CHEMICAL ENGINEERING TRANSACTIONS
VOL. 48, 2016
A publication of AIDIC
The Italian Association of Chemical Engineering
Online at www.aidic.it/cet
Guest Editors: Eddy de Rademaeker, Peter Schmelzer

ISBN 978-88-95608-39-6; **ISSN** 2283-9216
DOI: 10.3303/CET1648155

Issues In Assuring Human Controls in Layers-of-Defences Strategies

Ronald W. McLeod

Ron McLeod Ltd. 16, Waterside Avenue, GLASGOW, G77 6TJ, SCOTLAND, UK
ron@ronmcleod.com

The process industries place great reliance on layers-of-defenses, or barrier thinking, to protect against incidents. Human performance continues to be the single most widely relied on barrier: whether as a control in its own right, or in implementing, inspecting, maintaining and supporting engineered controls. Human error also continues to be a significant threat to the reliability of engineered and organisational defences. Many organisations struggle to know how to ensure the human defences they rely on are as robust as they reasonably can be when layers-of-defences strategies are being developed and implemented. Drawing on real-world incidents, this paper considers why organisations find it so difficult to know how to address human factors in their layers of defences strategies. Organisations can improve the strength of their human defences by being clearer about exactly what it is they expect and intend of human performance in their operations.

1. Introduction

The importance of having in place a number of layers of defences (also variously referred to as “controls”, “barriers”, and “protection layers”) to protect against the risk of major accident hazards is now virtually ubiquitous across industries with the potential for major accidents. Conceptually, the representation of layers-of-defenses is widely thought of in terms of James Reason’s ‘Swiss Cheese’ model (Reason, 2014), where accidents occur when ‘holes’ in protection layers (cheese slices) align.

For industrial operations, international standards, most prominently IEC 61508 (IEC, 2003) and 61511 (IEC, 2010), as well as a variety of sources of guidance on best practice (PSLG, 2010, CCPS, 2015) set out approaches to developing, analysing and validating layers-of-defences strategies. And there is a growing literature on experience and lessons learned with different techniques (see for example, Sklet et al, 2006, Chambers et al, 2009, Lewis and Smith, 2010).

Whichever approach is taken, the behaviour and performance of people continues to be relied on as probably the single most common form of control. One of the main findings of a review of the application of Layers-of-Protection-Analysis (LOPA) to the risk of overspill at fuel storage tanks in the UK, was that: “Human factors appear to dominate a number of initiating event (IE) frequencies and conditional modifier (CM) error probabilities in all the LOPA studies assessed in this work” (Chambers et al, 2009, p.2).

This paper considers some of the issues organisations that wish to formalise their layers-of-defenses strategies face in demonstrating that the human controls they rely on are as robust as they reasonably can be. The paper is based on three premises:

1. An issued layers-of-defenses strategy setting out the controls an organisation intends to have in place to defend against major accidents is a very important statement of intent. And if an organization chooses to rely on human performance as part of its strategy, then it must do everything it reasonably can to assure that human performance can be delivered consistently and reliably.
2. Rather than treating people solely as a factor that can lead to loss by defeating what would otherwise be safe and reliable operations, organisations should recognise that, most of the time, people are the reason operations so often go well despite the upsets and variability that is normal to complex industrial activities. That is recognised in the concept of Safety-II that has been adopted by Eurocontrol, who are responsible for air traffic management across much of European airspace (Eurocontrol, 2013, Hollnagel, 2014).

Please cite this article as: McLeod R., 2016, Issues in assuring human controls in layers-of-defences strategies, Chemical Engineering Transactions, 48, 925-930 DOI:10.3303/CET1648155

3. Because of that, both safety and productivity would benefit by being clearer about the role of people in operations, and what exactly it is organisations need and expect of people for the safe and reliable operation of their assets. And having achieved that clarity, focusing on what needs to be done to enable people to perform to the standards and to the level of consistency and reliability they need and expect, rather than concentrating solely on how people can defeat safety barriers.

