What can IT security professionals learn from safety literature?

From Safety Science to IT Security

Mate Soos

@msoos@post.lurk.org
www.msoos.org

Based on ideas and thoughts by Mario Platt, Sidney Dekker, Nancy Leveson, and others

What this talk is going to be about

Safety engineering has a LOT to give to us, IT security engineers.

I hope that by highlighting some of the issues safety engineers have thought about, I can help you be better IT security engineers

This talk is to pique your interest in this domain. You'll likely want to read about these yourself, and discover the ideas that you find most interesting & useful, and incorporate them into your work.

- Clearly, security is not quite safety – in safety, there is nobody maliciously trying to fill your mine with methane and ignite it.
- But there are miners who got tired of the methane warnings & turn them off
- There is management who set production targets impossible to achieve via tickboxes, without innovative ways to do it safely
- There are warnings signs, with experienced miners leaving your company due to safety concerns

And we have all been in this mine.

What this talk is NOT going to be about

What you have been doing, whatever it has been, has not been “wrong”. Instead, there may be things you can incorporate into your work from this talk to better your work, and improve your organization’s resiliency.

You need to find your own way, there is no one true way of doing IT security. Not even if you follow ISO/NIST to the letter ;)

A lot of time, effort, and money went into safety science, and many people died because of failures observed in this domain. Maybe they have something to teach us.

Scenario and its Classic Interpretation

Scenario: attacker got into company account via phishing attack, pivoted around infra, but was stopped because engineer noticed unusual behaviour in cluster. Sysadmin team disabled accounts.

Why did this happen? How do we fix it?

Classic Interpretation:

- User clicked bad email. Train the user.
- We lacked preventative measures against pivoting. Add preventative measures.
- We lacked detective measures, since no automated system notified us. Add detective measures.
- Sysadmin team had to manually fix stuff: we lacked reactive measures.

Down and In vs Up and Out

Sharp end (proximal cause), going “Down and In”:

- The user who clicked the email
- The sysadmin team that reacted
- The system that was compromised

Blunt end (distal cause), going “Up and Out”:

- Lack of priority for detective measures. What did we do instead? Was it worth the trade-off? What influenced that decision? Can we fix the influence?
- Systems built without involvement from security team, hence no preventative measures. Do they know we should be involved?
- Same phishing email was observed by other employees. Why didn’t they notify IT security?

Who/What do we Blame?

Most leaders will concentrate on the sharp end:

- You will get trainings
- You will get detective mechanisms
- You will buy new preventative measures

Notice: it’s always easy to blame the sharp end. They were/weren’t there, and they didn’t do the immediate, practical thing they should have done.

Instead of:

- Understanding why security was not involved in developing the the new system
- Running security war games together with the sysadmin team
- Building better relationships with the company to build trust & common understanding, so we’d get alerted to phishing attacks

Positive vs Negative Slack in the System

What most would concentrate on:

- The missing training of the user not to execute malware
- The missing automated systems to detect the malware
- The lack of adherence to rules: security should be part of all new systems built

What we tend to forget:

- The detection of the system by someone who clearly didn’t follow any process or procedure. Just noticed something unusual.
- Reaction of the sysadmin team, even though it’s not their duty, but they understood the gravity of the situation, and acted as best they could.

