



Resilience Engineering and Integrated Operations in the Petroleum Industry

Tempora mutantur, nos et mutamur in illis

"Times change, and we change with them"



Center for Integrated Operations in the Petroleum Industry



Design and print: Tapir Uttrykk, Trondheim, August 2010
Authors: Erik Hollnagel, Camilla K Tveiten, Eirik Albrechtsen
IO-center (SINTEF) report no: SINTEF A16331
ISBN 978-82-14-04901-5
Photo page 3, 7, 9 and 10: Statoil

Introduction

The petroleum industry is currently in a phase of development and implementation of new technology for oil and gas exploration and production. Advanced information technology and digital infrastructure enable new and more effective ways of working, but may also destabilize established work processes. The changes include using collaborative technology to connect distributed actors, using real-time data to monitor and manage operations across geographical and organisational borders, giving access to expert knowledge regardless of location, and a tighter integration of technology, data, disciplines, activities and organizations. Together this is known as Integrated Operations (IO).


The deployment of IO significantly changes established ways of working; this creates new possibilities as well as new risks. The industry must be prepared for both in order to maintain or improve current levels of safety and efficiency. The use of IO therefore requires an approach to safety management that can cope with the new challenges as well as the opportunities. Resilience Engineering has been developed to cope with increasingly complex socio-technical systems that often pose a challenge to established safety approaches. It provides a way to address the issues of emergent accidents and the often disproportionate consequences that may be the result of ever more complex technologies and integrated actors. Resilience Engineering methods are therefore going to be important tools in the management and assurance of safety in current and future IO systems.

This white paper provides an overview of Resilience Engineering and seeks to answer the following questions:

- What is Resilience Engineering?
- Why do we need Resilience Engineering in integrated operations?
- What is the significance of performance variability?
- How does Resilience Engineering work in practice?
- How does Resilience Engineering fit with other safety management approaches?
- How mature is Resilience Engineering?
- What is the added value of Resilience Engineering for integrated operations in petroleum production?



What is Resilience Engineering?



Resilience engineering was developed to address the new safety concerns that arise from the use of increasingly complex socio-technical systems, where performance depends on tightly coupled social and technical functions. Resilience Engineering, by its nature, has strong links with Human Factors, Control Theory, and Safety Engineering, and is based on the following premises:

1. Performance conditions are always underspecified because the work environment is too complex to describe in every detail and because it is never stable and fixed. Individuals and organisations must therefore adjust what they do to match current demands and resources. Because resources and time are finite, such adjustments will inevitably be approximate.
2. Many adverse events can be attributed to a breakdown or malfunctioning of components and normal system functions, but many cannot. The latter are better described as the result of unexpected combinations of performance adjustments or performance variability.
3. Safety management cannot be based exclusively on hindsight, or rely on error tabulation and the calculation of failure probabilities. Future occurrences may be due to a combination of performance variability that habitually is seen as irrelevant for safety. Safety management must therefore be proactive as well as reactive.
4. Safety cannot be isolated from the core (business) process, nor vice versa. Safety is the prerequisite for productivity, and productivity is the prerequisite for safety. Safety must therefore be achieved by improving the core processes rather than by constraining them.

Adopting this view creates a need for an approach that can represent the variability of normal system performance, and for methods that can use this to provide more compre-

hensive explanations of accidents as well as identify potential risks and opportunities. Resilience can be defined as the *intrinsic ability of a system to adjust its functioning prior to, during, or following changes and disturbances, so that it can sustain required operations under both expected and unexpected conditions*. The essential characteristic of a resilient system is the ability to adjust its functioning so that it can succeed in different – and difficult – situations. This implies four main aspects or dimensions, each representing an essential system ability:

- Knowing what to do, meaning how *to respond* to regular and irregular disruptions and disturbances either by implementing a prepared set of responses or by adjusting normal functioning. This is the ability to address the actual.
- Knowing what to look for, meaning how *to monitor* that which is or can become a threat or opportunity in the near term. The monitoring must cover both that which happens in the environment and that which happens in the system itself, i.e., its own performance. This is the ability to address the critical.
- Knowing what to expect, meaning how *to anticipate* developments, threats, and opportunities further into the future, such as potential changes, disruptions, pressures, and their consequences. This is the ability to address the potential.
- Knowing what has happened, meaning how *to learn* from experience, in particular how to learn the right lessons from the right experience – successes as well as failures. This is the ability to address the factual.

