An Insider Threat Framework– The SOFIT Ontology

Author: Tom Roberts

An Industrial Spy, Disgruntled Employee, Unhappy Customer and Business Partner and Recruiter walk into a virtual bar…. TOGETHER!

Just take a minute to imagine if this happened to your organisation…

The above may have sounded like the start of the world's worst joke, but it does occur. The significant risks of Insider Threat are still hard to quantify, monitor or prevent.

So, imagine the following scenario: You are the CISO of a large firm and your head of sales calls to inform you that the recent tender data your firm has been working on for three months, is now in the hands of your direct competition. You infer they will undercut you and seem to have valuable information about client projects and sales pipelines. You call the SOC and engage the security teams. You haven't been hacked by a nation state, but your data has been removed from your network by an insider threat.

There has been quite a lot of discussion about insider threat and the detection of it in the news recently. I became interested in this a few years back after reading a paper published in 2019 by Frank Greitzer (https://scholar.google.com/citations?user=ux7GO6sAAAAJ&hl=en) on forming a framework for detecting insider threat that modelled itself on elements provided by PERSEREC and others. It was called **S**ociotechnical and **O**rganisational **F**actors for **I**nsider **T**hreat Detection or SOFIT. It delved into how to categorise human based internal threat and calculate a metric to assess potential early warning indicators in behaviour as well as gaps within processes and procedures firms use that give rise to accidental internal threat, and not just focussing on the intentional or premeditative insider threat.

At first glance, and to many who read only parts of the wider picture and don't delve into the many other papers written about SOFIT, you may see aspects of science fiction and comparison to books like Minority Report by Philip K Dick. Some current predictive modelling does come with a slightly invasive side of human analysis and analytics. But this is only part of the picture and the wider ontology focusses on both the staff AND the business processes that staff adhere to.

It recognises that people are a vital part of business systems and that their failure can have wider impacts. Businesses should seek to formulate processes and procedures to monitor, assist, support and eventually protect themselves from intentional or accidental insider threat. Both myself and Mr Greitzer are here to stress it is not just about law and order and punishment. It is about creating a workplace that creates points of intervention and "***instead it argues for a proactive, comprehensive approach that seeks to help troubled staff find "offramps" to the critical pathway leading to the insider exploit. The objectives include individual and organizational health, safety and well-being, rather than solely focusing on finding and punishing bad guys. And the emphasis on contributing organizational factors is perhaps the most proactive aspect of the approach, since organizational changes that correct systemic organizational problems (such as toxic work environments, work-culture issues, workload stress, etc.) can be instrumental in preventing the growth of insider threats (i.e., by eliminating issues/environments that act as a 'breeding ground' for insider threats). This applies to both intentional/malicious and unintentional/accidental insider threats***." (Frank Greitzer 2021).

So, what is accidental insider threat? Well, it's the times when staff are either pushed to make a decision that negatively impacts the business, or, through a preventable accident, creates a security incident. Lack of resource, impending deadlines, JFDI culture and time saving short cuts are often the root causes of such insider threat and can have predictive warning signs which may be solved with cultural changes, process adaptation and policy adherence.

Examples of both accidental and intentional insider threat. Two tales about Alice, Bob, Charlie, and Dave

These can be quite complex or relatively simple scenarios and SOFIT can help in both. We will start with a firm that hasn't thought about SOFIT and has some cultural issues and large workforce spread across many hierarchical teams.

