



WHAT IS A RESILIENT SAFETY CULTURE

David G Broadbent
Safety Psychologist
Alliance House
12 Ken Tubman Drive
MAITLAND, NSW, 2320
AUSTRALIA

What is the first thing that happens after something bad happens in our workplaces?

Well, hopefully, the very first thing is that we focus on providing care and compassion to anybody who might have been affected by the “bad thing”, whatever it might be?

Once that has been achieved, then very soon after we find ourselves embarking on a process of investigation, in order to try and locate/allocate the contributory factor/s that we believe acted as primary drivers toward what has just happened?

In days past (and not that long ago), there would be a process to try and find a “root cause” or “route causes”. This is unfortunate, as this sort of language suggests that we are going to locate one, two, or three causes that resulted in the system failure (accident/ incident/disaster) – it is rarely that simple. If we did not go chasing the route cause, then we might go hunting for the “unsafe act” or the “unsafe condition”. This hunt is predicated on the work of Heinrich and his Domino Theory of Accident Causation – again truly unfortunate, as this is amongst the most simplistic means of thinking about system failure. Its very nature is about “ignoring” the system and trying to locate a single, or couple of factors, that led to the negative outcome. You might be seeing similarities with what has become known as behavioural based safety (BBS). More recently BBS is becoming a little less popular – thank goodness. The “new” focus is often around the language of “human error”. The human error guys are a little more enlightened in that they are, at least, willing to consider a less mechanistic understanding of causation. Nevertheless their focus still appears to be on “failure”, rather than “success”.

In my view, whilst the focus on human error and all the elements that may contribute to that error are important, it is very easy to be seduced into believing, that if we locate the error (and control for it) then the job is done. Nothing could be further from the truth. It certainly makes us feel “better” if we can say that we have “understood” what has happened, controlled for it, and it shall never happen again. Yet how often have we found that despite these best of intentions, the same sort of “stuff” just keeps happening? Indeed, much of the research evidence concerned with accident causation has found itself pointing in a somewhat different direction entirely. In almost every accident/disaster I have reviewed, the evidence seems to point more toward macro elements, rather than the more comfortable micro explanation of some poor bod or bods making an error (mistake).

There are a couple of guys who have confronted this type of question quite directly, and it is worth noting what they had to say:-

“You respect people, you listen to them, you work together. You don’t blame them. Maybe the **process** was not set up so well, so it was easy to make a mistake” – Gary Conviss, President, Toyota Motor Manufacturing.

“Human error is inevitable. We can never eliminate it. We can eliminate problems in the **system** that make it more likely to happen” – Liam Donaldson, WHO, World Health Alliance for Patient Safety.

We are seeing an increasing number of consulting organisations picking up on this intense focus on human error and designing interventions that are arguably allowing organisations to locate those error prone individuals – this should be very scary.

An exceptionally gifted workplace commentator, W Edwards Deming, referred to the performance “lottery”. His argument was that you could put almost any person, in an identical situation and the same error would be replicated. In other words the error was actually predicated by the system and not the poor bod who just happened to be in the wrong place at the wrong time. Throws the “error prone” individual theory out the window. Just another example of the snake oil that we have to be wary of in the safety space?



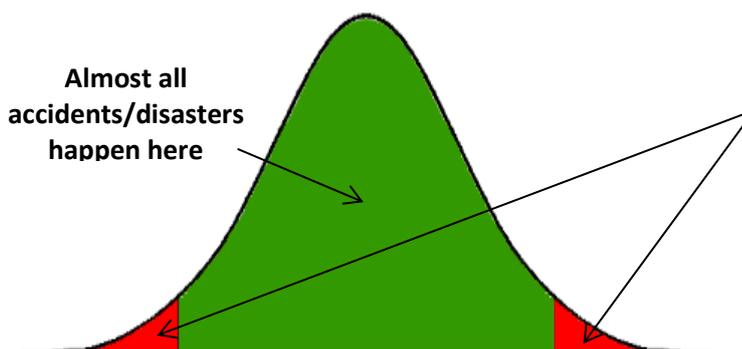
How often do we see those people at the pointy end of the operation being reprimanded, publicly vilified within the organization, even terminated from their position, all due to the making of an error that could have/would have happened to anybody. In contrast how often do we see the senior managers treated in a similar manner for their role in creating (or tolerating) a bad system. I would suggest there is a significant imbalance within this circumstance.

What is it then that sets the environment for these error states to actually come into being? If we were to take a robust look at some of the more recent workplace disasters the conclusions are remarkably similar. Despite the aforementioned focus on the error equation, the investigative outcomes almost universally point to something different, that’s if we are serious enough to look further back than the first or second system deviation? They point to organisational factors that have been key drivers for all that “bad stuff” happening – yes they are frequently organisational factors that set the “tone” for the human error type failure.

