The Emergence and Implications of Unconventional Security Controls

Jim Routh

Cybersecurity control frameworks, the foundation of security practices in any enterprise today, are becoming less significant with the evolving cyber threat landscape–driving a response towards innovation in control design and resulting in the deployment of unconventional controls. Control frameworks will remain essential, but they alone are no longer sufficient to avoid significant data loss from cyber breaches. In some respects, this represents an 180° change from how our cybersecurity professionals were trained over the past several decades.

Cybersecurity curriculums within the military services and in the public education system have grown significantly in recent years due to the increasing demand for cybersecurity professionals in private industry and government agencies. This is a generally a positive development, although the shortage of cyber skills in the market makes it difficult for the enterprise to attract and retain cyber talent. Some professionals entered cybersecurity through opportunistic means by taking advantage of the significant growth in demand for practitioners in industry. More and more are entering the field today after seeking out cybersecurity curriculums in college or by serving in various military branches with advanced cyber training. All of us learned the importance of security control frameworks as a foundation for any public or private enterprise seeking to manage risk effectively.

Security control frameworks remain core foundational components of cybersecurity programs, and I don't believe this will or should change. But I can't help acknowledging that when I first learned cyber security risk management techniques and practices, they were directly aligned with control standards from authoritative sources that represented the most maturity for enterprise adoption. The majority of the enterprise control standards for private industry in the past several decades were derived from

Jim Routh, CISM, CSSLP, CSO is the Chief Security Officer and leads the Global Security function for Aetna. He is the Chairman of the NH-ISAC Board. He serves on the Board of the National Cyber Security Alliance and is a member of the Advisory Board of the ClearSky Security Fund. Mr. Routh was formerly the Global Head of Application and Mobile Security for JP Morgan Chase. Prior to that he was the CISO for KPMG, DTCC and American Express.

Jim is the winner of the 2016 Security Alliance Award for Innovation, 2016 ISE Luminary Leadership Award, the Northeast and the 2014 North American Information Security Executive of the Year for Healthcare, the 2009 BITS Leadership Award sponsored by the financial industry in collaboration with NIST and the Department of Treasury.

authoritative sources (e.g. NIST 800-53, ISO-27001, FISMA, COBIT, and COSA). The maturity of the enterprise's security program was directly tied to the results of testing controls to determine if the enterprise practices were aligned with the control standards linked to authoritative sources, depending on the applied regulatory framework. Control standards (often referred to as policies) are documented and periodically tested by auditors or security assessors. The more stable the results from the testing of controls, the more mature the program. So as cyber professionals, we learned that changing business models, system architectures—and even the hiring and firing of people—all led to changes with direct implications for practices that evolve outside the alignment with controls and, therefore, opportunities for remediation and further testing.

The more consistent the business was with steady growth, the easier to prove cyber security maturity through alignment to the framework and consistently positive control testing results. Actual certifications (often conducted by third parties) resulting in the attestation of effective controls assures senior management and stakeholders about the resiliency of the security program. The underlying assumption was that the more change in control implementation, the less mature the program. In other words, if control standards changed continually, it was the result of an immature program that was "fixing" or remediating the practices to align with the control framework. In some cases, senior cybersecurity leaders that moved from one organization to another often increased the number of changes to control standards and practices as a direct result of the transformation of the program under their leadership. Once the new controls and practices were implemented, the program maturity took hold

(alignment of practices to control standards), and recertification or another assessment confirmed the improvement in resiliency. Ingrained in my thinking was that the number of changes to control standards was directly correlated to the maturity of the overall security program; more changes to controls meant less resiliency, while few changes meant maturity and higher program resiliency.

