FUNCTIONAL SAFETY NETS FROM A RESILIENCE ENGINEERING PERSPECTIVE

by Professor Erik Hollnagel

Originally, a safety net was a large net that could catch someone who accidentally fell from a height, such as the safety net used in a circus trapeze act or the safety nets used at many building sites ever since the construction of the Golden Gate Bridge in San Francisco (1933-1937). The purpose of such a physical safety net is to prevent harm when something or someone falls unexpectedly, either harm to the someone who is falling or harm to the someone who can be hit by the something that is falling.

Functional safety nets and loss of control

Today, the meaning of the term 'safety net' has been extended to describe not only *physical* safety nets but also *functional* safety nets, in the sense of the various ways in which a situation can be prevented from going out of control, or be saved if control has been temporarily lost. A loss of control can have serious consequences in two different ways. First, that it becomes impossible to ensure that an activity continues as intended: the future becomes more uncertain and neither safety nor productivity – or for that matter quality – can be effectively managed. Second, that the loss of control leads to a loss of life, time, and/or material or immaterial property.

From a resilience engineering perspective, the primary purpose of a functional safety net is, however, not simply to re-establish control but rather to dampen or delay unmanaged developments as a prerequisite to re-establishing control. Examples of functional safety nets are not limited to aviation but can be found in almost every line of activity. They range from a social or economic safety net in the case of unemployment or illness, over the collective experience that an organisation can fall back on when something happens, to the technical and non-technical competencies and experience that are ready for use to manage and stabilise irregular situations. A functional safety net can therefore be seen as a kind of active barrier that limits the consequences of a temporary loss of control.

A functional safety net involves a prepared systemic response that can be carried out either instantaneously or with very little delay. A functional safety net cannot serve its purpose if a response first has to be prepared or if the required resources first have to be activated – just as a physical safety net will fail to serve its purpose if it has to be installed prior to being used when the need arises. A functional safety net also primarily compensates for something that is missing in a situation – such as a specific practical or theoretical competence. The response therefore differs from a recovery action, which may take time to plan and activate and which may also be expected to work over longer periods of time.

In aviation, the term 'safety net' has been used to include also the automated systems that keep an eye on work and that intervene to help keep performance within safe limits, e.g., a TCAS. But in resilience engineering terms it would be

simpler, and more correct, to call these for automated safety (or protection) systems rather than safety nets, if for no other reason then because such systems are unable to learn on their own: they are designed but do not themselves develop. The functional safety nets I will now discuss will therefore exclude automated safety systems.

Today's socio-technical systems are often called complex, or even complex adaptive, systems (CS or CAS). Complex (adaptive) systems are partly intractable and must work in partly intractable environments where demands and resources may change when least expected. This makes it impossible fully to rely on a set of pre-defined responses. A functional safety net must continuously develop and improve its responses to prevent that the discrepancy between what it can do and what is needed becomes too large. And it must do so itself, rather than wait for some *deus ex machina* to bring it up to date.

Resilience engineering proposes that four fundamental abilities are required for a system's potential to perform in a resilient manner – or in short, for its resilience. The first is the ability to respond, the second the ability to monitor, the third the ability to learn, and the fourth the ability to anticipate. A functional safety net represents a subset of the ability to respond because it is only concerned with the responses to the potential or actual loss of control. The everyday functioning of a system clearly requires many other kinds of responses as well. The ability to respond, whether in the broad or the narrow sense, should, however, not be considered in isolation. Resilience engineering makes clear that the four abilities depend on each other and that they therefore must be seen together, as an integrated whole. Before we can begin to measure and manage a system's resilience potential, we must therefore first uncover and understand the ways in which each of the four abilities depends on the others.

In order to understand the ability to respond that is the essence of a functional safety net, we must find out what this ability depends on or requires as support. In other words, how does it depend on the other abilities – and possibly on other system functions?