It is my view that pretty much all accidents are “organisational accidents”. Regrettably, for most operations (irrespective of industry space) their journey toward finding themselves along this road is almost universally consistent.

Some look at the number of high frequency/low consequence events and become rabidly focused on metrics around LTIFR etc. Of even greater concern are those businesses who might point to the absence of any containment failures as a measure of their “success”. We might call these low frequency/high consequence events. In other words they didn’t “blow up”. Please believe me when I tell you that there is generally no predictive value in either of these approaches – yet we invest so many of our resources chasing these sacred cows. If I was to use the statistical

example of a Bell Curve I might agree that at the extreme tails of the curve there may be some predictive associations. The reality though is when we take a look at almost all the disasters that we know about – those organisations were not operating toward the tail of the curve, they were invariably operating right there within the green (average).



In other words, they were not some blatantly irresponsible organization. They were not some hierarchy full of psychopathic leaders. They were invariably just your run-of-the-mill organization that had managed to go a bit astray – and more importantly, not notice. It is fair to say that I am often asking myself, how could they not notice? Sadly, there are many answers to that question. Some of them might even be considered criminal by some – I am one of those.

What can we learn from disaster?



You may recall a recent train crash near Santiago de Compostela in Spain. Seventy nine (79) people lost their lives and a further one hundred and forty (140) suffered a wide range of injuries. An investigation was conducted and the key outcome was that the driver was charged with seventy nine (79) counts of homicide. There has been no argument presented with respect of the fact that the train was travelling at over twice the mandated speed for that particular curve. It has become known that the driver was not the only Spanish train driver who travelled at excessive speed. He was not the only Spanish train driver that posted photographs of a train speedometer on social media. In other words there is a strong case to be made that it was “known” that the Spanish high speed train network was even quicker than advertised. Does that mean that the safety culture operating within Renfe (the trains operator) was significantly different to any other operator of high speed trains. There is no evidence to suggest this was the case. Indeed, the current evidence was that Renfe was operating within the green zone presented previously. There was no significant data from either LTIFR style metrics, or “containment failure” history, to give any indication that Renfe was due a major train disaster.



We can then take a brief look at another horrific disaster. This one occurred in Turkey, at a town called Soma on May 13th 2014. On this occasion three hundred and one (301) workers lost their lives in the one accident. It seems that the “authorities” are not even able to agree on how many actually did lose their lives. Whilst the “official” figure might be as indicated above, there is also serious discussion that this figure could be closer to three hundred and forty (340). What message does that send? Now here’s a thing. The miners from Soma protested about the conditions of the local mines only several months previous. It drew enough attention that the Turkish opposition “demanded” that the safety of mining operations around Soma be investigated. The government of the day rejected that demand only days¹ prior to this disaster. Yet, it has been reported by the Turks themselves, that they lead the world in mining deaths per tonne of coal². That’s right Turkey, not China. That has to be scary.

Now the operator of the Soma mine in which three hundred and something miners lost their lives had not been directly identified as a place where something bad was

¹ April 29th 2014

² Arslanhan, S. & Cünedioğlu H. E. (2010), Turkish Economy Policies Research Foundation (TEPAV), Istanbul, Turkey

about to happen. Whilst we might find that a tad surprising, it seems that the operator was doing so, amongst a toxic culture of safety within and around them. Even despite the very evident (at least to us) markers of disaster that we might believe we can see, no one had predicted this was about to happen. Three hundred and something miners still went to work that day fully expecting to come home to their families at end of shift.

In short, we are organizationally very poor at being able to predict when this bad stuff is going to happen. Even when, as an organization, it is staring us in the face. I would go so far as to say we just can't, won't, or don't do it at all.

So I am hoping that we are seeing some greater clarity arising from the murky depths of disaster?

The place of Safety Culture in all of this?

Even though we might prefer to dodge the bullet, the question must be confronted at some point. What is it that seems to set some industries and/or organisations apart from the norm (average) when it comes to safety performance? It clearly is not an ability to predict when the "bad stuff" is going to happen. It is though, developing an intrinsic ability to do "stuff" in a better and different way.

Everything points to the prevailing existence of an effective and optimal "safety culture", and how it "lives" within an organisational structure. If we take an even wider brush to the image we are painting, then we also need to acknowledge that the safety cultures operating here can also be held hostage to the many socio-political and geo-political environments in which they function. The second example we have just explored makes that so powerfully, and distressingly, evident.



So we can no longer bury our heads in the sand. The evidence is clear. The safety culture that exists within our own operations, and our wider communities, is a critical factor for maintaining the "good" and minimizing the likelihood of confronting the "bad stuff".

