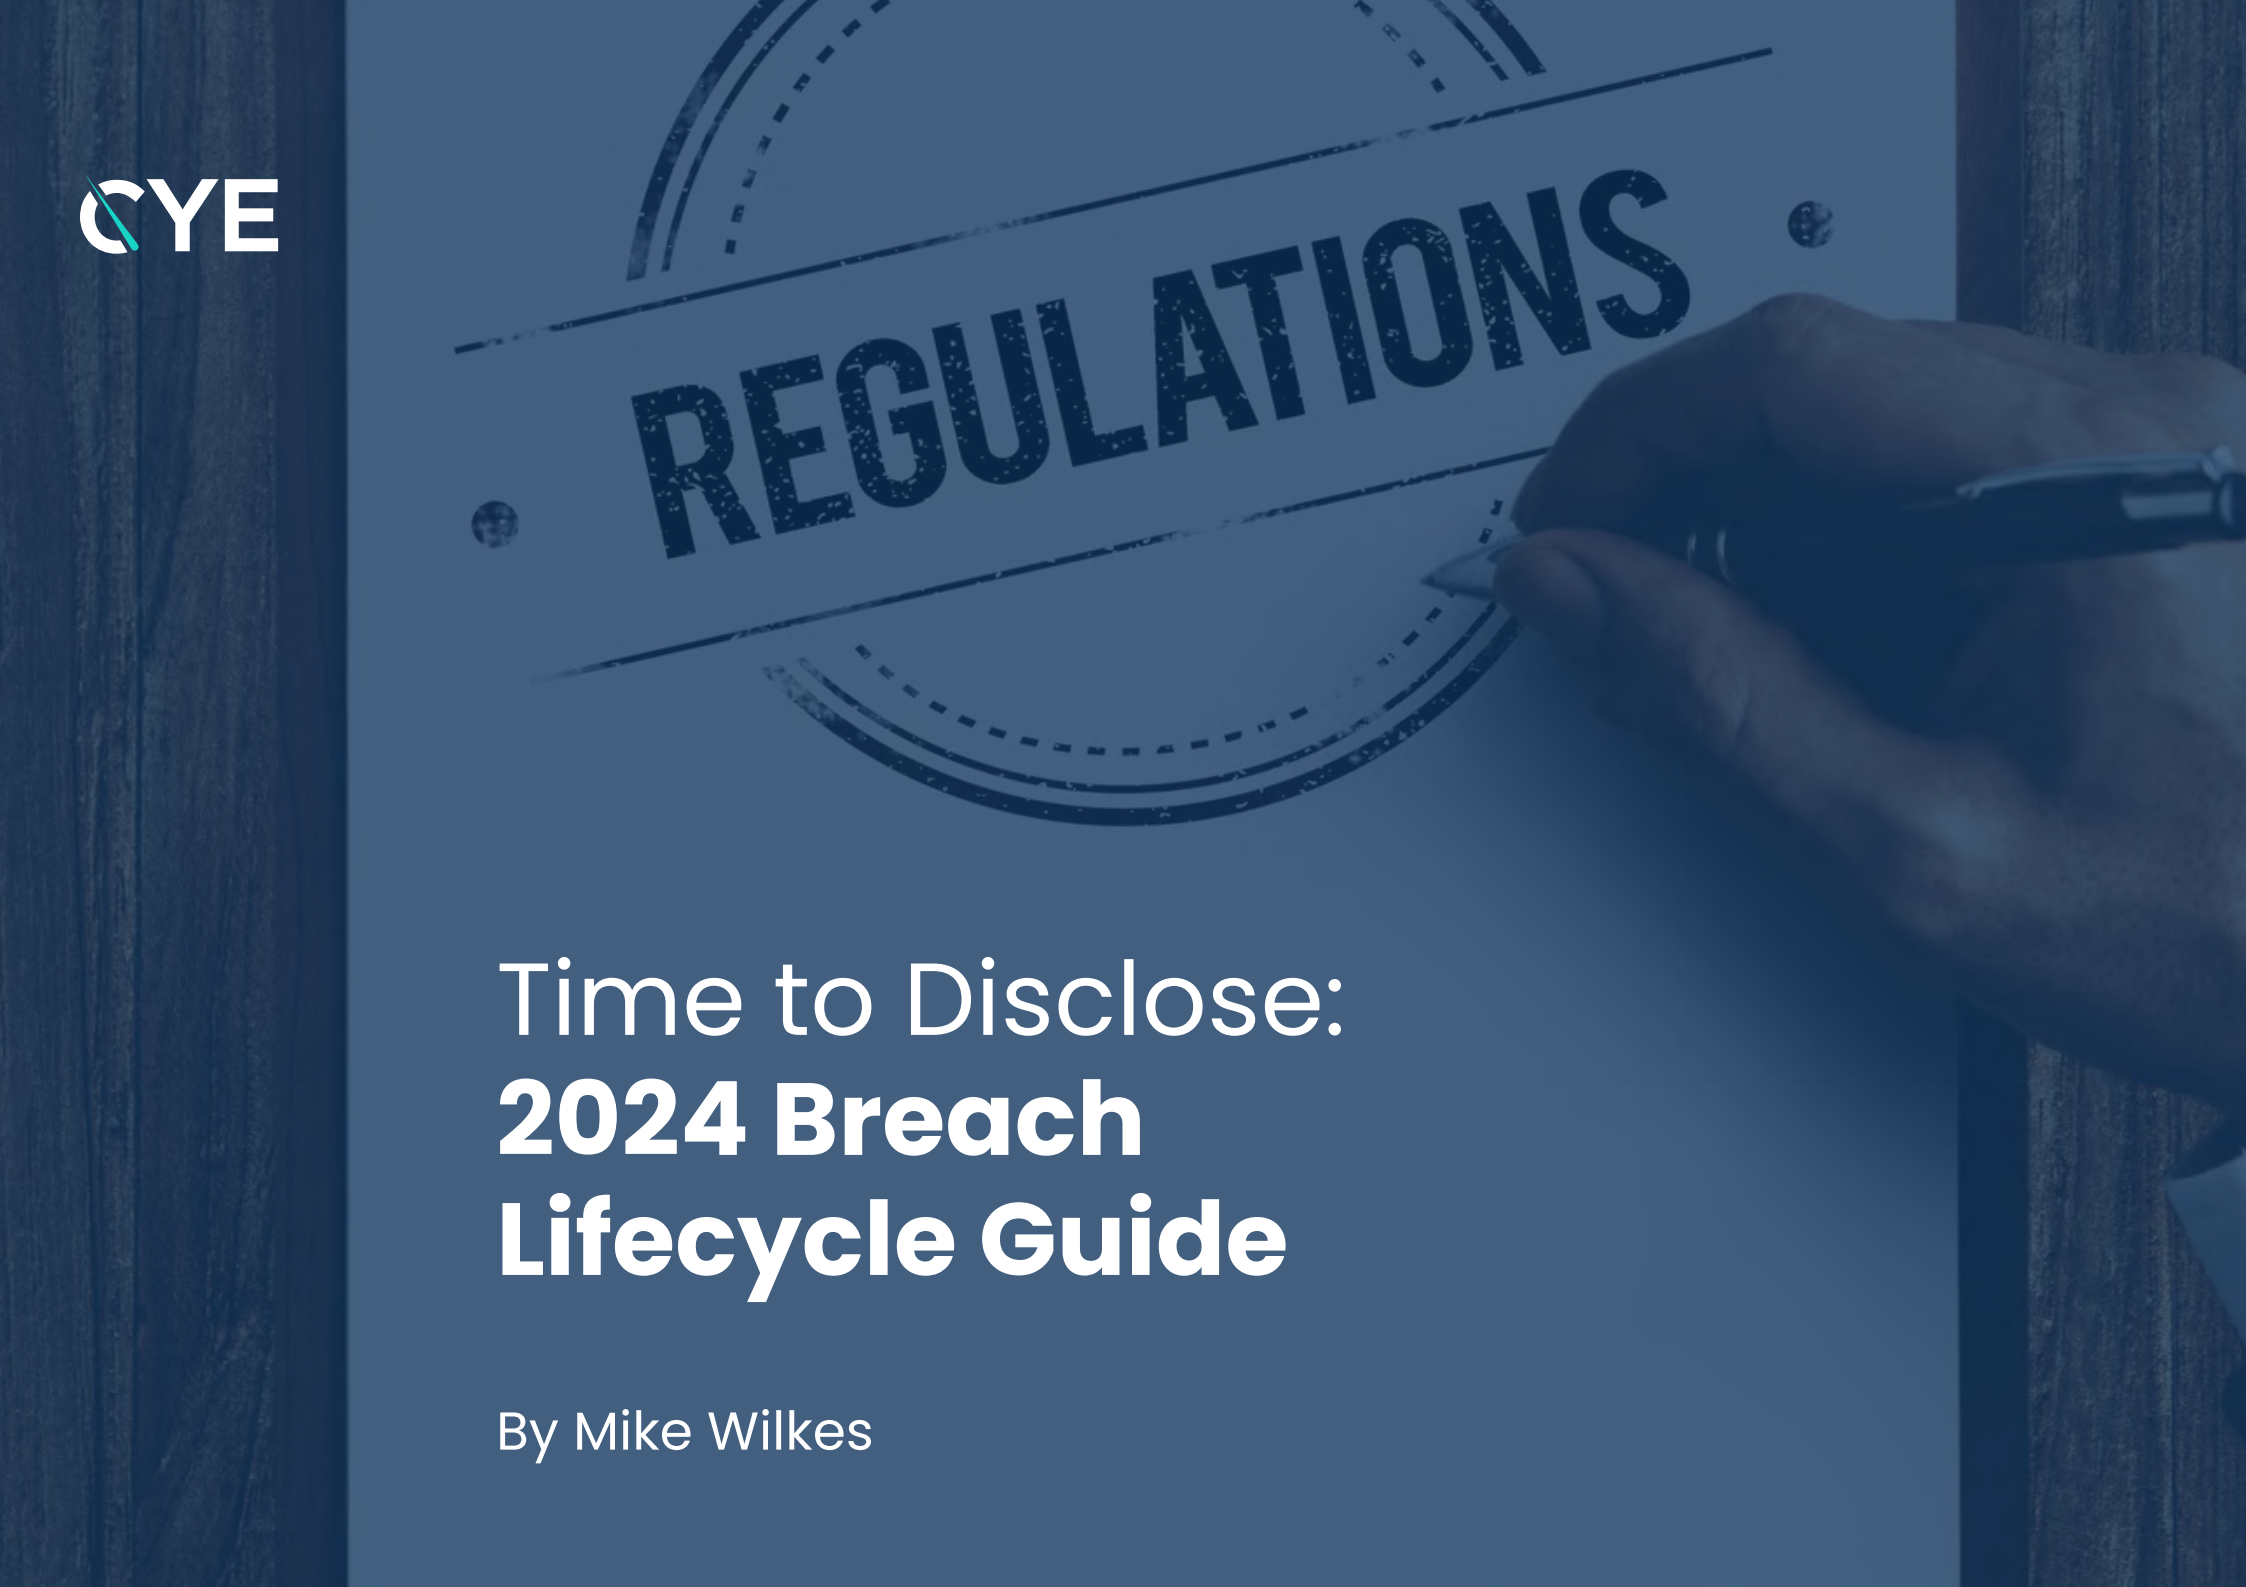

Time to Disclose: **2024 Breach Lifecycle Guide**

By Mike Wilkes

Introduction

As a Chief Information Security Officer (CISO), ensuring the availability, integrity and confidentiality of your organization's data and intellectual property is core to the mission. That includes protecting the data of your customers and partners. Breach disclosure requirements have been amped up several notches with recent changes in reporting requirements from the SEC, FTC, FCC and NYDFS. Some (or all) of these new requirements may affect you soon and some are already in effect. It would be dangerous to focus, for example, on just the new SEC reporting requirements. So let's see about how to get comfortable with them all, shall we?

If you operate a business in the US, you are already potentially beholden to 56 US states and territories breach notification requirements in terms of time to report and threshold for number of individual's information involved in an incident. Keeping up with each of these requirements is no small task. But it's also one that you do not have to perform alone. Remember that we are a community of practice and sharing our ideas and experiences makes us collectively resilient. At the very least, your chief legal officer should be working with you to make sure that your organization stays abreast of changes. Several states changed their privacy laws affecting breach disclosure requirements in January of 2023, for example.