To illustrate the importance of being clear about organisational expectations of people, consider the emphasis that organisations running safety critical operations give to the importance of people following procedures. When things go wrong, investigations commonly reach a conclusion that, if only people had followed procedures the incident would not have happened (see for example MMS, 2005). There are a number of implicit expectations underlying such a position: that the organisation has in place all of the procedures it needs; that they are sufficiently specific, accurate, clear and up to date; that they are accessible where and when they are needed; that the people expected to use them have the knowledge, skills and training to know what procedures to use in whatever context; and that they will actually recognise situations where procedures should be followed, identify the correct procedure, and be able to carry them out under the conditions that exist at the time. Such expectations can be far from reality, as many accident investigations have found.

More importantly, the expectation that if only people followed procedures, everything would function safely and reliably can be in direct conflict with the requirement that operations will be manned by competent and experienced people: especially when those people are “specialists” or “experts”. There is a fundamental conflict between the psychological characteristics of someone who is competent and has expertise, and someone who is expected simply to follow procedures. The reality is that competent, experienced people are usually able to perform to high levels, to understand what is happening, to know or work out what needs to happen to avoid or to recover from abnormal or unexpected situations, even when there are no procedures or the procedures that are available are out of date or technically incorrect. Somebody who is competent simply to follow procedures, is not the same as somebody who is competent to bring expert judgement and decision making to a situation or to cope with the many often unexpected variations in the way operations actually proceed on a daily basis. This has been demonstrated many times, perhaps most dramatically by Neil Armstrong in the final moments before he landed the moon landing craft Eagle on the moon in 1969, or by “Sully” Sullenberger when he landed his Airbus A320 aircraft on the Hudson river in 2009.

The point is not to suggest doing away with procedures: years of hard won experience has demonstrated the value of good procedures, checklists and working practices. The point is to emphasise the need to be clear about what is really expected of people, and to recognise the real contribution people make to the reasons why operations so often go right. To acknowledge that upsets, incidents and even disasters are often averted because motivated, competent people have diverged from standard procedures and normal practices. This is the difference between “Work-as-Imagined”, and “Work-as-Done”: “.everyone at the sharp end knows that it is only possible to work by continually adjusting what they do to the situation” (Hollnagel, 2014, p. 40)

A further reason why organisations find it so difficult to anticipate, and are so often surprised by situations where loss of human reliability occurs, is because expectations of how people are likely to behave are commonly based on a lack of understanding of some of the basic truths of human behaviour and performance. Among the most important are; a) human thought and performance is highly situated – they are strongly influenced by the situation or context as the individuals involved experience and believe it to be at the time; and b) it is part of human nature try to find the easier way of doing things, even if it is more risky. There is also a widespread lack of understanding of the operational importance of what psychologists refer to as two styles of thinking: System 1, or “fast” thinking, which is intuitive, rapid, prone to jump to conclusions, and subject to many sources of irrationality or bias. And System 2, or “slow” thinking, which is being slow, careful, evidence-based, doubting and rational. (See Kahneman, 2012 for a comprehensive review of the difference between fast and slow thinking, and of some of the characteristics and biases that can be associated with “slow” thinking). McLeod (2015, 2016), explored the implications of these two styles of thinking to operational risk assessment and decision making across all levels of organisations and demonstrated how they bring insight into why operators have behaved and took the actions they did in the events leading up to major accidents.

2. Requirements for robust controls

A number of criteria need to be met if any proposed defence is to be given credit as a control in a layer-of-defenses strategy. For process systems, IEC 61511 (IEC, 2010) requires not only that any Protection Layer must provide a minimum 100-fold reduction in risk with at least 90% availability, but it must meet four characteristics: i) it must be specific to a single potentially hazardous event (Specificity); ii) it must be independent of other protection layers (Independence); iii) it can be counted on to do what it was designed to do (Dependability); and, iv) it is capable of being audited (Auditability).