The four abilities provide a basis both for engineering and managing resilience. One result of that is the Resilience Analysis Grid (RAG), described later in this White Paper.

From safety management to Resilience Engineering

Risk governance and safety management have traditionally focused mainly on what can go wrong – and with good reason. Safety is therefore commonly defined and measured by the relative occurrence of unwanted outcomes. The set of possible outcomes can schematically be shown as in Figure 1, where the x-axis describes likelihood of occurrence, ranging from very low to very high, and the y-axis describes the value of the outcome, ranging from negative to positive. (The lower half is a simplified version of the frequently used risk matrix.)

The established approaches to risk and safety mainly focus on the things that go wrong, more specifically those areas in Figure 1 named disasters, accidents, and incidents – with occasional forays into ‘near misses.’ (The ‘mishaps’ region describes unwanted outcomes that in practice have been eliminated in well-functioning processes and organisations.) But there has traditionally been little or no focus on things that go right, despite the fact that these happen

far more often than things that go wrong. If, for instance, the probability of failure is 10^{-4} , then there will be 9,999 normal outcomes for every failure!

In contrast to safety management, resilience engineering covers the whole set of outcomes, i.e., things that go right as well as things that go wrong – with the possible exceptions of the areas of ‘serendipity’ and ‘good luck,’ where we still know too little to deal with them systematically. Safety is correspondingly defined as the ability to succeed under varying conditions. The aim of Resilience Engineering is not only to prevent things from going wrong, but also to ensure that things go right, i.e. to facilitate normal outcomes. Simply put, the more likely it is that something goes right, the less likely it is that it goes wrong. The added value of focusing on how to facilitate normal outcomes is that it reduces the traditional separation between safety and productivity.

