AIChE

Learning From Organizational Incidents: Resilience Engineering for High-Risk Process Environments

Stefanie Huber,[a] Ivette van Wijgerden,[b] Arjan de Witt,[b] and Sidney W.A. Dekker[c]

[a] Berlin Institute of Technology, Center of Human-Machine-Systems, Berlin, Germany; Stefanie.Huber@zmms.tu-berlin.de (for correspondence)

[b] Delft University of Technology, Faculty of Technology, Policy and Management, Delft, The Netherlands

[c] Lund University, Center for Complexity and Systems Thinking, Ljungbyhed, Sweden

Published online 18 December 2008 in Wiley InterScience (www.interscience.wiley.com). DOI 10.1002/prs.10286

For years, safety improvements have been made by evaluating incident reports and analyzing errors and violations. Current developments in safety science, however, challenge the idea that safety can meaningfully be seen as the absence of errors or other negatives. Instead, the question becomes whether a company is aware of positive ways in which people, at all level of the organization, contribute to the management and containment of the risks it actually faces. The question, too, is whether the organization has the adaptive capacity necessary to respond to the changing nature of risk as operations shift and evolve. This article presents the results of a resilience engineering safety audit conducted on a chemical company site. An interdisciplinary team of seven researchers carried out 4 days of field studies and interviews in several plants on this site. This company enjoyed an almost incident-free recent history but turned out to be ill-equiped to handle future risks and many well-known daily problems. Safety was often borrowed from to meet acute production goals. Organizational learning from incidents was fragmented into small organizational or production units without a company-wide learning. We conclude that improving safety performance hinges on an organization's dynamic capacity to reflect on and modify its models of risk as operations and insight into them evolve, for example, as they are embodied in safety procedures and policies. © 2008 American Institute of Chemical Engineers Process Saf Prog 28: 90–95, 2009

Keywords: chemical industry, incident reporting, accidents, hazards, anticipation, resilience engineering

INTRODUCTION

In recent years, ideas on how to improve safety in high-risk environments have followed a common approach. Suggestions are often made to improve safety by preventing human errors [1]. In this paradigm, safety is maintained by training, by disciplining of people, adding new procedures, and adding new safety automation. The system is considered to be basically safe because safety is something that is engineered into a system.

Lately, new insights into how safety is improved and maintained have revolutionized the whole discussion about safety in complex systems and have established a new paradigm for analyzing human error. In this new view, the emphasis is on the positive contribution of people at all levels of the organization, rather than solely on human errors [2,3]. Dealing with new procedures, new developments, and new process problems is a daily responsibility. This view acknowledges that the demands of high consequence/low probability events cannot always be handled by matching situational symptoms with scripts of coordinated action used in training. These scripts can help people (a) prioritize actions in the

face of time pressures and resource constraints, (b) to assign tasks, (c) organize roles, and (d) continually assess expectations. But the limits of and misplaced confidence in such preparations have been commented on before [4,5], and the literature has shed some light on the difficulty of processes of sensemaking in demand situations that lie beyond procedural reach [6].

Resilience engineering has been suggested as one answer to eliminating these limits. It is a good approach for handling systems that have to deal with dynamic and complex environments where daily routine no longer works and systems have to be dynamically stable and flexible rather than rigid [7–9]. Following this resilience engineering paradigm, safety cannot be engineered into a system as a property, but must emerge as a quality of the joint human-machine system [10].

Hence, it is not possible to find out if a system is safe or not by deconstructing it: "resilience engineering abandons the search for safety as a property, whether defined through adherence to standard rules, in error taxonomies, or in 'human error' counts" [10]. Furthermore, resilience engineering is a broader concept than existing error counting taxonomies or other models that include the "old view of human error" [11]. It is not only about determining the "probability that a given function or component would fail under specific circumstances" [10], but rather postulates that systems "must also be resilient and have the ability to recover from irregular variations, disruptions, and degradation of expected working conditions" [10]. Where old-view models focus on counting errors that happened in the past, resilience engineering concepts target present working conditions as well as the anticipation of future developments—or as Diamond [12] put it: "to anticipate a problem before it has arrived" [10,12].