While responding may be improved by monitoring, which enables timely responses, as well as anticipation, which supports the preparation of responses, the most significant dependence is clearly on the ability to learn. *The reason is simply that without learning, the responses will always remain the same*. But always responding in the same way is bound to be insufficient, unless the environment and the conditions of work are perfectly the stable. This may possibly be the case for some types of physical safety nets; but it will never be the case for functional safety nets, not even for systems that only change very slowly. And aviation is definitely not one of those.

Functional safety nets and organisational learning

Organisational learning is an issue where there are more theories and opinions than there are facts. But the basic idea is simply that organisations learn by encoding inferences from experience into routines that guide or support behaviour. If we consider the role and nature of functional safety nets, we can see that three types of learning may play a role. An organisation can learn from its own experience (direct or intra-organisational learning), from the experience of others (indirect or inter organisational learning), and by developing industry-wide conceptual frameworks or paradigms for interpreting practical experience (systemic learning). Direct and indirect learning are both relevant for functional safety nets.

Learning from own experience is direct and involves little delay, regardless of whether it is done by individuals, by groups, or by the organisations. Typical examples are the sharing of good habits, or even best practices, among colleagues or within a group or an organisation. Direct learning will usually be very specific to the organisation and the type of activity it performs. The advantage is that learning can be directly associated with specific situations or conditions. The disadvantage is that the specificity makes it difficult to generalise, in particular to other organisations.

In the case of direct learning, the time lag or delay between learning and use is short. Because the learning is specific to the organisation and/or some situations, the lessons learned will be readily available when the need arises. Since the safety nets are localised within the organisation they can also be maintained as part of everyday work.

While learning from own experience is valuable, it is inescapably limited. It is therefore important to learn also from other organisations that are involved in the same kind of activity or service, but probably less important to learn from completely different domains. This is the rationale for proposing industry-wide 'best practices' and for defining safety nets as collaborative, mutually-supporting activities to sustain safety within an industry. But while the experience of others may be useful, it suffers from being indirect rather than direct. No two organisations, such as two airlines or two ANSPs, work in exactly the same way or have exactly the same working conditions. The direct experience of one organisation therefore becomes the indirect experience of another, and must be interpreted or 'coded' in some way before this other organisation can use it.

In the case of indirect learning, there may also be a substantial time lag or delay between learning and use. The transmission mostly takes place by informal means, through talks among colleagues or via significant adverse events (though these are not the best to learn from), and therefore without systematic support from either organisation. The assimilation of the learning inevitably requires some form of 'tailoring' of the original responses to the new context. The indirect learning will not be immediately relevant or applicable by an organisation, but must be mediated in one way or another. This means that the readiness to respond is less than for direct learning. Indirect learning therefore has an associated cost that should be carefully considered when safety nets are built.

The Bottom Line

Functional safety nets are by their nature socio-technical rather than technical. They are not designed and fixed, but develop and change over time. They represent part of an organisation's ability to respond and their effectiveness depends on the ability of the overall system to learn. Organisations must therefore look for the best possible ways to ensure the learning on which the efficacy of the functional safety net depends. While individual organisations may find that a combination of direct and indirect learning is sufficient for the development and management of functional safety nets, there is also a need to encode or institutionalise such knowledge for even wider use. We often hear that we must learn from the good experiences of other industries. And strangely enough each industrial domain (e.g. nuclear, aviation, healthcare, off-shore, etc.) seems to believe that other domains are doing better and that one therefore should try to encapsulate or imitate the lessons learned there. But is the grass really greener on the other side of the fence?

PROFESSOR ERIK **HOLLNAGEL**

is Professor at the University of Southern Denmark (DK), Professor Emeritus at the University of Linköping (S). Professional interests: industrial safety, resilience engineering, patient safety, accident investigation, and modelling large-scale socio-technical systems. He has published widely and is the author/editor of 21 books, including five books on resilience engineering.
Erik also coordinates the Resilient Health Care net and the FRAMily.