RISKTEC ESSENTIALS

Quantitative Risk Assessment (QRA)

An introduction to the quantitative assessment of risks associated with high hazard facilities.

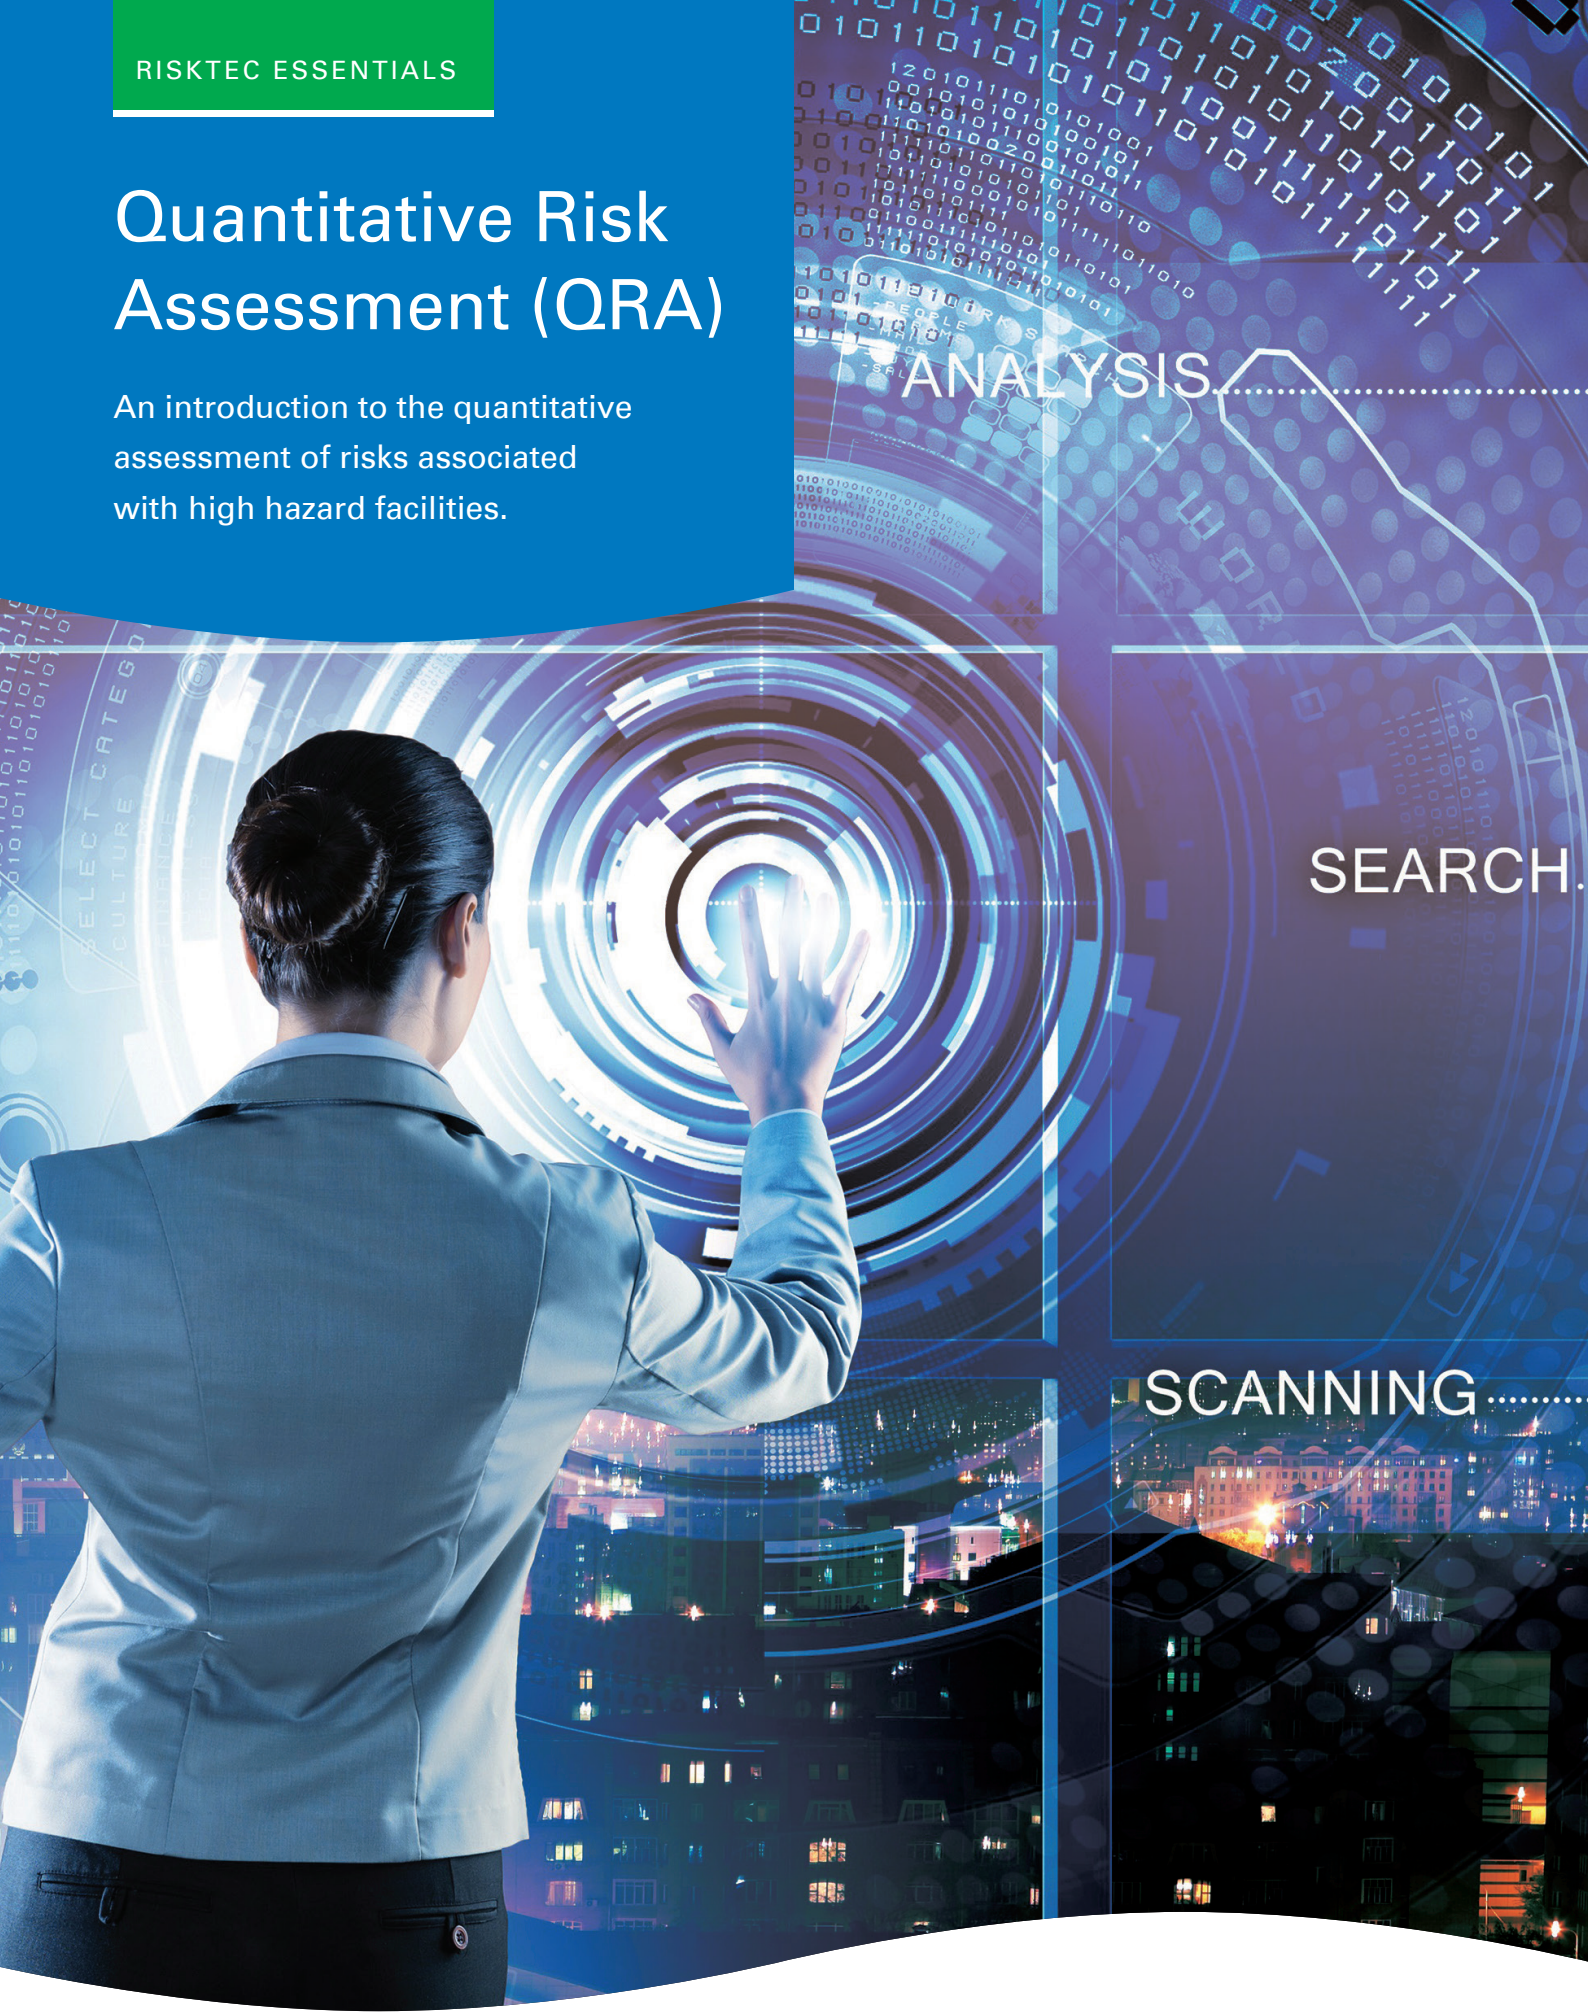

TÜVRheinland®
Risktec

risktec.tuv.com

Risktec is an established, independent and specialist risk management consulting and training company and is part of the TÜV Rheinland Group. At Risktec we believe in sharing our expertise and knowledge with our clients.

Contents

05 Foreword

06 QRA across major hazard industries

08 An introduction to QRA

10 An introduction to PSA

12 What's in a number?

14 QRA of multi-facility risks

16 Applying QRA more widely

18 Which QRA software?

20 Integrating CFD into QRA

22 Fire and gas detector mapping

24 Land use planning criteria for pipelines

26 Risk criteria

34 About Risktec

Foreword

Welcome to this volume of Risktec Essentials, which brings together a collection of short articles on quantifying risks in high-hazard sectors.

Many companies and regulators around the world require that safety risks are reduced to acceptable levels. This implies a need to quantify the level of risk. But how is this done? What are the limitations and pitfalls? And how do we compare the calculated risks with criteria? We hope Risktec Essentials provides some useful insights to help answer these questions.

Articles on other risk and safety management topics can be viewed at **risktec.tuv.com/knowledge-bank.aspx**

© Risktec Solutions Ltd. 2018

Quantitative Risk Assessment across major hazard industries

A BRIEF HISTORY OF QRA

The terms QRA (Quantitative Risk Assessment), PSA (Probabilistic Safety Assessment) and PRA (Probabilistic Risk Analysis) are used synonymously in different industries to describe various techniques for evaluating risk. Whilst quantification of risk for specific issues has been around for a long time, the grandfather of modern probabilistic assessment of the overall risk for an entire major hazard facility is generally accepted to be WASH-1400, commissioned by the US Nuclear Regulatory Commission in 1975. This quantified the safety risks associated with the operation of all electricity generating nuclear power plants in the US. The nuclear industry led the way, motivated by a desire to demonstrate that the actual risk was less than other industrial facilities and counter the public's perception that nuclear stations are very risky because the worst case consequences are potentially so catastrophic.

It is not surprising that the petro-chemical industry followed suit shortly after, since the toxic effects of large chemical releases can disperse many miles and affect large numbers of people in local towns and cities. Explosion effects can also be devastating. For example, an explosion in 1974 at the Flixborough chemical plant in the UK killed 28 people.

One of the first major QRAs for petrochemical installations was of the highly industrial area of Canvey Island near London, in 1978.