What I've learned recently is the opposite of what I learned decades ago: the number of changes to control standards today is actually an indicator of maturity, not immaturity. Unlike in the past, consistently changing control standards today is actually an indicator of resiliency in a program. Consistent changes to control standards or procedures indicate an active response to changes made by threat actor tactics, resulting in higher resiliency and greater maturity for a cybersecurity program. The fundamental difference is that the cyber threat landscape is changing more rapidly than any other time in our history (a trend likely to continue). In fact, the introduction of IoT in the marketplace is further accelerating the growth of the attack surface, and the growth in capturing consumer behavioral data is leading to a faster evolution of the cyber threat landscape. Essentially, when threat actors adjust their tactics (professional criminal and nation-state sponsored threat actors), it is most often due to either advances in controls by enterprises or new attack surfaces available from consumer product innovation. Cyber threat actors seek the most efficient way to achieve their objectives with the least amount of effort. If enterprises respond by consistently changing their controls, they can create friction for threat actors who adjust their tactics. Ensuring that an enterprise is a less attractive target is about as good as it gets for a CISO and is dependent on the nimbleness of adjusting controls.

Control frameworks will remain essential, but they alone are no longer sufficient to avoid significant data loss from cyber breaches.

This subtle shift, which changes the orientation of a CISO, does not mean control frameworks and testing controls are no longer valid means of measuring resiliency or program maturity. It simply means that testing controls against a framework is one data point representing a snapshot in time. It is an indicator of maturity and resiliency at a point in time. Another indicator of resiliency is how often control standards and procedures change in response to changes to the threats. The road to cyber program maturity will likely include the adoption of a set of control standards and a control framework. Aligning the framework with an authoritative source (or many) remains a part of the critical path to program maturity and remains an important component of a cybersecurity program. Security leaders need to recognize that the conventional controls defined within a framework alone will likely be inadequate to manage risk in a sustainable way. This is

not because the frameworks are no longer effective. The reality is that the threat landscape is more diverse and changes more rapidly for any framework to keep up with. Most meaningful changes to policy frameworks come about over time, as consensus among subject matter experts influence the need to update the standards. Risk frameworks with annual changes and updates are about as frequent as is practical. This pace of change, although admirable given the difficult work of codifying changes, is misaligned with the evolution of the threat landscape. As security practitioners, we have to evolve our practices driven by the changes in threat actor tactics. Keeping up with the changes to risk frameworks alone is insufficient, assuming we wish to keep our leadership roles.

Security control frameworks remain core foundational components of cybersecurity programs, and I don't believe this will or should change.

I went through a cycle over four years ago where I transformed a cybersecurity program from one based on regulatory compliance to one driven by risk and, specifically, changes to the threat landscape. I measured the number of control standard changes or adjustments made. In the early days of the transformation, control standards and procedures changed all the time. Daily changes were common. Over the three-year period, I assumed (incorrectly, it turns out) that the pace of changes introduced to control standards, procedures and practices would decrease dramatically. Today, the program is approaching its fifth year, and the average number of policy changes is one per day. We are converging the cyber and physical security programs which will result in more policy changes. When we change a control standard or, more frequently, a control procedure, it is triggered by a change in practices aligned with the new control requirement. Almost every control standard has several key performance indicators that measure the health of the process where the control is imbedded, and that is monitored frequently. One of the KPIs that carries more weight is how many changes are introduced (control standards, procedures, and the corresponding practices), and the average is one per day.

I've learned that a risk-driven security program needs to change security posture measured through the control standards and procedures at the same pace as threat actors who adjust tactics. This year, daily changes may be the right indicator of both maturity and resiliency, but next year it may be one and a half changes daily; the next year, two changes. It will never again be once a month or once a quarter or annually. I still remember the drudgery of changing the security policy document once a year and how I never thought there were significant changes made when I began my career in security. Today, significant changes happen every day in the policy, practices and measures of enterprise residual risk. We measure our enterprise risk trend daily and share it with

senior executives to help them understand what influences changes to risk. One of the most interesting aspects of this daily pace of change is that the majority of the changes in controls are in a category we call unconventional.