> *The first example has an intentional factor with a company that has good security but no early warning indicator framework. Alice notices Charlie has been working late recently and missing a few meetings. This ties in with a notification from IT that Charlie had installed a gambling app from the app store and, whilst it was not prohibited in policy, the policy will soon be updated and Charlie was asked to remove it, which he did, quite promptly.*
>
> *This prompts Alice into contacting HR. HR inform Alice that there are two recent incidents on file that were logged in relation to Charlie. One incident where Charlie swore at a co-worker; Dave, and was attempting to use a faux managerial excuse to try and make Dave supply Charlie with a database of client data for a project that did not yet require that data, and would not, for at least another three months. Dave was surprised at being asked to provide such a complex data set in a short period and even less enamoured with the approach in which Charlie was asking for it, by swearing and using demeaning language. Dave had logged a complaint against Charlie's language and general aggressive tone and behaviour.*
>
> *The next incident logged was in fact a grievance report by Charlie against Dave. This was logged almost the same time as Dave complained and Charlie had raised a managerial issue, regarding Dave's efficiency and work ethic, which Dave was unaware of. The incident led to a meeting with Dave's manager, which distressed Dave, as he felt he was following secure protocols and procedures. Dave is now more submissive and less likely to complain when asked by Charlie to provide information. Alice concludes it is office politics and a petty squabble and prioritises other problems.*
>
> *Next week, Charlie then approaches Dave again and asks for the data. When Dave cannot immediately comply due to workload, Charlie offers to assist and help create the data. As Charlie does not have access to the database, Charlie offers assistance, **if** Dave will share his credentials then Charlie can extract the data and save Dave some time. Dave isn't sure why Charlie had changed his behaviour so drastically from previous encounters, but Dave feels happy to halve his workload. Dave complies, as he doesn't want any more trouble, managerial comments or complaints made. Dave believes these complaints originated from Charlie, so Dave is reticent to rock the boat further and just wants an amicable workplace.*
>
> *Charlie now sees Dave as compliant and confides in Dave that he was recently passed over for promotion due to a poor yearly review and that Charlie needed to "show them how smart he is!". Dave thinks Charlie may even have been drinking during the day whilst in his home office. To add to this, Dave and others have noticed that Charlie isn't the happy go lucky person he once used to be and seems to have few nice things to say about work.*

A couple of weeks later Charlie hands in his notice. He's found another role and brags to others how much more he will be earning and how "valuable" he will be at the new firm. Charlie departs rapidly using unused holiday and there is no leaving party at Charlie's request.

Two weeks later Dave is called in the management office regarding an out of hours data extraction and backup of several databases Dave has access to, which happened over four weeks previously. Dave recalls no logins after hours and certainly no extraction of any data the firm is detailing. Dave has no idea what is happening. Dave is distressed and cannot recollect any of the mentioned events, as it was over a month ago, and the firm is considering dismissal as there is some evidence that a competitor now holds the data. Dave loves his job and never wanted a situation like this. He just wanted a quiet life, he follows orders and tries not to raise too many issues. Dave is aware the firm rarely gives second chances.

What Dave doesn't know, was that Charlie used his credentials to obtain the data Dave so helpfully provided access to, and quite a bit more. Charlie then used other means to extract this data and exfiltrate to a cloud-based server the firm does not control. Dave's record is tarnished, and Dave is now thinking of leaving as the accusations made during the investigations showed that he was under suspicion until a forensics team worked out the final exfiltration route and the probable culprit. Dave couldn't be fully exonerated, so Dave was deemed untrustworthy. His job prospects within the firm are now limited. He may even become disgruntled as time progresses.

Remember Alice, Dave, and Charlie – They will be referred to later.

There are also simpler scenarios where application of SOFIT can remove possible accidental insider threat.

In this instance we shall take an example of a company taking SOFIT into account and a relatively simple solution where SOFIT seeks to lower accidents caused by poor process or workers with resource issues:

Alice is Bob's line manager. She is aware that Bob's partner has recently given birth to their first child. Bob took his paternity leave, but health issues have meant his partner has had to remain in medical care longer than expected, which resulted in Bob undertaking duties and trips to hospital that they had not planned for. Bob's tiredness and personal stress levels are high. Alice recognises these as early warning indicators of accidental threat and offers him assistance by reducing Bob's workload and extending the part time cover that was applied during his paternity leave and insists that he does not work out of hours to make up time as these could be periods when attention to detail may not be high enough for the required role. This is not a poor reflection on Bob; it is merely the business protecting itself from accidental threats which Bob could inadvertently cause due to tiredness, stress, lack of focus or other immediate distractions. As this is not a permanent situation, the cost-effective method would be to assist Bob in the short term thereby maintaining his skills and lowering long term disruption. Alice's rapid intervention at an early stage helps lower Bob's long-term stress, improves his family engagement and personal relationship with the firm, lowers disruption caused by mistakes and errors, and hopefully reduces the time in which business as usual is established.

There are two parts to this example. The process the company follows which creates a framework in which to enact mitigation, and a human awareness framework that allows for staff to indicate their situation to management and for management to have the tools to look for variation of behaviour or indicators that could be early warning signs of intentional or accidental internal threat. SOFIT seeks to give indicators of accidental and intentional insider threat so companies can take mitigative action before the loss occurs.

The framework and how to apply it

Let us have a quick look at the SOFIT framework developed by Dr. Greitzer and colleagues (see footnote [1]. First, the individual factors branch of the hierarchy describes an individual's actions and psychological characteristics.

