# RAG – Resilience Analysis Grid

***Technical Document prepared by the Industrial Safety Chair, January 2009***

## Introduction

Resilience is today recognised as an important quality of an organisation or a system[1] and describes the systems ability to adjust its functioning prior to, during, or following changes and disturbances, so that it can sustain required operations even after a major mishap or in the presence of continuous stress. Resilience Engineering does not see failures as a breakdown or malfunctioning of normal system functions, but rather as the converse of the adjustments required to cope with the unpredictability of the real world. Resilience Engineering, as a practical discipline, looks for ways to enhance the ability of systems to *succeed* under varying conditions. This means, more specifically, the ability to *respond* effectively to disruptions or ongoing production and economic pressures, to *monitor* threats and revise risk models, to *anticipate* future threats, disruptions and other destabilizing conditions, and to *learn* from past events, to understand correctly both what happened and why.

### *Safety as a Quality*

A system is usually considered safe if the number of adverse outcomes is acceptably small.[2] Adverse outcomes are typically accidents and incidents, but may also include work time injury, work related illnesses, etc. Adverse outcomes are counted per *safety unit*, which can be characteristic operations or specific durations (cf., Amalberti, 2002). The level of safety is therefore measured by the number of such outcomes per *safety unit*, and the common interpretation is that higher safety corresponds to a smaller number of outcomes. One example of that is the following definition:

> Safety is the state in which the risk of harm to persons or of property damage is reduced to, and maintained at or below, an acceptable level through a continuing process of hazard identification and risk management. (ICAO, 2006).

Resilience engineering takes a broader look and defines safety as *the ability to succeed under varying conditions*. This clearly includes the traditional definition of safety, since it is a consequence of resilience that there will be fewer adverse outcomes. But resilience also has the system's ability to function broadly in focus. Resilience engineering is about the operations necessary for the system's continued existence and growth, hence address core

[1] In this technical note, the terms 'organisation' and 'system' will be used synonymously, even though they are not synonyms strictly speaking.

[2] It is significant that safety is defined in terms of outcomes rather than in terms of, e.g., events. This was noted by Heinrich (1959) who made a clear distinction between the accident as an event and the injury as the outcome.

business (productivity, quality, and effectiveness) as well as safety. This has consequences both for how safety is measured and for how safety is managed.

## *The Four Cornerstones of Resilience*

The key term in the definition of resilience is the ability of a system to adjust how the it functions. This makes clear that resilience is not just the ability to continue functioning in the presence of stress and disturbances. While the ability of a system to preserve and sustain its primary functions is important, this can be achieved by other and more traditional means. Continued functioning can for instance be achieved by isolating the system from the environment, or by making it impervious to exogenous disturbances. An example of that is the *defence-in-depth* principle, which means that there are multiple layers of barriers between the system and the environment in which it exists.

Adjustments can in principle take place either *after* something has happened (be reactive) or *before* something happens (be proactive). Reactive adjustments are by far the most common. For instance, if there is a major accident in a community, such as a large fire or an explosion, local hospitals will change their state of functioning and prepare for a rush of people that may have been hurt. Responding when something has happened may, however, be insufficient to guarantee the system's safety and survivability. One reason is that a system can only be ready to respond to a limited set of events or conditions, either because it only recognises a certain set of symptoms or because it only has the resources needed for some kinds of events but not for others – and usually only for a limited duration. Vivid examples of that can be found in everyday events, the most conspicuous case in recent years being the unfortunate lack of response by the Federal Emergence Management Agency (FEMA) to the Hurricane Katrina in 2005 (e.g., Comfort & Haase, 2006). In the world of business, the failure of Airbus company to recognise and effectively respond to the problems with the production of the A380 in the June 2006, and the later failure of the Boeing company to do the same with the production of the 787 in September 2007, suggest that limited readiness is not an unusual phenomenon at all.

The ability to adjust *after* something has happened relies on the experiences from past events, not only to establish a specific readiness but also to make decisions about structural or functional changes that may make the system better prepared for what can happen in the future. These changes are often directed at the causes, as determined by accident investigations, although such causes and explanations always must be seen relative to the accident models and the investigation methods that were used (Hollnagel, 2004; Hollnagel & Speziali, 2008).

Making adjustments prior to an event means that the system can change from a state of normal functioning to a state of heightened readiness *before* something happens. A state of readiness means that resources are allocated to match the needs of the expected event, that special functions are activated, and that defences are increased. A trivial example is to batten down the hatches to prepare for stormy weather, either literally or metaphorically. An everyday example from the world of aviation is to secure the seat belts before start and landing or during turbulence. In these cases, the criteria for changing from a normal state to a state of readiness are clear. In other cases it may be less obvious either because of a lack of experience or because the validity of indicators is questionable. An increased state of alertness should, of course, not last longer than necessary since it may consume resources that otherwise could be used for normal performance.

