How to design a safety organization: Test case for resilience engineering

David D Woods
The Ohio State University
372 PUBLICATIONS **15,313** CITATIONS

Chapter 19

How to Design a Safety Organization: Test Case for Resilience Engineering

David D. Woods

In the aftermath of the Columbia space shuttle accident (STS-107), the investigation board found evidence of an organizational accident as NASA failed to balance safety risks with intense production pressure (Gheman, 2003). Ironically a previous investigation examining a series of failures in Mars exploration missions also focused on breakdowns in organizational decision making in their recommendations (Stephenson et al., 2000). Both reports diagnosed a process where the pressure for production to be 'faster, better, cheaper', combined with poor feedback about eroding safety margins, led management inadvertently to accept riskier and riskier decisions.

Woods (2005a) links these accident analyses to patterns derived from previous results and argues that organizational accidents represent breakdowns in the processes that produce resilience. Balancing the competing demands for very high safety with real-time pressures for efficiency and production is very difficult. As pressure on acute efficiency and production goals intensifies, first, people working hard to cope with these pressures make decisions that consume or 'sacrifice' tasks related to chronic goals such as safety. As a result, safety margins begin to erode over time - buffering capacity decreases, system rigidity increases, the positioning of system performance relative to boundary conditions becomes more precarious (cf., Chapter 2). Second, when margins begin to erode as a natural response to production pressure, it is very difficult to see evidence of increasing or new risks. Processes that fragment information over organizational boundaries and that reduce cross-checks across diverse teams leave decision makers unable to recognize the big picture, that is, unable to reframe their situation assessment as evidence of a drift toward safety boundaries accumulates (until a failure occurs and with the benefit of hindsight the evidence of new dangers seems strong and unambiguous).

How do people detect that problems are emerging or changing when information is subtle, fragmented, incomplete or distributed across the different groups involved in production processes and in safety management? Many studies have shown how decision makers in evolving situations can get stuck in a single problem frame and miss or mis-interpret new information that should force re-evaluation and revision of the situation assessment (e.g., Johnson et al., 2001; Patterson et al. 2001). A recent synthesis of research on problem detection by professional decision makers (Klein et al., 2005) found that reframing is a critical but difficult skill. Reframing starts with noticing initial signs that call into question ongoing models, plans and routines. How do these discrepancies lead people to question the current frame? When do they become suspicious that the current interpretation of events is incomplete and perhaps incorrect?

The initial signs are always uncertain and open to other interpretations. These indicators easily can be missed or rationalized away rather than lead to questioning and revision of the current frame. For example, studies have shown that a skilled weather forecaster comes in to work searching for the problem of the day, which comprise the unsettled parts of the scene that will need to be closely monitored (Pliske et al., 2004). In other words, the expert adopts a highly suspicious stance to notice and pursue small discrepancies despite the workload pressures and attentional demands. Less-skilled forecasters are much more reactive given other demands and do not reserve time to pursue these small (usually unimportant) discrepancies. As this example indicates, factors related to expertise, workload, and attentional focus can all contribute to a tendency to become stuck in a single view or frame, even as evidence is accumulating that suggests alternate situation assessments (Klein et al., 2005).

A resilience perspective on accidents such as Columbia allows one to step away from linear causal analyses that become stuck on the proximal events in themselves, on red herrings such as human error, or vague 'root causes' such as communication. Major accidents like Columbia are late indicators of a system that became brittle over time, of a safety management process that could not see the increasing brittleness, and of safety management that was in no position to help line management respond to increasing brittleness. As a result, failures of safety management in the face of pressure to be 'faster, better, cheaper' reveal that more effective techniques should provide the ability:

- to revise and reframe the organization's assessment of the risks it faced and the effectiveness of its countermeasures against those risks as new evidence accumulates.
- to detect when safety margins are eroding over time (monitor operating points relative to boundaries as noted in Cook & Rasmussen, 2005), in particular, to monitor the organization's model of itself - the risk that the organization is choosing to operate nearer to safety boundaries than it realizes.
- to monitor risk continuously throughout the life-cycle of a system, so as to maintain a dynamic balance between safety and the often considerable pressures to meet production and efficiency goals.