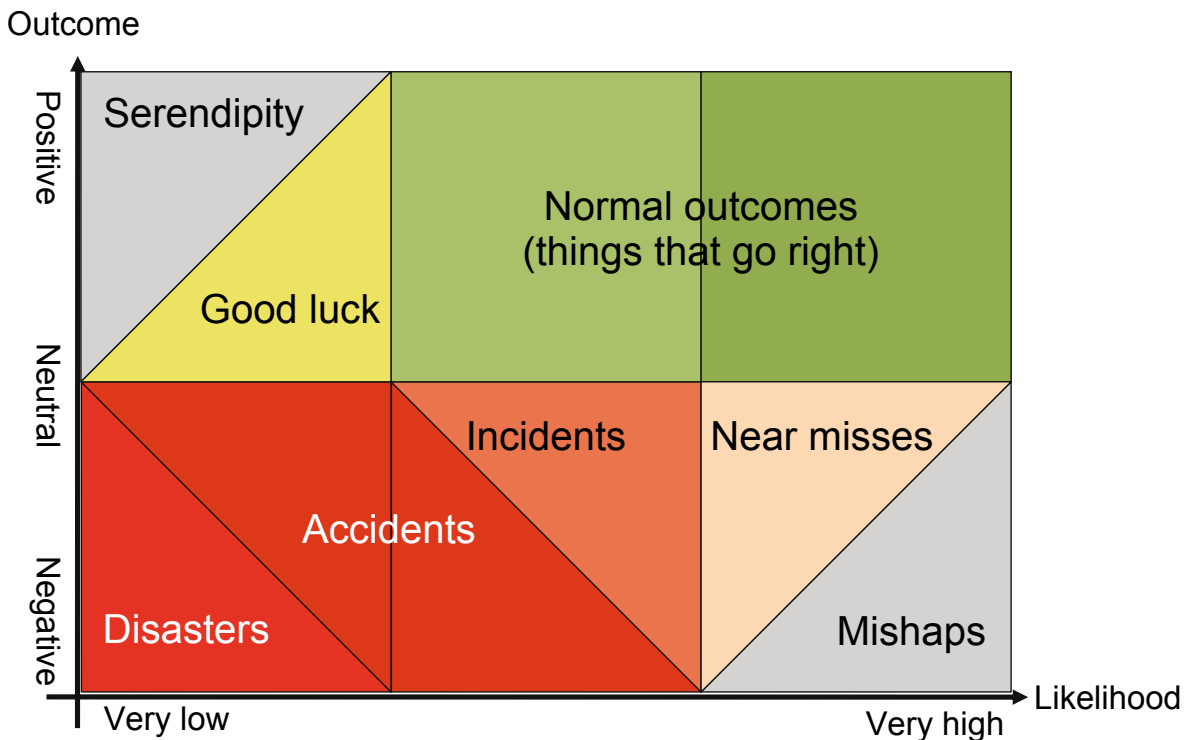


Figure 1 From safety management to resilience engineering.

Why do we need Resilience Engineering in integrated operations?

A simple answer to this question can be found by looking at the types of accidents that can occur in complex yet 'well-defended' systems, of which Integrated Operations is a good example. While incidents such as the Snorre A gas leak with hindsight can be explained by failure modes emerging from the weak interaction of multiple factors, few would have considered this situation credible ahead of time. Traditional thinking makes it difficult to describe and understand how multiple factors can come together in time and space to produce something as disastrous as a blow-out.

Today, most high risk systems such as petroleum production, have an abundance of safety 'nets,' barrier systems, safety management systems, safety assessment and assurance processes, and many are also improving their safety culture. These efforts add layers of safety to already safe systems, but also make them more complex, hence more difficult to understand and control. And while the multiple safety functions will further reduce the likelihood

of accidents, it also means that those that do slip through these 'nets' will be complex and multi-faceted. Future accidents may therefore be more due to coincidences among the variability of functions and human performance in different parts of the system, than to manifest failures and incorrect human actions.

Within the petroleum industry, as well as within other industrial domains, accident analysis and risk assessment methods are needed to deal with the problems coming from major accidents. Historically, methods have been developed in response to major technological developments or to cope with 'new' types of accidents. Figure 2 shows the distribution of some well-known methods used to address technical, human factors, and organisational issues, respectively. It is noteworthy that human factors methods came onto the scene after the accident at Three Miles Island in 1979, and that organisational methods were developed following the Chernobyl and Challenger accidents in 1986.

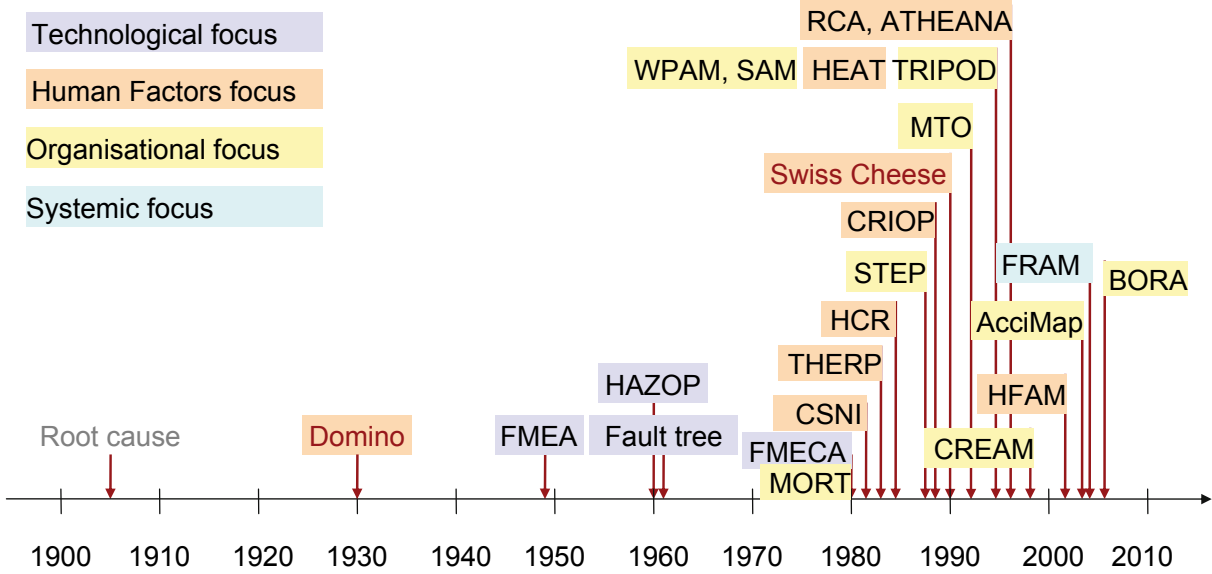


Figure 2 Methods used to address technical, human factors, and organisational issues in safety management