Senior managers within organisations shall often seek "guarantees" before they invest in what some might consider right brain intangible programs. Yet the return on investment (ROI) has been shown over and over again. If I bring us back to our look at disaster, what I can say is that in the many disaster investigation reports I have reviewed, I have never read a conclusion that said anything similar to:-

".....despite the presence of an excellent safety culture Aardvark Energy exploded
.....".

What I have consistently read is words to the effect of,

"..... Aardvark Energy showed little regard to maintain a culture of safety and this is felt to be a direct contributor to the events experienced and the lives lost.....".

I shall just draw attention to one more recent disaster, and the clear conclusions that were drawn from the investigation process. This is also from the world of rail. On July 6th 2013 a freight train actually became a "runaway". Seventy two (72) tanker cars loaded with crude oil crashed and exploded into flames. This occurred in the town of Lac-Mégantic (Quebec, Canada). The explosion and associated fires killed forty seven (47) people, destroyed fifty (50) buildings and burnt almost sixty (60) vehicles.



The Transportation Safety Board of Canada (TSBC) commenced an intensive investigation into the likely causes of this disaster. The final report revealed eighteen (18) different factors, including inconsistent oversight from the Canadian government, led to the accident.

Of direct relevance to this discussion were conclusions that bluntly described a lax corporate “culture of safety” at the Montreal, Maine and Atlantic Railway (MMA) as a primary contributing factor to the disaster.

Here is a highly relevant section drawn directly from the Investigative Report into this disaster.

The strength of an organisation’s *safety culture* starts at the top, and is characterized by proactive measures to eliminate or mitigate operational risks. MMA was generally reactive in addressing safety issues. Furthermore, there were significant gaps between MMA’s operating instructions and how work was actually conducted in day-to-day operations.

There were also other signs in MMA’s operations that were indicative of a weak organisational safety culture, such as:

- 🌻 MMA management’s acceptance of rail wear on the main track that was well beyond industry norms and their own track standards;
- 🌻 MMA managements’ tolerance of non-standard repairs (for example, to the locomotive engine and the QRB valve), which either subsequently failed, or did not return the parts to their proper operating condition;
- 🌻 The systemic practice of leaving unattended trains on the main track, and on a descending grade, at Nantes for several hours without in-depth defences to prevent an uncontrolled movement;
- 🌻 Crews and single-person train operators not always correctly applying CROR 112 and MMA’s instructions when securing trains at Nantes;
- 🌻 Inadequate company oversight to ensure the correct securement of trains at Nantes;
- 🌻 MMA’s inadequate recertification program and SPTO training that did not ensure that operating crews knew and understood the procedures for train securement;
- 🌻 The giving of extensions for competency cards by MMA management, in some cases for several months beyond the mandatory limit of 3 years;
- 🌻 The fact that only local corrective action resulted from recurring deficiencies identified during TC inspections of MMA track and operations; the systemic issues contributing to these deficiencies were not fully analysed by MMA, and thus persisted.

If instructions or rules are disregarded, and unsafe conditions and practices are allowed to persist, this leads to an increased acceptance of such situations. Deviations from the norm thus become the norm, and the likelihood of unsafe practices being reported and addressed is reduced. Although educational material about safety culture was provided to railway companies, safety culture was not formally assessed or documented within regulatory inspections or audits. MMA’s weak safety culture contributed to the continuation of unsafe conditions and practices, and compromised MMA’s ability to effectively manage safety.

An even more relevant conclusion from this Report³ is:-

“.....The cornerstone of a truly functioning Safety Management System (SMS) is an effective safety culture”, and notes that “an effective safety culture is one where past experience is not taken as a guarantee of future success and organisations are designed to be resilient in the face of unplanned events.”

³ Ross D et al, (2014), “Railway Investigation Report - R13D005”, Transportation Safety Board of Canada, 200 Promenade du Portage, Place du Centre, 4th floor, Gatineau, Quebec, Canada

The key take-away from this highly robust and detailed report was the acknowledgment that an optimal safety culture, within an organization, needs to be “resilient in the face of unplanned events”. In other words there is limited value in having what might be called a “fair weather” safety culture, something that seems to do well when nothing is happening. The real challenge is that when the seas get rough, the SS Safety Culture stays afloat. To do that the boat not only has to float, it has to be resilient enough to stay afloat, when it has rough weather coming at it from all sides.

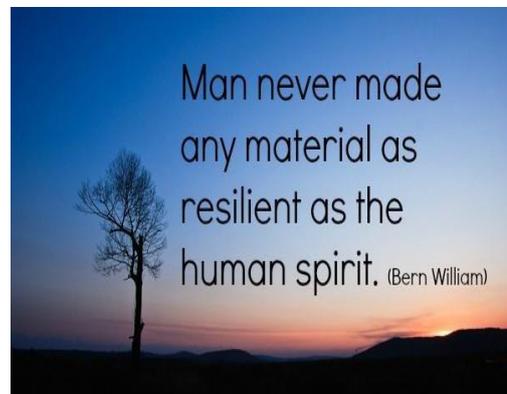


It is actually only relatively recently that there has been any real attention put toward the question of resilience and its application to safety culture.