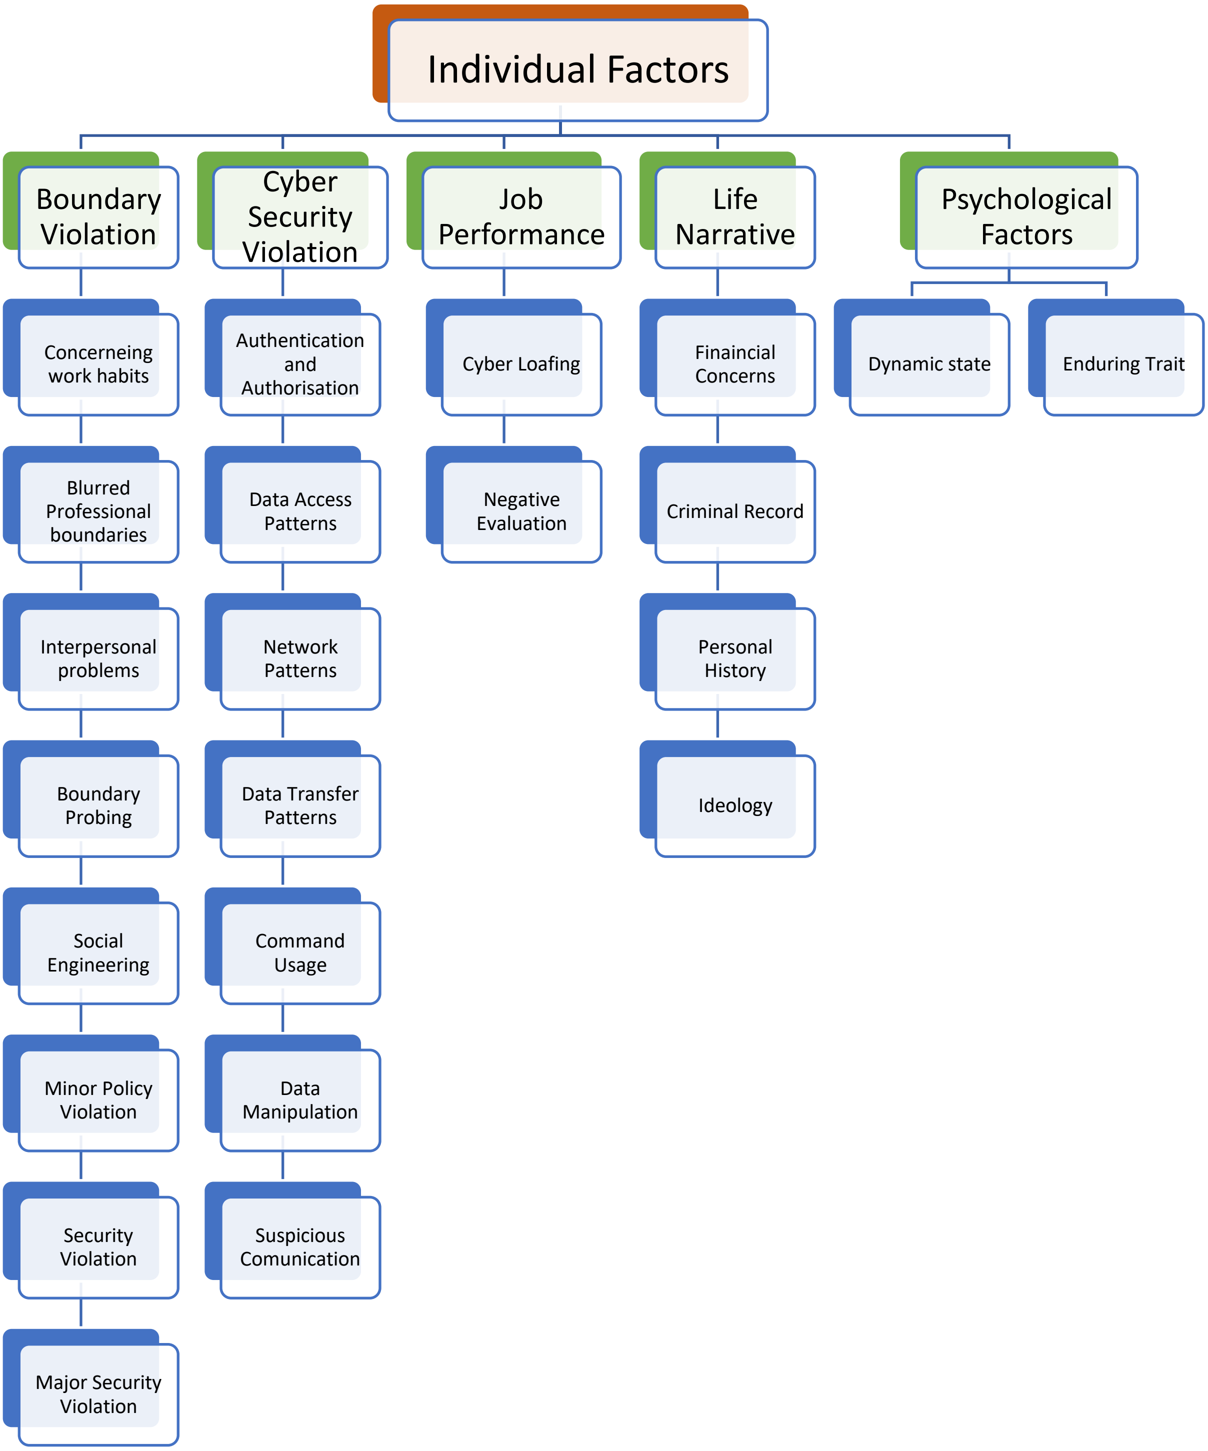

Figure 1 - SOFIT Individual Factors (Figure adapted from works of Dr. Greitzer) - SOFIT: Sociotechnical and Organizational Factors for Insider Threat | SOFIT: Sociotechnical and Organizational Factors for Insider Threat | IEEE Conference Publication | IEEE Xplore

But what about those Individual factors? They do seem on the face of it to be directed towards intentional threat, but don't let that fool you. They are about creating baselines for notification of changes in behaviour that might illicit red flags or warning signs to allow early intervention and prevention, rather than having to rely on reactionary forensics.

The diagram highlights there are many Individual Factors to consider. These include personality, and life narrative as baseline comparative values as well as more business-related areas around boundaries and work performance. All of these expand into metricised values and validation points, some of which can be monitored in pre-existing tools within Azure and other cloud-based systems to a limited degree.

They give businesses a heads up when people may be working out of pattern, or access areas of the network not normally associated with their role. It is fair to say the Machine Learning tools available are still being rapidly developed. Whilst some are more logging than predictive; they do create warning flags for review.

The second area is Organisational Factors. These are areas within a business that could allow insider threat. This framework can be used to highlight areas of processes where accidental or intentional insider threat can occur. This would include review of poor process, work culture or other organisational factors which can create an incubator of internal threat potential or obfuscate unintentional or intentional loss.