Methods must clearly be powerful enough to match the problems found in real-life applications. Since the complexity of industrial systems continues to increase, it is inevitable that established approaches to risk and safety at some time become unable to explain, predict, and prevent new types of accidents. Safety assessment methods have historically developed from technical methods, via human factors methods, to organisational methods. As part of that, the thinking about safety has developed from simple, linear models such as the Domino model, to complex linear models, such as the Swiss Cheese model. But in today's industrial systems the naturally occurring performance variability may often combine to produce effects that go beyond what current models can describe. In order to address these more complex phenomena, Resilience Engineering uses the principle of functional resonance to explain how the variability of dynamic performance adjustments can combine in ways that may lead to disproportionately large (non-linear) effects. Since resonance is something that happens on the level of the system as a whole, it is insufficient simply to combine or aggregate technical, human, and organisational factors.



What is the significance of performance variability?

The first premise of Resilience Engineering emphasises that performance variability is both inevitable and useful. Procedures and instructions are always incomplete, except for extremely simple situations. Following procedures and instructions to the letter will therefore either be inefficient or unsafe, or both. To compensate for this incompleteness, people (individually and collectively) and organisations habitually adjust what they do to match current demands, resources, and constraints. The ability to do so is at the heart of successful performance. But since information, resources, and time are finite, the adjustments will inevitably be approximate. The same performance variability that is the reason why things usually go right may therefore also be the reason why things sometimes go wrong.

The essence of a socio-technical system is that the conditions for successful performance – and conversely also for unsuccessful performance – depends on the interaction between social and technical factors. Socio-technical systems have since the 1980s become steadily more complex due to rampant technological and societal developments. As a result of these developments, safety methods must today address systems that are larger and more complex than the systems of yesteryear. Because there are many more details to consider; because some modes of operation may be incompletely known; because of tight cou-

plings among functions; and because systems may change faster than they can be described, the net result is that many systems, petroleum production operations included, are underspecified and therefore intractable. For these systems it is clearly not possible to prescribe tasks and actions in every detail, and performance must therefore be variable or flexible rather than rigid.


The distinction between tractable and intractable systems is useful to understand the challenges facing safety management systems today. Tractable systems can be completely described or specified, while intractable systems cannot. The differences between the two types of systems are summarised in Table 1.

Most established safety methods have been developed on the assumption that systems are tractable. As this assumption is no longer generally valid, there is a need to develop methods to deal with intractable systems. Resilience Engineering is one answer to this problem.

Table 2 shows the result of characterising different socio-technical systems on the dimensions of 'Coupling' and 'Tractability'. Existing methods are not well suited for systems in the upper right quadrant (intractable and tightly coupled), which makes this a primary focus for Resilience Engineering.

Table 1 Tractable and intractable systems

	Tractable system	Intractable system
Number of details	Descriptions are simple with few details	Descriptions are elaborate with many details
Comprehensibility	Principles of functioning are known	Principles of functioning are partly unknown
Stability	System does not change while being described	System changes before description is completed
Relation to other systems	Independence	Interdependence
Controllability	High, easy to control	Low, difficult to control



The picture has been modified. Interaction between humans and technology presupposes performance variability. Sometimes unintended interaction between normal variability produces unwanted outcomes.