The working definition of resilience can be made more detailed by pointing to four cornerstones of resilience, each representing an essential capability, cf., Figure 1:

- Knowing what to *do*, i.e., how to respond to regular and irregular disruptions and disturbances by adjusting normal functioning. This is the ability to address the *actual*.
- Knowing what to *look for*, i.e, how to monitor that which is or could become a threat 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*, i.e., how to anticipate developments and threats further into the future, such as potential disruptions, pressures, and their consequences. This is the address to address the *potential*.
- Knowing what *has happened*, i.e., how to learn from experience, in particular to learn the right lessons from the right experience. This is the ability to address the *factual.*

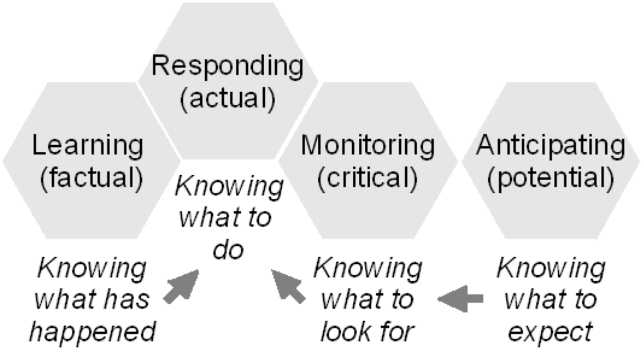

*Figure 1: The four cornerstones of resilience*

## *Basic Requirements to Manage Something*

In order to manage a system – indeed, in order to manage anything – three requirements must be met. First, it is necessary to know the *current* status or present position. Second, it is necessary to have a clear idea about what the *future* status or position should be. And third, it is necessary to know by which *means* an effective change from the present position to the future position can be made.

The problem of management, which in practice is equivalent to the problem of control, can in general be expressed as follows, where each of the three requirements corresponds to a specific question:

- How X are we? (Alternatively: how much of X do we have?)
- How X should we be? (Alternatively: how much of X should we have?)
- How do we become more / less X, or how do we get more / less of X?

The quality X can be many things, as the following two simple examples will show. In the case of safety, the three questions become:

- How safe is our operation or how safe is our department / company / primary process / etc.?
- How safe should our operation or system be, and when (by the end of this year / next year/ five years from now, and so on.)?
- Which means do we have that can make our operation or system safer, and how should we apply them?

Or, to take a more topical example, the questions for financial exposure could look like this:

- How exposed are we to the volatility of the stock market?
- How exposed should we be to the volatility of the stock market?
- How do we become less exposed to the volatility of the stock market?

(Notice that in the case of safety, the desired change is to *increase* something, where as in the case of financial exposure the desired change is to *reduce* something. The direction of change depends on how the second requirement is formulated.)

Process and safety management traditionally put much effort into meeting the first requirement, i.e., how something can be measured. But in order to be in control of something it is clearly necessary to know the direction in which change should be made, whether it should be *towards* something rather than *away* from something, and to have the means (tools, techniques) effectively to make the change. The present note will for practical reasons focus on the first requirement, but it is planned to address the two other requirements at a later time.

# The Measurement Problem

The first requirement raises the problem of measurements, specifically of performance measurements or performance indicators. This involves a number of commonly recognised issues, such as:

- Can the status be expressed by a single measurement or does it require a calculation, i.e., a combination of several measurements?
- Do the measurements represent lagging indicators or leading indicators, i.e., is the measurement of a condition in the past or is it indicative of a position in the future?
- Are the measurements reliable and valid?
- Are the measurements well-defined?
- Are the measurements objective or subjective, e.g., can they be made automatically and by technological means or do they rely on the judgement or opinions of people?

A measurement is normally understood as a quantity of something, for instance a value or a number, and is normally a scalar rather than a vector. But a number in itself is meaningless unless it can be set in relation to something or some context. In other words, quantities or numbers have to be interpreted. This can only be done by referring to a common understanding or a set of conventions. Consider, for instance, the following example:

| | |
|---|---|
| *52* | The number 52 taken by itself does not mean anything. Indeed, it could be a symbol as well as a number. |
| *52 kilos* | The number has now got a unit, so at least we know what the number represents, namely a weight of something. But we do not know whether it is 52 kilos of CO2 (which is how much a common car produces when driven for 350 km), or 52 kilos of cheese. |
| *Pierre weighs 52 kilos* | The added information, i.e., the name of a person, makes the number even more meaningful. We might even begin to make assumptions about whether this weight indicates a normal condition or an unusual condition. |
| *Pierre is 12 years old and weighs 52 kilos* | Because of the additional information, the number finally makes some sense. Even without being a physician we can say that the weight is not normal, and that Pierre possibly is obese. |

The final formulation does not only provide the necessary context to understand what the number means, but also invokes a frame of reference that can be used to plan what to do as a consequence of the measurement.[3]

## *Measurements of Safety*

As mentioned above, safety is usually defined as the absence of unwanted events. The level of safety is consequently measured by the number of such events per safety unit. As a simple example of that, consider the top five HSE[4] indicators used by the oil industry.

1. (Number of) fatal accidents
2. Total recordable injury frequency (TRIF)
3. Lost-time injury frequency (LTIF)
4. Serious HSE incident frequency (SIF)
5. Accidental oil spill (number and volume)

The European Technology Platform on Industrial Safety uses the same approach. The ETPIS does acknowledge that "Safety is ... a key factor for successful business and an inherent element of business performance," and then continues:

> ... (I)ndustrial safety performance will have progressively and measurably improved in terms of reduction of reportable accidents at work, occupational diseases, environmental incidents and accident-related production losses.

It is quite understandable that safety traditionally has focused on adverse outcomes, since these represent situations that any organisation would want to avoid. Adverse outcomes are also phenomena that by their very nature attract attention both in terms of their direct effects (loss of life, property, and money) and in terms of their indirect effects (disruption of functions and production, need of recovery operations, restoration, etc.).

Seen as measurements, adverse outcomes have two further advantages. The first is that they are countable, the second that they are (relatively) unambiguous.[5] They are also helpful towards meeting the second requirement, since it clearly is a step in the right direction to reduce the number of adverse outcomes, i.e., to reduce their number. But measurements or tallies of adverse outcomes are of little help to meet the third requirement, i.e., how effectively to make a change in the desired direction.

[3] This can, of course, involve not doing anything, if, e.g., the value was considered normal.