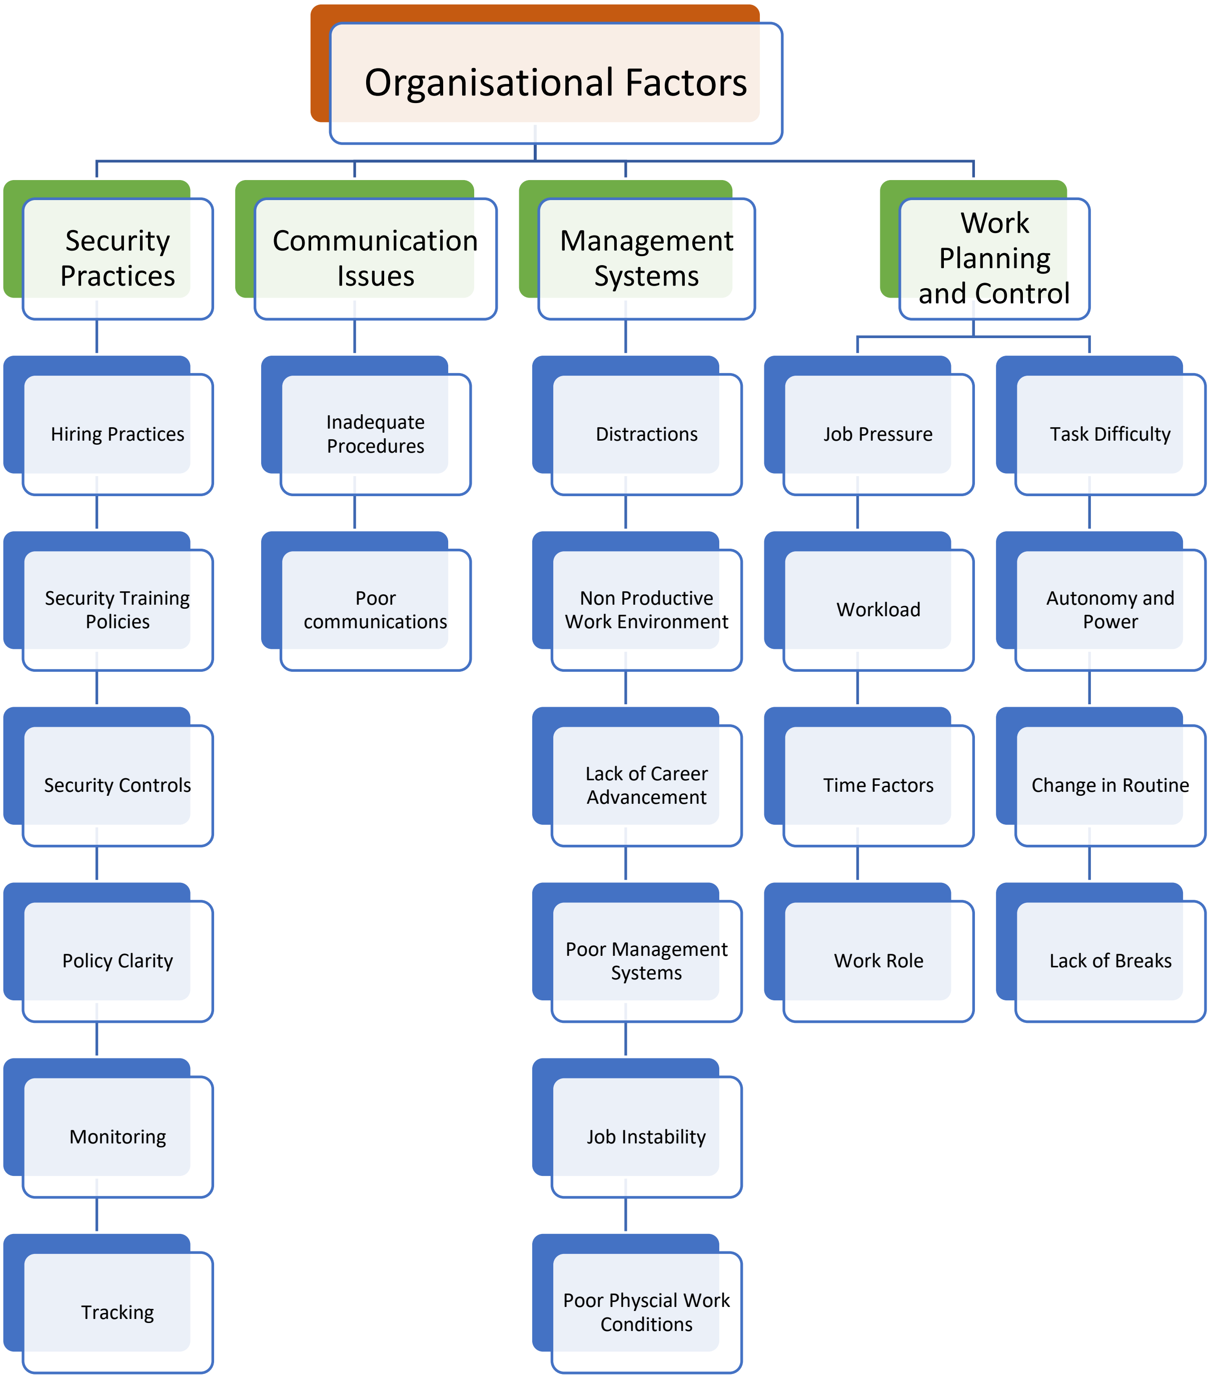

Figure 2 - SOFIT Organisational Factors (Figure adapted from works of Dr. Greitzer) - SOFIT: Sociotechnical and Organizational Factors for Insider Threat - SOFIT: Sociotechnical and Organizational Factors for Insider Threat | IEEE Conference Publication | IEEE Xplore

Gaps or imbalances in the above areas give rise to possible insider threat. For example. A lack of clear policy may encourage accidental abuse because the staff "didn't know".

A lack of perceived career advancement can give rise to feelings of disgruntlement and more importantly aspects of being “owed” something by the company for that lack of progression, giving possible precursors to a member of staff departing and taking data or intellectual property away with them, especially if they helped create it.