The resilience engineering approach is a new and fascinating one, because it de-emphasizes the old-view models of human error investigation, and it considers the system as a whole. It focuses on present and future safety, i.e. on proactive safety instead of reactive safety and hence does not concentrate solely on past errors [10]. It considers humans as an integral part of resilience and does not focus only on technical components or redundancy as the main elements for enhancing safety in systems.

All these ideas sound reasonable—but do they really work in practice? The aim of the research reported here is to test these ideas by conducting a resilience-engineering safety audit in a chemical company.

METHOD

In May 2007, a resilience engineering based safety audit was conducted in a chemical company (remains anonymous) with more than 300 employees. The company was founded more than 100 years ago and today runs more than a dozen production plants at a single site in Europe.

To maintain and enhance the plants' high-safety levels, a safety department was established some years ago and incident reporting has become a major way of keeping track of the plant's safety state. For confirmation, the safety manager was interested in (a) an external safety audit to obtain a more objective appraisal and (b) collecting ideas about further improvements or changes. An interdisciplinary, intercultural team of seven was formed to perform this audit.

A structured questionnaire was developed for interviews with operators and managers based on topics that are critical for organizational success and have turned out to be good indicators for resilient organizations [7,9]. Six dimensions were taken into account:

1. Top-level commitment (e.g. Do you think your boss appreciates your work?
2. Just culture (e.g. Do you feel comfortable reporting safety issues/problems to your boss?)
3. Learning culture (e.g. Do you feel the discussion about risk is kept alive in your company?)
4. Awareness and opacity (e.g. Do you know the major safety concerns the company has to deal with?)
5. Preparedness (e.g. Do you feel ahead of upcoming problems?)
6. Flexibility (e.g. Do you have any slack resources available to cope with sudden trouble?)

Each dimension was covered by at least four questions. General questions were included about the employees' positions, their daily work, their workload, and personal experience with accidents, as well as opener questions about the employees' feelings about the safety culture.

28 operators and 21 managers were interviewed during the audit period. Additionally, 2 days of field observation in various plants and with different operators were performed and background information was reviewed about the company, concerning organizational and industrial issues.

When analyzing the data a two-step approach was chosen. First, responses from both operators and managers were described and summarized in a report. As "data never speaks for itself" [13] the second step analyzed the interview and observation data. Patterns and regularities [14] were presented in the report, as well as higher-order categories for the findings. This second-step analysis was supported by literature and included proposals and interpretations about the investigated company. It aimed to answer questions such as: How could these phenomena have been generated? What underlying principles can be derived from the findings? What common principles do these results identify? This approach was chosen as research has shown that understanding the causes of failure (the so-called second stories) are crucial for learning in organizations [15,16].

For this article, the results of the comprehensive first step of the analysis will not be presented, but the findings and patterns that were discovered through the second step of our analysis will be described and discussed.

RESULTS AND DISCUSSION

Safety Versus Production

Operators and managers were convinced that safety and production are the two most important goals in the company and had to be met simultaneously. Company goals, often driven by external pressure to produce and to meet customer goals were internalized by almost all employees, and the consequence resulted in situations where an employee had to decide between safety or production, but could not achieve both at the same time. Usually, such incompatible goals arose at the organizational level and its interaction with the environment, but the actual managing of goal conflicts under uncertainty was pushed down into local operating units. There, the conflicts had to be negotiated and resolved in the form of thousands of daily decisions and trade-offs. Efforts were made by the operators and managers to deal with these conflicting goals simultaneously, but trade-offs were always the final resolution. Small incidents were neglected and often were considered as normal side-effects of daily work.

These daily trade-offs often resulted in operators not putting on their safety gear because this would "waste" some minutes in the production process. Hollnagel [17] uses the expression "efficiency-thoroughness trade-offs" (ETTO) for these types of situations. There people have to decide to go for safety and thoroughness or production and efficiency. "On the one hand people genuinely try to meet their (internalized) goals, i.e., they try to do what they are supposed to do—or at least what they believe is reasonable to do—and to be as thorough as they find it necessary. On the other hand they try to do this as efficiently as possible which means without spending any unnecessary effort or time to do it" [17]. Underlying organizational pressures and preferences were reproduced in what individual people did and valued (or undervalued), in a way that was invisible for the organization as a whole.