[4] HSE = Health, Safety, security and Environment

[5] Both advantages require that the outcomes are clearly defined in operational terms. In practice there may be large cultural and national differences in such definitions.

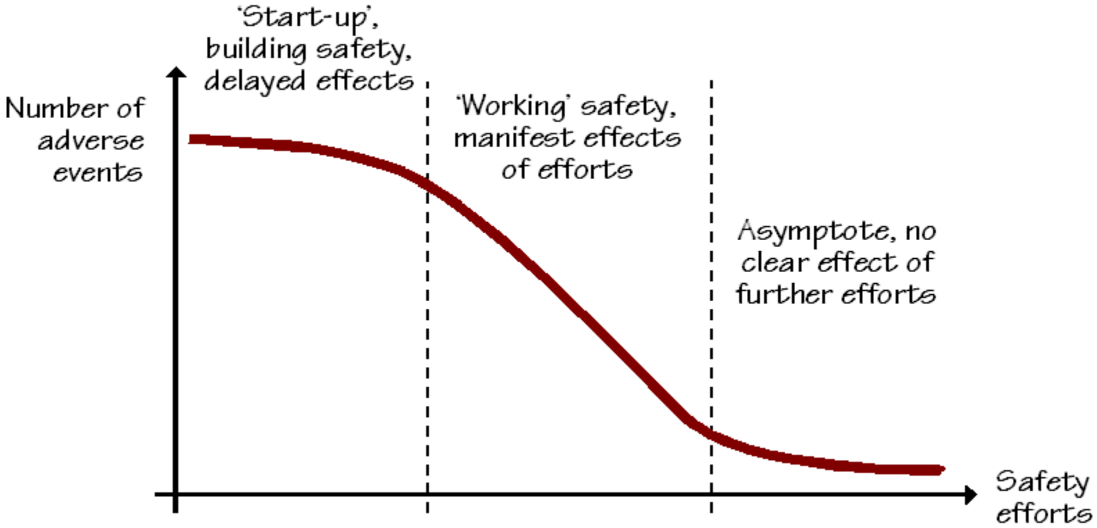

*Figure 2: Safety measured by adverse outcomes*

One problem with measuring safety by the number of adverse outcomes or unwanted events (accidents, incidents, etc.) is that it works best in the beginning, when safety is low, but not so well later, when safety is high. The reason is simply that if the number of adverse events has been effectively reduced, there is little to measure, hence little to show for the efforts made. Consider, for instance, the Strategic Research Agenda of ETPIS (2005). The explicit objective is "to achieve a 25% reduction in accidents by 2020 and to have programmes in place by 2020 to continue accident reduction at a rate of 5% per year or better." Assuming that these goals can be achieved, it does not require many calculations to realise that the proposed measurement at some time in the not too distant future will have reached an asymptote, hence will become useless.

## *Difference between Safety and Resilience*

The difference between safety and resilience makes a difference in how the two qualities can be measured. Safety is normally measured by the number of adverse events, where some examples have been provided above. This means that safety is measured by the product of a process, in this case the process of safety management. As any process, this will have two types of outcomes, the intended outcomes and the unintended outcomes. Whereas a normal feedback-controlled process used measurements of the intended outcomes as a basis for regulation, the safety management uses measurements of the unintended outcomes, i.e., the accidents, incidents, etc.[6]

Since resilience refers to the system's performance, the measurements should be of performance (process) rather than outcomes (products). This has two advantages. First, that

[6] From a regulation point of view this is probably not the most efficient way of controlling the process.

the measurements deal equally with successes and failures, or rather with the processes underlying successes and failures. This supports the Resilience Engineering view of failures as the flip side of successes, hence as being the outcomes of the same underlying processes. Second that the measurements are direct measurements of the process, and not indirect measurements of products. The traditional approach to safety requires an interpretation of what is measured (adverse outcomes) in terms of the possible underlying process – and more specifically the underlying process failure. Resilience makes that interpretation unnecessary by letting the measurement be of the process directly.[7]

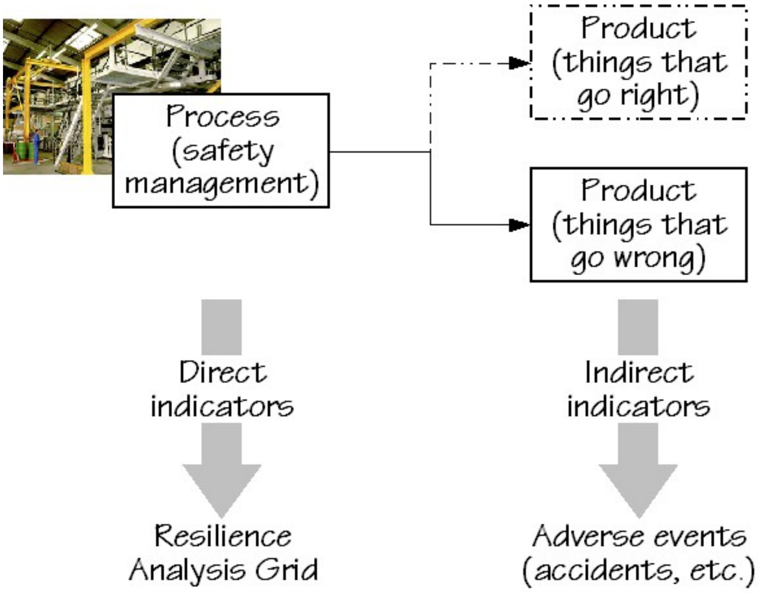
Process (safety management)
Product (things that go right)
Product (things that go wrong)
Direct indicators
Indirect indicators
Resilience Analysis Grid
Adverse events (accidents, etc.)