Flixborough explosion, 1974

BOX 1 - MISUSES OF QRA

- Using unrealistic or innacurate models and data
- Ignoring the uncertainties involved
- Manipulating the results to justify desired decisions
- Arguing everything is safe because a calculated risk level is lower than a numerical criterion
- Neglecting deterministic arguments

The UK's offshore oil and gas industry came relatively late to formal QRA of overall risks, prompted by the Piper Alpha disaster in 1988 in which 167 workers lost their lives. The rail industry in the UK also started formal QRAs in the early 1990's, against a background of train accidents, including the Clapham Junction crash in 1988 when three rush-hour trains collided, killing 34 people.

WHY CONDUCT QRA?

All industries were motivated to use QRA for much the same reasons – to provide insights into the nature of the facility that is being managed, to design defence in depth, to understand any constraints on operating the facility and any issues that require further investigation.

It is fair to say that in each industry, at various points in time, QRA has been misused, typically in efforts to 'prove' that calculated risk levels meet numerical risk acceptance criteria (see Box 1). This perhaps stems from the heavily engineering-biased culture within major hazard industries, where there is desire to have a precise answer represented by a number.

Whenever this has happened the industry has tended to redeem itself, getting back to a sensible use of QRA to help reduce risk by making better risk-informed decisions. In particular, the probabilistic approach of QRA can be extremely useful in demonstrating that a broad range of scenarios have been considered. This contrasts with more traditional deterministic approaches which are always left open to yet another "what if?" question and yet another study. A robust, well developed QRA can readily handle such questions and put the findings in the context of the total risk profile.

ARE THERE ANY DIFFERENCES IN QRA BETWEEN INDUSTRIES?

A formal QRA attempts to answer the questions:

1. What can go wrong?
2. How often does it happen?
3. How bad are the consequences?
4. Is the risk acceptable?

In addressing these questions, there are a number of subtle differences in the way the nuclear, petrochemical, oil and gas, and rail industries go about their QRAs. These differences tend to be shaped by the risk criteria, technology involved, the nature of the hazard itself, and whether or not the system being modelled is static.

Chernobyl nuclear accident, 1986

Risk criteria. The QRA must be able to generate risk results in the right form to allow comparison with the risk criteria set by the regulator or the operator. For example, where members of the public are exposed to toxic chemicals, the QRA may need to be

able to generate 'FN-curves', which are plots of the frequency F of events which may cause N or more fatalities. In contrast, the nuclear industry tends to require that the frequency of a given radiological release and associated doses are less than defined levels. On an offshore platform however, where only workers are exposed, the emphasis is more on the individual risk to personnel.

Technology involved. In the nuclear reactor industry, for example, a great deal of effort is expended in analysing the causes and frequency of initiating events because a lot of complex engineering has to fail before a reactor core can be damaged and radionuclides are released from the containment. The petrochemical and oil and gas industries, on the other hand, tend to start their QRAs with historical frequencies of gas or oil leaks because leaks are quite common and data are readily available. There is also less redundancy and diversity in safety systems to model than for nuclear power plants.

Nature of the hazard. Another difference relates to analysing the impact of the hazardous event. A fire or explosion in a chemical plant or a train derailment or collision both have the potential to cause fatalities immediately, within seconds of the event. Nuclear releases however may cause latent health effects, the extent of which may not be known until many years later. For example, the Chernobyl nuclear accident in 1986 caused 31 immediate deaths, but by 1991 some 7,000 clean-up workers were believed to have died and some estimates of the eventual death toll are as high as 75,000. Clearly, QRAs are sensitive to the models used to estimate the likelihood of fatality from the magnitude of the hazard.

Dynamic or static conditions. Unlike nuclear stations or chemical plants, which are located at a single site, trains can travel hundreds of miles through different local environments, picking up and putting down differing numbers of passengers. Rail industry QRAs are specifically designed to handle these transitory aspects.

There are other differences between industries as well, such as the way that fatalities are modelled during evacuation, or the extent to which frequencies and consequences are integrated into an overall risk picture, but what is evident is that the differences are not as great as one might first think – the level of detail and the focus of the analysis are shifted to enable QRA to help answer the specific questions unique to the industry.

SO WHAT IS THE SAME?

One thing that all major hazard industries do agree on, however, is the best uses of QRA (see Box 2).

Even so, QRA is not without its limitations (see Box 3). Because QRA involves lots of numbers it appears to be objective when in fact there are many judgements throughout the analysis. Some judgements are explicit but many more are hidden within data and methods. It is the role of QRA practitioners to interpret results in the context of the uncertainties inherent within the analysis.

In general, attitudes have changed from early scepticism, not helped by over-selling by QRA enthusiasts, to one of positive support for QRA approaches that provide clarity of focus to controlling hazards.

WHERE IS QRA TODAY?

Today, the nuclear industry is arguably taking an increasingly integrated approach to QRA, in that there is a more transparent combination of the frequency and consequences of events. It will be interesting to see how this plays out in light of the swelling interest in building new nuclear power stations.

Offshore UK, which is a mature operating environment, has seen the regulator pulling back somewhat from QRA and placing a much greater emphasis on the operational management of asset integrity. Nevertheless, internationally, there is an increasing demand for QRA and the value it can add, especially on major design projects.

The rail industry in the UK is using QRA more and more, applying it to increasingly complex situations to obtain a greater understanding for important risk-based decisions.

BOX 2 - BEST USES OF QRA

- Help to reduce risk by supporting risk-based decision-making
- Comparing options during the design phase or for modifications during operations
- Supporting the demonstration that risk levels are reduced As Low As Reasonably Practicable (ALARP)
- Defining requirements for land use planning
- Defining requirements for emergency response planning
- Not using QRA to support removing protective measures

BOX 3 - LIMITATIONS OF QRA

- Implies a level of accuracy that does not exist
- Can divert attention from precautionary or preventative measures
- Lack of reliable data in specific areas
- Often dominated by hardware issues and rigid handling of human factors
- Theoretical modelling may not reflect actual operations

CONCLUSION

Modern QRA has been around for over 40 years, led by the nuclear and onshore petrochemical industries, shortly followed by the offshore and rail industries. The differences in the focus and level of detail of QRA in each industry arise from the need to understand the critical risk issues unique to the industry.

But all industries agree that while QRA is not a panacea, it does help to make better risk-informed decisions, thus saving lives, protecting the environment, reducing economic loss and preserving the reputation of the associated organisation.

An introduction to QRA

Risk is defined as the combination of the probability of an event and the consequences of the event, but how do you go about undertaking a full Quantitative Risk Assessment (QRA) for an offshore oil and gas facility or an onshore petrochemical plant?

QRA is usually justified where there is:

- major personnel or environmental hazard potential; or
- significant economic implications; or
- a variety of risk trade-off decisions that need to be made.

WHAT IS QRA?

QRA is a technique used to systematically calculate the risks from hazardous events. It involves predicting the size of consequences associated with a hazard, and the frequency at which a release of the hazard may be expected to occur. These aspects are then combined in order to obtain numerical values for risk – usually risk of fatality.

QRA includes consideration of all identified hazardous events in order to quantify the overall risk levels. Similar hazardous events are often grouped and assessed together as bounding or representative events.

An array of third-party software packages exists for carrying out consequence modelling, frequency assessment or entire QRAs, but many of these calculations are also often done using spreadsheets.

QRA PROCESS

The main steps involved in a typical QRA study are shown in Figure 1. This also indicates how the outputs of each part feed into the next step of the process.

1. Identification of QRA scenarios

Hazards present on a facility will have been identified in a hazard study (e.g. HAZID) and the results are used as the starting point to identify the scenarios to include in the QRA. The hazard study will usually have ranked qualitatively the expected frequency and likely consequences of the identified hazards so that the QRA can focus on the significant hazards only.

All parts of the plant that contain a hazardous material (either toxic, flammable or both) should be included in the QRA. This could result in hundreds of different scenarios, so to simplify the analysis, the facility is split into sections (sometimes called isolatable sections) that contain similar materials under similar process conditions such as pressure and temperature.

2. Frequency assessment

To calculate the frequency of releases from an isolatable section involves counting the equipment items in each section (as shown

FIGURE 1 - QRA PROCESS

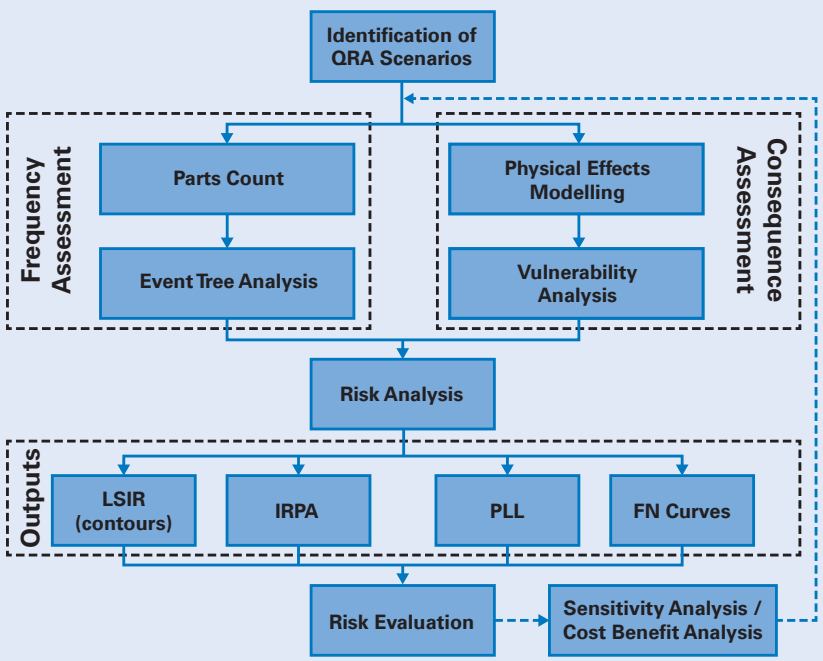

on the P&ID) and multiplying by equipment release frequencies from published databases. A range of hole sizes is required to obtain a spread of results and provide a realistic representation of the range of release sizes that could occur.

The many possible sequences that may result from these loss of containment releases are developed using event tree analysis, which considers probabilities for ignition and failure of protection systems.

The frequencies of non-process hazards such as ship collision or dropped objects are usually derived directly from incident data or using fault tree analysis.

3. Consequence assessment

A variety of models are available to estimate the consequences of the resulting fires, gas dispersion, explosions, etc. The vulnerability of people to these physical effects is determined in terms of probability of fatality using appropriate criteria. The consequence assessment will also identify potential escalation scenarios that may lead to further significant consequences.

A QRA may also consider the impact on the asset itself, the environment and the reputation of the company.

4. Risk analysis

The consequences and frequencies are then combined in an integrated QRA model to give numerical risk values. Offshore QRA is usually conducted using spreadsheets whereas onshore QRA is typically done with commercially available software.

Other non-process hazards also need to be analysed, such as personnel transport, occupational hazards, ship collision, aircraft impact and natural hazards. Each has its own specialist method for risk analysis.

The calculated risk values are summed for all possible outcomes and expressed in the required form. Offshore risk is usually expressed in terms of Individual Risk Per Annum (IRPA) and Potential Loss of Life (PLL), whereas onshore risk will typically calculate Location-Specific Individual Risk (LSIR) and FN-curves for societal risk considerations as well as IRPA.

5. Risk evaluation

The significance of the calculated risk levels is evaluated by comparing them with country and company risk acceptance criteria. The most significant contributors to the total risk are also identified to enable improvement measures to be targeted at those factors where the greatest risk reduction is likely to be gained.

Cost-Benefit Analysis (CBA) is a complementary tool to QRA and may be used to rank risk reduction options in order of cost-effectiveness.

QRAs have many sources of uncertainty, whether in the data, methods or assumptions, and these need to be understood through a targeted study that can inform the interpretation of the QRA results.

CONCLUSION

QRA is an established approach to understanding the risk levels associated with the operation of hazardous facilities and for providing insights into the main risk contributors and opportunities for risk reduction.

An introduction to PSA

Probabilistic Safety Assessment (PSA) provides an integrated and structured safety analysis for a nuclear facility that combines consideration of engineering design and operational features in a consistent framework. This established and systematic technique identifies all significant fault sequences which can lead to a radiological release and assesses their contribution to risk, on a best-estimate basis.

The operation of facilities such as a nuclear power plant involves:

- significant nuclear or radiological hazards; and
- complexity, with a high degree of redundancy and diversity of control and protection systems.

Typically, a PSA comprises combinations of event trees (modelling accident sequences) and fault trees (modelling system/equipment failures).

