DOI:10.1145/3418296

How to avoid insider cyber-attacks by creating a corporate culture that infuses trust.

BY ERIC GROSSE, FRED B. SCHNEIDER, AND LYNETTE L. MILLETT

Implementing Insider Defenses

CLASSICAL APPROACHES TO cyber-security—*isolation, monitoring, and the like*—are a good starting point for defending against attacks, regardless of perpetrator. But implementations of those approaches in hardware and/or software can invariably be circumvented by *insiders*, individuals who abuse privileges and access their trusted status affords. An organizational culture in which people and procedures are part of the system's defenses is thus necessary. Such a culture would instantiate classical approaches to cyber-security but implemented by people who follow administrative procedures. So, a careful look at a system's defenses finds that many of the same classical approaches reappear at each level. But the implementation at the lowest layers—structures we might term *insider defenses*—involves people.

People do not slavishly follow administrative procedures the way a computing system executes its programs. In addition, people are more prone than computing systems to making errors, and people can be distracted or fooled. Finally, because they can

be influenced by events both inside and outside of the workplace, people have very different kinds of vulnerabilities than computing systems. But people alter their behaviors in response to incentives and disincentives and, when empowered by organizational culture, they will (unlike computing systems) respond in reasonable ways to unusual or unanticipated circumstances. Thus, the use of people in a defense both offers benefits and brings different challenges than using hardware or software.

Those benefits and challenges are the focus of this article, which is informed by some recent discussions^a about best practices being employed at global IT companies and at the U.S. Department of Defense (DoD) for defense against insider attacks. The private sector and DoD are quite different in their willingness and ability to invest in defenses, in the consequences of successful attacks, and in the inclinations of their employees to tolerate strict workplace restrictions. Given those differences, two things we heard seemed striking and worth documenting for broader dissemination: How similar are the practices being used, and how these organizational structures and procedures to defend against insider threats can be seen as instanti-

a These discussions were facilitated by the National Academies Forum on Cyber Resilience.

» key insights

- Remarkably similar practices are being used at large-scale employers that range from DoD to the private sector, despite the significant differences in employee incentives, assets to be protected, and institutional cultures.
- People, processes, and culture are key to defending against insider attacks. Many of the processes will be implemented by people doing just what computer security mechanisms do: monitoring, replication of independent function, and logging.
- Carefully audited and constrained automation of a datacenter's system administration can replace most super-user access by people.

ating some classical approaches to cyber-security.

Assessing Risks from Insider Attacks

Risks from insider attacks will be part of any credible security assessment for an organization. In doing such an assessment, assets along with the protections they merit must be enumerated. That list is likely to include integrity and confidentiality of information about financial and customer data, confidentiality of intellectual property, integrity of system functionality, and availability of services. A risk assessment for insider attacks also requires determining which individuals and roles within the organization are being trusted and for what, as well as how those trust relationships are maintained and updated as roles change and as changes are made to the organizational structure itself. Part of this approach, articulated by Phil Venneables,⁹ then ‘a financial services CISO and Board Director at Goldman Sachs Bank, is to understand each role in the organization and the potential impact subverting that role could have; the aim would be to ensure no individual’s role has the potential for damage that exceeds the organization’s risk tolerance. Note that operational challenges of effective and comprehensive insider risk mitigation might delay deployment until other areas of an organization’s security program are mature but understanding and communicating the insider risk is nevertheless worthwhile.

Compromised insiders not only pose a threat to an organization’s assets but are a threat to organizational stability (for example, through personnel or organizational changes made in reaction to a compromise), mission

success (for example, when critical products fail to perform as expected), and customer satisfaction. Insider attacks also are an obvious vehicle for perpetrating supply chain attacks on an organization’s immediate or downstream customers. Moreover, certain functions and activities within an organization might be sensitive enough to warrant protection from inadvertent mistakes or accidents by even trustworthy employees. Many defenses against insider attacks can serve here as well.

Technical Controls

Classical technical controls and mechanisms play a key role in defending against insider attacks.

- ▶ Authorization prevents actions that compromise a security policy.
- ▶ Audit creates deterrence through accountability.
- ▶ Logging facilitates recovery after a security compromise.

Authorization is facilitated when the Principle of Least Privilege⁸ is followed, since access by individuals is then limited according to need, which might be characterized (and, therefore, validated) by using past activity, time, location, and role. So, role-based access control⁴ seems preferable, as should fine-grained privileges over coarse-grained ones.