Other bodies have identified similar requirements. In guidance to UK industry following the fire and explosion at the Buncefield fuel storage site in 2005, the UK's Process Safety Leadership Group (PSLG, 2009) requires only three criteria: that a valid protection layer needs to be Independent, Effective (which is essentially equivalent to Dependability) and Auditable. And the Center for Chemical Process Safety (CCPS, 2015), identifies seven "core attributes" that any Independent Protection Layer (IPL) included in a Layers-of-Protection-Analysis (LOPA) should have: Independence; Functionality; Integrity; Reliability; Auditability; Access Security and Management of Change. (For the present purpose, Functionality, Integrity and Reliability can all be considered as aspects of Effectiveness).

Some organisations have historically been reluctant to take credit for the risk reduction that can be achieved through a reliance on human performance when performing LOPA (Myers, 2013). Some still do not allow any control that relies on human performance for its operation to be considered an IPL at all. In common with other approaches, Myers' demonstrates how the potential for human failure can be quantified using Human Reliability Analysis (HRA) techniques, including taking account of aspects of the context of the expected human action, such as the available time. So one approach to demonstrating the robustness of human controls is to try to estimate the effectiveness of human performance by applying quantitative analysis of the likelihood of human error. While HRA is very widely used, there are concerns both about the validity of the data on which such estimates are based as well as the ways in which the results are used (see French et al (2009), for a review of concerns with HRA). An alternative, complementary approach is to challenge proposed controls and to be clearer about what exactly is intended and what is expected in order for controls that rely on human performance to have as high a likelihood of being effective as can reasonably be achieved.

McLeod (2015) has discussed in some depth what the requirements of Independence, Effectiveness and Assurance mean when they are applied to controls that rely on human performance. Based on application of the principles of Human Factors Engineering, and drawing on a detailed analysis of the investigation into the Buncefield explosion in 2005, McLeod explores the Human Factors challenges involved in meeting each of these requirements and demonstrates how each of them can be, and in the case of the Buncefield incident in fact were, defeated by real-world events.

Specificity

The requirement for controls to be specific to a single potentially hazardous event can be especially relevant to controls that rely on human performance. Indeed, because of the situational nature of human performance, it is only by being as specific as possible, identifying not only the judgments, decisions and actions that people are expected to take for a control to be effective, but putting them into a meaningful context that captures the individuals likely beliefs and expectations at the time they are expected to perform that the true robustness of human controls can really be tested.

LOPA analysis is scenario-based. It builds on a higher level review of operations and focuses on specific situations – known as "Initiating Events" (IEs) - that could lead to undesired consequences. Often, IEs are human actions or omissions. So in a LOPA analysis, treating human action or inaction as an IE is valid. Though as has been discussed earlier, it is not only a very limited way of looking at the role of people in operations – both when things go wrong and, much more frequently, when they go right – but also does not capture the situated nature of human performance.

In Bowtie Analysis however (Lewis et al, 2010, provide an introduction to Bowtie analysis and discuss lessons they have earned from applying the method) the opposite is the case – human performance cannot be either a Top Event or a Threat in Bowtie terms. This is widely misunderstood, and has led some organisations to invest effort into trying to build "human error bow-ties", where human error is modeled as either a specific threat or as a top event. To do so is to misunderstand the nature of human behavior and performance. Human error is not something that exists fully formed in the world, just waiting for an opportunity to strike. It is by nature situational. And, most usually, an "error" is only considered as such in hindsight, when judgements are being made about the consequences of human performance. In principle, and provided it met the necessary criteria, human performance could be counted as a control or a recovery measure, or as a control against their failure (i.e. a control against degradation factors). Though when human performance is being viewed as a risk, it can and should only be modelled as a degradation factor: it is something that in a specific situation can cause controls or recovery measures (and controls against their defeat by degradation factors) not to perform as intended or as expected.