PSA is a complementary approach to deterministic analysis, and aims to:

- assess overall safety against explicit or implicit standards or criteria;
- determine the balance of the design, identify importance of systems and sensitivity to change; and
- identify potential areas of improvement and constraints of safe operations.

BOX 1 - LEVELS OF PSA

Level 1: Takes a wide range of initiating events, develops accident sequences using systems modelling, and derives fault sequences to determine the frequency of plant damage.

Level 2: Takes the output from level 1 and examines accident progression, to consider release magnitudes and frequencies from loss of the containment function. Used for determining accident management strategies and identifying potential design weaknesses in reactor containment buildings.

Level 3: Takes the output from level 2 in order to determine the individual risk and wider (societal) consequences of accidents by considering the risks to the public from off-site releases. Used for emergency planning.

This is summarised in Figure 1.

LEVELS OF PSA

The extent of a PSA is defined in levels, with level 2 extending the level 1 analysis and level 3 developing the level 2 analysis (see Box 1).

International regulation typically requires a level 2 PSA that is plant-specific and considers all relevant operational states, covering fuel in the core, spent fuel pond and on-site storage and all relevant internal and external initiating events.

PSA PROCESS

So, what is involved in undertaking a PSA?

1. Identification of PSA scenarios

The foundation of a PSA is the output of the fault and hazard identification process which is typically a list of all faults and hazards within the scope of the PSA together with initiating faults and their causes, preventive, protective and mitigative safety systems. Whilst the PSA will typically assess a wider range of faults than the deterministic design basis analysis, it is usual to apply screening criteria for events of a very low frequency (typically less than 10^{-8} per year) or events which

FIGURE 1 - PSA LEVELS

Plant response to initiating event	→ Level 1	Frequency of plant damage state (/Yr)
Physical effects and containment response	→ Level 2	Frequency and amount of radiological release (Bq/Yr)
Environmental dispersion, pathways, radiological uptake, dose-effect relationships	→ Level 3	Individual and societal risk (/Yr)

lead to insignificant radiological consequences.

2. Accident sequence analysis

This stage models the behaviour of the facility for the chosen faults and hazards, considering all possible combinations of success or failure of the protection systems to perform the safety functions. This, in combination with underpinning plant physics models, will identify the fault sequences which correspond to failure to maintain the facility within safe limits.

3. Systems failure analysis

This stage models the combinations of failures within the various safety systems which would lead to overall system failure. Fault trees will include events typically corresponding to component failure, common cause failure, component unavailability during maintenance or test and operator errors.

4. Data

A key challenge in the development of a PSA model surrounds acquisition of suitable data for the estimation of the frequencies and probabilities in the model. Where plant-specific data are available this is preferable, however generic data may also be required. These data require suitable manipulation to ensure that frequencies and probabilities are appropriately calculated. New designs may have a reliance on inherent or passive safety, for which specific failure data derivation techniques may be required.

An approach to modelling 'common mode' failures and specific values for operator errors also needs to be developed.

5. Internal and external hazards

A key expectation for a modern PSA is explicit modelling of internal and external hazards, such as fire, flood, extreme environmental and seismic events. This relatively new field uses the PSA model as a basis for a vulnerability assessment based on, for example, the zonal location of equipment or the potential for induced failures based on derived fragility parameters for equipment and structures.

CONCLUSION

PSA is an established discipline and forms a key input into the safety assessment of nuclear power plant, however it offers much more than a means to generate risk levels. It can be used as a way of identifying design weaknesses and assessing improvements as well as giving real insights into the effects of internal and external hazards.

As with all safety assessment techniques, PSA continues to adapt for changing requirements, for example in the assessment of advanced reactor designs or in support of security assessments.

What's in a number? Myths and realities of PSA and QRA

In many major hazard industries there is a regulatory requirement that the risk posed to people from operating facilities is shown to be tolerable and as low as reasonably practicable.

But, how do you measure risk? There are several techniques available to estimate risk, ranging from simple, qualitative frequency-consequence matrices through to use of complex quantitative event tree, fault tree and consequence models. The precise technique used depends on the stage in the facility lifecycle, the complexity of its design and operation, and the potential consequences of any postulated accidents. In certain industries (such as nuclear, rail, oil and gas), acronyms like PSA (Probabilistic Safety Assessment) and QRA (Quantitative Risk Assessment) have become synonymous with the use of complex risk models and computer codes, accessible to a limited number of practitioners who often use mystifying language, but are

MYTH #1 - I CAN BELIEVE THE NUMERICAL ANSWER WITHOUT QUESTION

Reality #1 - PSA/QRA provides an estimate of risk based on a large number of assumptions and input data, some of which may be uncertain

Paraphrasing the old saying - there are lies, damn lies and PSA! The numerical risk evaluated by a PSA/QRA can be meaningless without an understanding of the purpose for which it was intended and the underpinning assumptions and uncertainties. Often data will be taken from other facilities, or even generic databases, which may or may not be directly applicable. Some of the phenomena considered in the analysis may not be well understood and, although PSA/QRA is expected to be best-estimate, a bounding assessment may need to be adopted.

MYTH #2 - PSA/QRA IS ONLY REQUIRED TO COMPLY WITH THE REQUIREMENTS OF THE REGULATOR

Reality #2 - If properly conducted, PSA/QRA can provide valuable insights into the strengths and weaknesses of a design and the way it is operated

At different stages of facility development PSA/QRA can be used to identify important design or operational issues, and can help focus or prioritise their resolution, as well as looking at potential solutions. It can be used to optimise maintenance activities and therefore minimise plant outages. Ultimately, PSA/QRA provides an input into demonstrating that the risk posed by planned operations is as low as reasonably practicable.

able to generate seemingly authoritative risk numbers. As a consequence, a number of PSA/QRA myths have flourished.

KEY TO SUCCESS

The key messages to take away are that PSA/QRA should:

- Be viewed as tools to aid in managing risk at all stages in a plant lifecycle.
- Be integrated with the design and operating processes.
- Be underpinned by an appropriate level of data.
- Not be considered an exact science.

Results should be used with a degree of caution and should be supported by qualitative understanding before informing decision-making, such as a design change.

The more complex the project, the more sophisticated the PSA/QRA is likely to be, involving a larger number of stakeholders. Equally, the higher the associated risk, or sensitivity to an increase in risk, the more robust and comprehensive the supporting evidence should be.

CONCLUSION

To some, PSA/QRA may seem like another legislative or corporate hurdle. However, in the right hands, it provides a very powerful tool that can be used to aid understanding and support the decision-making process during the design and operation of a facility.

MYTH #3 - THERE IS ONLY ONE ACCEPTED METHOD FOR CALCULATING RISK

Reality #3 - There are several valid methods available to assess risk

Given the potential expense, it is important to ensure that the intended purposes of a PSA/QRA are well understood so that the right level and type of analysis is undertaken. The quantitative method used at the preliminary design stage may not be what is appropriate or necessary to support detailed design or for that matter, operations. For relatively low hazard facilities, a simple evaluation of risk may well suffice.

MYTH #4 - THE PSA/QRA IS MY SAFETY CASE

Reality #4 - PSA/QRA forms part of the overall safety case and is not a catch-all that will ensure your design is acceptable

PSA/QRA complements, but is not a replacement for good, qualitative studies such as bowtie analysis, deterministic safety assessment or similar conservative, non-quantitative techniques that are used to assess faults and hazards and the suitability and sufficiency of controls. Neither should PSA/QRA be used to justify non-compliance with legal requirements.

Quantitative assessment of multi-facility risks

Operators of major hazard facilities are required to understand and manage the risks presented to their workforce, the general public and the environment. Complications can arise when multiple facilities are located in close proximity, as each will contribute to the risks at neighbouring facilities.

In some cases, such as large industrial cities in the Middle East, a large number of hydrocarbon processing facilities are co-located in the same super-complex. City-wide risks are important in such situations for land use planning, determining sites suitable for further process facilities or for worker camps.

For offshore developments of extensive oil and gas fields, there may be multiple platform hubs with drilling centres and processing trains, each with significant hazard ranges. Field-wide risks are important when assessing separation distances and scheduling construction teams, particularly in fields where the highly toxic hydrogen sulphide is present.

So how can we assess city or field-wide risks for a multi-facility development?

SINGLE QRA MODEL SOLUTION

The natural solution is to develop a conventional Quantitative Risk Assessment (QRA) model to cover the whole complex or field. This approach can present many technical challenges. Although each facility will usually already have its own QRA, these may have been developed with different rule sets and assumptions, so a process of rule set alignment with relevant stakeholders would be required initially. The QRAs may also have been developed in different custom or commercial software packages, further complicating the integration process.

The size of the final QRA model also needs to be considered, as a direct combination of individual facility QRAs could lead to a large and cumbersome model, taking significant computational

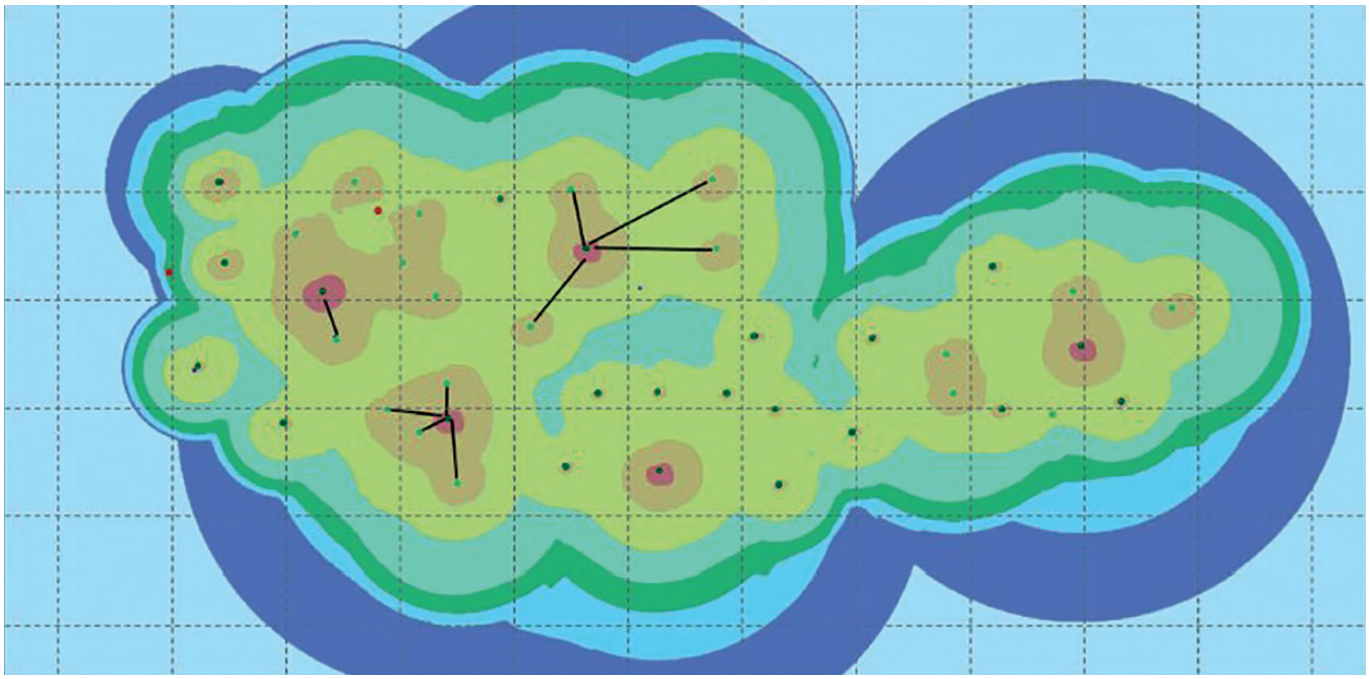

An example of risk contours for a multi-facility site

power to process - or even exceed the limitations of the software. A screening process to remove events with no potential for off-site impact would be required to mitigate this, on the understanding that ensuing results should only be used to analyse off-site impact. Nonetheless, it follows that developing an over-arching QRA model can involve substantial effort.

ALTERNATIVE SOLUTION

An alternate approach to consider before committing to an integrated model is to explore ways of exploiting the results of the existing QRAs. Bringing these together in a faster (albeit coarser) fashion has several advantages and can give useful results in a fraction of the time. One example is the use of the data behind the contour plots traditionally generated in onshore QRAs. Most QRA packages are able to export the underlying risk data generated from the model on a regular grid of points across a geographical area. These individual risk outputs can be combined on a common site-wide grid using simple translation and interpolation operations to produce risk contours spanning the full site.