Highly privileged administrator and operator accounts should be eschewed, which has led at least one global IT company to replace traditional system administrator root activities with automated systems that enable most data-center operations to run without the need for participation by individuals having root privileges. If fewer people have root privileges, then fewer people can abuse those privileges.

The choice between preventing action a priori versus deterrence through

accountability a posteriori often is dictated by the feasibility of monitoring, detection, and/or recovery. Can a compromise be detected in a timely way? Is recovery from a compromise possible? Of course, prevention should be preferred for actions that could lead to immediate and catastrophic outcomes.

Tamper-resistance of logs and backups is critical for implementations of accountability. For logs, tamper-resistance helps ensure attackers cannot easily change the logs to cover their tracks in order to avoid accountability. And for backups, tamper-resistance prevents attackers from installing system modifications that would perpetuate their presence after detection and restart. All accesses to sensitive data (see the sidebar “How Big Data and AI Complicate Things”) or functionality should be logged and attributed to an individual, if deterrence through accountability is intended. For programmatic access, log entries would indicate what program is running, what were its arguments, who is running the program, as well as who wrote and/or reviewed the program. Mechanisms that are guaranteed to intercept all requests and a strong form of authentication are thus a necessity for an implementation of logging.

Mechanisms to support deterrence through accountability can be further leveraged if the mechanisms also perform checks and/or pause to interrogate a developer or operator whenever sensitive or risky actions are initiated. For example, a pop-up message could inquire “Are you sure? Give a justification for your undertaking this action” whenever anyone is changing passwords, moving or removing large quantities of data, accessing highly sensitive information, and so on. Moreover, by requiring that a stylized form of justification (for example, bug identification references, support ticket, user involvement, requirements, and so on) be provided, the actor could be forced to reflect in a way that could head-off making an error. Obviously, automated tools can and should check logs after the fact to spotlight suspicious activities, such as actions involving too few individuals or where the explanation lacks detailed supporting references.

Physical security is an important element, not only to implement isolation

How Big Data and AI Complicate Insider Protections

Recent progress in data science and in machine learning means large datasets are more prevalent. Scientists and engineers using machine learning and other AI tools are taught to examine raw data and check for unspoken assumptions needed to validate models. However, an organization concerned with protecting against insider attacks cannot have staff exploring sensitive data at will. Automated tools to validate assumptions and models would be a way to obviate the need for insider access to that data. Such tools have not yet been developed, though.

but also for authorization and for deterrence through accountability. One prevalent scheme is to require people to badge-in (use an identification badge or other unique, auditable token to gain access to specific locations) and badge-out individually. First, it creates defense-in-depth if authorization within the computing system depends on the physical location from which a request is issued. Second, deterrence through accountability is strengthened if the physical security means attacks must be instigated from physical locations where the perpetrator might be observed.

Individuals as Monitors

Computing systems are not the only way to implement monitoring of an individual's behaviors. People can serve as monitors, too. We thus can see monitors as an underpinning for organizational structures where performing sensitive activities requires involvement by multiple individuals, observing each other. Common examples of such organizational structures that are in use today include:

- **Two-person integrity.** One person observes what a second person does. This structure, however, can be off-putting to staff, and it doubles the cost of the work being done.

- **Pilot and co-pilot or pair programming.** Two individuals are both active participants in the activity, reversing their roles periodically. Pair programming can be effective for some developers, but it can also impose costs and may pose workplace challenges for members of underrepresented or marginalized groups.

- **Maker/checker.** Widely employed within the financial industry, the *maker* creates a transaction and the *checker* approves it. High value transactions may require multiple checkers, although increasing the number of checkers can promote a climate where approval becomes a rubber stamp. In some implementations, a number of actions might be collected and the checker approves an aggregate rather than approving individual activities.

- **Poll-watching.** Used by U.S. election polling sites, this variation of maker/checker disaggregates critical sequential actions and has each performed by a different individual. In addition, inde-

pendent (uninvolved) individuals are engaged to ensure, either systematically or through spot-checking, that actions are undertaken properly.

- **Independent collaborators.** Tasks are accomplished by individuals who are sufficiently separated and independent that they are unable to collude, but sufficiently close to share a deep understanding of action and context. By selecting collaborators in a random fashion, we frustrate outsiders hoping to cultivate insiders for later compromise.