Table 2 Coupling and tractability

		Tractability (cf. Table 1)	
		High – system is difficult to describe and understand	Low – system is easy to describe and understand
Coupling	Tight – effect spread quickly, limited slack or substitutability of components or functions	Industries: Financial markets, IO, Nuclear power plants, air traffic management Methods: FRAM	Industries: Marine, power grids, railways Methods: Tripod, MTO
	Loose – effects spread slowly, buffers and flexible execution	Industries: Healthcare, public services, universities Methods: Practically none (!)	Industries: Manufacturing, mining, assembly lines Methods: FTA, HAZOP, FMECA, HERA

The reasons for performance variability


To predict how resonance may lead to accidents or new innovative ways of operations, we must be able to describe and model the characteristic variability of the system. A main source of performance variability is the underspecification of work, as described in the previous section. In addition to the underspecification, human performance can vary for several other reasons:

- Physiological and/or fundamental psychological factors (e.g., affecting perception and vigilance).
- Higher level psychological factors such as ingenuity, creativity, and adaptability.
- Organisational factors, as in meeting performance demands, stretching resources, substituting goals, etc.
- Social factors, as in meeting expectations of oneself, of colleagues, or of managers, complying with informal work standards, etc.
- Contextual factors, for instance if the workplace is unstable, too hot or cold, too noisy, too humid, etc. Likewise, if equipment is unreliable, if resources are unavailable or unpredictable, etc.
- Other factors such as the unpredictability of the domain, e.g., weather conditions, external disruptions, undocumented changes to the system, organisational hysteresis, technical glitches, etc.

The challenge for Resilience Engineering is to represent the variability of a system, such as integrated operation of petroleum production, in a way that makes it possible to identify what may affect performance either adversely or positively. This must overcome two obstacles:

- Because integrated operations of petroleum production is a complex and intractable system, its functions, their interactions, and potential variability, can only be specified approximately.
- In a time of increased demands of efficiency, planned changes to the fundamental infrastructure of operations and technical systems must be reconciled with strong but varying financial pressures and demands.

The operation environment feels the effect of the resulting instability and increasing variability as companies try to absorb the many changes whilst remaining safe and profitable. This makes it necessary for them to be flexible, to rely on human ingenuity and skill, and to manage performance variability rather than to constrain it. In other words, they must be resilient.



Resilience Engineering Methods

Resilience Engineering provides the conceptual basis for a new perspective on safety as well as for new methods. This section will introduce two such methods, the Functional Resonance Analysis Method (FRAM) and the Resilience Analysis Grid (RAG).

The Functional Resonance Analysis Method (FRAM)

FRAM has been developed to provide a practical and effective approach to describe and analyse the role of performance variability in socio-technical systems. As a method, FRAM has been used both for accident investigation and safety assessment, i.e. to understand what has gone wrong and to understand what may go wrong. The method is based on the following four principles:

- The equivalence of success and failures. Failures do not stand for a breakdown or malfunctioning of normal system functions, but rather represent the adaptations necessary to cope with the underspecification found in complex real-world systems.
- The principle of approximate adjustments. To get anything done people must adjust their performance to the current conditions. Because resources and time are finite, such adjustments will inevitably be approximate.
- The principle of emergence. Both failures and normal performance are emergent phenomena: neither can be attributed to or explained simply by referring to the (mal-) functions of specific components or parts.
- The principle of functional resonance. FRAM replaces the traditional cause-effect relation with resonance. This explains how the variability of a number of functions every now and then may resonate, i.e., reinforce each other, leading to excessive variability in one or more downstream functions. The consequences may spread through the system by means of tight couplings rather than separately identifiable cause-effect links.

The following shows what a FRAM analysis might look like for an integrated planning phase of a modification project scenario. To begin with, a FRAM analysis consists of five steps:

1. Define the purpose of the analysis, whether it is accident investigation (looking at past events) or safety assessment (looking at future events).
2. Identify and describe the functions that are necessary (and sufficient) for the intended (correct) performance to be produced (when 'things go right'). The functions can be assigned to either the set of foreground functions or the set of background functions. Characterise each function using the six basic aspects (Input, Output, Pre-conditions, Resources, Time, and Control, as depicted in Figure 3). Taken together, the functions must be sufficient to describe what should have happened, or should happen, (i.e. the normal or successful performance of a task or an activity).
3. Assess and evaluate the potential variability of each function. FRAM uses a distinction between foreground and background functions, which may all affect performance variability. Foreground functions are directly associated with the activity being modelled and may vary significantly during a scenario, while background functions refer to common conditions that may vary more slowly. Both sets of functions should be calibrated as far as possible using information extracted from accident databases.
4. Identify where functional resonance may emerge. This step finds the possible ways in which the variability of a function can spread through the system. In case of functional resonance, the combinations of this variability may lead to situations where the system loses its capability to safely manage variability.
5. The fifth and last step is the development of effective countermeasures. In addition to proposing barriers and defences, Resilience Engineering highlights the need to monitor and manage the system's performance. For monitoring, valid performance indicators can be derived from the FRAM description. Managing performance variability usually means dampening performance variability in order to prevent a loss of control, but can also be to sustain or amplify performance variability that leads to improved outcomes.