*Figure 3: Direct and indirect measures of safety*

## *Measurements of Resilience*

Since resilience is defined by the system's ability to adjust its functioning, it follows that a measure of resilience must be different from the traditional measures of safety. Since resilience refers to a quality rather than a quantity, something that the system *does* rather than something that the system *has*, it is not possible to point to any single or simple measurement. The solution is instead to consider the four capabilities that together define resilience, and on that basis develop a *Resilience Analysis Grid*, i.e., four sets of questions where the answers can be used to construct a resilience profile of a system or an organisation. The rest of the report will present an outline of such a *Resilience Analysis Grid.*

[7] From a methodological point of view it is, of course, only possible to measure a process by its products. But it makes a difference whether the 'products' are indicators of system states, which requires that the system is active, or whether they are of the outcomes as they exist independently of the process.

## *The ability to respond*

No system – be it an individual, a group, or an organisation – can sustain its functioning and continue to exist unless it is able to respond to what happens. The response must furthermore be both timely and effective so that it can bring about the desired outcome before it is too late. As described above, a resilient system responds by adjusting its functioning so that it better matches the new conditions. Other responses are to mitigate the effects of an adverse event, to prevent a further deterioration or spreading of effects, to restore the state that existed before the event or to resume the functioning that existed before, to change to stand-by conditions, etc.

In order to respond when something happens the system must be able to *detect* that something has happened. Second, it must be able to *identify* the event and *recognise* or *rate* it as being so serious that a response is necessary. Third, the system must know how to *respond* and be *capable* of responding, in particular it must have or be able to command the required resources *long enough* for the response to have an effect.

The detection that something has happened is not entirely passive but depends on what the system looks for – on what its pre-defined categories of critical events or threats are. If the system looks for the wrong events or threats it may either fail to recognise some threats (a *miss*) or respond to situations where a response was not actually needed (a *false alarm*). The former will leave the system vulnerable to unexpected events. The latter may be harmful both because the system may transition to a state that is not easily reversible, and because it wastes resources and reduces readiness.

Some events may be so obvious that they cannot be missed, yet without any response being ready – or even without a clear idea of what should be done. (The subprime crisis of 2007 was an example of that.) In such cases there may also be an urgency of the situation, i.e., an immediate pressure to act, which by its very nature limits the ability to consider what the proper response should be. Under such conditions the system may easily lose control by responding in an opportunistic or scrambled rather than in a more orderly mode (Hollnagel, 1998).

If an event or a threat is rated as serious, the response can either be to change to a state of readiness, or to take action in the concrete situation. In the first case, deciding that a state change is needed depends on a number of factors, cultural, organisational, and situational. The dilemma is nicely captured by the common definition of safety as the freedom from unacceptable risks, which forces the question of how large a risk is considered acceptable – and by whom. A common solution is to rely on probability calculations, and accept all risks where the probability is lower than some numerically defined limit (e.g., Amalberti, 2006). This, however, does not solve the problem of how the limit is set.[8]

The second case is how to decide whether a response should be activated in a given situation. As long as the activation of the response depends on technology, including software, the problem is in principle solvable. But in cases where the decision depends on humans, the problem is more difficult. Deciding whether to do something, and when to do it, depends to a

[8] It also requires that the risk can be accurately measured.

considerable extent on the competence of the people involved, and on the situation in which they find themselves (e.g., Dekker & Woods, 1999).

Finally, having the resources necessary for the chosen response is also essential. This is not only a question of having prepared resources, which really only makes sense for regular threats (cf., below), but also a question of whether the system is flexible enough to make the necessary resources available when needed.

| Analysis item (ability to respond) | Score or evaluation |
|---|---|
| **Event list**: What are the events for which the system has a prepared response? | |
| **Background**: How were these events selected (experience, expertise, risk assessment, etc.? | |
| **Relevance**: When was the list created? How often is it revised? On which basis is it revised? | |
| **Threshold**: When is a response activated? What is the triggering criterion or threshold? Is the criterion absolute or does it depend on internal / external factors? | |
| **Response list:** How was the specific type of response decided? How is it ascertained that it is adequate? (Empirically, or based on analyses or models?) | |
| **Speed**: How fast is full response capability available? | |
| **Duration**: For how long can a 100% effective response be sustained? | |
| **Stop rule**: What is the criterion for returning to a “normal” state? | |
| **Response capability**: How many resources are allocated to the response readiness (people, materials)? How many are exclusive for the response potential? | |
| **Verification**: How is the readiness to respond maintained? How is the readiness to respond verified? | |

## *The ability to monitor (keeping an eye on critical developments)*

A resilient system must be able flexibly to monitor what is going on, including its own performance. The ability to monitor enables the system to cope with that which could become critical in the near term. In order for the monitoring to be flexible, its basis must be assessed and revised from time to time.

As argued above, it is in practice only possible for a system to be ready to respond to regular threats, or even just to a subset of these. It is nevertheless a potential risk if the readiness to respond is limited to too small a number of events or conditions. The solution is to monitor for things that may become critical, and use that to change from a state of normal operation to a state of readiness when the conditions indicate that a crisis, disturbances, or failure is imminent. Such as two-step approach will be more cost effective. If a system can make itself ready when something is going to happen, rather than remain in a state of readiness more or less permanently, then resources may be freed for more productive purposes. The difficulty is, of course, to be able to decide that something may go wrong so early that there is sufficient time to change to a state of readiness. It is also necessary that the identification of the impending event is so reliable that preparations are not made in vain. There will of course always be situations that completely defy both preparations and monitoring – the dreaded unexampled events – but more can be done to reduce their number and frequency of occurrence than established safety practices allow.