As there are no risk calculations being performed, there is no dependence on specialist QRA packages and processing is fast. The relative locations of facilities can be quickly updated, with changes to the risk profile visualised almost instantly without long simulation times, giving a real benefit in situations where a multitude of potential layouts are being assessed. In this case, although the detailed contribution from specific events is lost, the overall benefits are realised for a fraction of the cost of an integrated QRA.

CONCLUSION

Developing QRAs spanning multiple major hazard facilities can be a challenging and time consuming process. Coarser results-based strategies have the potential to shortcut this, allowing city or field-wide risk profiles to be generated in a much shorter timescale.

Applying QRA more widely

In recent years, there has been increasing interest in extending the scope of Quantitative Risk Assessment (QRA) beyond the risk to people to look at areas such as environmental damage, economic impact and the effect on reputation. High profile accidents with widespread consequences, such as Deepwater Horizon, Buncefield and Fukushima, have left many organisations with a desire to understand better their exposure across the whole spectrum of potential risks. But how easy in practice is it to apply QRA more widely?

TRADITIONAL SAFETY QRA

QRA in the oil and gas industry, for example, focuses on the risk to workers and the general public from major hazards such as fires, explosions and toxic gas release. This process involves identifying the hazards, evaluating the frequency of the various hazardous events and undertaking consequence analysis to estimate the magnitude and effects of the resultant fire, explosion or gas cloud.

Geographical information is captured, including the location of the hazardous events and the number and distribution of people. This information and supporting analysis of the hazard progression (taking into account detection, isolation and ignition, for instance) are combined with the vulnerability of people to each hazard to calculate the risk to people.

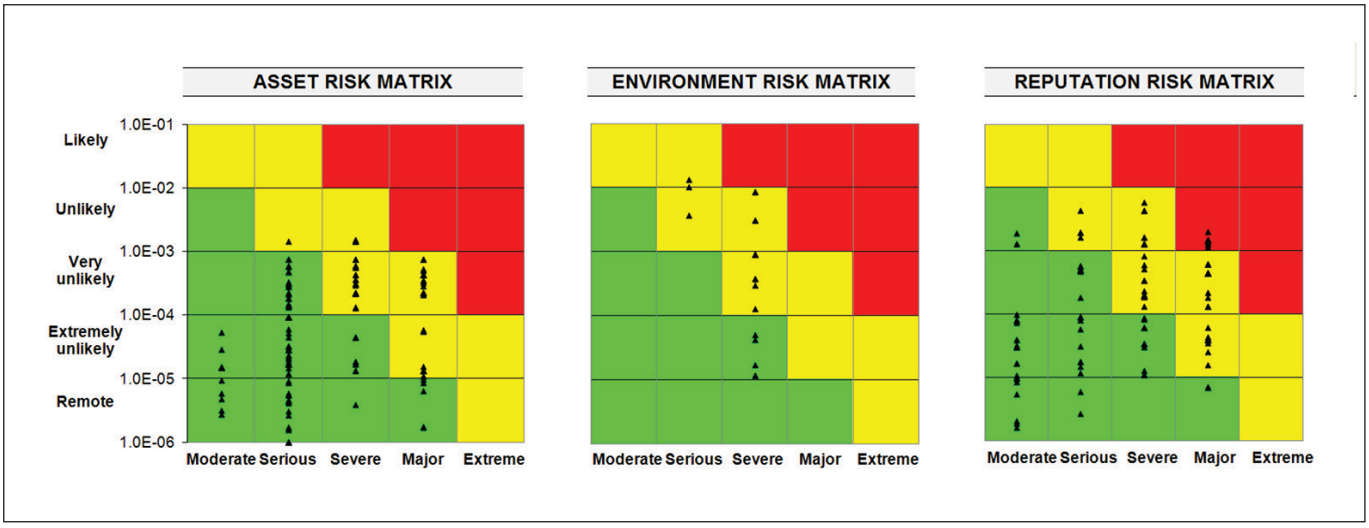

FIGURE 1 – Semi-probabilistic presentation of QRA results

WIDER EFFECTS

However, hazards may also have other negative effects beyond harming people. Liquid spills may cause harm to the environment, whereas fires and explosions can damage assets and infrastructure. These may lead to lost revenue, regulatory penalties, compensation to third parties, as well as damaging the reputation of the company involved. The information in QRA models can be extended to quantify some of these additional risks.

Harm to the environment is normally associated with releases of hydrocarbon or other chemicals either into the sea or onshore where it flows into water courses or permeates into the ground. The volume of release can often be estimated from the process data used in the conventional consequence modelling (release rate, duration and the volume of the isolated inventory). In practice, all potential sources of release would be screened first to determine whether they would reach the environment. Whilst quantifying clean-up costs is feasible, measuring the harm to the environment is more subjective and is perhaps best achieved using a number of discretely defined, qualitative categories (for example, see Figure 1).

Damage to assets and infrastructure depends on a combination of magnitude (overpressure or radiation) and in the case of fires, the duration, which may be limited by isolation and depressurisation. For onshore and offshore facilities it is usually straightforward to estimate the repair or rebuild cost. Lost production or processing revenues are sometimes a simple function of the outage period, though in many cases production is actually deferred rather than lost. However, oil and gas blowouts need to factor in the cost of bringing the well under control, which can be very high especially if a relief well needs to be drilled.

Regulatory penalties extending to loss of operating licence, compensation to neighbours and the public, and reputation issues are difficult to quantify, but it is usually possible to assign a qualitative indication of the harm, which can be presented as a risk matrix (similar to Figure 1).

CONCLUSION

The analysis of event frequency, event progression and consequences developed in traditional safety QRAs provides a sound platform from which to develop a wider picture of risk that can naturally include environmental, asset and economic factors.

Which QRA software?

Risktec completed a comprehensive survey of software currently available for undertaking Quantitative Risk Assessment (QRA) for onshore and offshore oil and gas facilities. The key requirement was that the software had to be available to users under licence, with full user support. This immediately removed from the search any "in-house" tools developed by consultants.

From an initial list of over 80 tools, only a handful of software products were found that could undertake full QRA. Also, offshore and onshore QRA tools tend to be packaged separately, reflecting the different characteristics that need to be modelled, e.g. offshore evacuation, or onshore far field impact on the public. What is clear is that there is no single best tool designed for both offshore and onshore QRA.

KEY FINDINGS

- There are no commercially available tools for coarse QRA at concept selection stage, but some consultants have in-house models.
- There is no single fully integrated offshore tool. In practice, most companies develop bespoke, installation-specific, linked spreadsheet models – see Fig 1.
- Onshore is better served and software products are generally well used and accepted. Non-hydrocarbon/ chemical risks (e.g. transport) still need to be quantified off-line, though they tend to be less critical onshore than offshore.
- A handful of products stand out as technical leaders – see Fig 2.

KEY SELECTION CRITERIA

Key factors to consider when selecting QRA software include:

- **Scope** – what exactly do you want to model and in how much detail? Can the software meet your requirements or will you be overwhelmed by the functionality?
- **Repeatability and transparency** – are the methods, rule sets and data visible and traceable?
- **Cost** – how much will licences, training, in-house time and external consultants cost over the long-run?
- **Integration** – how easy will it be to integrate the processes for managing the software and assessments into your company's management system?

RISKTEC VERDICT

Risktec survey – the choice is limited and there is plenty of scope for improvement in the software currently on the market.

Risktec choice – spreadsheets for offshore, either Safeti Onshore or Riskcurves for the more complex onshore studies because users have the flexibility to enter results from their preferred physical effects tool.

Don't be fooled by – good looks. Users want flexibility and transparency in methods, rule sets and data.

CONCLUSION

Users need to consider very carefully their requirements before selecting specific software. Often, using one of the onshore products is the best way to proceed. But the complexities of modelling offshore risks mean that most organisations develop their own spreadsheet models to utilise the methods, assumptions and data they understand to an appropriate level of detail. Organisations with multiple facilities, who want a flexible but more robust approach than spreadsheets, have an alternative cost-effective option: to develop their own bespoke model making use of Microsoft development tools or equivalent.

FIGURE 1 - INTEGRATED QRA MODELS VERSUS SPREADSHEET MODELS

Integrated QRA models *note 5*		Spreadsheet models	
Advantages	**Disadvantages**	**Advantages**	**Disadvantages**
• Inclusion of many models in a common computing environment • Models validated against experiment • Software quality assured by supplier • Technical support from software supplier • Available 'off-the-shelf' enabling early start of work • Recognised and generally accepted within the industry	• Difficulty of use and understanding – onerous user training and familiarity requirements (but decent results require complex modelling) • Lack of control and flexibility – user unable to modify software (can be an advantage) • Lack of transparency – hidden assumptions and calculation methods, 'black box' (requires high quality technical user manual) • High initial and ongoing costs (licences)	• Relatively easy to understand • Lower user training requirements and easier user familiarisation • Good spreadsheet models provide transparent calculations and assumptions • Better control – user able to develop spreadsheet model to level of detail required (flexibility of calculation and presentation) • Lower external cost (but man-hour time can be expensive)	• Prone to errors by the analyst • Can be personal to analyst and difficult to update by others without errors (requires careful QA) • Macro programming can be difficult to check • More time consuming to demonstrate validation • Perception – less sophisticated (when reverse is often true)

FIGURE 2 - LEADING QRA SOFTWARE

Offshore QRA	Onshore QRA *note 4*
• Safeti Offshore	• Safeti Onshore *note 1* • Riskcurves *note 2* • Shepherd *note 3*

Note 1 Incorporates PHAST physical effects tool.

Note 2 Incorporates Effects consequence modelling software which itself includes Damage software.

Note 3 Only available for use on Shell projects. Incorporates FRED physical effects tool.

Note 4 There are some new entrants such as HAMS-GPS, though these do not appear to be widely used.

Note 5 'Integrated' means that most calculations are done on-line within the software rather than off-line by other tools.

Integrating CFD into QRA

It would seem obvious that Computational Fluid Dynamics (CFD) studies of physical effects such as explosion overpressure or toxic cloud dispersion should be an integral part of Quantitative Risk Assessment (QRA) modelling. But in practice this is often not the case. What is important is to get the interface right at an early stage.

CFD FOR DESIGN

CFD studies are often performed to assess the dispersion of flammable or toxic releases, or overpressures resulting from explosions.

Traditionally, the output of such studies is a report containing 2D and 3D graphical presentations of results, sometimes accompanied by short videos showing how the dispersion or overpressures develop over time. The results may also be combined with event frequencies to derive 'exceedence plots' which indicate the likelihood of overpressures exceeding certain magnitudes.

These studies are generally targeted at designers with a view to improving understanding and informing the design, e.g. for developing a blast wall design to protect against a 1 in 1,000 year explosion.

QRA is often used to assess probabilistically the risks associated with a facility, including those due to explosions.

In principle, CFD should enable more refined QRA modelling of explosions, but often the CFD output is of limited benefit due to a range of factors, including, for example:

- The large number of scenarios that need to be considered in the QRA, compared to the limited number of design cases modelled by CFD.
- Extracting information from a CFD report can often be difficult and time consuming, and relies on manual estimates from figures or interpolation from data tables.

Ironically, this means that non-CFD consequence modelling software may have to be used instead, noting that this approach is unable to take specific account of the facility geometry.

In summary, the traditional CFD approach produces a fairly standard set of outputs which is not suited to the development needs of QRA models.

