Preface

Resilience Engineering in a Nutshell

Erik Hollnagel

Since the publishing of *Resilience Engineering: Concepts and Precepts* (Hollnagel, Woods & Leveson, 2006), there has been considerable interest for what this approach to system safety really means and how resilience engineering can be brought into practice. At times the curiosity has been tempered by reasonable scepticism and doubts whether resilience engineering really is something new. This is, of course, a welcome challenge, and one that fortunately can be answered by an unqualified *Yes* and *No*. It can be answered in the positive because resilience engineering does offer a different approach to system safety as this book, and the previous, demonstrates. And it can be answered in the negative because resilience engineering does not require that methods and techniques that have been developed across industries over several decades must be discarded. Instead, it makes sense to try to retain as many of them as reasonable, with the proviso that they must be looked at anew and therefore possibly used in a way that may differ from what has traditionally been the case.

Another way of elaborating the answer is to note that resilience engineering differs more in the perspective it provides on safety, than in the methods and practical approaches that are used to address real-life problems. Resilience engineering makes it clear that failures and successes are closely related phenomena and not incompatible opposites. Whereas established safety approaches hold that the transition from a safe to an unsafe state is tantamount to the failure of some component or subsystem and therefore focus on what has gone or might go wrong, resilience engineering proposes that:

> ... an unsafe state may arise because system adjustments are insufficient or inappropriate rather than because something fails. In this view failure is the flip side of success, and

therefore a normal phenomenon." (Hollnagel, 2006)

Since both failures and successes are the outcome of normal performance variability, safety cannot be achieved by constraining or eliminating that. Instead, it is necessary to study both successes and failures, and to find ways to reinforce the variability that leads to successes as well as dampen the variability that leads to adverse outcomes. This is relatively easy in the world of reactive safety management, which is concerned with how to respond after something has happened – although it generally looks as failures only. But effective safety management cannot be based on a reactive approach alone. (Nor, in fact, can effective business management.) For control or management to be truly effective, it is necessary also to make corrections or changes in anticipation of what may happen. The problem with that is, of course, that the future always is uncertain, hence that actions taken to prevent – and sometimes to ensure – a possible outcome, are never guaranteed to succeed. In that sense the management of safety and the prevention of risk cannot be done without also taking some risk. For the good of the organisation's survival it is necessary to accept the chance – or risk – that something may happened and to invest efforts (time, money, work, resources) either trying to make it happen (if it is positive), or trying to prevent it from happening (if it is negative).

Because performance variability is both normal and necessary, safety must be achieved by controlling performance variability rather than by constraining it. In agreement with this principle, a resilient system is defined by its ability effectively to adjust its functioning *prior to* or *following* changes and disturbances so that it can continue its functioning after a disruption or a major mishap, and in the presence of continuous stresses. The quality of resilience can be defined more precisely by pointing to following four essential abilities that a system or an organisation must have.

- The ability to respond to various disturbances and to regular and irregular threats. It is not enough to have a ready-made set of responses at hand, since actual situations often differ from what was expected or imagined – with the possible exception of routine

normal operation. The organisation must be able to apply the prepared response such that it matches the current conditions both in terms of needs and in terms of resources. Relative to the three types of threats proposed by Westrum (2006), this is the ability to deal with the regular threats. The responses enable the organisation to cope with the *actual*.

- The ability to monitor flexibly what is going on, including the system's own performance. The flexibility means that the basis for monitoring must be assessed from time to time, to avoid being trapped by routine and habits. The monitoring enables the organisation to cope with that which is, or could become, *critical* in the near term.
- The ability to anticipate disruptions, pressures, and their consequences. This means looking beyond the current situation and the near future, to consider what may happen in the medium- to long-term. In terms of the three types of threats proposed by Westrum (*op. cit.*), this is the ability to deal with the irregular threats, possibly even the unexampled events. The anticipation enables the organisation to cope with the *potential*.
- Finally, the ability to learn from experience. This sounds rather straightforward, but a concrete solution requires careful consideration of which data to learn from, when to learn, and how the learning should show itself in the organisation – as changes to procedures, changes to roles and functions, or changes to the organisation itself. The learning enables the organisation to cope with the *factual*.

Whereas classical safety management mostly looks at the *actual*, resilience engineering equally tries to look at the *factual*, the *critical*, and the *potential*. Another difference is that traditional safety management focuses on the system or organisation as a whole, rather than on safety alone. This is a natural consequence of taking performance variability rather than performance failures as a starting point. Resilience is achieved both by damping variability that may lead to adverse events and by reinforcing variability that may have positive outcomes. An increased availability and reliability of functioning on all levels will therefore not only improve safety but also enhance *control*, hence the

ability to *predict*, *plan*, and *produce*. Just as failures are the flip side of successes, so is safety the flip side of productivity. You can't have one without the other!

About this book

The chapters that follow have been selected to demonstrate the developing practice of resilience engineering. Many of the papers are based on presentations made at the Second Resilience Engineering Symposium that was held 8-10 November 2007 in Juan-les-Pins, France. In each case the authors were requested to elaborate on their presentations, taking the discussions during the symposium into account. (The complete proceedings from this symposium are available for download at http://www.resilience-engineering.org.) In addition, a number of papers were solicited to complement the state of resilience engineering anno 2006. Among these are the last six chapters, which all address the same event, namely the unintended overexposure of a patient during radiation treatment at the Beatson Oncology Centre, Glasgow in January 2006. The common objective of these chapters is not so much to criticise the official investigation of the event, as to demonstrate what can be learned by adopting a resilience engineering perspective.

Good resilience engineering produces a system that can adapt. Resiliency can be built into any system, and it offers a lens to look at critical areas like cybersecurity and operations. Here are some examples: If the system adapts by acquiring servers from a different region when all the servers in its zone fail, it has been successfully engineered. If the system adapts by taking the next best CPU when the cloud provider cancels providing the present CPU, the system has been successfully engineered.Â If a document were to suddenly disappear from a computer in a hard drive crash, that disappearance would be a failure of the system. That failure would strike a user as oddâ€”something that should never occur. Resiliency in systems has become something we all expect. Building resilience will allow you to ride the changing tides of business and keep your calm anchor at the centre. How we feel internally influences the external world we create, so first of all hereâ€™s a brief exercise to help you achieve that calm centre. Start to think about a time in your life when you felt more confident, visualise it just for a moment and connect with the feeling of being there right now.Â Resilient business owners are people who are committed to their lives and their goals, and they have a compelling reason to get out of bed in the morning. Think about what your compelling reason is and carry this thought with you. Having a clear identity and purpose helps you build greater business resilience. In a nutshell, resilience engineering involves incorporating engineering practices to ensure that a system functions properly under the presence of external and internal disturbances and uncertainties. This Special Issue invites professionals, experts, researchers, and practitioners of resilience engineering to share their knowledge and findings that emphasize the efficient, economic, and sustainable development of our community. This Special Issue focuses on the revolutionary engineering concepts, sophisticated engineering practical solutions, and innovative engineering frameworks that help t