Monitoring normally looks for certain conditions or relies on certain indicators. These are by definition called leading indicators, because they indicate what may happen *before* it happens. Everyday life is replete with examples, as the indicators for the weather tomorrow or for the coming winter (or summer). In the case of the weather there are good leading indicators because we have an accurate understanding of the phenomenon, i.e., of how the (weather) system functions. In other cases, and particularly in safety related cases, we only have weak or incomplete descriptions of what goes on and therefore have no effective way of proposing or defining valid leading indicators. Because of this, most systems rely on lagging indicators instead, such as accident statistics. While many lagging indicators have a reasonable face validity, they are only known with a delay that often may be quite considerable (e.g., annual statistics). The dilemma of lagging indicators is that while the likelihood of success increases the smaller the lag is (because early interventions are more effective than late ones), the validity or certainty of the indicator increases the longer the lag (or sampling period) is.

| Analysis item (ability to monitor) | Score or evaluation |
|---|---|
| **Indicator list**: How have the indicators been defined? (By analysis, by tradition, by industry consensus, by the regulator, by international standards, etc.) | |
| **Relevance**: How often is the list of indicators revised, and on what basis? | |
| **Indicator type**: How many of the indicators are leading, and how many are lagging? | |
| **Validity**: For leading indicators, how is their validity established? | |
| **Delay**: For lagging indicators, how long is the lag? | |
| **Measurement type**: What is the nature of the “measurements”? Qualitative or quantitative? (If quantitative, what kind of scaling is used?) | |
| **Measurement frequency**: How often are the measurements made? (Continuously, regularly, every now and then?) | |

| Analysis item (ability to monitor) | Score or evaluation |
|---|---|
| **Analysis**: What is the delay between measurement and analysis/interpretation? How many of the measurements are directly meaningful and how many require analysis of some kind? | |
| **Stability**: Are the effects measured transient or permanent? | |
| **Organisational support**: Is there a regular inspection scheme or schedule? Is it properly resourced? | |

## *The ability to anticipate (looking for future threats and opportunities)*

While looking for what may go wrong in the immediate future generally makes sense, it may be less obvious that there is an advantage to look at the more distant future. The difference between monitoring and looking ahead is both that the time horizons are different (short versus long), and also that it is done in different ways. In monitoring, a set of pre-defined cues or indicators are checked to see if they change in a way that may demand a response. In looking for the potential, the goal is to identify possible future events, conditions, or state changes – internal or external to the system – that may threaten the system's ability to function. While monitoring tries to keep an eye on the regular threats, looking for the potential tries to identify the most likely irregular threats.

While risk assessment already does look for the potential, it is constrained because it relies on representations of linear combinations of discrete events, such as event trees and fault trees. Established risk assessment methods are developed for tractable systems where the principles of functioning are known, where descriptions do not contain too many details, where descriptions can be made relatively quickly, and where the system does not change while the description is being made (Hollnagel, 2008). For such systems it may be acceptable to look for the failure potential in simple combinations of discrete events or linear extrapolations of the past. Many present day systems of major interest for industrial safety are unfortunately not like that. This means that the principles of functioning are only partly or incompletely known, that the description is elaborate and contains many details, that it takes a long time to make, and that the system therefore changes while the description is made. In consequence of that there will never be a complete description of the system and it is therefore ill advised to rely on established risk assessment methods.

Looking for the potential requires requisite imagination or the ability to imagine key aspects of the future (Westrum, 1993). As described by Adamski & Westrum (2003), requisite imagination is needed to know from which direction trouble is likely to arrive and to explore those factors that can affect outcomes in future contexts. The relevance of doing that is unfortunately not always accepted, since it requires resources that could have been used for, e.g., production.

Even if the possibility that something could go wrong is acknowledged, thinking about the potential is fraught with difficulties. Many studies have, for instance, shown that human thinking makes use of a number of simplifying heuristics such as representativeness, recency, and anchoring (Tversky & Kahneman, 1974). While these may improve efficiency in normal working conditions, they severely restrict the more open-minded thinking that is necessary to look at the possible. Looking for the potential is also difficult because it requires a disciplined combination of individual or collective imagination. It can also be costly, both because it cannot be hurried but must take its time and because it deals with something that may happen so far into the future that benefits are rather uncertain. Relatively few systems therefore allocate sufficient resources to look at the potential. However, a truly resilient system realises the need at least to do something.

| Analysis item (ability to anticipate) | Score or evaluation |
|---|---|
| **Expertise**: What kind of expertise is relied upon to look into the future? (In-house, outsourced?) | |
| **Frequency**: How often are future threat and opportunities assessed? | |
| **Communication**: How are the expectations about future events communicated or shared within the organisation? | |
| **Strategy**: Does the organisation have a clearly formulated ‘model of the future’? | |
| **Model**: Is the model explicit or implicit? Qualitative or quantitative? | |
| **Time horizon**: How far ahead does the organisation plan? Is the time horizon different for business and safety? | |
| **Acceptability**: Which risks are considered acceptable and which unacceptable? On which basis? | |
| **Aetiology**: What is the assumed nature of future threats?<br>• Same as previous threats/accidents?<br>• Combination/extrapolation of known accidents / incidents?<br>• Completely novel threats? | |
| **Culture**: Is risk awareness part of the organisational culture? | |

## *The ability to learn (finding and making use of the right experience)*

A resilient system must be able to learn from experience. Although this is mentioned last, it is in many ways the basis for the ability to respond, to monitor, and to look ahead. To learn from experience sounds rather straightforward and few safety managers, administrators, or regulators will disagree with that. Yet if it is to be done in an efficient and systematic manner, it requires careful planning and ample resources. The effectiveness of learning depends on what the basis for the learning is, i.e., which events or experiences are taken into account; on how the events are analysed and understood; and on when and how often the learning takes place.