Conventional controls are well established within risk frameworks and clearly defined. In addition, the audit testing procedures are mature, well established, taught to others and repetitive. When external auditors test for identity and access management controls today, the methods and techniques used for sampling and testing control effectiveness are based on decades of practical experience and are well documented. Auditor skill level is measured and quantified through certifications and ongoing education (see ISACA.org) for industry, including The American Institute for Certified Public Accountants (AICPA) certifications. Auditor opinions matter, but the methodologies used are mature and established as effective.

What I've learned recently is the opposite of what I learned decades ago: the number of changes to control standards today is actually an indicator of maturity, not immaturity.

Unconventional controls are not easily identified within the most commonly used risk frameworks and represent innovation, either in the technology capability being applied or in the techniques applied by the security practitioner. Unconventional controls often result in either a new control standard or, at a minimum, new control procedures. Here is an example of an unconventional control standard and its implications.

Conventional controls for monitoring and controlling access for privileged users (those with the entitlement rights to add or delete accounts like domain or server administrators) are well established in all control frameworks, as are auditing practices related to monitoring privilege users. We do not use conventional controls for privilege user management, a more important control objective given the fact that all cyber incidents involving data exfiltration required some kind of breach or bypass of privilege user rights. This puts more of a premium on controlling for the misuse of privilege user rights or credentials being used by a threat actor to exfiltrate sensitive data. Instead, we use behavioral risk models to create patterns of use for every person or account with privilege access for a temporary period of time. The user patterns are derived from four sources of data:

1. Entitlement data
2. Web browsing data from the web proxy
3. Email usage data from the data loss prevention log data
4. Physical access data

These patterns or models are used in real time to identify anomalies or behavioral events that don't match a pattern. The pattern matching creates a risk score for the privileged user based on all of the attributes collected and analyzed in real time. We define specific risk score thresholds established for the types of privileges provisioned so when an anomalistic event or series of events breaks a behavioral pattern; the risk threshold determines one of two required actions. If the behavior score is within a specific range of tolerance, an email is automatically generated to the privilege user's leader asking them to confirm that anomalistic event is reasonable or not. If the leader response to the email is no (big red button), then the security operations center is notified to begin intrusive monitoring. If the behavioral risk score is high and above the threshold, then the specific entitlements are revoked for the privileged user automatically with no human intervention. The security operations center is notified as is the leader of the privileged user. Essentially the effectiveness of the primary control is tied to the behavioral risk model. The more data on the behavior of the user, the better performing the model is.

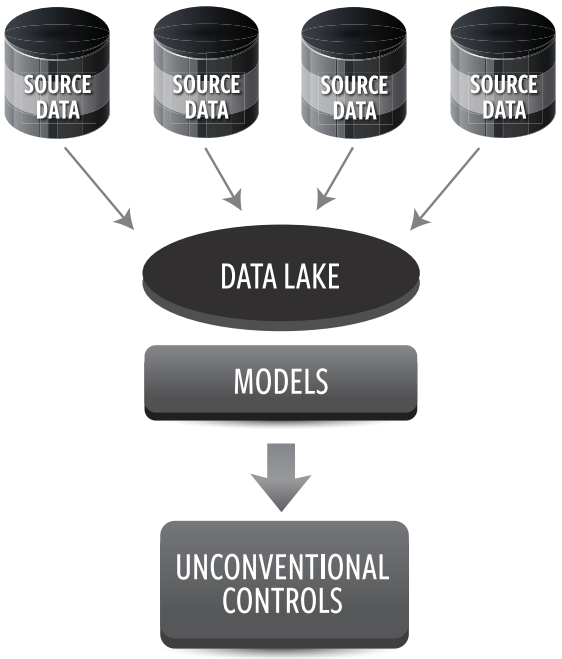

One of the first questions I typically receive when describing this model for privilege user management is about the accuracy of the models. The answer is that for over four thousand users with privilege for a specific period of time, we typically get a handful of anomalies a day and a large majority of the alerts received are benign or modifications that go back into the models.