These and other organizational structures that embody *distributed trust* (so named because our trust in the aggregate exceeds our trust in its constituents) ultimately depend on assumptions about their participants. Typically, we assume a significant fraction will be trustworthy and that participants exhibit independence from each other—the cost to compromise N of them is N times the cost of compromising a single one, and the probability of collusion among multiple participants is small. There also will be an assumption that only mandated procedures are followed.³ Role-based access control and other privilege assignments can help ensure that no participant in an organizational structure usurps authority and engages in actions on behalf of another. Monitoring by an independent party (typically, management) also can help check any assumptions required for an organizational structure.

Professionals in the public and private sector alike expect to be treated with respect, and constant monitoring can negatively affect their morale. The physical presence of another person necessarily creates a loss of privacy in the workplace, although expectations of workplace privacy vary by industry, sector, and type of work. Loss of privacy due to monitoring sometimes can be mitigated by also providing isolated space and time for private, solo work. In addition, an organizational culture that is explicit and public about assigning high priority to security and privacy of sensitive data (whether it's customer data or mission-critical intelligence analysis) helps staff accept that workplace privacy might not always be possible and that organizational procedures do not reflect any individual's trustworthiness.

Individuals who serve as monitors

of current actions or as analyzers of logs listing past actions could become inured to false positives. Informal discussions with managers from government and the private sector alike suggest that the risk of such burnout can be avoided if this kind of activity is limited to approximately one-third of an individual's time. And secondary benefits do accrue from having individuals check their own actions. Daily system reports of unusual actions can serve as a reminder and a training tool about actions considered suspect. However, besides training individuals about what actions are sensitive, this practice does risk training bad actors about how to avoid triggering alerts. Insiders are anyway likely to be experts in whatever automation the system they work with and likely also in how to defeat it.

Monitoring and other organizational structures for minimizing insider risk come with costs. Senior leadership must be prepared to make allowance for lower productivity, for the additional resources that will be needed, and for lowered work force morale. Even with the help of automated systems (for example, that only expose safe interfaces and thus employ prevention to block attacks), burdens imposed by security enforcement may contribute to the decision of valued staff to leave. Fortunately, security fatigue is usually given as a secondary, rather than primary, reason for these departures. And some employees—depending on their roles and responsibilities—even welcome monitoring and other tools that help reduce human mistakes or that provide means for defending against allegations of malfeasance.

Ideally, additional costs incurred to reduce the risk of insider attacks would not be a competitive disadvantage in the corporate sector. But in most sectors, today, they are. And significant, revealed insider attacks have been rare enough that the market has not incentivized expenditures for suitable defenses. Legal standards and/or regulatory approaches would be one way to a level playing field where the market is not responsive.

Organizations as Monitors

Insiders include anyone who has access (even if unintended) to sensitive information and/or operations: em-

ployees, contractors, and friends. These different classes of individuals respond to different incentives which, in turn, affords different opportunities for defenses and requires different approaches to assessing threats.

In all cases, predicting when greater scrutiny of an individual may be needed can contain costs connected with reducing insider attacks. In addition, a model of an organization's business processes can be used to identify parts of the process that insider attacks are likely to target.¹

Previous work in combining psychosocial data with cybersecurity efforts provides a path forward for identifying individuals who might warrant additional monitoring.^{2,5,7} And in government and intelligence organizations, there is typically a small, intensely supervised group that integrates human resources and technical indicators to monitor the workforce. Private-sector best practices from the Securities Industry and Financial Markets Association similarly recommend the deployment of an institutionalized insider-threat team. However, some surveillance practices are not allowed in all jurisdictions where a multinational corporation might operate.^b

Whether an individual becomes a threat is often correlated with signals from system-implemented prevention and monitoring as well as from information about non-technical activities. Staff turnover (incoming or outgoing) is one event where greater attention is typically justified. Experience has shown that theft of information is more likely to occur when an individual is preparing to leave the organization. So, monitoring indicators of staff dissatisfaction can help to anticipate such exfiltrations of confidential information. Another time to be vigilant is just after hiring—it can take time for newcomers to absorb the culture of an organization. Finally, it is wise to plan for exceptional events that require changes to a trustworthy person's normal behavior. These events

Organizations about to implement new security policies and procedures can benefit by first identifying staff whose work demands high levels of security and reliability, since they are likely to embrace the transition.

can range from dealing with climate and weather emergencies to operating in a country or region that suddenly finds itself in a violent conflict. Security and risk minimization processes must not be so compliance-oriented that they cannot handle complexity and extraordinary circumstances.