Now, we shall briefly discuss some of those operational security practices. Many will not be new, such as hiring practices; vetting and clearances or background checks and validation of an applicant CV. Some are regularly adopted, like security training and policies. Many are deemed part of IT teams… and this is where it may start to break down a bit. Policy clarity, and how it is applied technically can be two different things, and monitoring and tracking are often achieved but sometimes not reviewed with internal threat in mind, depending on the size of the firm or the resources it has available.

Track and monitoring are often sensitive subjects too. No one wants to feel like they work in a prison but there are procedures for monitoring and tracking that staff may forget can help them. It may help prove someone did not do something, as much as it may help prove they did. In today’s world of ever-growing ransomware and 0-day exploits being used in wild abandon, monitoring and tracking are the key tools to determining the how, when why and what of a breach. Even if it is after a breach, lessons can be learned from insider threat frameworks that help close gaps for threat actors to exploit.

Communication Issues, Management Systems, and Work Planning and Control align heavily to old school ERP, MRP and process flow models that many companies are aware of. The common security principles of segregation of duties, role-based access control, oversight and validation between connected processes are all there too. Management interface, empowerment, guidance, accountability as well as staff engagement are intermixed with these to establish a more secure culture and more proactively secure workforce.

Noticing the early warning signs

The following conceptual diagram illustrates how the combination of technical and behavioral factors may provide a more complete picture that can reveal early warning signs of potential insider threat risk.

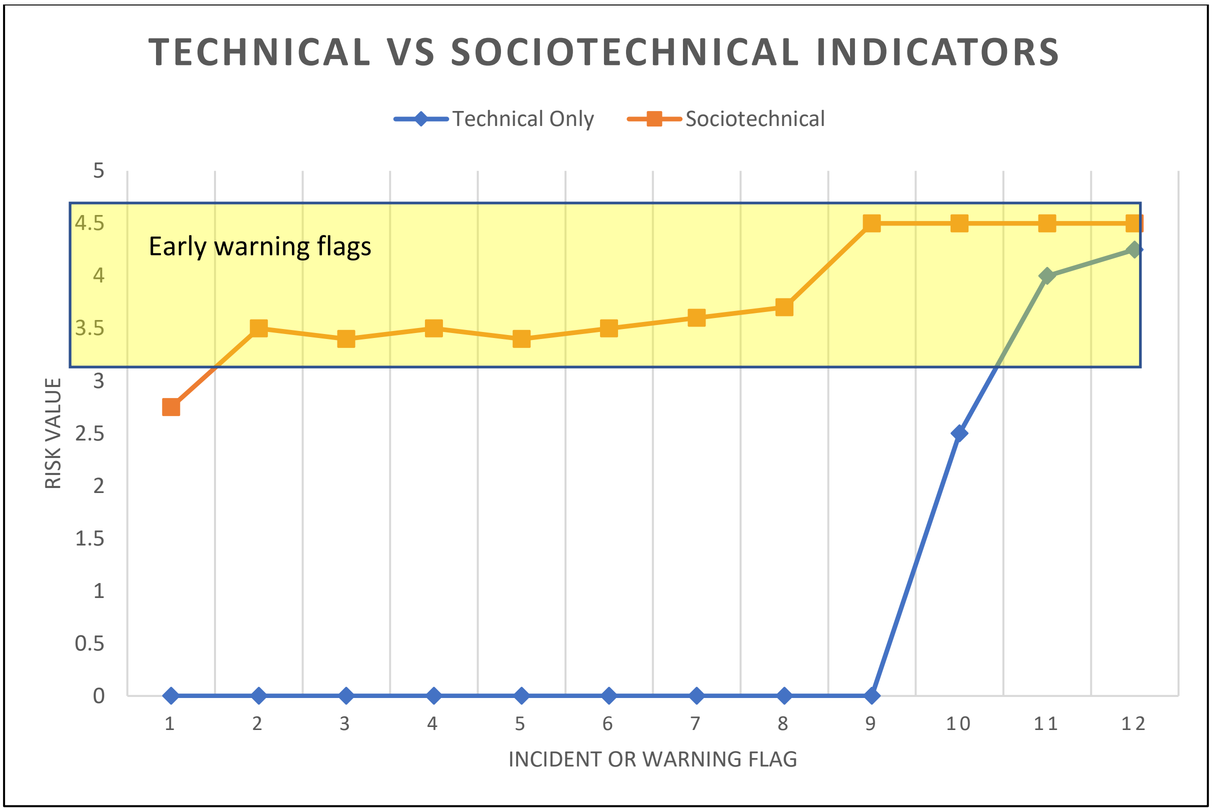

Figure 3 - [illustration adapted from the works of Dr. Greitzer.