One of the biggest implications for the implementation of this type of unconventional control is not in the implementation effort itself (which was relatively easy to do) but in the auditing of the effectiveness of this unconventional control. This privilege user management (PAM) control represents a growing trend of applying models to real-time access management, a trend that is accelerating as security practitioners build and implement better machine learning capabilities. This trend makes the job of the auditor more complicated and challenging, requiring the testing of models to determine their control effectiveness. Unfortunately, there is no body of data or techniques for testing models established over the decades of control testing in practice. The use of an unconventional control, like this one for PAM, requires new testing procedures and techniques for the auditors.

I'll provide another example of an unconventional control that has significant implications for both security architectures and auditing procedures. In this case, this unconventional control has implications for any IT or security professional designing applications now and in the future, and the trigger event is directly related to changes made by threat actors to bypass access controls for web and mobile applications. Threat actors today have access to billions of credentials (user ids and password combinations) harvested from security breaches and posted to public sites along with Social Security Numbers (SSNs) available in Dark Web forums. The result is that binary authentication, a one-time event at the front end of a user interaction with a web or mobile application, is becoming obsolete. Passwords (including one-time generated passwords) are becoming less effective as an access control since this control is based on the difficulty of the threat actor getting access to the factor or factors. The number of successful login events today in any large enterprise that is actually someone with credentials from a legitimate user is increasing due to the availability of the credential and SSN information to criminals. The net effect is an evolving obsolescence of passwords as a primary user access control, and that has significant implications for how we design application architectures going forward.

> Ensuring that an enterprise is a less attractive target is about as good as it gets for a CISO and is dependent on the nimbleness of adjusting controls.

A number of security professionals are now moving beyond passwords to deploy behavioral models using many attributes of online behavior and to create a pattern for each user across mobile and web channels. The behavioral attributes are collected during account registration and refined with more online account usage, creating a risk score to which the application can react in real-time so the actual authentication event is integrated into the user lifecycle of the application, rather than occurring one time at the beginning of the

lifecycle. The sensitivity of the application allows the application to respond to the behavioral authentication risk score and provide the level of access commensurate with the risk score at any point in the user's experience with the application. The number of architectural constructs that change in this kind of behavioral authentication is significant for the security professional, the application designer, the developer and the auditor.

This is another case of the use of an unconventional control in response to changes in threat actor tactics that has significant implications for determining the effectiveness of the control—and conventional risk frameworks do not offer much in the way of guidance for the auditor. Security professionals need to evolve control standards and procedures in response to shifts in threat actor tactics, which means enterprises must change how they build and deploy technology architecture and create more challenges for the auditors dealing with model-driven controls in real-time that are clearly key controls (heavily relied on for risk management). Sometimes the enabling technology available from an early-stage company offers game-changing capability for the enterprise. We are implementing an authentication model using behavioral risk scoring from patterns that also enables us to make adjustments to authentication controls without changing application code, saving millions of dollars every year. Changing authentication controls quickly provides the enterprise with more resiliency to respond to changes in threat actor tactics, avoids the need for developers to write or change application code every time we make an adjustment to an authentication control, and saves on operating costs. This is another positive outcome for pursuing unconventional controls.

A few years ago, I hired a chief data scientist, formerly with the National Security Agency, exclusively for the security program. His contribution to raising the skill level in data analytics for security professionals has been instrumental in our ability to deploy unconventional controls in response to changes in the threat landscape. What I had no idea of at the time was that his deep skills in data science would be so important to helping auditors figure out how to test unconventional and model-drive controls throughout the enterprise going forward.

Security professionals who understand that risk-driven programs are essential to improving resiliency for the enterprise are reaching beyond conventional control frameworks and creating unconventional controls enabled by models. These unconventional controls have the potential for mimicking another highly resilient system called the human immune system. Antibodies responding to threats automatically are essential to the human immune system and models driving unconventional controls are becoming more essential to the enterprise.