CFD FOR DESIGN AND QRA

Better integration of CFD studies and QRA can be achieved by specifying the outputs required by the QRA in advance and having raw data from CFD simulations delivered in electronic formats that can be interrogated easily. In the past, this has been hampered by the significant volume of data involved, but modern data storage and transfer arrangements mean this is no longer such an issue. With the

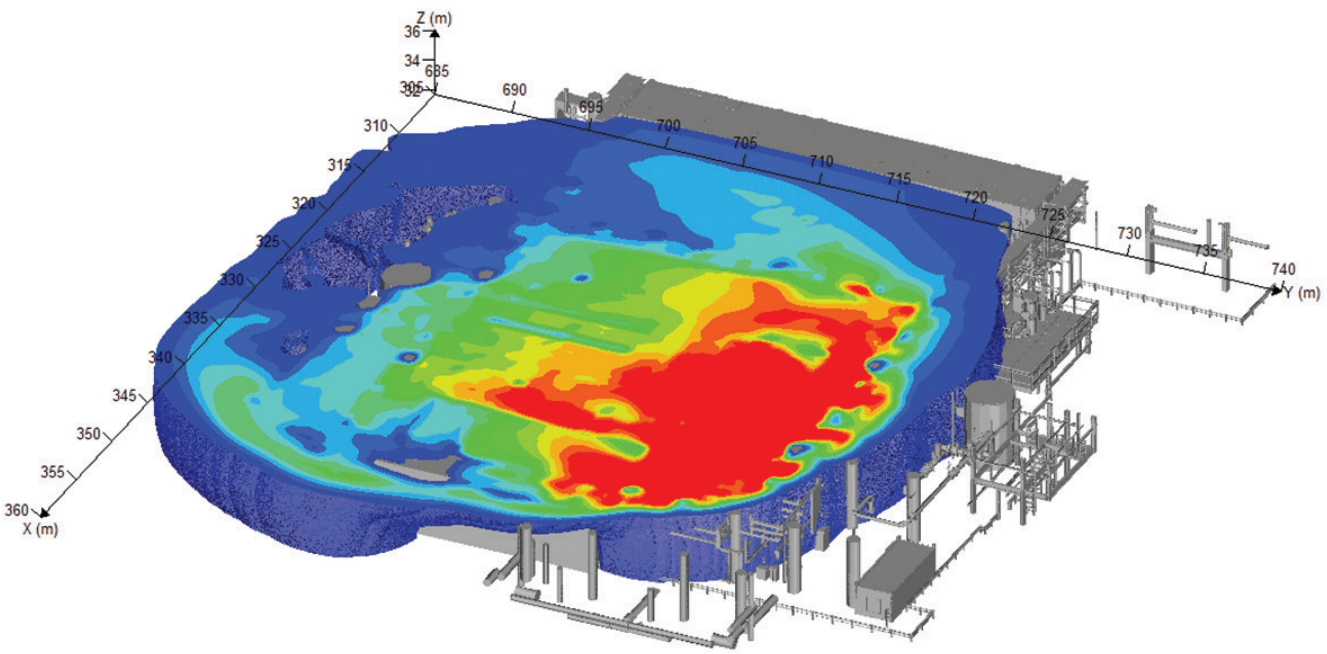

Typical CPD explosion modelling output

whole dataset available, the inherent limitations of the data published in the CFD report no longer applies.

Making it clear from the start what results are needed for QRA can also reduce the requirement for additional CFD analysis runs. For example, by specifying a grid of overpressure results across the facility rather than just at specific design elements means that it can also be used by the QRA to answer much wider questions relating to the effect of overpressures at different locations (an issue that would traditionally involve time consuming iterations with the CFD analyst).

Recent experience tells us that it is possible to obtain targeted outputs of raw data from CFD studies without significantly affecting cost, provided they are specified in advance.

OPTIMISING QRA

Making the most of CFD data depends on building the QRA model so that CFD output can be imported directly, without the need for post-processing. Achieving this relies on a consistent specification and format between the QRA and CFD. On occasion it is also possible to automate this data transfer process. Overall, this makes things much more efficient and enables more straightforward updates to the QRA following revisions to the CFD.

Ultimately, this should result in a QRA that is informed more explicitly by CFD results, enabling a more representative evaluation and understanding of risk, and its sensitivity to potential explosions.

CONCLUSION

As the use of CFD to support QRA continues to increase, getting the interface right at an early stage can reap real benefits in cost-effectiveness, as well as producing a better understanding of risk and, by extension, more focused risk mitigation.

Making the most of fire and gas detector mapping

Fixed fire and gas detection systems in processing facilities typically ensure that risk mitigation systems such as isolation, blowdown and active fire protection are activated in the event of a hazardous event. A well-designed system provides an appropriate level of detector redundancy to guard against false trips and detector faults. Fire and gas detector mapping studies provide an objective analysis of detector layouts to support the design process and optimise the number of detectors needed to meet coverage targets.

MAPPING OR MODELLING?

Modelling gas detection following a leak is a difficult task, since gas dispersion depends on a large number of variables such as process conditions, hole size, release position and direction, ventilation conditions, impingement, etc. Probabilistic dispersion studies using Computational Fluid Dynamics (CFD) can assess the likelihood of cloud formation across a process area to identify favourable detector locations. However, such analysis is time consuming and expensive, inevitably meaning that only a sample of the variables involved can be considered.

On the other hand, a mapping study avoids this degree of complexity by considering a reference cloud or fire of fixed dimension. For example, a maximum tolerable flammable cloud of 5m diameter at its Lower Flammable Limit (LFL) is often used in offshore environments, based on research indicating the onset of damaging explosion overpressures from clouds of this size.

A map and coverage statistics are generated by considering the number of gas detectors that would alarm as the position of this cloud is moved across a detection zone, and the detector layout is tuned to ensure predetermined spatial coverage goals are met. A similar approach is applied to fire detection, where the ability to detect a reference fire size is assessed. There is no consideration of likelihood in this approach, with the reference fire or cloud treated equally likely at each position in the detection zone.

INTEGRATION WITH OTHER ASSESSMENTS

Whilst good detection coverage levels are claimed in safety cases, demonstration of adequate coverage is not normally demanded by regulators. As such, fire and gas detection mapping tends to be viewed as a stand-alone study, separate from the traditional set of fire and explosion studies supporting the safety case. There are many advantages to taking a more holistic approach, though.

Studies such as Fire And Explosion Risk Assessment (FERA) and gas

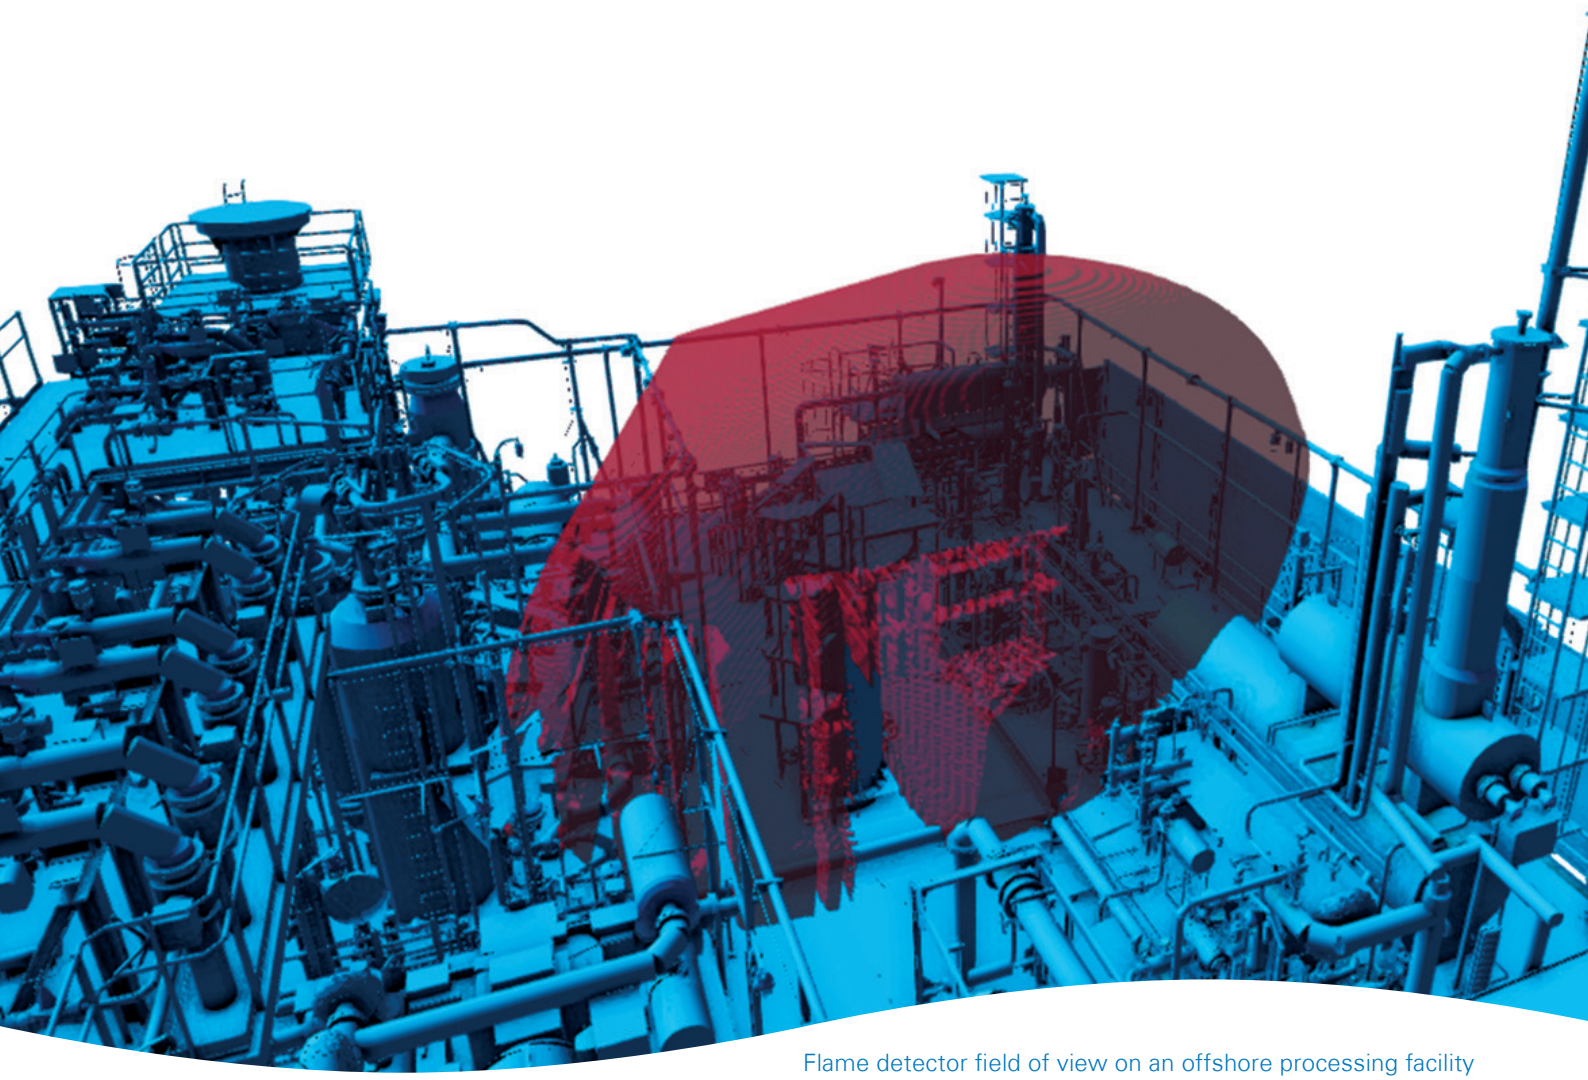

Flame detector field of view on an offshore processing facility

dispersion assessment provide a comprehensive analysis of fire, explosion, flammable and toxic gas events across a facility, identifying what hazardous materials exist, their location, consequences, durations and potential for escalation. This can provide crucial information for mapping.

For example, detection zones can be selected based on identified toxic and flammable hazard sources; flame detector fields of view can be calibrated according to the radiant intensity of fires in each area; and dispersion distances to detection levels (high alarm set points are typically 40%- 60% LFL) may be used to define the detection distance for the reference cloud. Additionally, scenarios with significant escalation potential can be identified from the FERA and extra importance placed on detection in these areas.

Detecting events clearly reduces risk and high detection probability is usually claimed in a QRA, yet failing to detect is often based on the reliability of detectors rather than the ability of the detection system to actually detect an event.

Coverage levels from the mapping study (e.g. >2 detectors in high alarm) can be used to estimate the minimum detection probability for many scenarios considered in the QRA. This allows for a more refined evaluation of escalation frequencies and the associated risk to personnel and plant.

CONCLUSION

Fire and gas detector mapping is becoming more commonplace for oil and gas facilities, supported by sophisticated software tools. Ideally such studies should be undertaken in conjunction with related assessments since they can provide valuable insights into associated safety claims. In this sense, they can contribute to a more thorough understanding of the installed hazard protection and ultimately lead to improved safety through better risk-informed design.

An introduction to land use planning criteria for pipelines

In practice, the segregation of hazardous industries and populated areas is not always achievable. Urban expansion and industrial development can often lead to an increased pressure to site hazardous industries, such as refineries, chemical plants, pipeline networks, etc., adjacent to vulnerable populations like residential areas, schools, hospitals and shops.

The risk to offsite populations from major accidents arising from the release of hazardous substances can be managed through the application of criteria for Land Use Planning (LUP), which are designed to aid planning decisions.

Pipelines transporting hazardous products present unique challenges to LUP criteria compared to fixed facilities. For example, without a site security fence, they can be accidentally or deliberately damaged; it may not be immediately apparent to the operator that a release from the pipeline has occurred; and emergency response may not be available along the length of pipeline, which in extreme cases can cross international borders.