IT Security as a Hierarchical Control System

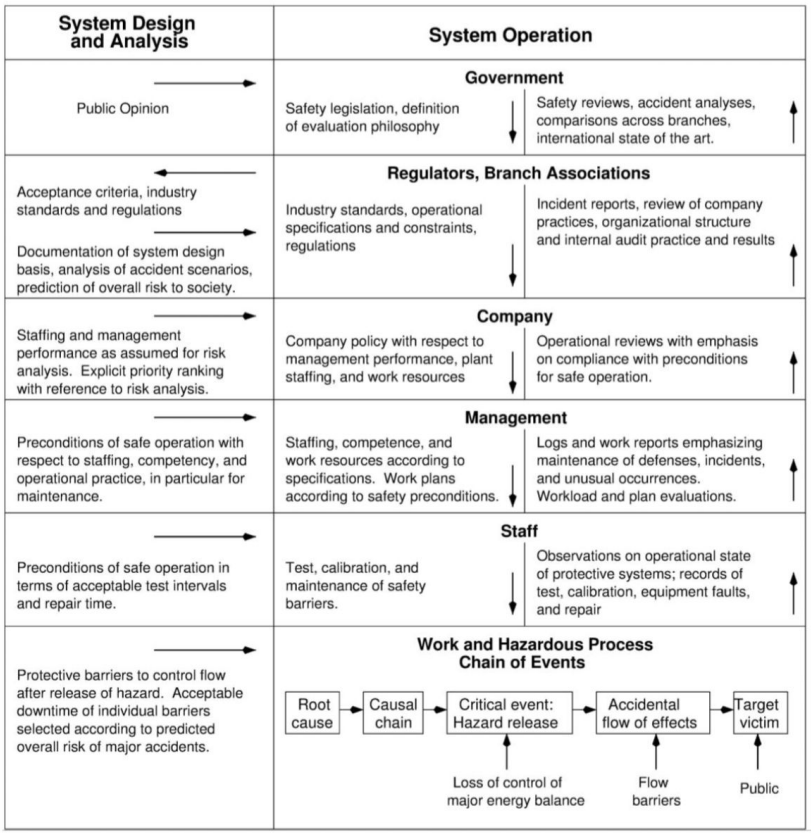

Figure 4: Hierarchical Model of Safety Control (Rasmussen, 1997)

- Each level observes the behaviour of the level below and pushes it with a control input in some direction
- Where are you on this chart?
- Where do you want to be?

Notice: there is no “right” place to be. The only mistake to make is to forget that there are layers

Rasmussen: Risk management in a dynamic society: A modeling problem (1997)

Adaptation of front-line work to a centralized control mode of safety

- **Plan**: Existing strategies, plans, roles, requirements and process that should be applied to activities. To an insider, the expectations and understanding of work never match the reality of what it takes, and how work gets done.
- **Fluency**: Well adapted activity that smooths over contradictions and challenges to make things work. To an outsider, the work seems well coordinated which hides the difficulties that they had to work around to make things work.
- **Discounting**: Problems and issues with front-line work are discounted by management and safety professionals if they are outside of work as imagined.
- **Double binds**: Managers, front-line workers and safety professionals face irreconcilable decisions between two simultaneously necessary but incompatible choices. Neither decision resolves the other issue.
- **Role retreat**: Front-line workers retreat to just performing their role as defined: 'work to role'. This undermines collaboration, especially when things are difficult.
- **Covert work systems:** Work as done is hidden from outsiders due to the fear that it will be stopped or changed. The greater the gap between work as done and work as imagined, the greater the effort that goes into keeping the shadow work systems underground. Work as done has the illusion of alignment with work as imagined through teams dutifully meeting outside expectations through surface compliance activity (i.e. tick and flick, lip-service).

Provan et al: Safety II professionals: How resilience engineering can transform safety practice (2020)

Centralized Control Mode, Safety Personnel Perspective

Safety personnel:

- Support the task-based identification of hazards
- Facilitate the identification and assessment of system level hazard
- Develop controls for tasks and processes
- Monitor controls proactively (e.g. inspections) and reactively (e.g. incident investigation)
- Provide safety incident and compliance reporting to line management and regulators
- Support line management decision-making and arbitrate between stakeholders as necessary
- (Promote an 'authority to stop work' for safety across the frontline workforce)
- Develop and promote safety culture improvement programs

Provan et al: Safety II professionals: How resilience engineering can transform safety practice (2020)

Guided Adaptability Mode

Safety personnel:

- Explore everyday work. Understand the way the organisation is currently operating and where resilience and brittleness is present
- Support local practices and guide adaptations for safety.
- Reduce goal conflict and negotiate redistribution of resources. Monitor goal conflict and create action to alleviate it.
- Facilitate information flows and coordinate action: create mechanisms to transfer information and coordinate action across organisational boundaries.
- Generate future operational scenarios. Utilise current understanding of the organisation to predict future conditions.
- Facilitate Sacrifice Judgements. Support the understanding of trade-off decisions and the resolution of acute goal conflict.
- Facilitate Learning. Create organisational change based on current conditions and future scenarios.

Provan et al: Safety II professionals: How resilience engineering can transform safety practice (2020)

Setters of Requirements vs Providers of Capability

“Are you writing uncontextualised policy documents that no one knows exist in the organisation, and then holding employees accountable when those requirements set with poor understanding of operational context aren’t met?” [Mario Platt]

Maybe we could be provisioners of capability as opposed to requirement setters?