Incentivizing Trustworthy Behavior

Organizations with mature security cultures learn to treat their staff well (including staff not in security-specific organizations) while remaining slightly paranoid about damage that staff might inflict. A government intelligence agency will necessarily have different approaches and incentives in place than a private company. Insider defenses at government agencies can benefit from security clearances, background checks, and criminal penalties for disclosure of protected information. And the same general principles hold for public sector, national security, and private companies. In all, staff and senior leadership must understand that polite questions or requests for clarification are not rude and reinforce behavior when difficult cases are handled well. Existing training for compliance with the Foreign Corrupt Practices Act may already teach staff how to deal with some of these problems.^c

Organizations about to implement new security policies and procedures can benefit by first identifying staff whose work demands high levels of security and reliability, since they are likely to embrace the transition. It is also wise to find staff whose work might be negatively affected or inconvenienced by new security measures, since they might need additional persuasion.

Establishing a baseline of trust and appropriate openness helps ensure that all staff are inclined to share any concerns, and that can often ease the way. In government contexts, whistleblower protections can help. Here, reported concerns should not be ignored, but overreacting could discourage borderline reports. Be care-

^b An *Insider Threat Best Practices Guide* (<https://www.sifma.org/resources/general/best-practices-for-insider-threats/>) produced by the Securities Industry and Financial Markets Association cautions that “using such traits to profile insiders carries some degree of legal risk, particularly in EU member states where automated decision-making based on such profiles is restricted.”

^c The Foreign Corrupt Practices Act is a law aimed at preventing companies and their senior leadership from paying bribes to foreign officials in order to assist a business deal and also imposes accounting transparency guidelines. The law applies to publicly traded companies.

ful that new security policies and procedures do not put staff into personally untenable positions, for instance, by ignoring local laws. For companies with staff located around the world, best practices for security increasingly have come to depend on location, citizenship status, secret laws on law enforcement access, border inspections of devices, the fragmentation of the internet, and sometimes even coercion of family.

To trust an individual is to assert you believe you can predict how that individual will behave in various contexts. Of course, people surprise even themselves when left alone and confronted with extraordinary circumstances. Mature security organizations recognize this fact and know that collaborative efforts with shared goals are likely to produce better results than imposing controls on creative individuals (who might simply be motivated to show how those controls can be defeated).

Staff inevitably must respond to competing demands—productivity and efficiency on the one hand versus diligence, care, and security on the other. When checks or safeguards are put in place, especially those seen as an impediment to efficiency or productivity, leadership should expect creative and amusing workarounds. One example is the use of a single password for all accounts, which is easy to generate but increases the damage from succumbing to a phishing attack.

Discussion

There already has been a good deal of research on insider threats. That literature was recently surveyed, quite thoroughly, by Homoliak et al.,⁶ who populate a taxonomy with the goal of systematizing knowledge and research in the area. Some of that prior work relates to our focus, by exploring behavioral frameworks and models, how organizations might put these to use, and psychological and social theory related to insider threat. But most prior work on the insider threat concerns technical aspects: defining what constitutes an insider^d or an insider

attack,⁹ formulating security policies to defend against those attacks, designing technical means for enforcing those security policies, and creating datasets for testing mechanisms designed to detect insider attacks.

The focus of this article is non-technical means because technical means are invariably subject to compromise by insider abuses. By establishing the right culture and imposing administrative procedures, thereby enlisting trustworthy insiders to the cause, a defense in depth is achieved and a more comprehensive solution results. Rather than propose new administrative procedures, this article focuses on existing administrative procedures in use for defending against insider attacks. Specifically, we reported on practices at one large global IT company and one large DoD security organization who had not previously talked with each other about those practices. Considering the different incentives of employees there and the different kinds of assets being protected, we found it striking to see such overlap in methods that had been independently devised and deployed.

Finally, although the main thesis of this article concerns policy and administrative approaches, our discussions revealed that organizational culture and personal integrity are what matter most for building an organization that minimizes insider risk. Leadership support and buy-in is required, since mitigations can be costly. And a culture of trust and collaboration is necessary. No collection of safeguards will be sufficient to overcome a culture that is not security conscious or that lacks rigorous engineering practices. Developing and sustaining such a culture is incumbent on senior leadership.