INTERNATIONAL APPROACHES

LUP restrictions in proximity to transmission pipelines are regulated in one of three ways, depending on the jurisdiction:

1. Deterministic, e.g. USA and Canada.
2. Risk-based, e.g. Australia.
3. Combined deterministic and risk-based, e.g. UK, Singapore and Netherlands.

As an example, a deterministic criterion could be a development exclusion zone of 30m either side of the pipeline, or a requirement for public consultation within a distance of 200m.

FIGURE 1 - LUP MANAGEMENT APPROACH

	Deterministic	Risk/Goal Based
Pros	Provides absolute clarity of what is required of operators (and how to comply)	Performance goals allow more focus on the outcomes rather than the methods of achieving compliance
	Consistency of safeguarding reduces ambiguity and debate amongst safety professionals and regulators	Freedom to use different safeguarding solutions foster innovation and can lead to cost-savings
	Specific protection perceived as desirable by the regulators and other stakeholders can be directly specified	Specific outcomes desired by the public and stakeholders can be directly required
Cons	Requirements may fail to anticipate all circumstances	Requires more analysis and documentation to verify compliance
	May not encourage innovation with respect to safeguarding	Requires a well-trained, resourced and active regulator and operator
	May limit the operator's willingness to go beyond compliance	Public and other stakeholders may not have enough trust in government and industry to ensure risk goals are met

An example of a risk criterion could be restrictions on certain types of development, e.g. schools, in a zone on both sides of the pipeline where the individual risk of fatality is greater than $1x10^{-6}$ per year.

The advantages and disadvantages of the deterministic and risk-based approaches are summarised in Figure 1.

The deterministic approach, whilst simple and relatively easy to implement, may be overly pessimistic in nature and result in the unnecessary restriction of developments. That said, in specific circumstances the deterministic approach may be less conservative than a risk-based approach – for example, in toxic releases, where toxic clouds may extend to significant distances before they are diluted to safe exposure limits.

Rapid population growth and urbanisation may prompt the consideration of a risk-based approach since this potentially facilitates a more efficient use of land in proximity to pipelines. However, the success of risk-based approaches depends crucially on the use of appropriate data, assumptions and methods and the uncertainty inherent within key variables.

A combined deterministic and risk-based approach would appear to offer the best of both worlds: risk-based criteria tend to ensure that the solution is not overly conservative, while fixed distance exclusion zones tend to ensure a precautionary approach is taken where risk results may be uncertain.

CONCLUSION

Many developed countries around the world use LUP criteria to manage the location of new industrial developments and the encroachment of urban development near to existing hazardous facilities. Pipelines pose some unique challenges, not least when they bridge entire countries.

Some criteria are deterministic only, whilst others are solely risk-based. The most robust criteria tend to combine both deterministic and risk-based elements, enabling a balanced approach to safety.

Risk criteria – When is low enough good enough?

Criteria are used to help decide whether the risk associated with a project or activity is low enough to proceed. The key question then, is when is low enough good enough?

INTRODUCTION

No industrial activity is entirely free from risk and so many companies and regulators around the world require that safety risks are reduced to acceptable levels. The key question then is what level of risk is considered to be low enough? A subsidiary question is also what risk are we talking about, individual risk or societal risk? This article attempts to answer these questions.

Why have risk criteria?

Risk criteria are standards used to translate numerical risk estimates, e.g. risk of fatality of 10^{-7} per year, as produced by a Quantitative Risk Assessment (QRA), into value judgements such as 'negligible risk' that can then be set against other value judgements such as 'high economic benefit' in the decision-making process (Ref. 1).

Put more simply, criteria are used to help decide whether the risk associated with a project or activity is low enough to proceed.

A framework for risk criteria

The most common and flexible framework used for risk criteria divides risks into three bands (Ref. 2):

- An **unacceptable region**, where risks are intolerable except in extraordinary circumstances, and risk reduction measures are essential.
- A middle band, or **tolerable if ALARP region**, where risk reduction measures are desirable, but may not be implemented if their cost is disproportionate to the benefit achieved.
- A **broadly acceptable region**, where no further risk reduction measures are normally needed.

This framework is shown in Figure 1.

Figure 1 – Framework for tolerability of risk

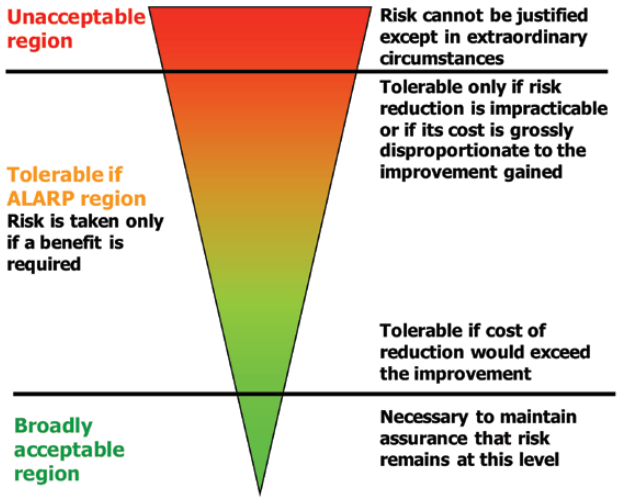

To define these bands, two levels of risk criteria are required; a maximum tolerable criterion above which the risk is intolerable and a broadly acceptable criterion below which the risk is insignificant.

Risk measures for loss of life

Risks to people may be expressed in two main forms:

1. **Individual risk** – the risk experienced by an individual person.
2. **Societal (or group) risk** – the risk experienced by the whole group of people exposed to the hazard. Where the people exposed are members of the public, the term societal risk is often used. Where workers are isolated and members of the public are unlikely to be affected, the term group risk is often used. Here, the term societal risk is used to encompass both public and worker risk.

By way of illustration, the maximum tolerable (upper) criterion and the broadly acceptable (lower) criterion in use in the oil, gas and petrochemical industries are now described, firstly for individual risk and then for societal risk.

INDIVIDUAL RISK

Individual risk criteria are intended to show that workers or members of the public are not exposed to excessive risk. They are independent of the number of people exposed and hence may be applied to a broad range of activities.

Individual risk is calculated by identifying all sources of fatality risk to a given individual, deriving the contribution from each source and then summing these to give the overall risk. For typical oil, gas and petrochemical workers the primary sources of risk are:

- Occupational, e.g. slips and falls, drowning
- Transport, e.g. road traffic accidents, air transport accidents
- Hydrocarbon related, e.g. loss of containment leading to toxic releases, fires or explosions

What are the levels of the upper and lower criteria for individual risk?

Individual risk criteria are most commonly expressed in the form of Individual Risk Per Annum (IRPA). Today, the following IRPA values for these criteria are generally regarded internationally as applicable for hazardous industries:

Table 1 – Individual risk criteria

	Workers	Members of Public
Maximum tolerable criterion	10^{-3} per yr	10^{-4} per yr
Broadly acceptable criterion	10^{-6} per yr	10^{-6} per yr

An IPRA of 10^{-3} per year was first used by the UK HSE as the maximum tolerable criterion because it approximated to the risk experienced by high risk groups in mining, quarrying, demolition and deep sea fishing (Ref. 3). As such, it would appear quite lenient for offshore and onshore oil, gas and petrochemical facilities. This is borne out by typical risk levels on offshore installations (which generally have higher risk levels than onshore facilities) shown in Figure 2 (Ref. 4).

Figure 2 – Risks on representative UK offshore installations

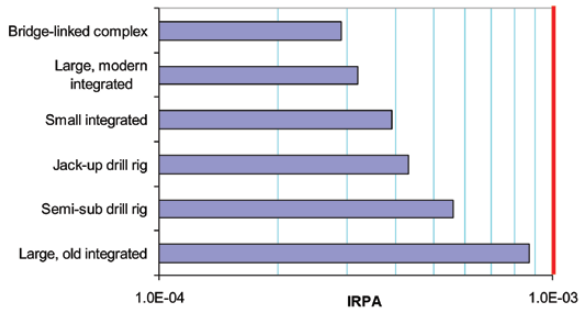

Comparing risks is not a straight-forward task, but nevertheless Table 2 shows how many risky activities an individual would need to undertake in one year to reach an IRPA of 10^{-3} per year (derived from Ref. 2).

Table 2 – Risky activities

Activity	Number of activities in one year that equals an IRPA of 10^{-3} per year
Hang-gliding	116 flights
Surgical anaesthesia	185 operations
Scuba diving	200 dives
Rock climbing	320 climbs

This further illustrates that 10^{-3} per year is actually quite high. However, in practice, few modern facilities with proactive risk reduction strategies have risk levels approaching 10^{-3} per year. This tends to be recognised in company risk tolerability standards, where benchmark design targets are often set for new facilities in the region of $3x10^{-4}$ to $1x10^{-4}$ per year. Risk levels are rarely ever insignificant, i.e. less than 10^{-6} per year and therefore tend to lie in the middle band of the risk tolerability framework, sometimes referred to as the ALARP region.

SUMMARY – INDIVIDUAL RISK CRITERIA

The maximum tolerable IRPA criteria of 10^{-3} per year for site workers and 10^{-4} per year for members of the public are generally accepted internationally. However, these levels are rather lenient for most facilities and companies often set more stringent criteria as much as 10 times lower for new designs.

SOCIETAL RISK

All fatal accidents are a cause for regret, but society generally tends to be more concerned about multiple fatalities in a single event. Whilst such low-frequency high-consequence events might represent a very small risk to an individual, they may be seen as unacceptable when a large number of people are exposed. Such incidents can significantly impact shareholder value and, in some cases, the company never recovers (Ref. 5).

The concept of societal risk is illustrated in Figure 3. Situations A and B have equal individuals risk levels (IR and IR') but B has a larger societal risk (SR) because more people are exposed (Ref. 6). If the individual risk levels are acceptable, when is the societal risk not acceptable?

Criteria may be defined to limit the risk of major accidents and help target societal risk reduction measures such as restrictions on concurrent activities or land use, enhanced engineered safeguards, and improved building siting or protection.

Figure 3 – Societal risk

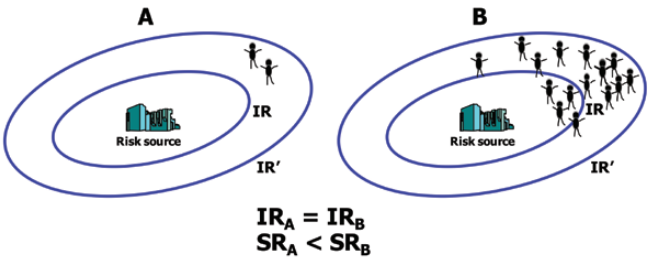

FN-diagram

A common form of presenting risk tolerability criteria for societal risk is on an FN-diagram, where two criteria lines divide the space into three regions – where risk is intolerable, where it is broadly acceptable and where it requires further assessment and risk reduction as far as is reasonably practicable, as shown in Figure 4. This is the same framework for risk tolerability shown in Figure 1 earlier.

FN-criteria are not without their drawbacks but they are undoubtedly helpful when used in context. They clearly show the relationship between frequency and size of accident. A steep criterion slope also builds in multiple fatality aversion and favours design concepts with lower potential for large fatality events. The pros and cons of FN-criteria are summarised in Table 3.

Figure 4 – Illustrative FN-criteria

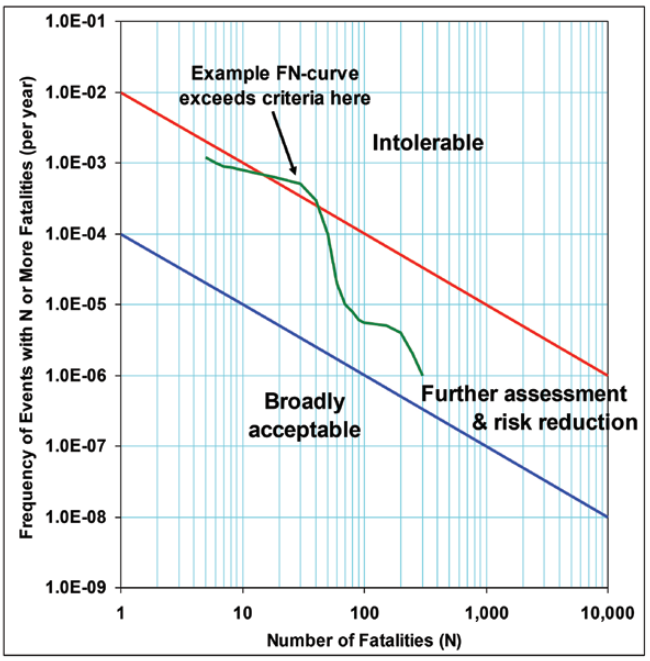

Table 3 – Pros and cons of FN-criteria

Pros	Cons
Clearly shows relationship between frequency and size of accident	Cumulative expression makes it difficult to interpret, especially by non-risk specialists
Allows judgement on relative importance of different sizes of accident	Can be awkward to derive
Slope steeper than -1 provides explicit consideration of multiple fatality aversion and favours concepts with lower potential for large fatality events	May be difficult to use if criterion is exceeded in one area but otherwise is well below
Allows company to manage overall risk exposure from portfolio of all existing and future facilities	Much debate about location of criterion lines

What are the levels of the upper and lower criteria for societal risk?