Independence

Achieving true independence of controls in terms of their reliance on human performance can be a major challenge. Controls frequently rely on the same individual or team, so any factor – workload, fatigue, distraction, lack of training, etc. – that defeats one control will often have the potential to defeat others controls in the same threat line. And even when controls are assigned to different individuals, organisational factors – leadership messages emphasizing production over safety, rewarding individuals for delivery even when

corners have been cut, poorly thought out incentive schemes that reward unsafe behavior in annual bonus schemes or the way contractors are incentivized in their contract arrangements – can all lead to the defeat of many controls that depend on people. The independence that can be achieved by relying on other people to cross-check somebody's work may not be as effective as is widely assumed. This was recognised as far back as 1983 in the classic work by Swain and Guttman that has provided the basis for most approaches to quantifying human reliability since (Swain and Guttman, 1983) and has also been reiterated by the UK Process Safety Leadership Group (PSLG, 2009). Unfortunately, those and similar warnings are frequently over-looked when decisions are made about how to assure human performance in operations.

Effectiveness

Making judgements about whether any control is likely to be Effective (i.e. that each control, on its own, should be capable of preventing an event from leading to an undesirable consequence) needs a great deal more information than is typically produced when a Process Hazard Analysis, HAZOP, Bowtie Analysis, LOPA or other form of risk analysis to support a layers-of-defenses strategy is carried out. In particular, to make judgements about the likely effectiveness of controls that rely on human performance, means being clear about exactly what is *intended*, and what is *expected* of human performance for the control to be effective.

Intentions are things that can reasonably be expected to be within the scope of influence of the team that assesses the risks and develops and approves a set of controls to mitigate those risks. Intentions will often relate to aspects of the design of the work environment or equipment interfaces. A control that relies on someone opening or closing a valve brings with it the clear intention not only that the individual will be act on the right valve (or even the right aircraft, see AAIB, 2015), but that the human interface to the valve will be designed and labelled such that the chances of design-induced human error are reduced to ALARP. Similarly, relying on an alarm as a control brings with it numerous implicit intentions about the design quality of that alarm and the environment in which it will be presented.

Expectations, by contrast, can be defined as things the analysis team cannot usually be expected to be responsible for, but that it needs to assume will be true for a human performance based control to be effective. McLeod (2015) has demonstrated how the challenges of Intentions and Expectations can be applied to four types of controls related to a tank filling operation: having a plan for a fuel transfer; pro-active operator monitoring; control room alarms; and a fully automatic, independent high-level shut-off.

To demonstrate some of the insight that can come from examining the implicit intentions and expectations about human performance that underpin assumptions about control effectiveness, the remainder of this paper will consider two incidents that occurred in process systems.

3. Example Incidents

Incident A: Over-pressurised pig launcher

The incident: A team was preparing for a pipeline inspection using an in-line inspection tool (known as a 'pig'). The team believed that the pipeline valves were open and began pumping nitrogen from a truck to purge the line. However, the valve between the pig-launcher and the pipeline was actually closed preventing nitrogen from entering the pipeline. The truck included a pressure trip set at 6000 psi although the Maximum Allowable Working Pressure (MAWP) of the pig launcher was 350 psi (i.e. the truck was capable of supplying nitrogen at a much high pressure than the launcher was designed to withstand). The pig launcher was not equipped with a pressure relief valve. When pressure was applied, the 100 psi gauge on the pig launcher almost instantaneously swung to the zero position. The team at the pig launcher mistakenly read the gauge as indicating there was no nitrogen flowing from the nitrogen truck and called for more pressure. The pressure release happened within two minutes of the call from increased nitrogen flow.

So what kind of controls might have been in place that should have prevented the pig launcher from being over-pressurised? The obvious ones would include a pressure trip and relief valves. More importantly, what sort of expectations might the organization involved reasonably have held about those controls? Here are two suggestions that can reasonably be assumed from the incident report:

- The job would have been planned and a safety review conducted before starting work;
- Operators would be aware of the risks and exercise caution. If they had any doubt they would stop.

It must have been assumed that holding a safety review prior to starting work would identify the risks, and ensure everyone involved knew what was involved in ensuring they were controlled. Safety reviews prior to starting work are widely relied across many industries. A safety review had indeed been carried out (the control was in place), but it was not effective in recognising or mitigating the over-pressurisation risk. So why did it fail in this case? Was there something unusual about the way the review was held on that day, or the engagement of the people involved at the time? Are safety reviews usually effective as a means of raising awareness and ensuring risks are under control? Or was this only one incident among many where such

safety reviews have failed to identify and provide the expected control over significant risks? Are they, in fact, as effective as is widely assumed?