So what is Resilience?

Dealing with change or loss is an inevitable part of life, none of us can avoid this – as much as some of us might try. At one point or another each of us experiences varying degrees of “bad stuff”. Some of this “stuff” might be relatively minor (missing a meeting due to long delays on the transport system), while others can be personally disastrous on a much greater scale (floods, earthquakes, relationship breakdown, etc.). How we deal with each with these sorts of challenges often plays a pivotal role in not only the outcome, but also the long-term health of ourselves, and those around us.

Have you ever given some thought as to why some people seem to remain calm in the face of disaster, while others seem to trip over themselves? Those people that are able to keep their cool have what psychologists call “resilience”, or an enhanced ability to deal (cope) with problems and challenges that get in their way. Resilient people are able to draw upon skills and strengths to cope and recover from problems and challenges, which may include job loss, financial problems, illness, natural disasters, medical emergencies, divorce or the death of a loved one.



Those who lack this resilience may end up becoming damaged by such experiences. They may end up ruminating on their “unfair” circumstances, and use unhealthy coping mechanisms to deal with such challenges e.g. alcohol, shopping, eating, gambling etc. Generally, these individuals are significantly slower to recover from these challenges, and may end up experiencing more long term psychological dysfunction as a consequence. In the world of clinical psychology the provision of resiliency based coping strategies has become increasingly utilized within treatment planning.

Resilience does not eliminate life's challenges. Instead, it gives people the strategies and stamina to tackle problems head on, overcome adversity and move on with their lives.

Even in the face of events that seem utterly unimaginable to many, some people are able to marshal the strength to not only survive, but to prosper. It is these examples of resilience that many of us are in awe of. We often hear people who have witnessed or heard of such resilience say things like “aren’t they incredible, they must have been born that way”. Not true, by the way. Some people may seem naturally resilient, though when you look closer they have often been exposed, in their formative years, to circumstances and support systems that have encouraged the development of resilience. Certainly the evidence from the world of clinical psychology is that resilience is actually learnable.

What about Resilience and Safety Culture?

It is also pretty clear that some organisations/industries appear to handle the “bad stuff” better than others.

It just so happens that there has been some quite recent quantum leaps with regard to how we might think about safety. The two most well-known are associated with the concepts of “resilience” and “reliability”.

If the definition of Resilience is an ability to deal (cope) with life’s challenges, then it stands to reason that the same definition might be applied to our organization/s. After all, whilst an organization might well be a living entity in itself – is it not also made up of many other living entities (us)? Thus when we are aiming to have a “resilient safety culture”, our target is an organization that is able to be adaptive in the face of significant challenge/s. Not unlike the outcomes we strive for in the individual, we shall be seeking an ability to recognize the “bad stuff” as we are being confronted by it, and apply those self-same adaptive techniques in such a way as to ameliorate the “bad stuff” altogether. If that is not possible, and let’s face it, sometimes it just sneaks up on us, then those self- same strategies may well support our organisations (and ourselves) to come through the “bad stuff” in a more adaptive, healthier way.

Now the positive influence of developing a “resilient safety culture” has also been clearly expressed by the Western Australian Department of Mines & Petroleum⁴. In their Safety Culture Spectrum document they make the unequivocal statement “all operations should aspire to be resilient”. Absolutely spot on.

This whole area of resilience, and how it can add enormous value, has received lots of serious attention over the past several years. To say that resilience, and its application to our safety world is relatively recent, is probably a bit of an understatement. The key event was the International Symposium on Resilience Engineering held in Sweden in late 2004. The result was unanimous agreement that this was a hugely promising field of endeavour, which could add enormous value to the field of organisational performance. A book was subsequently published in 2006 that captured some of the initial ideas about the importance of resilience – that book has become a best-seller.

That subject is what has become known as Resilience Engineering (RE). Whilst we see the word resilience here, there are applications that go way beyond the earlier exploration of individual resilience. Whilst we also noted previously that organisational resilience is aimed at assisting an organization recover from all that “bad stuff”, Resilience Engineering adds something extra. It adds a level of understanding around “tolerance” of the system. What this means is “how much can a safety culture ‘bend’ before it breaks”. Other words that might seem appropriate could be “flexible”, and/or

⁴ 2013, Safety Culture Spectrum: Resources Safety Matters, Vol 1. No 1, p.27

“adaptive”. The most prolific author in this space is a Swedish guy named Eric Hollnagel.

In his first book⁵ on the subject he stated:-

"The essence of resilience is therefore the intrinsic ability of an organization (system) to maintain or regain a dynamically stable state, which allows it to continue operations after a major mishap and/or in the presence of a continuous stress."