The organizational reforms proposed by the Columbia Accident Investigation Board try to meet these criteria which makes this accident report the first to recommend a resilience strategy as a fundamental mechanism to prevent future failures.

Dilemmas of Safety Organizations

Using a resilience approach to safety, I provided some input to the Columbia Accident Investigation Board (CAIB) which seemed consistent with the Board's own analysis and recommendation directions. Later Congress, as NASA's supervisor, wanted to check on NASA's plans to implement the CAIB's recommendations, especially the modifications to NASA's safety office. Congressional staffers asked several people to comment on the changes. As background I circulated a draft of my input to the board (what later evolved into Woods, 2005a). The staffers were very interested in this perspective, but to my surprise asked a simple and challenging question - how does one design a safety organization to meet these criteria? I was caught completely off guard, but immediately recognized the centrality of the question. Resilience engineering, if it is a meaningful and practical advance in safety management, should be able to specify the design of safety organizations as a work-a-day part of the organization's activities.

The staffers' question put me on the spot. As always when confronted with a conceptual surprise my mind shifted to a diagnostic search mode: why is the job of a safety organization hard? The resilience paradigm suggested organizations needed a mechanism that question the organization's own model of the risks it faces and the countermeasures deployed. Such a 'fresh' or outside perspective is necessary for reframing in cognitive systems in general. A review and reassessment was necessary to help the organization find places where it has underestimated the potential for trouble and revise its approach to create safety. A quasi-independent group is needed to do this -- independent enough to question the normal organizational decision-making but involved enough to have a finger on the pulse of the organization (keeping statistics from afar is not enough to accomplish this).

Why is developing and maintaining this questioning role difficult and unstable? Because organizations are always under production pressure (though sometimes the pressure on these acute goals can be stronger or weaker), the dilemma for safety organizations is the problem of 'cold water and an empty gun.' Safety organizations, if they assess the organization's own models of how it is achieving safety, raise questions which stop progress on production goals - the 'cold water.' Yet when line organizations ask for help on how to address the factors that are eroding or reducing resilience, while still being realistic and responsive to the ever-present production constraints, the safety organization has little to contribute - the 'empty gun.' As a result, the safety organization fails to better balance the safety/production trade-off in the long run and tends to be shunted aside. In the short run and following a failure, the safety organization is emboldened to raise safety issues (sacrifice production goals), but as time flows on, the memory of the previous failure fades, production pressures dominate, and the drift processes operate unchecked (as has happened in NASA before Challenger and before Columbia, and can happen again).

From the point of view of managing resilience, a safety organization should monitor and dynamically re-balance the trade-off of production pressure and risk. The safety organization should see 'holes' in the organization's decision processes, reframe assessments of how risky the organization has been acting, to question the organizations assumptions about how it creates safety. How could a safety organization be designed to meet these ambitious goals since these are rather difficult cognitive functions to support in any distributed systems? Even worse, in order to avoid the trap of 'cold water and empty guns,' I was in effect asking the leadership of an organization to authorize and independently fund a separate group whose role was to question those leaders' decisions and priorities.

And then, if the safety organization was authorized and provided with an independent set of significant resources, it was committed to offer positive action plans sensitive to the limited resources and larger pressures imposed from outside. To accomplish this requires a means for safety management to escape the fundamental paradox of production/safety conflicts: safety investments are most important when least affordable. It is precisely at points of intensifying production pressure and higher organizational tempo that extra investments are required in sources of resilience to keep production/safety tradeoffs from sliding out-of-balance. What does Resilience Engineering offer as guidance to better balance this trade off?

The 4 'I's' of Safety Organizations: Independent, Involved, Informed, and Informative

At this point I had used a resilience perspective to provide common ground for an exchange on the dilemmas of safety organizations. But I was still on the spot and the staffers were insistent, how can safety organizations be designed to cope with these dilemmas? How did successful organizations confront these dilemmas?

To help organizations balance safety/production tradeoffs, a safety organization needs the resources and authority to achieve independence, to be involved, informed and informative. My response was that safety organizations are successful when they:

- provide an independent voice that challenges conventional assumptions about safety risks within senior management,
- have constructive involvement in targeted but everyday organizational decision making (for example, ownership of technical standards, waiver granting, readiness reviews, and anomaly definition).