Yet people also took their ability to reconcile the irreconcilable as a source of considerable professional pride. It was seen as a strong sign of their expertise. The general faith in, and commitment to the company, and its advertised high concern for safety initially seemed overwhelming. "Safety is always first" people often said at all levels, but often contradict it decisively with their own actions not much later. However, further questions about the main goal or goals in the organizations produced responses to keep the production running, not only from operators but also from middle-managers. They doubted that safety was unconditionally put first by their superiors. Most of them felt that other goals such as production, costs, or delivery were at least equal. But how can an organization claim that safety always comes first and then insist on keeping production running? Organizations often resort to "conceptual integration, or plainly put, doublespeak" [18].

Digging deeper, we experienced critical voices that talked about daily trade-offs that were made between safety and production.

> Safety is important when you are not busy, but if there is production pressure safety is not important, production is put first.[1]

Daily trade-offs between safety and production occur in most companies that depend on producing and selling their goods. It is a common problem in the process industry, where the production constitutes the right to exist [2]. But as one can imagine, it is not possible to be safe and efficient to the same great extent in a given time [10] or as Hollnagel put it: "If anything is unreasonable, it is the requirement to be efficient and thorough at the same time" [17]. To avoid such situations, sacrificing decisions have to be made to deal with these problems—situations where safety is put first and the pressure on throughput and efficiency goals is relaxed [19]. Management plays a major role in making sacrificing decisions. Whom else, if not the middle and higher management, could mandate that safety is worth more than production? But why is it so difficult to stand up and stop the production in favor of safety?

The Internalization of External Pressure

One answer is the internalization of external pressure. There are a lot of goals that operators and managers have to face everyday. Safety is never the only goal in systems; often there are economic pressures, cost goals, production goals, and so forth [2]. Many operators have internalized these multiple goals and try to achieve all of them simultaneously with their best efforts. This was illustrated by a story told by one of the operators in the audited company:

> Usually you have to [...] monitor the process every 15 min. If you do it only after 16 min, you have to fill a report and the 'paper error' goes down to the office. To avoid that, you usually write down that you checked it on time. Sometimes you have to work for two production lines at the same time and then it's not possible to perform the checks in time. So you look afterwards what the temperature was some minutes before and you write it down (e.g., 5 min too late). Because you have to do things simultaneously and you have to do things right - these are two goals that sometimes conflict.[1]

One of the managers said that "the higher in hierarchy you are, the more goals you have. You have fewer goals if you go down the hierarchy, because you act more locally"[1]. But actually the opposite is the case. Goals are usually cascaded down the hierarchy, which is known as the management by objectives principle. If a superior has the goal to produce more the next year, then his subordinates have to execute this goal. One manager told us about the pressure he has and stated that he has to "improve every working process and be cost-efficient and at

[1]Direct quotes from operators' and managers' interviews, resilience engineering based safety audit, 2007.

the same time produce and work safely"[1]. Exactly the same goals and pressures were faced by the operators, because "institutional pressures are reproduced, or perhaps really manifested, in what individual people do" [11]. There is no difference on your hierarchy level: "Safety is first with maximum production"[1]. Keeping that in mind when asked for the most valued goal of their superiors it is not surprising that operators internalize these often conflicting goals and make them their own.

This phenomenon is also known as the macro-micro-level connection. Company goals that are established at upper hierarchy levels are reflected in the thinking and behavior of operators at the lower end of the hierarchy. The "competitive environment (competition, scarce resources, and norms) [. . .] generates pressure on organizations to violate laws and rules in order to attain goals" [20]. There is also a "wider cultural belief system [. . .] of society: capitalism and competition are 'the'economic way; concerns with cost, production goals, and efficiency dominate industries" [20]. This is recognized by employees and they complain that "it would be nice not to be on the stock market (as a company) to get rid of the quarterly and end of year reports—because of the high stress that goes with the reports, more incidents, and mistakes are bound to happen"[1] or that "money is of course no. 1"[1]. One can clearly understand how company goals are passed down to managers and operators.

This study showed that it is very difficult to put safety first. To be successful, every level in the plant and company must recognize the hazards of these external pressures and seriously internalize safety first.