Or consider the expectation that the operators would be aware of the risks associated with pig launchers, will exercise caution, and will stop the job if they are in any doubt. The evidence suggested that this expectation was not met. The individual(s) who read the pressure gauge, concluded that there was no pressure in the pig launcher and called for the flow of nitrogen to be increased cannot have been in any doubt. They did not mis-believe the pressure gauge: they believed it was reading zero flow, when in reality the pressure was already too high for it to provide a reading. But there can have been no doubt involved, and therefore no need to exercise caution and stop and question what they were about to do. It has the characteristics of thinking and decision making dominated by what psychologists refer to as System 1 thinking.

Incident B: Pipeline corrosion

Organisations have to believe that at the time a risk analysis is performed their analysis teams are capable of identifying the ways operations can be put at risk as a consequence of human performance - which can be many years before an event actually takes place. Because of the situational nature of human performance, and the broad range of factors that influence how any individual perceives, understands and experiences the situation they are in at the time they assess a situation, make decisions and act, that belief can be extremely optimistic. The following incident illustrates how some of the “hard truths of human performance” (McLeod, 2015) can play out in the real world to defeat unrealistic expectations about how people will behave and perform in the real world.

The incident: A fire occurred when a sample supply line was breached due to corrosion in a carbon steel line allowing release of hydrocarbons to atmosphere under high pressure and at high temperature. A change in sampling procedures had been introduced that required the valves on the sample line to be accessed more frequently than had been anticipated during design. The valves were located up to 15 feet above grade and were known to be difficult to operate. With the increased sampling frequency, the practice became to leave the valves open, avoiding difficulties operating the valves. As a result, the piping became more exposed to the corrosive environment in the process stream, accelerating the rate of corrosion and leading to the breach.

At the heart of this incident was an operational change that increased the frequency of operating a valve that was both inaccessible and difficult to work with. That change led operators to adopt a working practice that defeated the premises of the defences-in-depth strategy. Allow, for the purpose of this discussion, that access to the valves had been optimized during design in accordance with a Valve Criticality Analysis reflecting the anticipated frequency of use and criticality of the valves. One of the hard truths of human performance is that if working life is made unnecessarily difficult, people will find the easy way to do things (even if it is more risky) (see McLeod, 2015). Would a PHA, HAZOP or LOPA analysis be capable of recognizing that a change in sampling policy could trigger that hard truth, leading operators to adopt a working practice that defeated key assumptions behind the controls? Would a change in sampling policy be subject to a formal Management of Change procedure? One that not only recognized the implications for operators working with inaccessible and difficult to use valves, but that considered the possible implications on human behavior if the change made work more difficult? As Myers has pointed out: “..unintended or poorly managed changes may be a more important consideration than those due to malicious intent”, (Myers ,2013, p.535).

The piping design would certainly have allowed for the anticipated rates of corrosion based on the anticipated exposure to the process materials. And the corrosion inspection plan would have been based on the anticipated exposure rate, assuming that the valves would only be open during sampling: a significantly lower exposure to the process stream than had been introduced.

Even for a relatively simple incident such as this, there are major challenges facing any team tasked with identifying the human performance issue that could contribute to this incident *a priori*. And it is clear that a simple bowtie or LOPA analysis cannot capture what the organization really would have intended and expected should have prevented this incident. While it is easy, with hindsight, to see how the holes in the Swiss cheese lined up in this situation, that is very far from the case when trying to do an analysis proactively, and to identify the likely situations that could realistically give rise to a threat scenario. Indeed, it might be thought unrealistic to have expected any analysis team to be capable of coming up with a pro-active analysis that actually reflected what did in reality happen.

The question is whether it is reasonable to expect any risk analysis team to imagine how the combination of events that actually did occur in this incident might coincide. And, even if they had the imagination, is it likely that they would give it credibility as being a realistic combination of events that needed controls in place? It is far more likely that operators would simply be assumed to be competent, and expected to follow procedures.