- actively generate information about how the organization is actually operating and the vectors of change that influence how it will operate (informed).
- use information about weaknesses in the organization and the gap between work as imagined and work as practiced in the organization to reframe and direct interventions (informative).

These four 'I's' provide a simple mnemonic that concisely captures the difficulty in designing a safety organization: these 4 requirements are in conflict! At best the relationship between the safety organization and senior/line management will be one of constructive tension. Safety organizations must achieve independence enough to question the normal organizational decision making, provide a 'fresh' point of view, and help the parent organization discover its own blind spots. Challenging conventional assumptions of senior management limits the voice as fresh views bring unwelcome information and seem to distract from making definitive decisions or building support for current management plans. Inevitably, there will be periods where senior management tries to dominate the safety organization. The design of the organizational dynamics needs to provide the safety organization the tools to resist these predictable episodes by providing funding directly and independent from headquarters. Similarly, to achieve independence, the safety leadership team needs to be chosen and accountable outside of the normal chain of command.

Safety organizations must be involved in enough everyday organizational activities to have a finger on the pulse of the organization and to be seen as a constructive participant in the organization's activities and decisions that affect the balance across safety and production goals. In general, safety organizations are at great risk of becoming information-limited as they can be shunted aside from real organizational decisions, kept at a distance from the actual work processes, and kept busy tabulating irrelevant counts when their activities are seen as a threat by line or by upper management (for example, the 'cold water' problem). Simply by being positioned to have a voice at the top can leave the safety organization quite disconnected from operations and exacerbate information limits. By being informed, the safety organization can be informative, and the strongest test of this criterion is the ability to identify targets for investments to enhance aspects of resilience and to prioritize across these targets of opportunity. To be constructive, a safety organization needs to control a significant set of resources and have the authority to decide how to invest these resources to help line organizations increase resilience and enhance safety while accommodating production goals. For example, the safety organization could decide to invest and develop new anomaly response training and rehearsal programs when it detects holes in organizational decision making processes. Involvement, balanced with independence, allows the safety organization to provide technical expertise and enhance coordination across the normal chain of command. In other words, the involvement focuses on creating effective overlap across different organizational units (even though such overlap can be seen as inefficient when the organization is under severe cost pressure).

Balancing the '4 I's' means that a safety organization is more than an arm's length tabulator, does more than compile a trail of paperwork showing the organization meets requirements of 'safety' as defined by regulators or accreditors, is more than a cheerleader for past safety records, and more than a cost center that occasionally slows down normal production processes. Being involved and informed requires connections to the character and difficulties of operations (the evolving nature of technical work as captured e.g., in the studies in Nemeth, Cook & Woods, 2004). Being independent and informative requires a voice that is relevant and heard at the senior management level. By achieving each pair and making them mutually reinforcing, safety management becomes a proactive part of the normal conduct of the organization.

The safety organization's mission then is to monitor the organization's resilience including the ability to make targeted investments to restore resilience and reduce brittleness. In reaching for the '4 I's', the safety organization functions as a critical monitor of the gap between work as imagined and work as practiced and generates tactics to reduce that gap. As a result, the safety organization becomes a contributor to all of the organization's goals - by enhancing resilience both safety and production are balanced and advance together as new capabilities arise and as the organization faces new pressures.

Safety as Analogous to Polycentric Management of Common Pool Resources

The analysis above and the '4 I's' as a potential solution to the challenge case parallels analyses of how complex systems avoid the tragedy of the commons (Ostrom, 1990; 1999). The tragedy of the commons concerns shared physical resources (among the most studied examples of common pools are fisheries management and water resources for irrigation). The tragedy of the commons is a name for a baseline adaptive dynamic whereby the actors, by acting rationally in the short term to generate a return in a competitive environment, deplete or destroy the common resource on which they depend in the long run. In the usual description of the dynamic, participants are trapped in an adaptive cycle that inexorably overuses the common resource; thus, from a larger systems view the local actions of groups are counterproductive and lead them to destroy their livelihood or way of life in the long run.