In learning from experience it is important to separate what is *easy* to learn from what is *meaningful* to learn. Experience is often couched in terms of the number or frequency of occurrence of some event or other, usually ones that are negative (accidents, incidents, loss time, etc.). But counting is not the same as learning. In order for a measure to be useful, it must be meaningful, hence refer to a principle, a model, or some kind of conceptual basis. While compiling extensive accident statistics may seem impressive it does not mean that the system actually learns anything. Knowing how many accidents have occurred says nothing about why they have occurred, nor anything about the many situations when accidents did not occur. And without knowing *why* accidents occur, as well as knowing why they do *not* occur, it is impossible to propose effective ways to improve safety.

| Analysis item (ability to learn) | Score or evaluation |
|---|---|
| **Selection criteria**: Which events are investigated and which are not? How is the selection made? Who makes the selection? | |
| **Learning basis**: Does the organisation try to learn form successes as well as from failures? | |
| **Classification**: How are events described? How are data collected and categories? | |
| **Formalisation**: Are there any formal procedures for investigation and learning? | |
| **Training**: Is there any formal training or organisational support for investigation and learning? | |
| **Learning style**: Is learning a continuous or discrete (event-driven) activity? | |
| **Resources**: How many resources are allocated to investigation and learning? Are they adequate? | |
| **Delay**: What is the delay in reporting and learning? How are the outcomes communicated within and without the organisation? | |
| **Learning target**: On which level does the learning take effect? (individual, collective, organisational) | |

## Rating Resilience

The four sets of items described above constitute the *Resilience Analysis Grid* (RAG). In order for the RAG to be useful as a tool, it is necessary that the answer to each item can be rated or assigned a value. This can be done in a number of fashions. As a beginning, it is proposed that the rating is done according to the following scale:

- Excellent – the system/organisation meets and exceeds the criteria for the required capability.

- Satisfactory – the system/organisation fully meets all reasonable criteria for the required capability
- Acceptable – the system/organisation meets the nominal criteria for the required capability.
- Unacceptable – the system/organisation does not meet the nominal criteria for the required capability.
- Deficient – there is insufficient capability to provide the required capability.
- Missing – there is no capability to provide the required capability or information is not available.

Consider, for instance, the first item for 'ability to respond': 'What are the events for which the system has a prepared response.' The rating can be made using the following guideline:

- Excellent – the system/organisation can respond to all imaginable events.
- Satisfactory – the system/organisation can respond to all reasonable events.
- Acceptable – the system/organisation can respond to events that occur frequently, or events that are defined by a regulator.
- Unacceptable – the system/organisation can only respond to some of the events that occur frequently.
- Deficient – the system/organisation can only respond to the most critical events.
- Missing – there are no prepared responses or information is not available.

Similar guidelines to evaluate or rate responses can be developed for the other items. Once the rating has been done for the 10 items that characterise the 'ability to respond,' it is possible to show the ratings graphically using the star-diagram shown in Figure 4. The advantage of this style of graphical representation is that it uses of a regular polygon to show the result of the evaluation. Any irregularities in the shape of the polygon will be easy to detect, and provide a clear signature of how well the system/organisation rates in regard to the ability to respond.

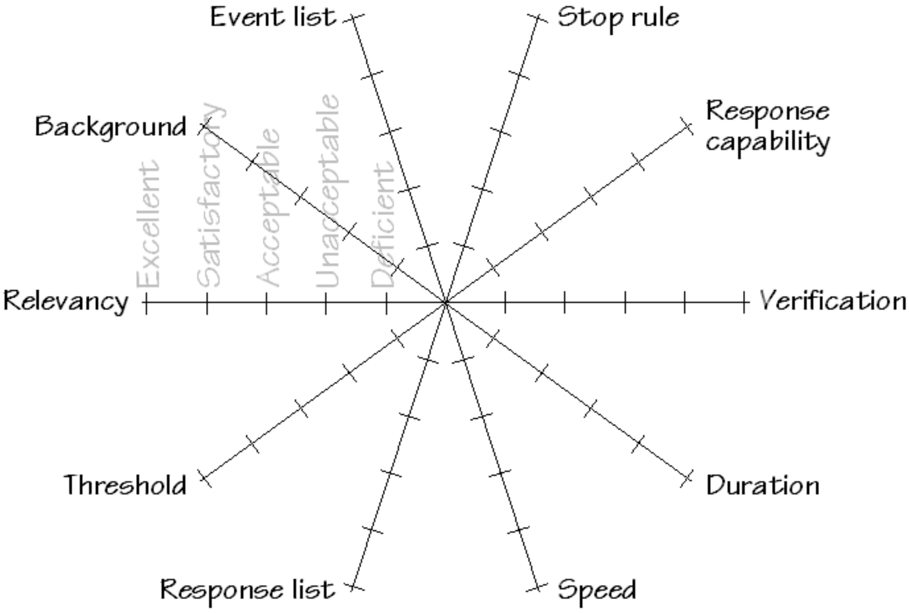

*Figure 4: Star diagram for 'ability to respond'*

Similar star-diagrams can be developed for the other abilities. In order to provide an overall impression of the resilience of an organisation, it is necessary to specify a way to collapse each star diagram, or rather the ratings of the individual items, to a single rating. When that is done, it is possible to show a star diagram for the four components of resilience, i.e., the abilities to respond, monitor, anticipate, and learn. Figure 5 shows what a RAG star diagram would look like for an organisation, where the combined ratings for all of the four components were ‘acceptable.’

