 10^{-5}.

Steps 1 to 3 – As above

1. Establish safety requirements
2. Establish suitability for service
3. Failure mode and effects analysis

Step 4 – Estimate the performance achieved

1. What SIL could be achieved by a single channel architecture?

The overall undetected dangerous failure rate of a single channel is ≈ 2,800 + 200 FITS ≈ 3,000 FITS

≈ 0.026 pa, $MTBF_D$ ≈ 40 y

The contribution to *PFD* from the logic solver subsystem is < 10^{-5} and can be neglected.

The effect of limited proof test coverage factor can be estimated using $T = PTC.T_1 + (1 - PTC).T_2$

For the sensors, *T* ≈ 0.95 x 1 + 0.05 x 6 ≈ 1.25

For the valves, *T* ≈ 0.96 x 1 + 0.04 x 8 ≈ 1.25

The simplified method gives:

PFD ≈ 0.026 pa x 1.25 y / 2 ≈ 0.016

RRF ≈ 2 x 40 y/1.25 y ≈ 60

This is towards the upper end of the SIL 1 range, not enough for SIL 2.

The full IEC 61508-6 equations give a similar result:

Sensor *PFD*	≈ 1.2×10^{-3}
Logic solver *PFD*	≈ 1×10^{-5}
Valve *PFD*	≈ 1.5×10^{-2}
Total *PFD*	≈ 1.6×10^{-2} corresponding to *RRF* ≈ 60.

Calculations based on Markov models will also give similar results.

2. What SIL could be achieved fault tolerant dual channel architecture?

For a fault tolerant dual channel architecture (1 out of 2 voting, or '1oo2'), assuming that $\beta = 0.1$:

The simplified method gives:

$PFD \approx 2/3 \times 0.1 \times 0.026$ pa x 1.25 y / 2 ≈ 0.018

$RRF \approx 3/2 \times 40$ y/(0.1 x 1.25 y) ≈ 450

The full 1oo2 equations given in IEC 61508-6 produce a slightly lower estimate for *PFD*:

Sensor *PFD*	≈ 1.2×10^{-4}
Logic solver *PFD*	≈ 10^{-5}
Valve *PFD*	≈ 1.8×10^{-3}
Total *PFD*	≈ 1.9×10^{-3}, corresponding to *RRF* ≈ 520.

Both models estimate the SIL in the upper mid SIL 2 range.

The estimate from the simplified model is 12% more conservative than the detailed calculation in this example.

Step 5 – Propose an architecture to achieve SIL 3

The performance needs to be improved by a factor of at least 2 to achieve the minimum *RRF* required for SIL 3. Improving the performance by a factor of at least 3 would provide some margin for uncertainty.

Again, the simplified model clearly shows that there are only 3 main factors that drive the *RRF*: $MTBF_{DU}$, T, and β.

$$RRF \approx 3/2.MTBF_{DU}/(\beta.T)$$

As before in the first example, the performance of a dual channel 1oo2 architecture can be improved to SIL 3 by any of these strategies:

- Reducing the average test interval
- Reducing the failure rate through reliability centred maintenance
- Reducing the rate of undetected failures through diagnostic functions
- Reducing the likelihood of common cause failures.

References

This summary of simple design methods follows on from these previous I&E Systems publications:

New Approach to SIL Verification (2018)

The Myth of Proof Testing and Mission Time (2020)

Dealing with Uncertainty (2022)

High demand and continuous mode safety functions compared with low demand mode (2023)

The MooN SIF model (2023)

A simplified MooN model for safety functions (2024)

Creative Commons Licence

The document was prepared by Mirek Generowicz and Blake Merritt of I&E Systems Pty Ltd.