Often referred to as "left of the bang" the graph highlights a potential internal threat actor and the personal traits they display to colleagues and co-workers vs. the technical only indicators they may present to network logs. The risk is the Y axis and the indicators are along the X axis.

Let us say that the following indicators were shown, or were highlighted, as part of an evaluation, and cast your mind back to Alice, Charlie and Dave and recall the timeline and how this ties in with the early warning indicators:

1. Big Ego – maybe this is normal – Several groups such as sales often rely on such behaviours. But are they maybe more egotistical than normal? Or maybe their outlook has changed, and their ego is drastically diminished or subdued and out of character?
2. Manipulative behaviour – Rarely a good thing and can lead to stress in other workers and disruption in work forces. But maybe this isn't manipulation and is a complaint to encourage employees to improve output. This may not be destructive, yet.
3. Callousness to co-workers – Not always intentional. Some people have personality types that mean they are not as emotionally aware of others and may say things that offend without actual intent to abuse.
4. Verbal Abuse – We are starting to cross some lines here. This might even turn into an HR issue.

5. Intimidating behaviours – Is this person a bully? Have they got other problems that are causing them to lash out?
6. Threats of retaliation – Dissatisfied at a poor management review and then discussions with colleagues where use of language may speak of "making the firm pay" or someone stating they are "taking what they deserve" can be early warnings.
7. Excessive absence or lack of engagement with other co-workers – Maybe deciding to apply for another role? Maybe an Illness or personal issues.
8. Overly critical – Nothing is right for this person and nothing will make it better. This might be their tipping point. Is this where they decide to leave and maybe take something with them?
9. Access/Policy Violation – They test their physical or network limits and access. Nothing overly overt and just seeing what is available or looking for low hanging fruit. Is this just an accident or a process flaw?
10. Working unusual hours – They believe that monitoring is not happening and that out of hours means out of sight. Or maybe this is a shift pattern and perfectly normal.
11. Access to prohibited materials – It might be client lists, code sets, financial documents. Whatever they feel will give them value in their next role. Do they have a need to know and a right to access?
12. IP theft or loss – they find a means of exfiltrating the data via physical or cloud-based means or even just emailed to a web-based email account.

Any one indicator would not be proof of intent and may even go unnoticed and as the examples show without enough data any single indicator may be completely innocent, or have some other factor that underlies the problem. Many or all of them are far more indicative and may tell a story, or highlight behaviours, that are intended to obfuscate the actual intent.

A company might only know they lost data at points 11 or 12 or even after it happened and sometime later. Many proactive firms might take preventative steps at points 7 or 9. SOFIT seeks to show that maybe the time to take steps to correct behaviour, or even disciplinary action, was at points 5 or 6. Also, the time to have gentle words or encourage improved engagement might be at 2 to 4. It may also be a point to ask; does your work culture allow for behaviours like these to go unchecked or unrecorded.

The issue is that hindsight is 20/20 as they say, and as a result, SOFIT seeks to give early warning indicators based on science and wide population sampling that can provide early warning signs to take less aggressive and thus, potentially, less disruptive action to prevent internal threat.

Issues or Early Indicators

There are even some potentially obvious indicators of compromise that can be early flags that HR or IT may already consider. They can be:

- Large amounts of stored leave or excessive unused vacation – shows a lack of wishing to leave work, it may even hide a control mechanism to prevent detection. It is also poor work life balance and should be discouraged.
- Consistent first in and last out of the office and is not a member of the board or the owner. Again, this may be a sign of access into areas or activities they don't wish others to witness. It, again, may also be indicative of a poor work life balance and should be discouraged in regard to mental wellbeing for all staff.
- Significant life change – sudden wealth, sudden change in marital status. Death of a spouse or a child. Many are not indicators of ill intent but may be precursors to behavioural changes that are early warning signs. Help and assistance rather than correction and penalties may be needed.
- Layoffs, redundancies – These will often cause a fight or flight response and may see sudden changes in behaviours that can be indicators of who is likely to leave with files full of paperwork or a USB drive full of valuable data.
- Passed over for promotion, rejected for a raise – may cause feelings of anger or disgruntlement. How this person is engaged with before, during and after will be key in preventing them turning from unhappy - to angry or vengeful.
- Disciplinary action – Again how this is handled and how the person reacts are key indicators of how they might approach any departure or improvement.
- Increased number of logins from remote or unknown IP's, logging in at odd times or out of hours on a regular basis or even just once if the IP address comes from an untrusted IP range or country.
- Using other people's logins – sharing passwords is bad practice and should be discouraged for several reasons. Finding someone abusing such passwords is rarely beneficial and if this is a business process that drives it - then understand it can hide both intentional and unintentional threats.
- Change in website visitation behaviour – betting, recruitment, web email, cloud storage, and a variety of others may be purely innocent but may also be behaviour indicators that might be warning signs.
- Export of large data volumes or significantly increased traffic – It may be related to a project, it might be export of data.