4. Conclusions

Some people find it difficult to accept the argument that industry needs to treat the “hard-truths” of human behavior and performance seriously in the decisions they make and the actions they take. They believe people who are trained and competent, are assessed as being fit for work, properly motivated and working in a culture that places a high value on safety and compliance, should in some way be capable of overcoming not only poorly designed working environments and procedures, but very powerful instincts of human nature. Though the reality is that we are all human, and are all subject to these powerful motivations.

As the Nobel prize winning psychologist Daniel Kahneman puts it, “..people will eventually gravitate to the least demanding course of action…laziness is built deep into our nature” (Kahneman, 2012, p. 35). That is not to suggest, when incidents do occur, that people should not be held responsible for violating what are clear expectations and procedures. But it does mean that, when layers-of-defenses strategies are being developed, they should be subject to robust challenge about whether the expectations of human behavior and performance that are relied on for those barriers to achieve an acceptable level of effectiveness and reliability are realistic and reasonable. And when changes to operational practices are being implemented, the potential consequences on human behaviour need to be considered carefully.

An issued layers of defenses strategy, in whatever form it takes, is a significant statement of intent on behalf of the organization that prepares and approves it. Organisations can make improvement in assuring the human controls they rely on for safety and production are as robust as they reasonably can be by being clearer about exactly what is the role of people in their systems. When controls that rely on human performance are relied on, they should be subject to robust challenge about what exactly is intended and what is expected for those controls to do what is needed of them.

References

AAIB., 2015, Report on the accident to Airbus A39-131, G-EUOE London Heathrow Airport 24 May 2013. Aircraft Accident Report 1/2015, Air Accidents Investigation Branch.

CCPS., 2015, Guidelines for initiating events and independent protection layers in layer of protection analysis. Center for Chemical Process Safety, Wiley & sons, New Jersey.

Chambers C., Willday J., Turner S., 2009, A review of layers-of-protection analysis (LOPA) of overfill of fuel storage sites. HSE Books.

Eurocontrol., 2013, From Safety-I to Safety–II: A White Paper. Eurocontrol. http://www.skybrary.aero/bookshelf/books/2437.pdf Accessed 11.09.15.

French S., Bedford T., Pollard S., Soane EC., 2011, Human Reliability Analysis: A Critique and Review for Managers. Safety Science, 49, 753–763.

Hollnagel E., 2014, Safety-I and Safety-II: The Past and Future of Safety Management. Ashgate, Farnham.

IEC., 2003, Functional safety of electrical/electronic/programmable electronic safety-related systems. IEC 61508. International Electrotechnical Commission.

IEC., 2010, Functional safety – Safety Instrumented systems for the process industry sector. IEC 61511, International Electrotechnical Commission.

Kahneman D., 2012, Thinking, Fast and Slow. Allen Lane, London.

Lewis S., Smith., K., 2010, Lessons learned from real world application of the bow-tie method. American Institute of Chemical Engineers 6th Global Congress on Process Safety, San Antonio, Texas.

McLeod R.W., 2015, Designing for Human Reliability: Human Factors Engineering for the Oil, Gas and Process Industries. Gulf Professional Publishing, Oxford.

McLeod, R. W. 2016, The impact of styles of thinking and cognitive bias on how people assess risk and make real world decisions. Oil and Gas Facilities (In Press). SPE Paper no: PFC-1115-0010.

MMS., 2005, Human Engineering Factors Result in Increasing Number of Riser Disconnects. Safety Alert No. 231. US Department of the Interior Minerals Management Service, Gulf of Mexico OCS Region.

Myers P.M., 2013, Layer of protection analysis – Quantifying human performance in initiating evcents and independent protection layers. J. Loss Prevention in the Process Industries, 26, 534-546.

PSLG., 2009, Safety and environmental standards for fuel storage sites. Process Safety Leadership Group, HSE Books.

Swain, A.D., Guttman, H.E., 1983, Handbook of human reliability analysis with emphasis on nuclear power plant applications. Final Report. NUREG/CR-1278, USNRC