Organizational analyses of accidents like Columbia seem to put production/safety tradeoffs in a parallel position to tragedies of the commons. Despite organizations' attempts to design operations for high safety and the large costs of failures in money and in lives, line managers under production pressures make decisions that gradually eat away at safety margins undermining the larger common goal of safety. In other words, maybe safety can be thought of as an abstract common pool resource analogous to a fishery. Or, alternatively, dilemmas that arise in managing physical common pool resources are a specific example of a general type of goal conflict where different groups are differentially responsible and affected by different sub-goals, even though there is one or only a couple of commonly held over-arching goals (Woods et al., 1994, Chapter 4).

Developing the analogy further, the standard view of how to manage common pool resources is to create a higher level of organization responsible for the resource over its entire range and over longer periods of time. This organization then needs authority to compel individuals or local groups to modify their behavior sacrificing short term return and autonomy in order for the higher level organization to analyze and plan behaviors that sustain or grow the resource over the long term - a command organization. Safety management theory often seems to make similar assumptions and propose similar responses, i.e., a command structure is needed from regulators to companies or from management to line operations that takes a broader view and compels workers and line managers to modify behavior for a long term common good.

Ostrom (1999) reviews the empirical results on how people actually manage common pool resources and finds the standard view unsupported by the evidence. Basically, she found that overuse by local actors is not inevitable and that command style relationships across levels of organizations do not work well. Instead, she finds from research on co-adaptive systems that common pool resources can be effectively managed through polycentric governance systems. Polycentric systems provide for multiple levels of governance with overlapping authority in a dynamic balance but where there is no single governance center which directs or 'commands' unilaterally. Her synthesis of research identifies a variety of conditions and properties for polycentric management of common resources (such as cross-communication, shared norms, trust, and reciprocity; Ostrom, 2003).

The proposed '4 I's' of safety organization design can then be seen as additional policy guidance for how to build effective polycentric management to balance multiple interacting goals. Achieving a dynamic balance across multiple centers of governance - some closer to the basic processes but with narrower field of view and scope of action and others farther removed but with larger fields of view and scopes of action, would seem to require a quasi-independent, intersecting organization that can cross connect these different levels of organization to be both informed and informative. By being outside a nominal chain of command, such groups can question and help revise assessments as evidence and situations change, as well as intervene with targeted investments to help resolve short term dilemmas (independent and involved).

Recent research on distributed cooperative systems made possible by new computer technology also seems to support the analogy, for example studies of the change to 'free flight' in managing the national air transport system support and extend Ostrom's findings (see Smith et al., 2004). The tools that have proved necessary to make collaboration work between air carriers and FAA authorities given new capabilities for communication at a distance and given the demands for adaptive behavior as congestion and weather change also provide other ideas for the design of polycentric management systems. Similarly, studies of how military organizations delegate authority to adapt plans to surprising situations provide lessons that also can be applied to guide polycentric management (e.g., Woods & Shattuck, 2000).

The analogy suggests that findings from managing physical common pool resources and findings from how goal conflicts between safety versus production are resolved (Woods et al., 1994, chapter 4) may converge and mutually reinforce or stimulate each other. For example common pool research may benefit from examining the reframing processes which are central to the resilience approach to safety under different management structures.

Summary

Organizations in the future will balance the goals of both high productivity and ultra-high safety given the uncertainty of changing risks and certainty of continued pressure for efficient and high performance. This organization will be able to (a) find places where the organization itself has missed or underestimated the potential for trouble and revise its approach to create safety, (b) recognize when the side effects of production pressure may be increasing safety risks and, (c) develop the means to make targeted investments at the very time when the organization is most squeezed on resources and time.

To carry out this dynamic balancing act, a new safety organization will emerge – designed and empowered to be independent, involved, informed, and informative. The safety organization will use the tools of Resilience Engineering to monitor for 'holes' in organizational decision-making and to detect when the organization is moving closer to failure boundaries than it is aware. Together these processes will create foresight about the changing patterns of risk before failure and harm occur.

Acknowledgements

This work was supported in part by grant NNA04CK45A from NASA Ames Research Center to develop resilience engineering concepts for managing organizational risk. I particularly thank the congressional staffers who provided an opportunity to review NASA's post-Columbia

reform plans and who challenged the concepts for achieving resilience. The ideas here benefited greatly from the inputs, reviews, and suggestions of my colleagues Geoff Mumford and Emily Patterson. The remaining gaps are my own.