An example of what this may look like, applied to an integrated planning phase scenario, is in the following. The analysis of system functions (Step 2) produced the following list:

- Identify problem or opportunity
- Screen for solution
- Define solution/concept
- Estimate cost
- Execute (for pre engineering)
- Pre – engineer the chosen solution

By characterising each function using the six aspects described in Step 2 above (cf. Figure 3), the following instantiation of the model was produced.

The Resilience Analysis Grid (RAG)

The RAG has been developed as an easy-to-use way of obtaining the resilience profile of an organisation. The purpose of the RAG is to produce a relative rather than an absolute rating or evaluation of the resilience of an organisation. The RAG is designed to be used repeatedly, and the importance is the changes in the resilience profile between consecutive ratings, rather than the absolute ratings. This can, for instance, be used to show the effects of changes made, hence support resilience management.

The basis of the RAG are the four abilities of a resilient organisation mentioned above, namely the ability to respond, to monitor, to anticipate, and to learn. Starting with

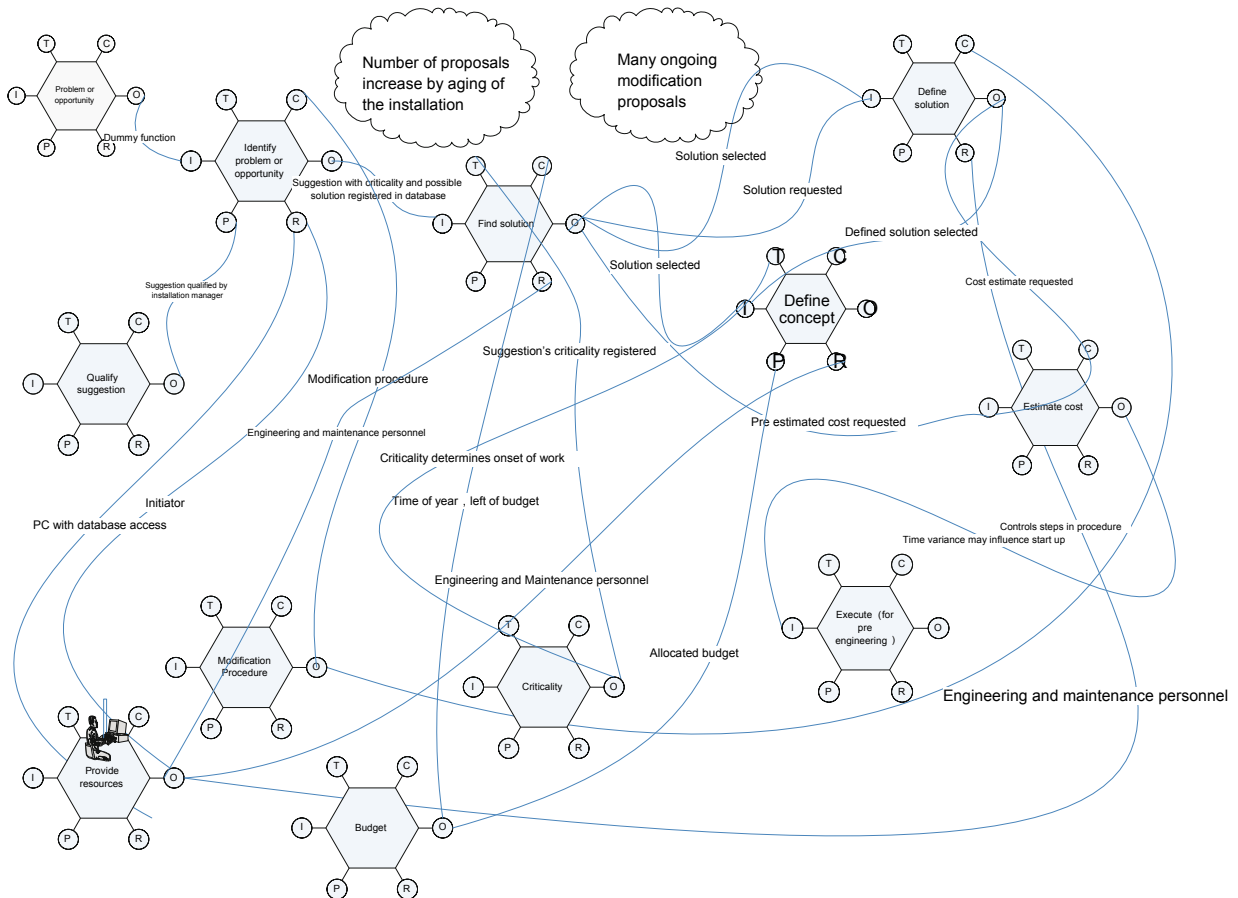


Figure 3 Instantiation of a FRAM model of a planning phase of a modification project.

these abilities it is possible to develop sets of more specific questions (cf. Table 3). The questions should be tailored to match the characteristics of the organisation, and of the particular aspect of functioning that is in focus. (The specific questions can, of course, also be used as a basis for thinking about how to improve the resilience of the organisation.) These questions can then be answered by the people involved in the work, from operators to managers.

To illustrate that, Table 3 shows an example of a set of questions developed to look at the integrated planning process in an oil company¹.

In this example, the questions are answered simply by ticking the appropriate box. This makes it possible to get a quick rating of the current state of resilience (in this case vis-a-vis integrated planning). The ease of getting the ratings is important for the use of the RAG as a basis for resilience management. The result of the rating can be shown