Mario Platt: “The future of Infosec is interdisciplinary and integrated” [link]

Safety-I vs Safety-II (Hollnagel)

Safety I	**Safety II**
Learn from Errors	Learn from successes
Safety defined by absence	Safety defined by presence
Ractive approach	Proactive approach
Understanding what goes wrong	Understanding what goes right
Accident causation models	Repeat what goes right
Avoidance of errors	Enforce successful behaviours
Reducing losses	Create new processes based on successful behaviour

Hollnagel: A Tale of Two Safeties (2013) [link]. Also: https://safety4sea.com/cm-safety-i-vs-safety-ii-an-overview

Drift Into Failure

- We have been running with risky system X for Y years
- Therefore, risky system X is fine, how about we stretch it further?

Classic example is the Challenger accident. But e.g. downing of black hawks over Northern Iraq in 1994 was a similar incident.

- Notice that risk is evaluated by people
- People form a group, and reinforce each others' beliefs
- Decision making is influenced by group dynamics
- People come and leave the organization. The "extended norm" is "the norm" for the new-hires.

Have you noticed that new employees are sometimes astonished at some of the risks you are running with?

Dekker: Drift Into Failure (2011)

A Side Note on Accepting Risk

Have you thought about “risk to whom”? Is that a column in your risk sheet? Who is represented in the columns?

- 1st party victims? (e.g employee)
- 2nd party victims? (e.g. customer)
- 3rd party victims? (e.g. people who your customers talked about in their DMs)

… or only the business?

Note that the people accepting the risks are not the only ones suffering the consequences. It may be the ethically correct choice to keep this in mind when accepting risks for people who likely didn’t consent to be put at risk.

High Reliability Organizations (HROs)

A relatively new trend in health, but used widely in mining, oil & gas, aeronautics, i.e. in high-risk organizations.

Wikipedia: “A high reliability organization (HRO) is an organization that has succeeded in avoiding catastrophes in an environment where normal accidents can be expected due to risk factors and complexity.”

Characteristics of HROs:

- Preoccupation with Failure: Process Failures are Addressed Immediately and Completely
- Reluctance to Simplify: Complex Problems Get Complex Solutions
- Sensitivity to Operations: Every Voice Matters
- Commitment to Resilience: Recovery is Swift
- Deference to Expertise: Experts are Trusted (rather than authority)

Safety Currently vs Safety Differently

Safety Currently	**Safety Differently**
People are a problem to control	People are the solution
Tell them what to do	Ask them what they need
Count success by absence of negative	Count positive capacities

Dekker's talk at DevOps Summit [link]

Any questions?

Mastodon: @msoos@post.lurk.org

Email: mate.soos@gmail.com

Blog: www.msoos.org

GitHub: @msoos

References:

- Mario Platt: Security Differently - Insights from RE and Safety (2022)
- Mario Platt's blog at: https://www.securitydifferently.com/
- Provan, Woods , Dekker, Rae: Safety II professionals: How resilience engineering can transform safety practice (2020) [link]
- Sidney Dekker: The Safety Anarchist (2018)
- Nancy G Leveson: Rasmussen's Legacy: A Paradigm Change in Engineering for Safety (2016) [link]
- Sidney Dekker: The Field Guide to Understanding 'Human Error' (2014)
- Hollnagel: A Tale of Two Safeties (2013)
- Sidney Dekker: Drift Into Failure (2011)
- Snook: Friendly Fire: The Accidental Shootdown of U.S. Black Hawks over Northern Iraq (2000)
- Rasmussen: Risk management in a dynamic society: A modeling problem (1997)
- Diane Vaughan::The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA (1997)
- Charles Perrow: Normal Accidents (1985)
- Barry A. Turner: Man Made Disasters (1978)

On Empire Building

“Paperwork begets paperwork” is what Dekker termed “Bureaucratic entrepeneurialism” [1]. We fear the consequences of curtailing security functions. Combined with the promise of future useful work and reminders of past successes, it helps to perpetuate their existence and serve as justifications for more of it.

When this is combined with security leaders interested in Empire-building, it’s a surefire recipe for long term disaster.

- Leads to bureaucratic accountability as the central factor of providing security to the organization
- Leads to institutionalisation and legitimatisation of counting negatives (non-compliances, deviations, incidents). Most of the language revolves around deficit and control and that we need more of it.

When was the last time you counted the positive in your system?

Nancy G Leveson: Rasmussen’s Legacy: A Paradigm Change in Engineering for Safety (2016)
[1] Dekker: Safety Differently: Human Factors for a New Era, Second Edition Safety Differently (p. 32)