When it comes to safe operations within an organization we could not ask for more.

A definition of a resilient safety culture can be made clearer by giving some real thought as to what makes resilient performance possible. Since resilient performance is theoretically possible for most, if not all, systems, the answer must therefore be independent of any specific industry/profession etc. It must actually be more of a “state of mind”, or in other words “how we do things around here” – I hope you might have heard that somewhere before?

Resilience engineering has therefore proposed the following four (4) basic abilities:

- 🌟 The ability to respond. Knowing what to do, or being able to respond to regular and irregular changes, disturbances, and opportunities by activating prepared actions or by adjusting a current mode of functioning.
- 🌟 The ability to monitor. Knowing what to look for, or being able to monitor that which is or could seriously affect the system’s performance in the near term – positively or negatively. The monitoring must cover the system’s own performance as well as what happens in the environment.
- 🌟 The ability to learn. Knowing what has happened, or being able to learn from experience, in particular to learn the *right* lessons from the *right* experience.
- 🌟 The ability to anticipate. Knowing what to expect, or being able to anticipate developments further into the future, such as potential disruptions, novel demands or constraints, new opportunities, or changing operating conditions.

These four (4) basic abilities as posited by the Resilience Engineering community go a long way to allowing us to create/evolve resilient safety cultures.

Resilient Safety Cultures and High Reliability Organizing.

We cannot have a serious conversation about resilient safety cultures and not pay attention to the seminal work of Karl Weick and Kathleen Sutcliffe⁶.

What we find when we look at the influence of High Reliability Organizing (HRO) on the safety equation, is a very intuitive system that when responsibly applied cannot help but make our systems safer, and in a sustained way – leading to a “better way of doing things around here”.

HRO consists of five (5) hallmarks. These are those areas of focus that, if applied, shall almost osmotically contribute to a sustained and resilient safety culture. The five (5) hallmarks of HRO are:

1. Preoccupation with Failure

Everyone is focused on errors and near-misses, learning from them and figuring out how to prevent them from happening again. In a resilient safety culture that embraces RE and HRO we see error as a *fait accompli*. It is always present, so attention to detail is crucial. Locating potential systems failures is everyone’s responsibility and is actively encouraged and supported by leadership.

⁵ Hollnagel, E., Woods, D. D., & Leveson, N. (Eds.). (2006). “Resilience Engineering: Concepts and Precepts”, Burlington, VT: Ashgate.

⁶ Weick, K.E. & Sutcliffe, K.M. (2007). “Managing the Unexpected: Resilient Performance in an Age of Uncertainty”, Jossey Bass:Wiley, New York.

2. Reluctance to Simplify Interpretations

Requires constantly asking the “why” question and inviting others with diverse experience to express their own opinions. The underlying belief here is that the more you’re immersed in something, the harder it is for you to objectively observe and question things that need questioning. In the world of safety we might sometimes refer to this as risk blindness. Whilst there is nothing intrinsically wrong with trying to make a process “simpler”, there is an enormous risk that in doing so you remove some of the “redundancies” that have been placed within that system. This has become a more genuine threat as the thinking behind “lean manufacturing” begins to impact safety thinking. So the position is to always be “reluctant”. Does not mean you don’t do it, just means your only do so in a very thorough and informed manner.

3. Sensitivity to Operations (an HRO Distinguishing Characteristic)

Everyone needs to be constantly aware of how processes and systems affect the organization. In a resilient safety culture, each and every employee pays close attention to operations, and maintains an acute awareness as to what is or isn’t working. There are no assumptions here.

This steady concentration on processes leads to observations that inform decision-making and new operational initiatives. Most people just “know” when the temperature has fallen by as little as one degree – it just feels “different”. Whilst there is a sensitivity borne of experience and training, there can also be a sensitivity borne of instinct – it just “feels off”. In a resilient safety culture we respect that “sensitivity” no matter where it has come from.