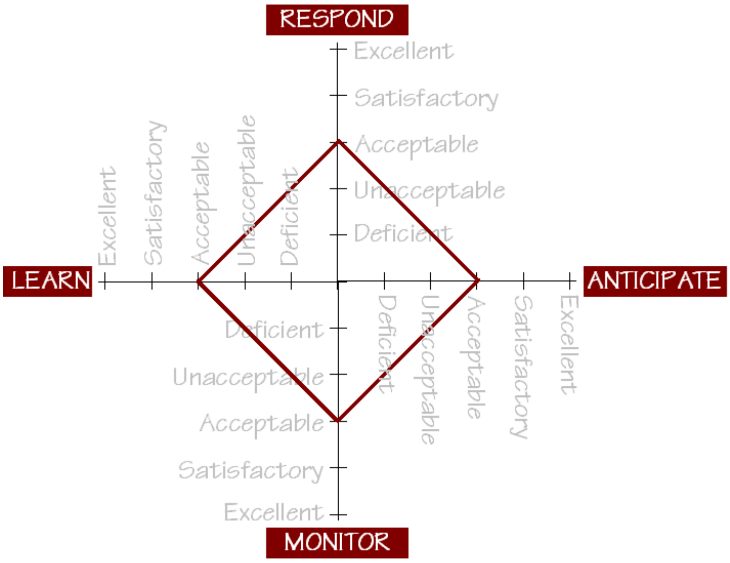

*Figure 5: RAG star diagram (showing acceptable ratings)*

The next two figures illustrate how other ratings would appear. Figure 6 shows the RAG for a organisation that does well in terms of the ability to respond and monitor, but which fails in terms of the ability to anticipate and learn. While such an organisation may be safe in the short run, it is not resilient.

Figure 7 shows what a star diagram would like for a resilient organisation. The RAG reveals that the organisation does very well in terms of the ability to respond and monitor. But in addition it puts efforts into looking at the potential ('anticipate') and has an acceptable ability to learn. The diagram also suggests how the resilience can be improved.

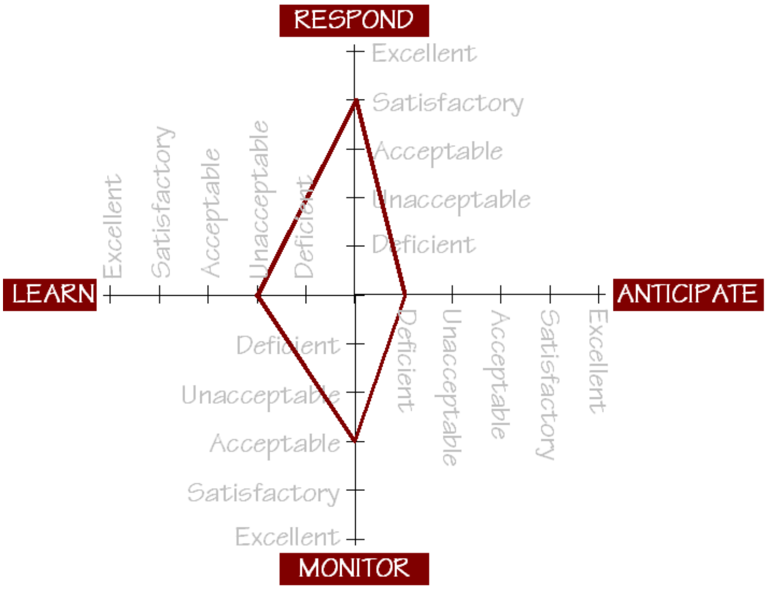

*Figure 6: RAG star diagram (show lack of resilience)*

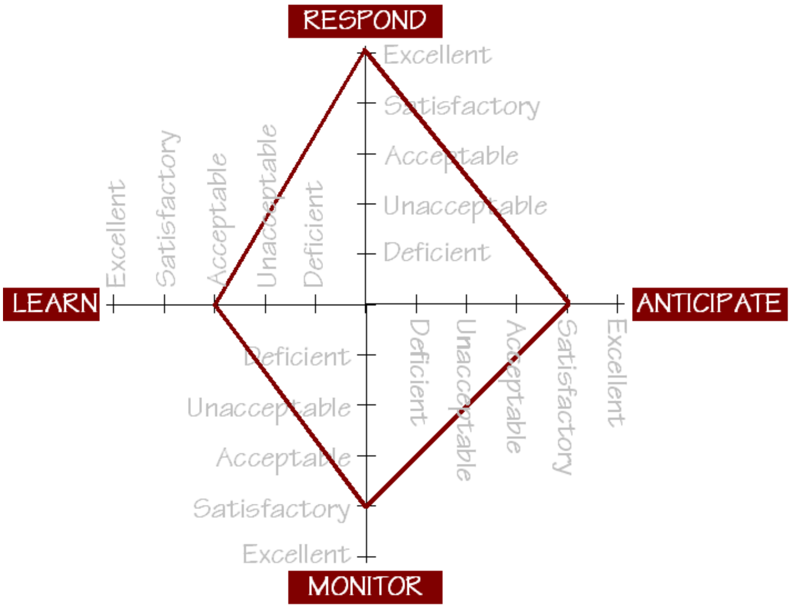

*Figure 7: RAG star diagram for resilient organisation*

# Summary

The *Resilience Analysis Grid* presented here is not proposed as a final tool that can be used directly. It is rather intended as a basis from which a more specific grid – or set of questions – can be developed. The questions must clearly be relevant for the organisation where they are intended to be used, and may therefore require clarification and reformulation.

The note has outline the principles for how the evaluations can be rated, and the star digram similarly suggest what constitutes an acceptable score. The star diagram is not in itself a measure of resilience, but is a compact representation of how the various items have been rated. It is also a process measure rather than a product measure, i.e., it shows the current state of things – the current level of resilience and of how well the organisation does on each of the four main capabilities.