The Normalization of Daily Risk

Asking for a rough estimate of annual incidents and accidents, the employees on average stated that big accidents happen zero to five times a year, but almost everyone stated that small accidents or incidents happen all the time. It was interesting that most operators rated events such as acid in the eyes or a skin burn only as a minor incident. Operators seemed to be so used to the frequency of such events that they did not consider them to be "a big thing":

> I have just experienced some small incidents. Scalding from hot steam or hot water are typical incidents that often happen. I was involved in a burning incident.[1]

Crucial for the managerial awareness of this normalization of incidents is that these little incidents are usually not reported and therefore management does not get to know them, especially their frequency. Operators did not consider these daily events as incidents and believed they were naturally associated with normal work. The boundaries of what is an incident are redefined because of the mere frequency of events; incidents become normalized:

> If the fire alarm goes off nothing happens. No one runs out of the factory, because we are so used to it. The fire alarm goes off about once a month because of steam leaks (e.g., out of some pipelines) and then the smoke detector reacts.[1]

Not only is daily risk normalized, but also the deviance from procedures or manuals are normalized; i.e. procedures and manuals functioned as guidelines instead of fixed and mandatory rules. Taking into account that operators at the investigated company had about 400 overall procedures and additionally 10–30 local procedures in every plant, where "some are in good shape, others are in bad shape"[1], it was not surprising that some of them were used as mere guidelines or suggestions. Sometimes the sheer number of procedures or manuals made it impossible to follow them. During the audit, one operator who was interviewed had to run three production lines at the same time without a colleague. It was impossible to follow the manuals as required. Furthermore, experienced operators often claimed to know which situations were dangerous and which were not.

Are these people ignoring good rules, or are the rules bad and unsuited to the demands of real work? There is a deeper, more complex dynamic where real practice is continually adrift from official written guidance, settling at times, unsettled, and shifting at others. Deviance from original rules becomes normalized; nonconformity becomes routine [20]. Digging deeper and asking for reasons for the deviances was because their managers did not have enough time to fix the problems. One of the managers explained: "I am a little behind with changing procedures in some plants, because I have too little time"[1]. This lagging behind could be one reason for the gap between procedure and practice. Another reason could be that "instructions do not cover all issues—sometimes they are very general, sometimes they are too detailed and you can't follow them"[1]. Again, people have to adapt these procedures locally to create safety, because there is always a gap between a written rule and an actual practice. This distance needs to be bridged; the gap must be closed, and the only thing that can close it is human interpretation and application. It is often more fruitful to look at such gap-closing as a creation of safety than at the "violation" of some procedure.

The challenge is to make the gaps visible and provide a basis for learning and adaptation where necessary. When operators change the procedures or manuals and nothing happens then they presume the situation is safe. But as the procedures change little by little, the situation could become unsafe and an accident could occur. For progress on safety, organizations must monitor and understand the reasons behind the gap between procedures and practice. Additionally, organizations must develop ways to support people's skill at judging when and how to adapt. To prevent this problem with procedure-changes, the company must establish and use a formal management of change (MOC) process. Procedures must not be changed unless this MOC process is used. The MOC process should require a technical review by a team to make sure that the recommended change in the procedure is safe. The MOC process should also include (a) checks to make sure

that the latest and approved procedure is used, (b) a corresponding document containing dates and approval signatures for the latest approved procedure, and (c) a file with old procedures and MOCs. Of course, the MOC as a process itself should be kept up to date too, as its ways and strategies for dealing with change problems may lag behind the developing nature of risk in the plant.

Closing the Loop of Learning

Not analyzing how the company learns from failures was a major cause of dissatisfaction among the operators. Learning appeared to be based on individual experience rather than on a structured learning process that covered the whole company. As some operators pointed out, they wanted to be better informed about failures in other plants. This statement was not surprising, because most information the employees received was only put on the intranet and not actively distributed, for example, by a supervisor in a weekly meeting.

Operators really did want to know what was going on in other plants. "Weick [6] refers to 'collective mindfulness' as a characteristic of highly reliable organizations: collective mindfulness includes the fact that safety issues and concerns are widely distributed throughout the organization at all levels" [9]. It is ineffective to use only the intranet to inform operators about other plants—the chance to create such a collective mindfulness relies on the intranet as a single source. Hence, people had to search for information on the intranet, and they found this difficult and not user-friendly. Another problem was that employees wanted to know more about a failure than was reported; people were only informed by the intranet that an accident had happened—but the cause of a failure and what management did to prevent a repeat was not reported. Many employees said that when a big accident occurs in other plants, only an email will be sent. Putting information on the intranet or sending out emails is not enough. You need to involve people and heighten their awareness that such an accident could happen to them also.