Unfortunately, unlike individual risk criteria, there are no single 'one-size-fits-all' criteria for societal risk in use by operators and regulators in the major hazard industries world-wide. Indeed, the variation in regulatory criteria is very wide, as shown by the upper tolerability criterion lines in Figure 5, which span a factor of over 100. The Dutch criterion is so restrictive that it raises a question about its merits.

Multiple-fatality accidents can significantly impact shareholder value and, in some cases, the company never recovers

Figure 5 – Regulatory upper tolerability FN-criteria

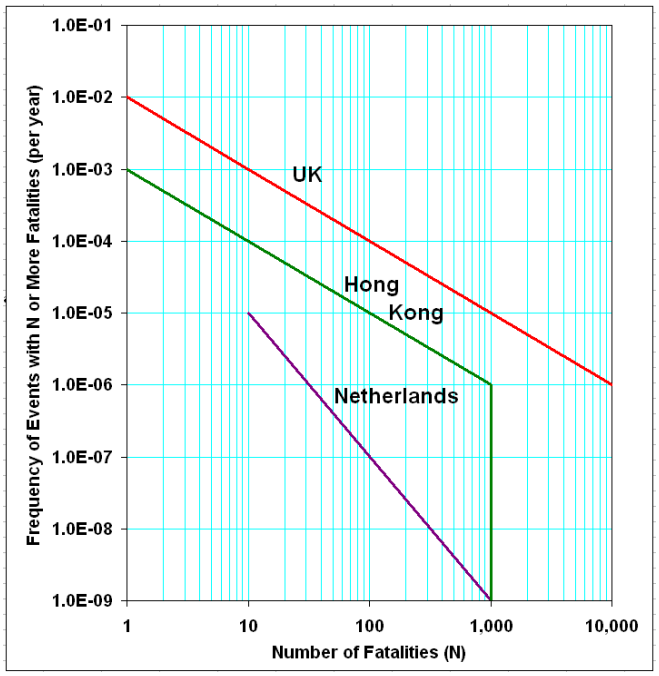

Figure 6 – Operator upper tolerability FN-criteria

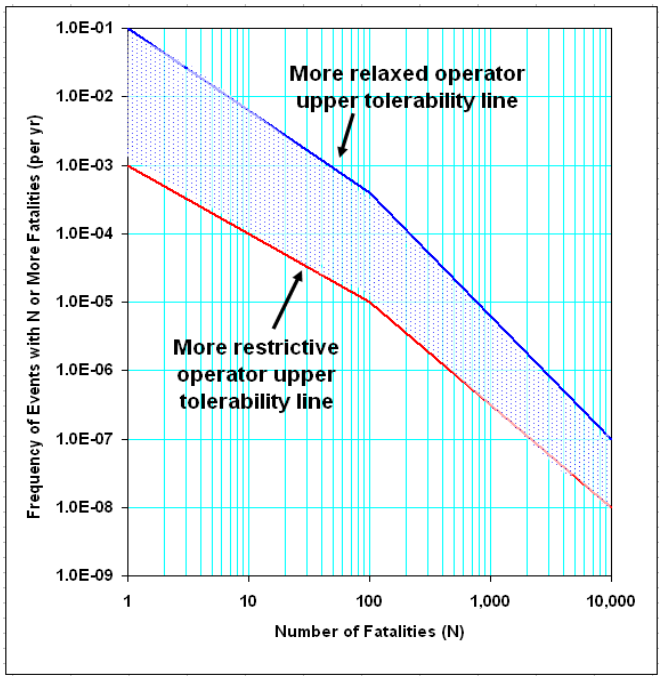

A number of operators, including US corporations, do have their own FN-criteria or guidelines for their facilities but many others, including European corporations, do not. A review of a relatively small sample of international operators' FN-criteria shows a similarly wide variation to that seen in regulatory criteria, Figure 6.

For a company operating in regions where there are no regulatory criteria to meet, the choice of criteria to help decision-making therefore largely comes down to one of company values, i.e. the perceptions of the stakeholders directly affected by the decision and the values of the company in terms of its safety commitment and reputation.

When is too big too often?

Expressed from a dispassionate business perspective, the company must decide how frequently large-fatality accidents would need to occur before the company's survival is put severely at risk due to the adverse reaction of shareholders, the regulator, media and public.

To illustrate this point, assume for example that a company believes that its future survival would be severely threatened if an accident causing 10 or more fatalities occurred more regularly than once every 10 years across all of its facilities, and if an accident causing 100 or more fatalities occurred more regularly than once every 300 years. A straight line can then be drawn between these two points and extrapolated to higher values of N. The slope of the line will determine whether large fatality aversion is included or not; in this example the slope is -1.5, implying a relatively high aversion.

Furthermore, if the company operated 30 facilities, it might decide to allocate its risk evenly between each facility. The resulting company upper criterion is shown in Figure 7 below (red line), together with the single facility criterion line if the company operated 30 facilities (blue line).

Figure 7 – Illustrative company-wide FN-criteria

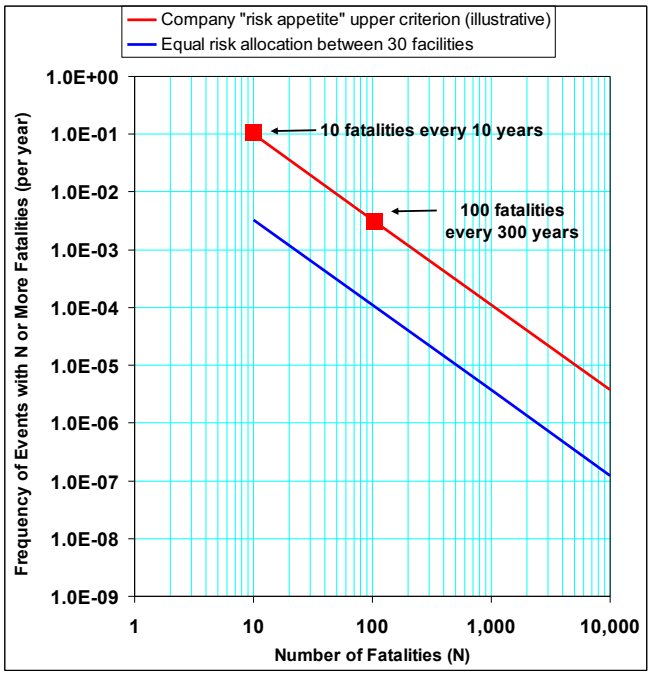

An FN-curve for a single facility that lies on the 'wrong side' of the facility criterion line would use up an excessive proportion of the company risk appetite. In practice, the criterion line may be a reporting line, above which a higher level of corporate scrutiny would be applied. Only then can the decision be made by senior corporate management to proceed with the project or continue existing operations.

This example assumes the company allocates the total risk evenly between each of its 30 facilities. A variation on this is to allocate the total risk in proportion to the size of facility. Whilst the approach has some merit (it recognises larger facilities generate greater economic value), it has one major disadvantage in that it may be misused to keep facilities on the 'right side' of the criteria.

PLL criteria as an alternative to FN-criteria

The other main measure for societal risk is the annual fatality rate, where the frequency and number of fatalities are combined into a Potential Loss of Life (PLL), which is a convenient one-dimensional measure of the total number of expected fatalities.

PLL is well suited for comparing alternative solutions for the same facility, is relatively easy to understand for non-risk specialists and must be calculated to be able to derive the cost-effectiveness of risk reduction options. However, because no information is provided on the relationship between frequency and size of the accident, it is difficult to draw meaningful conclusions from completely different facilities, and it often favours the concept that has the lowest manning level.

As such, there is little benefit to be gained in limiting PLL by explicit criteria. It is extremely rare (unknown) for organisations to have such limits.

Risk contour criteria as an alternative

There are some other 'surrogate' measures of risk which do not explicitly show the relationship between frequency and consequence, but nevertheless provide a proxy for group risk in that useful inferences can be drawn to protect against large fatality events.

Risk contours are amongst the most common, where iso-risk contour plots represent the geographical variation of the risk for a hypothetical individual who is positioned at a particular location for 24 hours per day, 365 days per year. This is also known as Location-Specific Individual Risk (LSIR).

Although there is no consideration of the total number of expected fatalities or of any explicit aversion to low-frequency high-consequence events, an approach of lowering the risk contour criteria with distance away from a facility reflects an attempt to do this. Risk contour criteria tend to be used for land use planning purposes, with the local planning authority left to enforce land use controls. For example, the Major Industrial Accidents Council of

Canada (MIACC) recommends individual risk levels for use in respect to hazardous substances including the risk contributions from all sources, with the inner zone criteria of LSIR from 10^{-4} to 10^{-5}, middle zone 10^{-5} to 10^{-6} and outer zone beyond 10^{-6} per year. Restrictions are placed on activities or structures within the various zones, as shown in Figure 8.

Figure 8 – Canadian risk contour criteria

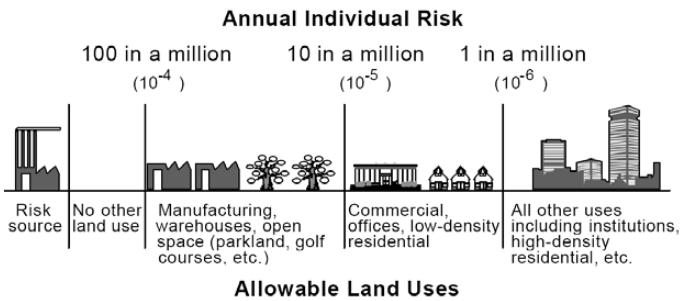

The guidelines are considered to be realistic in terms of existing practices of risk management and levels of risk. They are also compatible with criteria that have been selected and implemented in other industries and other countries.

FN, PLL or risk contour criteria?

In the absence of regulatory FN-criteria, some international operators have set their own FN-criteria but other operators believe there are simply too many issues associated with defining the upper and lower criterion lines. The preferred, quantitative, way for such operators of comparing risk reduction options in design and layout is through determining the change in PLL and the change in risk contour profiles.

SUMMARY – SOCIETAL RISK CRITERIA

In the absence of regulatory criteria, the choice of societal risk criteria largely comes down to the company's values. But whatever criteria are selected, they need to be workable in practice – if they are too severe or too relaxed they will lose their usefulness – and should be based on a sound assessment of current good practice in industry.

ALARP

Whether considering individual risk or societal risk, safety risks need to be reduced to levels that are As Low As Reasonably Practicable, or 'ALARP'. The region that lies between unacceptably high and negligible risk levels is sometimes referred to as the 'ALARP region'.

If risk is in the ALARP region is it ALARP?

No, this is a common misconception. Even if a level of risk for a 'baseline case' has been judged to be in this ALARP region it is still necessary to consider introducing further risk reduction measures to drive the remaining, or 'residual', risk downwards.

The ALARP level is reached when the time, trouble and cost of further reduction measures become unreasonably disproportionate to the additional risk reduction obtained.

When does the ALARP principle apply?

Risk can be reduced by avoidance, adopting an alternative approach, or increasing the number and effectiveness of controls.

At the concept stage of a new project there is the greatest opportunity to achieve the lowest residual risk by considering alternative options, e.g. for an offshore oilfield development, options may range from fixed legged platforms to floating production vessels to subsea facilities.

Once the concept is selected and the early design progresses, the attention shifts to considering alternative layout and system options to optimise inherent safety. In the detailed design phase, the focus moves on to examining alternative options for improving safety systems.

During operations, the attention is on collecting feedback, improving procedures and personnel competence, and managing change to maintain the residual risk at an ALARP level. However, with advances in technology, what is ALARP today may not be ALARP tomorrow, so periodic reviews will be necessary.

How is ALARP demonstrated?

The definition of ALARP implies there is a mathematical formula to wield at the problem, and it is true that there is one.