None of these will be a single issue that will indicate an insider threat. There is no "one weird trick" and the tools are many and varied and expand into facial recognition, body language and facial expression understanding. Remember, not all insider threat is intentional. You may flag data exfiltration as part of the exercise only to realise it's part of a flawed client contact method, or communications channel, and not someone intentionally trying to exfiltrate it.

Conclusion

SOFIT is based on both social and technical science and many papers and studies over large data sets, both commercial and social. Whilst still in its development, and being expanded on by others, it has potential to allow firms to recognise employee stress factors earlier and thus help prevent incidence of internal threat both accidental and intentional.

This rather long explanation is just the tip of a very large and fascinating topic which is starting to integrate with other elements such as word and language analysis, body language, facial reading or FACS, personality baselining, functional process, staff engagement, as well as machine learning and AI techniques to highlight changes and give indicators. Remember, we said indicators and NOT proof, but it is data which may help with early warning and lowering the incidence and impact of insider threat. This is just an introduction, and the subject matter is wide and varied in both personal and technical fields. The field is moving at a fast pace, and with outstanding rates of change and adaption.

If you think you need to start down the path of insider threat detection, then here are the first steps.

- Consider threats from insiders and business partners in enterprise-wide risk assessments.
- Clearly document and consistently enforce policies and controls.
- Institute periodic security awareness training for all employees which includes personality factors and indicators.
- Monitor and respond to suspicious or disruptive behaviour, beginning with the hiring process.
- Track and secure the physical environment. This includes removable media and non-trackable storage, backup or hard copy.
- Implement strict password and account management policies and practices.
- Enforce separation of duties and least privilege.
- Consider both accidental and intentional insider threats in the software and hardware development life cycle, as well as the business-as-usual procedures.
- Use extra caution with system administrators and technical or privileged users.
- Implement auditable system change controls.
- Log, monitor, and audit employee online actions; especially regarding privileged access, commands, or applications.
- Use layered defence against remote attacks.
- Deactivate all computer access directly following or at the time of termination.
- Implement secure backup and recovery processes.
- Develop an insider incident response plan. Include lessons learned and improvement to close gaps.
- Develop an insider threat prevention plan.

SOFIT provides a new sphere of analysis for understanding insider threats in advance. It enables Organisations to monitor their workforce and business partners to identify potential insider threats proactively and provides the opportunity to address them, before the risks becomes a reality, or the individual is even aware!

It's fascinating new stuff and you may already have some tools available within your estate!

[1] There are multiple papers by Frank Greitzer. The following all highlight or discuss the SOFIT ontology and are worth a read if you are interested in the topic and wish to explore further.

a) Greitzer, FL, J Purl, PJ Sticha, MC Yu, & J Lee. (2021). Use of Expert Judgments to Inform Bayesian Models of Insider Threat Risk. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA), 12(2), 3-47. June 2021*. DOI:10.22667/JOWUA.2021.06.30.003 https://dx.doi.org/10.22667/JOWUA.2021.06.30.003

b) Greitzer, FL, J Purl, YM Leong, & PJ Sticha (2019). Positioning your organization to respond to insider threats. *IEEE Engineering Management Review, 47(2),* 75-83. https://ieeexplore.ieee.org/document/8704879

c) Greitzer, FL, J Purl, YM Leong & DE Becker. (2018). SOFIT: Sociotechnical and Organizational Factors for Insider Threat. IEEE Security and Privacy Workshops (SPW), Workshop on Research for Insider Threat (WRIT), San Francisco, CA, May 24, 2018, pp. 197-206. DOI: 10.1109/SPW.2018.00035 http://conferences.computer.org/sp/2018/Resources/spw/2018/SOFITSociotechnicalandOrganizationalFact.pdf

Diagrams and illustrations reproduced with kind permission by Frank L Greitzer PhD. www.psyberanalytix.com

(Greitzer et al., 2016, 2018, 2019 , 2021)