Learning from failure was somehow organized by collecting incident reports, by placing memos about accidents on the intranet, and by establishing a safety department. But when asking "how do you learn from failures?" there seemed to be no active and structured way of distributing information through the whole company except for the intranet. Most of the operators and managers stated that they keep things in mind and try to use their knowledge if they get in the same situation again. But leaving the responsibility for corporate learning to single individuals who "try to remember"[1] is not the way to create a structured, holistic approach, where people are given room to share their best practices and their knowledge.

Distributing information about failures should not solely hinge on the intranets and computers—more personal communication is essential here.

CONCLUSION

The study reported here suggests that we should be wary about thinking that safety can be engineered into a system by extensive error counting, or by removing or tinkering with small elements that seem to be unreliable (e.g. individuals, procedures, and equipment). The capability of a system to adjust to the constantly changing environment could be a more important predictor of future safety. Ways to do this include having effective safety meetings and identifying gaps between work as imagined (or specified) and work as actually done [15,19]. Production pressures and similarly company goals play a role in this process.

In our study, managers admitted that they needed more time to know what was really going on in their own plants; managers complained about having to produce reports for annual and quarterly results and also having the obligation to meet production goals. We also saw operators voicing concerns about the same safety problems for years without success; operators gave up because "management never listens,"[1] and operators were forced to function quickly and safely at the same time. They were pushed into inadvertently taking procedural short-cuts to meet the requirements of parallel production processes. Finally, we encountered operators working 16-h shifts for some days in a row to meet production and manning constraints.

Although these circumstances are common in the chemical industry, we found little evidence of conscious discussions about how this company could decide to relax the pressure on throughput and put safety first. Such sacrificing decisions are considered to be a major contributor to safety in high-risk environments [21]. Counting negative events (errors, violations, and incidents) does not always reduce risk. Incident reporting can help organizations get an overview of accidents and incidents. However, every incident is not always recorded, e.g., when operators do not report incidents that happen daily because they are considered to be normal. Moreover, past recorded accidents and incidents do not predict future mishaps, because the systems are already very safe, "the next accident has never been seen before. It may involve a series of already seen micro incidents, although most have been deemed inconsequential for safety" [22]. Taking this into account, incident reporting and error analyzing alone cannot raise safety to a higher level. Unexpected events cannot always (if ever) be predicted from the previously gathered data about things that went wrong. Instead, resilience engineering suggests that a company must recognize, adapt to, and absorb challenges that fall outside the scope of its design and historical experiences. This is a capacity that necessarily relies on more than knowing what happened before.

The especially valued lessons learned from this audit study include:

- Safety must be first, but there are always uncertainties and goal conflicts that make this very difficult in practice.
- All incidents must be reported and analyzed, but it can be very difficult for managers and operators alike to agree on what counts as an "incident." Furthermore, analyzing an incident is not the same as learning from it; for this a whole suite of follow-up activities is necessary.

- It would be nice to say that operating procedures must be used as specified, and only changed after an MOC, but in reality there is always a gap between written guidance and actual practice. The real challenge for an organization is to be sensitive to this gap, to find out where and why it exists and resist judging the operator for not following the procedures as specified, as reasons for this may lie buried more deeply in the organization or operation.
- Plants need sufficient operators and managers to operate the plant safely, but definitions of "sufficient" are often negotiable and based on incomplete evidence.
- Person to person safety meetings are needed and intranet/computer communication regarding safety should be discouraged or solely used as complementary source of information, for only then are there real opportunities for sharing narratives about risk that people can use for vicarious learning.
- Operations, engineers, and managers need to constantly adapt to a changing environment—This is the key factor to a resilient organization.

ACKNOWLEDGMENTS

The authors thank the research team, especially Travis Haigler, Roberta Usher, and LaToya Hall for their contribution, and most importantly the managers and operators of the anonymous chemical company for their extraordinary openness and patience while being observed and asked over and over again.