Acknowledgments. Thanks to members of the National Academies of Sciences, Engineering, and Medicine Forum on Cyber Resilience for participating in discussions that contributed to this article. Special thanks to panelists at a July 2019 meeting: H. Adkins, J. Dhanraj, M. Gregory, A. Stubbs, N. Ziring. Thanks also to Max Poletto, Susan Landau, Bob Blakley, and Brad Martin. Finally, we appreciate the helpful feedback from reviewers of our initial submission.

The Forum on Cyber Resilience is sponsored by National Science Foundation under award number CNS-14194917, the National Institute of Standards and Technology under award number 60NANB16D311, and the Office of the Director of National Intelligence under award number 10004154. Schneider is supported in part by AFOSR grant F9550-19-1-0264, and NSF grant 1642120.

References

1. Bishop, M., Conboy, H.M., Phan, H., Simidchieva, B.I., Avrunin, G.S., Clarke, L.A., Osterweil, L.J., and Peisert, S. Insider threat identification by process analysis. In *Proceedings of the Security and Privacy Workshops*. IEEE, Los Alamitos, CA, 251–264.
2. Bishop, M., Engle, S., Frincke, D.A., Gates, C., Greitzer, F.L., Peisert, S., and Whalen, S. A risk management approach to the 'Insider Threat'. *Insider Threats in Cyber Security: Advances in Information Security*. C. Probst, J. Hunker, D. Gollmann, M. Bishop, eds. Springer, Boston, MA, 115–137.
3. Clark, D.D. and Wilson, D.R. A comparison of commercial and military computer policies. In *Proceedings of the 1978 IEEE Symp. on Security and Privacy*. IEEE Computer Society Press, 184–194.
4. Ferraio, D.F. and Kuhn, D.R. Role-based access controls. In *Proceedings of 15th Nat. Computer Security Conf. National Institute of Standards and Technology, National Computer Security Center*, (Oct. 1992), 554–593.
5. Greitzer, F.L. and Frincke, D.A. Combining traditional cyber security audit data with psychosocial data: Towards predictive modeling for insider threat mitigation. *Insider Threats in Cyber Security: Advances in Information Security*. C.W. Probst, J. Hunker, D. Gollmann, and M. Bishop, eds. Jan 2010, Springer Science+Business Media LLC, New York, NY, 85–113.
6. Homoliak, I., Toffalini, F., Guarnizo, J., Elovici, Y. and Ochoa, M. Insight into insiders and IT: A survey of insider threat taxonomies, analysis, modeling, and countermeasures. *ACM Comput. Surv.* 52, 2, (Mar. 2019), Article 30; <https://dl.acm.org/doi/pdf/10.1145/3303771>
7. Kandias, M., Mylonas, A., Virvilis, N., Theoharidou, M. and Gritzalis, D. An insider threat prediction mode. In *Proceedings of the 7th Intern/ Conf. Trust, Privacy and Security in Digital Business*. S. Katsikas, J. Lopez, and M. Soriano, eds. Springer-Verlag, Berlin, Germany, 2010, 26–37.
8. Saltzer, J.H. Protection and the control of information sharing in multics. *Comm ACM* 17, 7 (July 1974), 388–402.
9. Venables, P. Insider Threat Risk—Blast Radius Perspective. (Dec. 1, 2019); <https://www.philvenables.com/post/insider-threat-risk-blast-radius-perspective>

Eric Grosse (grosse@gmail.com) is a private consultant in Los Altos, CA, USA.

Fred B. Schneider (fbs@cs.cornell.edu) is the Samuel B. Eckert Professor of Computer Science at Cornell University, Ithaca, NY, USA.

Lynette I. Millett (lmillett@nas.edu) is Senior Program Manager of the Computer Science and Telecommunications Board at the National Academies of Sciences, Engineering, and Medicine, Washington, D.C., USA.

© 2021 ACM 0001-07/21/5

Watch the authors discuss this work in the exclusive *Communications* video. <https://cacm.acm.org/videos/insider-defenses>

^d Beebe and Chang,² for example, suggest expanding the definition of an insider to include technologies within a system that have access and whose outputs are trusted by other machines and humans.