4. Commitment to Resilience

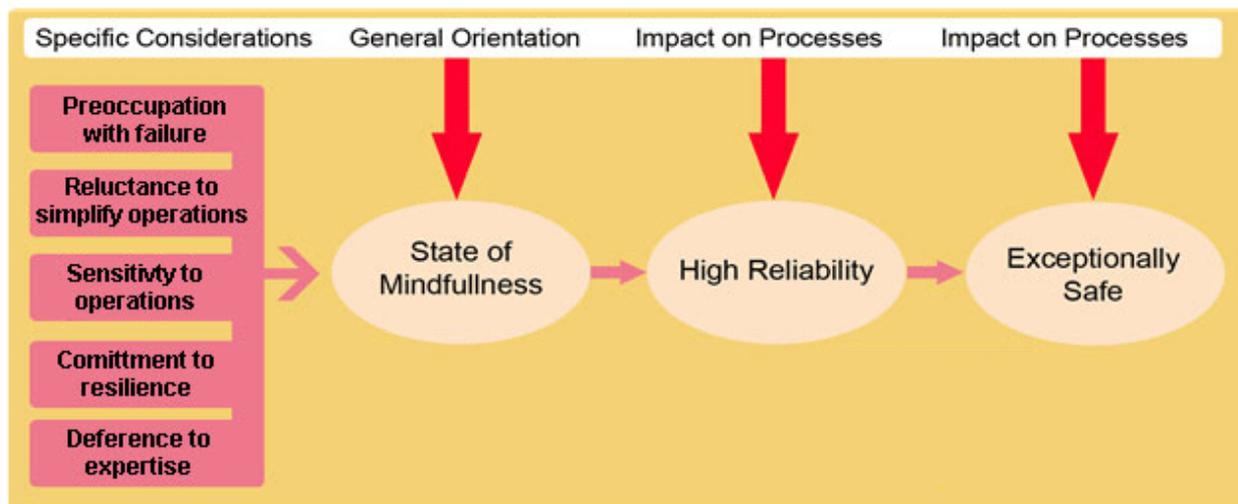
Having a commitment to resilience means being prepared for whatever. In a resilient safety culture we accept that “bad stuff” is going to happen. The difference is that we are actually prepared for it. When it does happen we have drills and practices in place that allow both the organization (and its people) to feel secure in the knowledge that the path to return and recovery is “in place”. This means that a resilient safety culture shall train, train, and train some more. It shall train its people in those areas such as personal stress management/resiliency skills, organisational trauma recovery, etc.

5. Deference to Expertise

In a resilient safety culture there shall be an overt recognition as to where the expertise lies. More often than not, it means recognizing that those closest to the “frontline” are the experts, and providing them with the skill, resources, and delegation to make decisions when a critical issue arises shall often result in quicker mitigation of harm.

In a resilient safety culture, senior leaders are conducting frequent SAFE-T-WALKS to reinforce safety behaviours and find and fix critical safety issues. These senior leaders are generally quite naïve when it comes to recognizing safety issues. They are reliant upon sourcing this information from the “experts”. This means they must create a climate of openness and sharing in their communications with operational staff.

When senior leaders are in operational review meetings, they should “lead with their ears”. It is all about listening intently to all that is going on around them, and using their own position within the organization to remove speed bumps that might be getting in the way of resilient performance.



Putting it all together – SAFETY I and SAFETY II

For a long time I have been writing, and saying, whenever I have had the chance, that we have just not been “getting it” when it comes to safety. In a very real way we have been trying to measure the success of the profession of safety, by actually measuring its failures. Whenever we are using reductionist metrics, what we are actually doing is measuring failure, and not success. So when we hear a mining executive express genuine pride when their operations only killed sixteen (16) people last year, that is a clear measure of failure. It does nothing to encourage resilience. Similarly when we see safety professionals (who should know better), standing around at a conference, espousing how good their safety programs are because they have had a twenty-five percent (25%) reduction in their LTIFR number, they are actually reporting on their failures. In short any metric that is being used, which has as a goal a result approaching zero, is what we would call a reductionist metric. The contrast is what we might refer to as an expansionist metric. The latter are those things where we actually want to see the numbers increasing, not reducing. Examples of expansionist metrics might be:-

- ☀ The number of safety walk-arounds being conducted by leadership,
- ☀ The number of safety observations,
- ☀ The percentage of close-outs with regard to safety improvements,

Hollnagel⁷ puts these frustrations into an excellent format in his discussion of SAFETY I and SAFETY II. He defines SAFETY I thinking as being primarily reactive. This has probably been the way it has been for time immemorial. That is quite understandable. It is much easier to measure something that “has been” rather than what “might be”.

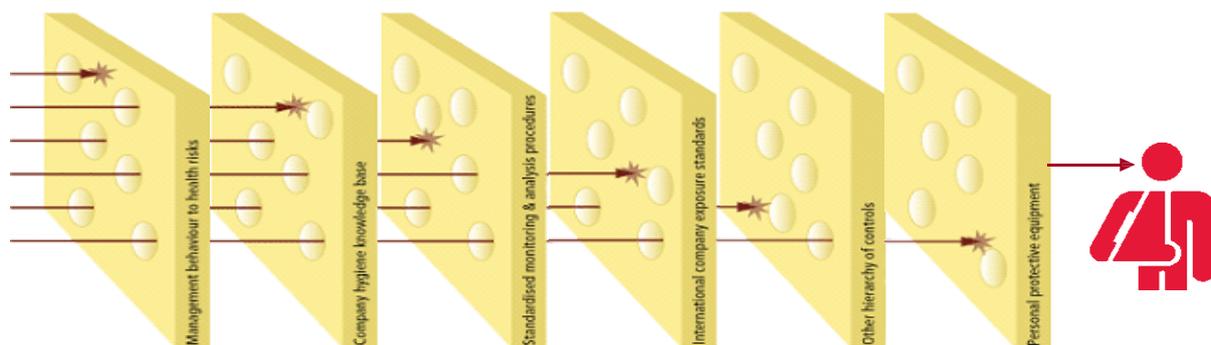
The reason we use the term “reactive” is because SAFETY 1 thinking is based on responding to something that has gone wrong, or has been identified as a risk, i.e. something that might go wrong. The standard SAFETY I response involves looking for ways to eliminate the cause/s, or to control the risks, either by finding the causes and eliminating them, or by improving options for identification and recovery. Reactive safety management (SAFETY I) centres around “causality”.

