Three Deviations from Established Cybersecurity Standard Practice

I was appointed the first CISO of American Express specifically as a result of compliance with GLBA, so I owe much to the established cybersecurity standards and practices that have evolved over time. My practical experience as a cybersecurity practitioner over several decades has shown me three deviations from standard practice that I have applied consistently and that I firmly believe all cybersecurity practitioners should consider. I will share each of them and encourage you to decide for yourself how practical it is to deviate from well-established practice and standards.

Risk Acceptance

Many established standards and authoritative sources in the U.S. and globally reference cyber-risk acceptance. They enable a CISO or cybersecurity professional to point out the risks of a specific decision, configuration, or technology deployment and require the business owner to accept that risk if they decide to move forward. This effectively transfers the risk from the cybersecurity professional to the business owner requesting the technology.

This practice (risk acceptance) has negative side effects that I experienced in my second CISO role. The majority of business leaders who accept the risk of a specific deployment view it as a permanent exemption from established standards. Ironically, those same leaders who initially accepted the risk are often the first to point fingers at others when negative consequences arise. I also learned that these risk-acceptance documents—formal emails or memoranda of understanding—were kept only by the CISO and were rarely referred to again. Often the exemption lasted for many years and was ultimately forgotten.

I therefore suggest an alternative practice that I used successfully: a cyber-risk policy-exception process that references a specific control standard and identifies the business-owner of the exception. The policy exception expires after a designated period determined by the control-standard owner, not to exceed 18 months. All exceptions are reviewed 30 days before the expiration date to determine whether the business need still exists. Every control exception has a designated owner who is responsible for managing the risk of that exception. Because this differs from traditional risk acceptance, I never used the term “risk acceptance.” Instead, I referred to the owner of the exception as the responsible party for managing the risk. Auditors can review all

policy exceptions, see who the exception owners are and determine if follow up steps were taken. Every policy/control standard exception allows the policy owner to require a secondary remediation or risk management task to be the responsibility of the exception owner. This is part of the risk management process.

The rationale for this process change is that it avoids the notion that risk exemptions are permanent and never revisited. It clarifies who owns the control risk for a specific timeframe and forces a periodic review of the exception. The eGRC tool we used was modified to identify policies, control objectives, and control standards. The control-standard exception process was added to capture the dates, business rationale, and owner of each exception. It tracked the review timeframe and decisions made by control-standard owners. This process was validated by internal audit, external auditors, and regulators. I still cringe when cyber leaders use the term "risk acceptance," even though the term appears in many regulatory requirements.

Control Design

The second topic adds to cyber-risk professional practices. I believe every cybersecurity professional must recognize that new control design is essential when technology fundamentally evolves. I often hear colleagues claim that existing policies and control objectives are perfectly reasonable for dealing with disruptive technology evolution. In my view, that belief is far from reality. When technology evolves, the attack surface is extended and new controls are necessary.

Consider Enterprise Identity-and-Access-Management (IAM) systems. Over the decades, auditors and regulators have scrutinized IAM processes, procedures, and controls because they govern humans controlling and managing access to systems. Agentic AI represents a massive technological shift that dramatically expands the attack surface for threat actors. Most identities are now non-human, making the traditional human-centric IAM controls obsolete. Pretending that existing IAM controls are sufficient to support the evolution to agentic AI is absurd and irresponsible. Simply identifying agents in production is a challenge using conventional IAM controls.

Cybersecurity professionals need to master the policy and control frameworks that are well established for any enterprise size, but that mastery alone is insufficient to address technology evolution. Professionals must also apply their technical skills to control design. This means creating new control standards and capabilities when technology evolves. Today and going forward, cybersecurity control design includes building and deploying agents that enforce control standards. Designing these new capabilities helps shrink the expanded attack surface created by emerging technologies. Claiming that

existing controls, policies, and regulations are adequate in the face of massively disruptive technology evolution is irresponsible. Designing new controls is part of every cybersecurity professional's responsibility, requiring an investment in understanding emerging technologies. This leads to the third practice, which is becoming obsolete primarily for large enterprises and, ultimately, for all enterprises regardless of size.

Feeding Log Data to Humans

Well established cybersecurity practices have long been based on the principle of sharing and organizing data so that a human can digest it, add context, and take specific action. This applies to security-operations analysts, identity-access-management administrators, threat-detection and response analysts, threat-intelligence analysts, GRC analysts, third-party-risk administrators, and others.

When new cyber capabilities are developed by vendors (often early-stage), the CISO typically wants someone who monitors screens all day to evaluate whether the new capability might be beneficial. The concept of a "single pane of glass" emerged from the need to feed new information to a human expert for action. Ironically, the expert asked to consider the new capability is often set in their ways with existing tools and may be reluctant to change.

I advocate a fundamentally different model, accelerated by the proliferation of generative-AI capabilities for cybersecurity professionals. We should strive to create a **digital immune system** that ingests data from multiple sources, applies reasoning to correlate events, and measures pattern deviation with a score. The score enables the establishment of thresholds that trigger automated workflows without requiring human input. Administrators become analysts and architects, adjusting a system that operates in milliseconds and can be tuned by tweaking models. The digital immune system operates without humans in the middle of every transaction, reducing latency and cost. Humans, however, still drive better outcomes by overseeing the system, interpreting high-level insights, and adjusting risk-management strategies as needed.

If you would like to learn more about the concept of a digital immune system, see: https://www.icitech.org/post/the-digital-immune-system-how-ai-can-outpace-cyber-threats