The *Resilience Analysis Grid* can be used useful to determine the current state of resilience of an organisation, and also to define the future position or objective. This does not follow from the questions themselves, but from the way in which they are evaluated. The *Resilience Analysis Grid*, is also useful to meet the third requirement, because the questions are derived from an underlying theory of what resilience is.

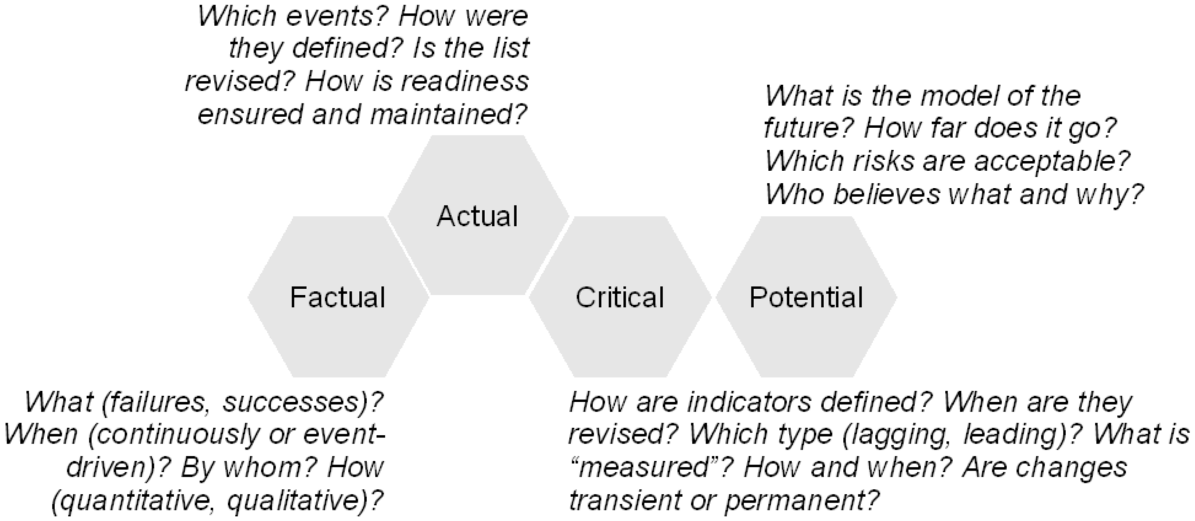
Which events? How were they defined? Is the list revised? How is readiness ensured and maintained?
What is the model of the future? How far does it go? Which risks are acceptable? Who believes what and why?
Actual
Factual
Critical
Potential
What (failures, successes)? When (continuously or event-driven)? By whom? How (quantitative, qualitative)?
How are indicators defined? When are they revised? Which type (lagging, leading)? What is “measured”? How and when? Are changes transient or permanent?

*Figure 8: Implementation questions for resilience engineering*

Resilience engineering does not prescribe a certain balance or proportion among the four qualities. But it does make clear that it is necessary for an organisation to address each of these qualities to some extent, in order to be resilient. This can be illustrated by the shape of the polygon in the star diagrams, provided appropriate rating rules and weights have been developed. All organisations traditionally put some effort into the ability to respond to the actual. Many also put some effort into the ability to learn from the factual, although it often is in a very stereotypical manner. Fewer organisations make a sustained effort to monitor the critical, particularly if there has been a long period of stability. And very few organisations put any serious effort into the ability to anticipate the potential.

## References

Adamski, A. & Westrum, R. (2003). Requisite imagination. The fine art of anticipating what might go wrong. In E. Hollnagel (Ed.), *Handbook of cognitive task design* (pp. 193-220). Mahwah, NJ: Lawrence Erlbaum Associates.

Amalberti, R. (2002), Revisiting safety and human factors paradigms to meet the safety challenges of ultra complex and safe systems. In B. Wilpert & B. Fahlbruch (Eds.), *System safety: Challenges and pitfalls of intervention.* Pergamon.

Comfort, L. K. & Haase, T. W. (2006). Communication, coherence and collective action: the impact of Hurricane Katrina on communications infrastructure. *Public Works Management & Policy*, 11(1), 6-16.

Dekker, S. W. A. & Woods, D. D. (1999). To intervene or not to intervene: The dilemma of management by exception. *Cognition, Technology & Work*, *1*(2), 86-96.

European Technology Platform on Industrial Safety (ETPIS) (2005). *Safety for Sustainable European Industry Growth: Strategic Research Agenda*. www.industrialsafety-tp.org

Hollnagel, E. (1998). *Cognitive reliability and error analysis method*. London, UK: Elsevier.

Hollnagel, E. (2004). *Barriers and accident prevention*. Aldershot, UK: Ashgate.

Hollnagel, E. (2008). *From protection to resilience: Changing views on how to achieve safety*. Proceedings of the 8th International Symposium of the Australian Aviation Psychology Association, April 8-11, Sydney, Australia.

Hollnagel, E. & Speziali, J. (2008). *Study on developments in accident investigation methods: A survey of the "state-of-the-art"* (SKI 2008:50). Stockholm, Sweden: Swedish Nuclear Inspectorate.

ICAO (International Civil Aviation Organisation), (2006). *Safety management manual*. Montreal, Canada: Document sales unit.

Tversky, A. & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*, 1124-1131.

Westrum, R. (1993). Cultures with requisite imagination. In J. A. Wise, V. D. Hopkin & P. Stager (Eds.), *Verification ad validation of complex systems: Human factors issues*. Berlin: Springer-Verlag.