⁷ Hollnagel, E., (2014), “Safety-I and Safety-II: The past and future of safety management”, Ashgate Publishing, Burlington:VT, United States

1. Bad stuff (accidents, incidents) happens when something goes wrong.
2. Bad stuff therefore has causes, which can be found and treated.

The problem confronted is that this is a toxic circuitous journey that requires the “bad stuff” to happen (or nearly happen, i.e. at least be discernible) before we recognize the need to react.

A key understanding here is that our work systems, and the system failures that occur within them, are actually not linear. SAFETY I thinking is grounded within that linear progression. Consider for a moment the well-respected accident causation model put forward by James Reason (BTW, this one is also respected by me). Essentially it is based upon a view that all accidents occur as a consequence of system energy being allowed to travel from one point to another without being challenged. The image below is a representation of what has become known as the “Swiss Cheese” model of accident causation.



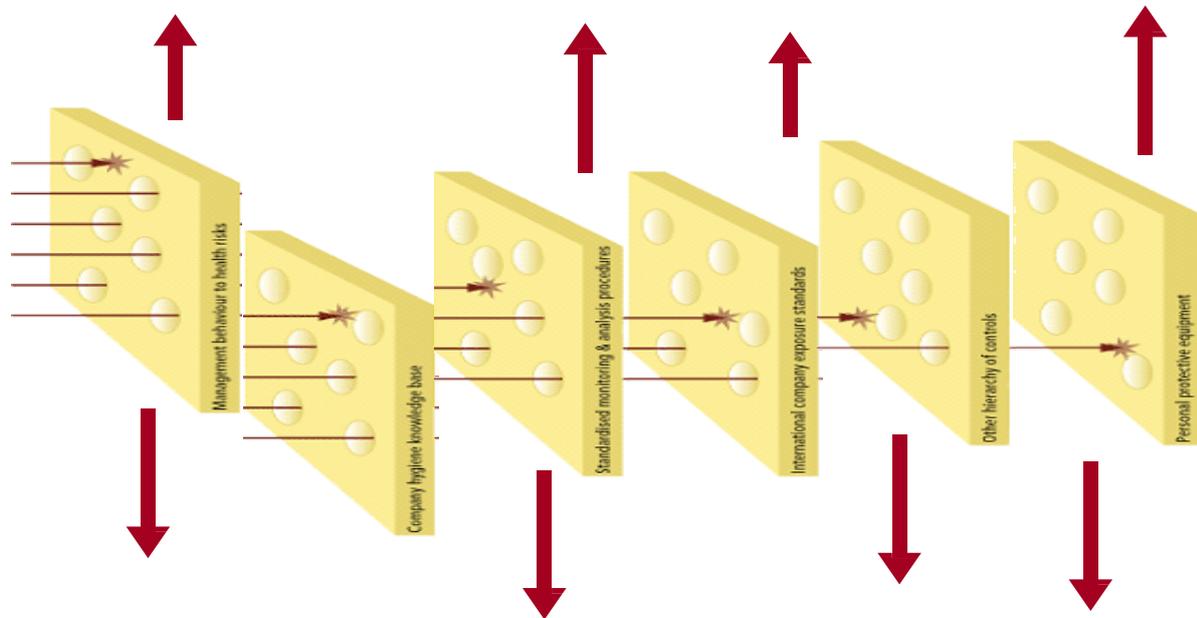
In short, each time you have a system failure, you investigate etc. and apply a new “barrier” to get in the way of that system energy, or you may try and plug up a hole in a slice of a pre-existing barrier. Experience has shown, over many years, no matter how many barriers you put in place “bad stuff” still happens, and sometimes it is really “bad”. That is because the entire understanding around the Swiss Cheese model is premised upon the predictability of causation. Yet, we have shown, time and time again, that “bad stuff” does not fit that predictability model. If it did, we would have controlled for it, and it would not happen.