Breach Notifications	
without unreasonable delay	AK, AR, CA, HI, IL, IA, KS, KY, ME, MA, MI, MS, MO, MT, NE, NV, NH, NJ, NY, NC, ND, OK, PA, SC, TX, UT, VA, WV, WY, DC, PR
10 day breach notification	AL
14 day breach notification	VT
24 hour breach notification	GA, ID, IN
30 day breach notification	CO, FL
45 day breach notification	MD, NM, OH, OR, RI, TN, WA, WI
48 day breach notification	MN
60 day breach notification	DE, LA, SD
90 day breach notification	CT

Table: Matrix of Breach Notifications

Definitions

Before going any further with this guide, it is important to clarify a few terms that are used when talking about security incidents and breaches. The definition of a breach is actually not defined consistently. There is also an intentionally vague approach to defining a material cybersecurity incident or material breach event. And on top of that, some companies (and their lawyers) have elected to use terms like data spillage and data leak in order to avoid invoking the phrase "data breach."

So what are we to do about all this imprecision of language and terminology? For the moment, we have to live with it and be careful about generally accepted terms and strive to be as explicit as possible about when we are talking about:

- **Event Logging** - not all logs are interesting for infosec teams, but you must have logs
- **Monitoring** - interesting logs should be monitored by algorithms and by humans
- **Alerting** - specific events in logs should trigger alerts (humans can trigger alerts too)
- **Escalation** - a process by which alerts are elevated to a higher level of awareness
- **Security Event** - indication that the infosec team should be contacted
- **Incident** - formal process begins around an event that has been escalated
- **Breach** - unauthorized access has been demonstrated or is suspected to have occurred
- **Disclosable Breach** - conditions around the breach merit additional awareness

Accepting the validity of the above differences in terminology and definition, breach disclosure is not an event, it is a lifecycle with several phases and transitions. The breach lifecycle has always been there in our work as infosec professionals, but in 2024, the number of legally disclosable breaches is expected to skyrocket from hundreds to tens of thousands.

It was CISCO's former CEO John Chambers who once said, "There are only two types of organizations: those that have been hacked and those that don't know it yet!" The mantra that everyone should be familiar with by now is that it's not a matter of if you will be breached, but only a matter of when. We must assume compromise in order to operate our infosec programs with full awareness that the bad actors are already in our environment.

Progress MOVEit

Example of a legally disclosable breach: MOVEit file transfer software breach

As of December 2023, the number of impacted organizations was 2,686 and impacted individuals was > 90 million.

10 Guiding Steps for Breach Discovery and Disclosure

Discovering and disclosing a security breach is a critical aspect of maintaining trust and transparency. Before drilling down into where the new cybersecurity reporting requirements are driving a significant change in process, let's quickly refresh ourselves on the matter at hand with a nice 10-step guide to breach discovery and disclosure:

1 Establish an Incident Response Team

- Designate key personnel responsible for managing security incidents.
- Ensure team members have defined roles and responsibilities.
- Create and deliver role-based training specific to team responsibilities.

2 Develop a Security Incident Response Plan (SIRP)

- Create a comprehensive SIRP outlining steps to be taken in the event of a security breach.
- Update the SIRP at least annually and have it signed by the CISO and their boss.
- Define communication protocols and roles within the organization.

3 Continuous Monitoring

- Implement robust monitoring systems to detect unusual activities.
- Utilize intrusion detection and prevention systems to identify potential breaches.
- Third-party service providers, vendors, and partners must also be monitored.

4 Network Segmentation

- Segment the network to limit the lateral movement of attackers and blast radius.
- Isolate critical systems to minimize the impact of a breach.
- Configure not only inbound firewall rules, but also outbound rules too.

5 Regular Security Audits and Assessments

- Conduct periodic security audits and vulnerability assessments.
- Stay proactive in identifying and addressing potential weaknesses.
- Use threat intelligence to inform your breach scenarios with threat actor emulation.

6 Anomaly Detection and User Behavior Analytics

- Deploy tools for anomaly detection and user behavior analytics to identify unusual patterns.
- Monitor privileged user activities for any suspicious behavior.
- Create a deception capability with honeypots and canary tokens for early warning insights.

7 Data Classification and Encryption

- Classify sensitive data and implement encryption to protect it.
- Ensure that data remains unreadable without proper decryption keys.
- Inventory your encryption libraries and cipher suites in preparation for post-quantum.