in different ways, for instance as the star diagram below (Figure 4; the values are randomly assigned for the purpose of illustration). This rendering gives a quick overview of how the various abilities were rated, and is particularly useful for comparing repeated ratings, as when following a development (change in policy, use of new tools, change in staffing, etc.)

How does Resilience Engineering fit with other safety management approaches?

Resilience Engineering is an alternative to established safety approaches. It provides a new perspective on safety assessment and management and offers practical methods such as FRAM and RAG to complement existing tools. Adopting a Resilience Engineering view does not require that existing practices are discarded wholesale. But it does mean that they are looked at in a different way, which in turn may change how they are applied, as well as the way in which their results are interpreted.

How mature is Resilience Engineering?

At present, the use of Resilience Engineering in the petroleum production industry is at the feasibility stage of development. However, it seems evident that Resilience Engineering is a suitable approach for managing safety in the complex socio-technical environment of Integrated Operations. Within the IO Center, Resilience Engineering has been the theoretical background for many of the projects within Program 4. Further test cases will be conducted to develop and mature the approach to make it an additional tool for safety cases.

In other domains, such as air traffic management, general aviation, and healthcare, Resilience Engineering has been used longer and has provided a substantial body of knowledge and experience. This has been documented in a number of books and marked by several international symposia and workshops, as well as some commercial applications, e.g. maintenance of heavy machinery, risk modelling of helicopter safety, and improving patient safety.



Figure 4 Resilience rating system

1 Apneseth, K (2010) "Resilience in integrated planning" Master thesis at the Norwegian University of Science and Technology

Table 3 Set of questions to evaluate abilities of a resilient integrated planning process ¹