A pretty confronting example of that observation is the number of refinery fires in the United States since the Texas City disaster. They have not significantly reduced. That is despite huge investments in process safety programs etc. In the same industry space it has also been shown that the number of fatalities has not significantly reduced since Texas City either. OSHA have even acknowledged they are not even “sure” how many fatalities have actually occurred within the industry – it could even be more than the recorded number. Now this is the United States saying that, not some third world country with questionable data collection methodologies. That just has to be plain scary. This one example, and there are many others, shows that “safety” as we would like to think of it, is not significantly improving in these areas.

If we want to consider an example outside of the more traditional workplace type scenarios, give some thought to road safety. More specifically let’s consider the number of cyclists killed in road accidents for the last ten (10) years⁸. I have paid some special attention to this figure as in 2008 I almost became one of those statistics. Anyway in 2008, the magic number was twenty eight (28). This was

⁸ Department of Infrastructure & Regional Development, (2014), “Road Deaths Australia - Annual Summaries”, Australian Government Publishing Service, Canberra: AUSTRALIA

obviously a very “safe” year. The reason I say this is that the year prior (2007) there had been forty one (41) deaths and the year following (2009) the number was essentially the same, thirty one (31) deaths. If we stretch out the period of review we see that in 2004 there were forty three (43) deaths and in 2013 there were fifty (50) deaths. In short the number of cyclists dying on Australian roads has actually significantly increased over this ten (10) year period. I am sure that all the governments of Australia would be loudly spruiking the investments they have made in trying to reduce these numbers. Clearly with no success – at least if we use the traditional SAFETY I measurement of failure methodology. Now back to work.



The previous image is a representation which pretty clearly shows what can happen to a safety system. I would argue strongly that this is what happens within our safety management systems all the time. We know that organisational safety systems are entropic in nature. This means that left to their own devices they shall inevitably fail. They also suffer greatly from “variation” – this is defined by the red arrows. It is this “variation” that allows for highly protected environments such as petrochemical and aviation etc. to still suffer from all sorts of disasters. What is actually happening is that this “variation” is permanently present in all of our safety management systems. It is just that the degree of “variation” varies from one environment to another. The Swiss Cheese model only becomes more useful; if we can see it operating through this third dimension (variation). Failure to acknowledge the influence of “variation” means that the only solution to terminate the linear causality sequence, is to just add more and more complexity to the system. This means just more and more barriers. We know that this just does not work!

It also means that if this way of describing accident causation is going to be helpful, we need to do everything we can to minimise or extinguish that “variation”. In other words ensure those barriers consistently line up – that way the barrier model (Swiss Cheese) might just work?

SAFETY II is a different concept altogether. It is based on a fundamentally different promise. That being that systems (including safety systems) work because the people within them are able to adjust to match the requirements being asked of them.

People can also detect and correct when something goes wrong, or when it is about to go wrong, hence intervene before the situation becomes seriously worsened. In the HRO discussion this “skill” comes from that “sensitivity to operations” coupled with a “pre-occupation toward failure”. Others describe this as deliberately encouraging a permanent state of un-ease. If you like a permanent position of being wary as to what is happening around you.

The result of all this is a different form of “variability”, not in the negative sense where “variability” is seen as a deviation from some procedure (as depicted in the image presented above), but in the positive sense in that “variability” represents the dynamic adjustments that are consistently applied in the interests of safety and productivity.

One of the dilemmas we often confront is the commentary from some organisations who shall suggest that “machines don’t make mistakes, people do”. Whilst I would fundamentally beg to differ, now is not the time to embark upon that crusade. What I will do is offer the following observation. While the reliability of “machines” may be pretty high, organisations frequently trade-off completeness for productivity. In addition we know the quality of labour varies widely. When confronted with these daily challenges (and I have yet to meet a senior business leader who does not acknowledge the reality of this observation), people are clearly an asset rather than a liability and their ability to adjust what they do to the conditions is a strength rather than a threat.

So What is a Resilient Safety Culture?

It is my hope that the preceding discussion demonstrates what we now consider to be the power and influence of developing a mind-set of resilience within an organization. The Resilient Safety Culture is one that has the ability to recognize when both “bad stuff” and “good stuff” is happening and respond accordingly.

In a resilient safety culture variability is actually encouraged. The trick though is to ensure that any such variability is what we shall call “positive variance”. In other words, members of resilient safety cultures have the ability to recognize when the system is moving in sub-optimal direction, and then put in personal and organisational processes to bring things back to alignment. It is when everything is aligned that more traditional accident causation models (Swiss Cheese) have a better chance of being able to deliver

that which the organization aspires to. I would go so far as to suggest that the new SAFETY II thinking, when properly applied throughout an organization, goes a long way to actually making the original model proposed by James Reason far more effective. In other words, if the “positive variance” surrounds the underlying organisational behavioural patterns, the negative variation that has been shown to contribute to almost all accidents/disasters, is significantly reduced. In essence, there shall be so much “good stuff” happening throughout the organization, that the “bad stuff” has so little room to move.