Having selected a range of possible risk reduction options, a QRA can be re-run for each option to identify the associated reduction in risk. Combining this improvement with the total cost of each option enables the options to be ranked in order of cost-effectiveness, using a Cost-Benefit Analysis (CBA). The Implied Cost of Averting a Fatality (ICAF) is expressed in terms of $ per statistical fatality averted and comprises the following generally annualised elements:

$$ICAF = \frac{\textit{Net cost of option}}{\textit{Potential saving of life}}$$

where *Net cost of option* =
Cost of option – Reduction in loss of assets & production

This calculation takes account of the fact that measures to reduce risk to people are also likely to reduce the potential damage to assets and loss of production.

The derived ICAF values for the proposed options may then be ranked and compared against company standards for ICAF. The typical ICAF value used by the UK offshore industry is around £6,000,000, i.e. in simplistic terms a measure that costs less than £6,000,000 and saves a life over the lifetime of an installation is reasonably practicable, whilst one that costs significantly more than £6,000,000 is grossly disproportionate and therefore is not justified. The UK HSE considers this to be the minimum level for the application of CBA in the offshore industry (Ref. 7).

In reality there is no simple cut-off and often a band of ICAF values is applied, as illustrated in Table 4.

Table 4 – ICAF guidelines

ICAF (US$)	Guidance
<$10,000	**Highly effective** Always implement
$10,000 - $100,000	**Effective** Always implement
$100,000 - $1,000,000	**Effective** Implement unless risk is negligible
$1,000,000 - $10,000,000	**Consider** Effective if individual risk levels are high
$10,000,000 - $100,000,000	**Consider** At high risk levels or when there are other benefits
>$100,000,000	**Ineffective** Cost grossly disproportionate

Discussion on this subject can be emotive and care must be taken to provide an explanation as to why it is necessary to venture into this seemingly sensitive area of option evaluation. However, experience is that derivation of ICAF achieves not only a ranking of improvement options but also provides a spur to the creative development of yet safer and more economic options.

So is ALARP demonstrated by QRA and CBA?

This is another common misconception. QRA and CBA are inexact and a high variability in results is often seen. This variability can arise from poor standards in performing the study, e.g. omitting hazards or making calculation errors, as well as genuine uncertainty in data and modelling methods. The use of numerical estimates of risk, by themselves, can be misleading and can result in decisions that either do not meet adequate levels of safety, or overestimate the real risks.

The 'formula approach' therefore should be used very cautiously and only in support of qualitative or engineering arguments. In general an approach that uses information from engineering and operational analysis, supplemented where appropriate by QRA, will lead to more robust decisions. The steps to follow are shown in Figure 9.

Figure 9 – ALARP process

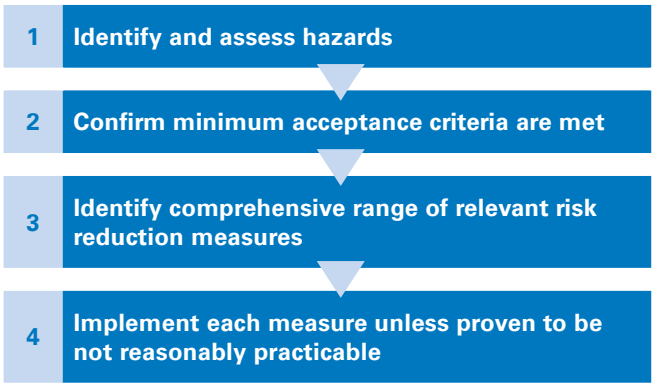

The critical step of this process is step 3, the need to identify a comprehensive range of relevant risk reduction measures. They should be based on modern good practice and be targeted at the largest risk contributors. This is typically achieved through 'brainstorming' workshops to identify technically feasible improvements that may:

- Eliminate the hazard
- Reduce the exposure of personnel to the hazard
- Reduce the frequency of occurrence
- Mitigate the consequences if the event does occur
- Improve evacuation if control is lost

Risk levels are only ALARP once every measure identified during step 3 has either been implemented or proven to be not reasonably practicable. It is surprising how many people need reminding that risk levels will remain the same, or even increase, until real improvements are fully

implemented. A formal risk assessment can generate a large number of recommendations and they need to be properly managed.

What tools are available to help demonstrate ALARP?

The tools available for demonstrating risks are reduced to levels that are ALARP are illustrated in Figure 10 (Ref. 8):

Figure 10 – ALARP tools

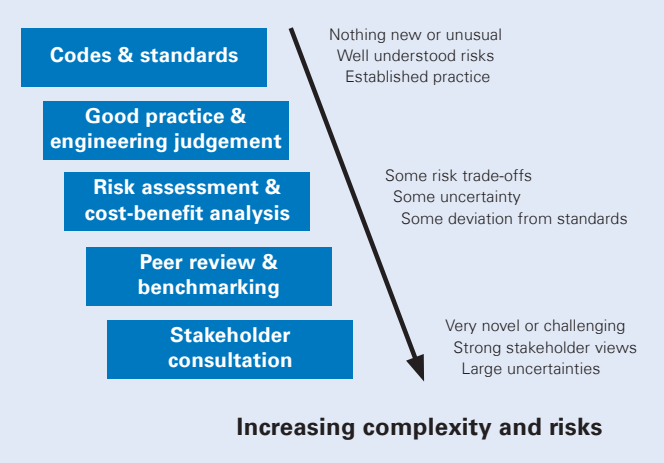

In general, the more complex the project, the more complex the decisions and the more sophisticated the tools required. Also, the higher the risk, the more comprehensive and robust the ALARP assessment needs to be.

For example, in many common engineering situations, what is reasonably practicable may be determined simply by reference to the relevant code or current practice. The majority of decision making will usually fall into this category. The codes and standards capture the lessons from past experience and try to reflect best use of current technology and understanding.

In other cases, optioneering together with the use of risk assessment and cost-benefit analysis may be appropriate. A risk based approach can go some way towards addressing situations where, for example, there is high complexity, high costs, conflicting risks and uncertainty. It can provide a clearer picture of the decision implications and the pros and cons of the various decision options.

There may also be the need to take into account the views and concerns of those stakeholders affected by the decision. Their perception of the risks and benefits may be different from that analysed, affecting what they believe to be reasonably practicable as a solution. What one organisation may deem as the appropriate solution to manage the risks may be different from another organisation and in excess of that required by regulation.

SUMMARY – ALARP ASSESSMENT

In practice ALARP decision making amounts to taking a balanced view and reaching a defensible consensus on prioritised improvements. A convincing ALARP demonstration lies in the documented consideration of improvement options, both implemented and discounted, at a level of detail appropriate to the facility lifecycle and magnitude of risk.

CONCLUSION

There is a high degree of commonality in individual risk criteria internationally, but societal risk criteria show a large variation. QRA is inexact and any quantitative criteria should be seen as guidelines. The risks associated with most facilities lie in the middle band of the risk tolerability framework – the 'ALARP region' – and require qualitative and sometimes quantitative demonstration of risk reduction to ALARP levels. In practice, this amounts to taking a balanced view and reaching a defensible consensus amongst stakeholders.

References:

1. **CMPT 1999.** A Guide to Quantitative Assessment for Offshore Installations, Publication 99/100a, The Centre for Marine and Petroleum Technology (CMPT), 1999.
2. **HSE 2001.** Reducing Risks, Protecting People – HSE's Decision-Making Process, 2001.
3. **HSE 1992.** The Tolerability of Risk from Nuclear Power Stations, 1992.
4. **HSE 1995.** Update of the UKCS Risk Overview, OTH 94-458, 1995.
5. **Sedgwick 2002.** The Impact of Catastrophes on Shareholder Value, A Research Report Sponsored by Sedgwick Group, Rory F. Knight & Deborah J. Pretty, The Oxford Executive Research Briefings, 2002.
6. **Jonkman 2003.** An Overview of Quantitative Risk Measures for Loss of Life and Economic Damage, Journal of Hazardous Materials A99 (2003) 1–30, Jonkman SN, van Gelder PHAJM, Vrijling JK, 2003.
7. **HSE 2006.** Offshore Installations (Safety Case) regulations 2005 Regulation 12 – Demonstrating Compliance with the Relevant Statutory Provisions, HSE Offshore Information Sheet No. 2/2006.
8. **UKOOA 1999.** Industry Guidelines on a Framework for Risk Related Decision Support, UKOOA, Issue 1, May 1999.

About Risktec

Risktec is an independent and specialist risk management consulting and training company. We help clients to manage health, safety, security, environmental (HSSE) and business risk in sectors where the impact of loss is high.

OUR SERVICES ENCOMPASS:

Consulting

Specialist risk management services, delivering packaged and proportionate solutions to help reduce and manage risk.

Learning

Online and classroom training and postgraduate education to help develop competent risk management professionals.

Resourcing

Specialist risk, HSSE and engineering associates to work at client locations to help fill resource and skills shortages.

Inspection

Industrial and vendor inspections and assessments to ensure asset integrity and mitigate project risks.

Consulting

Our experience ranges from delivering small self-contained work packages to managing complex multi-disciplinary projects with a large number of stakeholders.

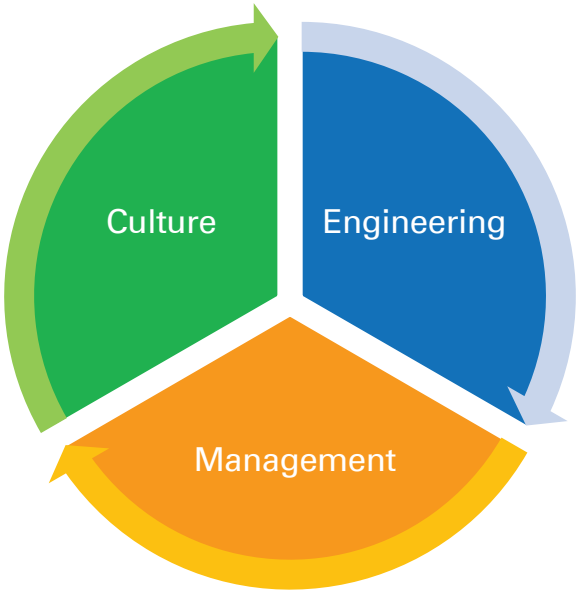

Our services recognise that controlling risk requires understanding engineered and technological systems, management systems and organisational, cultural and behavioural factors.

ENGINEERING
Identifying, analysing, evaluating and reducing the risks associated with facilities, operations and equipment to acceptable levels.

MANAGEMENT
Identifying, developing and implementing effective policies and procedures to maintain control of risks and minimise loss.

CULTURE
Accelerating cultural and behavioural improvement, and ensuring a solid foundation for building sustainable improvements in risk control.

Learning

We provide a unique training and education service, from a single training course to a Risktec professional qualification or a tailored master's programme in Risk and Safety Management, all developed and taught by our experienced consultants. Our courses encompass the breadth and depth of our consulting services.

- Postgraduate Certificate, Diploma or Master's Degree (MSc) in Risk and Safety Management
- Degree Apprenticeship in Risk and Safety Management
- Risktec Professional Qualification (RPQ) in Risk and Safety Management
- Training courses from single modules to multi-year programmes for corporate clients
- Game-based learning
- Computer-based training
- Delivery via face-to-face, distance or blended learning
- Accredited by professional engineering institutions and industry bodies
- Our whole approach is flexible to meet client needs

Resourcing

We provide resource to support our clients' activities by working at their main offices, project locations or industrial sites, anywhere in the world. The support is delivered by our professional resourcing business, ASTEC, which has access to a huge pool of professional associates.

We provide associates who:

- Are well known to us.
- Are suitably qualified and bring the required specific skills and experience.
- Have many years' experience and hence can make an immediate and positive impact on projects.
- Can be supported by work packages from consultants in our own offices.

Inspection

We provide a risk-based programme to focus inspections where they are most needed, to mitigate project, safety, environment, production and regulatory risks.

- Inspection strategies and workscopes
- Site inspections including non-destructive testing
- Integrity assessment
- Weld repair solutions
- Component life extension
- Failure assessment
- Third party equipment inspection
- Quality assurance / quality control
- Vendor capability assessments

TÜV Rheinland

As part of the TÜV Rheinland Group we have access to a very large range of services via the group's 20,000 employees in over 65 countries worldwide, including:

Testing, inspection and certification services to ensure the safety, reliability and regulatory compliance of assets and components throughout their lifecycle; as well as technical consulting and training to energy, industrial, transportation, products and healthcare sectors.

Images © Shutterstock

TÜVRheinland®
Risktec

risktec.tuv.com
enquiries@risktec.tuv.com