8 Communication Strategy

- Develop a communication plan that includes internal and external stakeholders.
- Define the criteria for determining when a breach warrants disclosure to customers, service providers, partners, law enforcement and regulatory agencies.

9 Legal Compliance

- Understand and comply with relevant data breach notification laws and regulations.
- Work closely with outside legal counsel to ensure the organization's response aligns with legal requirements.

10 Post-Incident Review and Improvement

- Conduct a thorough post-incident review to analyze the breach and the effectiveness of the detection and response.
- Use insights gained to update and improve incident response processes and security measures.

Remember, speed is crucial in breach discovery and disclosure. Some threat actors can pivot within 18 minutes of their compromise of “patient zero.” The sooner you detect and respond to a security incident, the better you can minimize its impact. Transparency and communication are key components in maintaining trust with stakeholders during and after a breach. Regularly update and practice your incident response plan to ensure your team is well-prepared for any security eventuality.

The Role of CRQ

Organizations face a growing need to quantify and manage cyber risks effectively and strategically. Cyber risk quantification solutions play a pivotal role in this process, providing a structured approach to assessing, measuring, and prioritizing potential risks. By integrating these solutions into your security strategy, you gain valuable insights into the financial impact of potential breaches, enabling more informed decision-making and resource allocation. This proactive approach not only strengthens your organization's resilience but also enhances your ability to detect, respond to, and disclose security incidents with precision and speed, ultimately safeguarding trust.

Disclosing in Time in 2024

This is what you came for… this is the bit that needs your full attention. Question: Where in the typical six-phase incident response lifecycle does breach disclosure occur exactly? That's the decision that authors and approvers of SIRPs in every business need to make.

1. **Preparation**
2. **Detection**
3. **Containment**
4. **Eradication**
5. **Recovery**
6. **Lessons learned**

There are really only two options: between steps 2 and 3 or between steps 3 and 4. Pretty much all of the new regulations are ruling out it happening anywhere later in the lifecycle. Let's call the first option "trigger happy" and the second option "cool hand" (yes, you can infer my preference here given the fact that I believe "tactical restraint" is lacking in a lot of organizations). A trigger-happy breach disclosure scenario is perfectly fine for an organization to adopt. But keep in mind that historical analysis of breaches reveals that the transition from the detection phase to the containment phase is often rushed.

So while it may be true that your disclosure committee has almost all of the details needed to publish at that point, a new and salient detail is likely to arise. Reinfection with ransomware attacks occurs not only when containment is declared prematurely, but also when there are multiple threat actors targeting a company with different intents. It is not uncommon for third-party DFIR teams to discover two or three threat actors on a system when called in to investigate. Russian malware is looking for and trying to eject Chinese malware from the infected system and vice versa.

A great example of this is the inadvertent discovery by Dutch chip maker NXP that they had been breached for over two years. They would still be unaware of their breach had it not been for researchers investigating a breach at another Dutch company, Transavia, who noticed data access patterns indicating that the same threat actors had traffic routed to the NXP headquarters. This exact story about a very patient and methodical APT performing a "low and slow" attack ("low" meaning staying below the radar of the intrusion detection tools and "slow" meaning the temptation to deploy ransomware or other payloads is resisted in favor of a longer-term objective) is also a strong endorsement for the enactment of these new disclosure rules. Our greatest strength as defenders comes from threat intelligence sharing and the collective resilience that is the result of ISAC participation and other communities.

Ownership Discussion

The more interesting discussions that need to happen, however, are not around sequencing the disclosure into your SIRP, but rather center around the topics of ownership and uncertainty. A SOC Tier 1 Analyst is not going to declare an incident; they are going to escalate it to a Tier 2 Analyst who is also not likely to be declaring an event an incident. The discussion of exactly who is going to make the decision of declaring a breach (or a disclosable breach) must be worked out and made explicit. And the criteria by which materiality for a breach is declared must be reasonably transparent, repeatable, and defensible. For companies regulated by the SEC, your next **Form 10-K** will need to include exactly that.

Essentially every SIRP needs to have a big red button labeled "disclose this incident" with some good automation in place to make it relatively easy for the person or persons charged with pushing that button to quickly acquire sufficient context to make that decision. A sophisticated asset metadata tagging system is critical for identifying and classifying certain types of assets (S3 buckets, databases, APIs, etc.) to aid in this effort. Can you tag specific infrastructure components that will always qualify as disclosable?