LITERATURE CITED

1. S.W.A. Dekker, The re-invention of human error, Technical Report 2002-01. 2002. Available at: **http://www.lusa.lu.se/upload/Trafikflyghogskolan/TR2002-01_ReInventionofHumanError.pdf**. Last accessed on July 12, 2007.
2. S.W.A. Dekker, The Field Guide to Human Error Investigations, Ashgate Publishing Co., Aldershot, 2002.
3. D.D. Woods and R.I. Cook, Nine steps to move forward from error, Cognit TechnWork 4 (2002), 137–144.
4. L.A. Suchman, Plans and Situated Actions: The Problem of Human-Machine Communication, Cambridge University Press, Cambridge, 1987.
5. P.C. Wright and J. McCarthy, "Analysis of procedure following as concerned work," Handbook of Cognitive Task Design, E. Hollnagel, (Editor), Lawrence Erlbaum Associates, Mahwah, NJ (2003), pp. 679–700.
6. K.E. Weick, The collapse of sensemaking in organizations, Admin Sci Q 38 (1993), 628–652.
7. A. Hale, F. Guldenmund, and L. Goossens, "Auditing resilience in risk control and safety management systems," Resilience Engineering: Concepts and Precepts, E. Hollnagel, D.D. Woods, and N. Leveson, (Editors), Ashgate Publishing Co., Aldershot (2006), pp. 289–314.
8. E. Hollnagel, "Resilience—the challenge of the unstable," Resilience Engineering: Concepts and Precepts, E. Hollnagel, D.D. Woods, and N. Leveson, (Editors), Ashgate Publishing Co., Aldershot (2006), pp. 9–18.
9. J. Wreathall, "Properties of resilient organizations: An initial view," Resilience Engineering: Concepts and Precepts, E. Hollnagel, D.D. Woods, and N. Leveson, (Editors), Ashgate Publishing Co., Aldershot (2006), pp. 275–286.
10. E. Hollnagel and D.D. Woods, "Epilogue: Resilience engineering precepts," Resilience Engineering: Concepts and Precepts, E. Hollnagel, D.D. Woods, and N. Leveson, (Editors), Ashgate Publishing Co., Aldershot (2006), pp. 347–358.
11. S.W.A. Dekker, Ten Questions About Human Error: A New View of Human Factors and System Safety, Lawrence Erlbaum Associates, Mahwah, NJ, 2005.
12. J. Diamond, Collapse. How Societies Choose to Fail or Survive, Allen Lane, London, 2005.
13. C.L. Bosk, Forgive and Remember: Managing Medical Failure, University of Chicago Press, Chicago, IL, 2003.
14. J.W. Creswell, Qualitative Inquiry and Research Design: Choosing Among Five Traditions, SAGE Publications, Thousand Oaks, CA, 1998.
15. S.W.A. Dekker and T. Laursen, From punitive action to confidential reporting, Patient Saf Qual Healthcare 5 (2007), 50–56.
16. D.D. Woods, L.J. Johannesen, R.I. Cook, and N.B. Sarter, Behind Human Error: Cognitive Systems, Computers and Hindsight, CSERIAC, Columbus, Ohio, 1994.
17. E. Hollnagel, Barriers and Accident Prevention, Ashgate Publishing Co., Aldershot, 2004.
18. D. Dörner, The Logic of Failure: Recognizing and Avoiding Error in Complex Situations, Perseus Books, Cambridge, MA, 1989.
19. S.W.A. Dekker, "Resilience engineering: chronicling the emergence of confused consensus," Resilience Engineering: Concepts and Precepts, E. Hollnagel, D.D. Woods, and N. Leveson, (Editors), Ashgate Publishing Co., Aldershot (2006), pp. 77–94.
20. D. Vaughan, The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA, Chicago University Press, Chicago, IL, 1996.
21. D.D. Woods, "Essential characteristics of resilience," Resilience Engineering: Concepts and Precepts, E. Hollnagel, D.D. Woods, and N. Leveson, (Editors), Ashgate Publishing Co., Aldershot (2006), pp. 21–34.
22. R. Amalberti, "Optimum System Safety and Optimum System resilience: agonistic or antagonistic concepts," Resilience Engineering: Concepts and Precepts, E. Hollnagel, D.D. Woods, and N. Leveson, (Editors), Ashgate Publishing Co., Aldershot (2006), pp. 253–274.