	Excellent	Satisfactory	Acceptable	Unacceptable	Deficient	Missing
Respond						
1. The integrated plan is continuously updated to reflect the varying needs of the installation.						
2. Active short-term plans are rescheduled when a certain threshold for risk on the activities are reached.						
3. If there are problems in execution of activities, the activities can be reprioritized and/or replaced.						
4. Our planners are experienced and understand the problems that may occur in the execution of activities.						
Anticipate						
5. External factors are taken into consideration in medium and long-term planning, even if that may mean that planned activity will not be executed as intended.						
6. For the planning we have developed performance indicators that give us direct/indirect information about future changes in risk level at our installations.						
7. The ICT-tools used visualize our future activity and activity conflicts in a satisfactory way.						
8. Most of the future activity conflicts and problems in execution are anticipated by the planning tools						
Monitor						
9. The ICT-tools used in the integrated planning, makes it easy to identify different data sources.						
10. External factors that can lead to problems in the execution of activities are taken into consideration in the short-term planning.						
11. In the integrated planning, changes in risk levels are taken into consideration when future activities are considered.						
12. The ICT-tools used to prioritize our activities gives updated and clear information.						
Learn						
13. There is a well-functioning two-way communication between the offshore- and onshore organization during planning						
14. There is a well-functioning performance measurement system for how the integrated planning process works.						

Conclusion

This White Paper has presented the main concepts and practical principles of Resilience Engineering, a developing field which will be important for safety in the future. IO is among the socio-technical systems where rapid changes and developments exceed the ability of established safety assessment approaches to address all the issues and identify all the risks. In complex socio-technical systems, things may go wrong in the absence of manifest failures and malfunctions, and outcomes may often be disproportionately large. To safeguard against such developments, it is necessary to have tools and methods that can deal with the underspecification and tight couplings of complex, highly interactive systems such as IO. Resilience Engineering offers a conceptual and methodological basis for achieving that goal. The research and development efforts will continue and experiences will be documented and disseminated to demonstrate the added value of this way of thinking. Special emphasis will be put on case studies and guidance on how a smooth integration with conventional safety approaches can be accomplished.

Further reading

Albrechtsen, E (ed.) (2010). Essays on socio-technical vulnerabilities and strategies of control in Integrated Operations. SINTEF report A14732. Available at www.sintef.no/rio

Hollnagel, E., Woods, D. D. & Leveson, N. (2006). Resilience engineering: Concepts and precepts. Aldershot, UK: Ashgate publishers.

Hollnagel, E, Nemeth, C.P., Dekker, S (2008) Resilience Engineering volume 1. Remaining Sensitive to the possibility of failure. Aldershot, UK: Ashgate publishers.

Hollnagel, E. (2009). The ETTO principle: Efficiency-thoroughness trade-off. Why things that go right sometimes go wrong. Aldershot, UK: Ashgate publishers.

Nemeth, C.P., Hollnagel, E, Dekker, S (2009) Resilience Engineering volume 2. Preparation and restoration. Aldershot, UK: Ashgate publishers.

Hollnagel, E., Paries, J., Woods, D. D. & Wreathall, J. (eds.) (2011). Resilience engineering in practice: A guidebook. Farnham, UK: Ashgate.

Glossary

Instantiation: A concrete version of the model, where detailed knowledge of the scenario defines how the aspects of the functions are coupled or 'linked.' A model may thus give rise to different instantiations, depending on the scenario details.

Intractable: A system which cannot be described in every detail and where the functioning therefore is not completely understood. Intractable systems are only partly predictable.

Performance variability: The ways in which individual and collective performances are adjusted to match current demands and resources, in order to ensure that things go right.

Resilience: The intrinsic ability of a system to adjust its functioning prior to, during, or following changes and disturbances, so that it can sustain required operations under both expected and unexpected conditions.

Resonance: A principle that explains how disproportionate large consequences can arise from seemingly small variations in performance and conditions.

Serendipity: The making of happy and unexpected discoveries by accident or when looking for something else; such as discovery.

Center for Integrated Operations in the Petroleum Industry

The center for Integrated Operations in the Petroleum Industry (IO Center) conducts research, innovation and education within the IO field, to promote accelerated production, increased oil recovery, reduced operating costs and enhanced safety and environmental standards. The center was established by NTNU, SINTEF and IFE, in collaboration with major international oil companies and suppliers and is financed by 12 industrial partners and the Norwegian Research Council.

For more information, visit www.ntnu.edu/iocenter/

IO center point of contact: eirik.albrechtsen@sintef.no



Center for Integrated Operations in the Petroleum Industry