As we clearly saw play out with the LastPass breach and, more recently, the Storm-0558 breach, the narrative of what is known about the details can take quite some time to unfold. The first responders show up and assess the situation and make their recommendations based on available information. A press release is published with all of the standard language around "ongoing investigation" and engagement of a world-class forensics firm in response to a "sophisticated attack." But then maybe a month or more after that disclosure we learn that the situation was worse than originally thought. It also often turns out that the attackers were not so sophisticated, but rather that the company was naive and immature in their defense posture.

Take the example of the new SEC Form 8-K disclosure requirements in which the following information must be supplied within four days of a breach being declared material:

- **nature**
- **scope**
- **timing of the incident**
- **financial impact**
- **operational impact**

Both the trigger-happy disclosure approach and the cool hand disclosure approach need to operate with incomplete knowledge of the facts. When the nature of the breach is an insider risk scenario, the "who decides" question might land with the HR team (in consultation with the legal and infosec teams of course). When the nature of the breach is a third-party supply chain attack, the "who decides" question might land with legal.

The ownership discussion for breach disclosure in a nutshell is: Who gets to make what decisions, who is accountable for those decisions, and how do you make those decisions in a time of crisis?

Conclusion

There is work to be done. If you're reading this, then you are probably not confident that you and your organization are 100% ready for complying with these new cybersecurity regulations on breach disclosure and reporting. Maybe part of you is hoping that they will be overturned or nullified. To be honest, there is a non-trivial chance that this might happen for the new FTC and SEC rules. But I'm sure that another part of you is accepting the fact that these are not horribly onerous and unheard-of requirements. Many of them are solid best practices that large enterprises are already doing. And if you're a small company, additional time has been built into all of the requirements to provide for a longer period within which to comply.

I would advise everyone, however, to resist the temptation to think that these breach requirements are not in scope for your company because you are not a financial services organization (NYDFS), not publicly traded (SEC) or not a telecommunications provider (FCC). You are going to be required to understand and contribute to any and all of those previous kinds of organizations if you are a provider of services to them since that makes you one of their third-party providers. You are part of their supply chain, and the rules make it clear that third-party breach notification is in scope for those companies.

Lastly, a clarifying question can be used when thinking about when is the time to disclose: Why disclose? Which of these four cybersecurity outcomes do the new disclosure rules impact/help?

- **Reduce likelihood of a breach? No. Disclosure is a post-breach event.**
- **Detect a breach faster? No. Detection occurs before disclosure.**
- **Ensure compliance? No. Setting a one-day or four-day rule merely shifts the compliance goalposts.**
- **Defend the brand? No. Earlier disclosures won't improve the brand image with investors.**

The intent of these changes to the breach lifecycle is to improve the ability of the market to manage risk and reward. By creating new risks to businesses (lawsuits, fines, and corrective action plans) that are currently not disclosing their breaches and attacks and compromises, the regulators are desiring to introduce a healthy feedback loop into the system. Whether this is what emerges or not will, of course, remain to be seen.

Mike Wilkes has built, transformed and protected companies such as SecurityScorecard, ASCAP, Marvel, AQR Capital, CME Group, Sony, Macy's, as well as European banks and airlines. A graduate of Stanford University and author of a book for Cisco Press in 2002, he is a featured speaker at security conferences for Black Hat, SANS, GovWare and Gartner and is an adjunct professor at NYU teaching graduate-level courses to CISOs and aspiring CISOs.

Would you like to learn more about how CYE can help you quantify your cyber risk? Contact us.

About CYE

CYE's optimized cyber risk quantification platform and expert guidance transform the way organizations manage cybersecurity. Using AI, machine learning, and innovative technology, CYE visualizes attack routes, quantifies, mitigates, and communicates cyber risk, and matures organizational cybersecurity posture. In doing so, CYE provides clear and relevant insights that empower companies to make effective cybersecurity decisions. The company serves organizations in multiple industries globally. Founded in 2012, with headquarters in Israel and operations around the world, CYE is funded by EQT Private Equity and 83North. Visit us at cyesec.com.

© 2024 CYE, All rights reserved. www.cyesec.com | info@cyesec.com