Healthcare Cybersecurity

W. ANDREW H. GANTT III, EDITOR

ABA
AMERICAN BAR ASSOCIATION
Health Law Section

CHAPTER 3

How to Prepare for and Respond to Cybersecurity Attacks

Laura Ferguson and Anthony Hess

I. Cybersecurity and Compliance Requirements of Healthcare Entities

Maintaining robust HIPAA compliance practices is a significant challenge for healthcare entities (covered entities) and service providers of such entities (business associates), due to increasing reliance on electronic health records in the healthcare industry and such records becoming more frequent targets of cybersecurity attacks. Covered entities and business associates are facing increased enforcement activity by the U.S. Department of Health and Human Services (HHS), Office for Civil Rights (OCR) for privacy and security violations, including breaches. In the past few years, covered entities have paid various settlement amounts to OCR for cybersecurity attacks in which HHS determined the covered entity failed to maintain compliance with HIPAA, from as low as $100,000[1] to a record-breaking penalty amount in the Anthem breach of $16 million.[2] These settlements demonstrate the importance of developing and implementing effective HIPAA compliance policies and procedures, including procedures to comply with the Breach Notification Rule. In order to do so, covered entities and their business associates must navigate HIPAA's privacy and security requirements and be able to effectively identify, address, and mitigate HIPAA risks and breaches. In addition to HIPAA's requirements, covered entities and business associates need to identify other applicable privacy and security laws that are not preempted by HIPAA that may apply to the individually identifiable health information they maintain, so as to be prepared to respond appropriately when an attack occurs.

[1] *See* the Press Release and Resolution Agreement for the Medical Informatics Engineering, Inc. breach at https://www.hhs.gov/hipaa/for-professionals/compliance-enforcement/agreements/mie/index.html.

[2] *See* the Press Release and Resolution Agreement for the Anthem breach at https://www.hhs.gov/hipaa/for-professionals/compliance-enforcement/agreements/anthem/index.html.

In October 2018, HHS announced its largest settlement agreement highlighting cybersecurity attacks as one of HHS's top security concerns for HIPAA-covered entities and business associates. Anthem's $16 million settlement was a result of an advanced, persistent threat attack, in which cyber attackers gained access to the IT system via an undetected continuous and targeted cyber attack for the purpose of extracting data. During its investigation, Anthem determined that the hackers got into the system using spear phishing e-mails and at least one employee responded to the attack. During the time in which the attackers were in Anthem's system, the attackers stole approximately 79 million individuals' electronic protected health information (PHI). HHS stated in its press release announcing the settlement agreement that Anthem failed to (1) implement appropriate measures for detecting hackers, (2) conduct an enterprise-wide risk assessment, (3) implement sufficient information system activity review procedures, (4) identify and respond to suspected or known security incidents, and (5) implement adequate minimum access controls to prevent cyber attackers from accessing sensitive electronic PHI.

In May 2019, OCR entered into a $100,000 settlement agreement with Medical Informatics Engineering, Inc. (MIE) related to a breach in which the cyber attackers were able to access MIE's system using a compromised user's identification and password, giving them access to electronic PHI of approximately 3.5 million people. OCR determined that the company had failed to conduct a comprehensive risk assessment prior to the breach, as required by HIPAA, to assess the potential risks and vulnerabilities to the confidentiality, integrity, and availability of electronic PHI.

OCR has consistently communicated that covered entities and business associates must carefully consider the security of IT resources, ensure that PHI available through online access is secured through access controls, and that an enterprise-wide security risk assessment is routinely conducted. Based on the number of settlement agreements in this area and OCR expectations that are consistently communicated regarding cybersecurity, entities must also adopt—and follow—HIPAA security policies and procedures and risk management policies that effectively identify and mitigate risks related to electronic PHI. One of the key considerations is how to prepare and respond to cybersecurity attacks.

A. Focus on "Reasonable" Processes for Attack Prevention and Preparation Strategies

It is important that entities focus on reasonable internal and external processes to attempt to prevent a security incident and prepare for a breach. What this means is that for healthcare entities focused on preparation, mitigation, and recovery from potential cyber attacks, it is imperative that preparations be scaled to the size of the entity. Entities should also consider investing in technology solutions that will aid them in identifying and detecting cyber threats before the attackers have the opportunity to wreak their damage.[3] A large number of manual tasks can quickly overwhelm security teams. By adopting tools and technology to automate these repetitive tasks, security teams can more quickly detect, gather data on, and rapidly respond to high-risk incidents across the enterprise. As an example, Optiv research finds a 96% reduction in the average time to triage an alert after implementing automated workflows.[4] Lastly, it is also imperative with modern cyber threats to develop a comprehensive cyber extortion strategy. This topic is addressed later in this chapter.

To determine what is a reasonable process, looking to the large amount of guidance issued by HHS and OCR on cybersecurity risks and HIPAA compliance is important. Items such as the *Fact Sheet on Ransomware and HIPAA*[5] and quarterly Cyber-Security newsletters aim to help HIPAA-covered entities and business associates remain in compliance with the HIPAA Security Rule by identifying emerging or prevalent issues, and highlighting best practices to safeguard PHI. These newsletters, available on the HHS website, have addressed hot topics in cybersecurity such as cyber extortion, phishing, cloud computing and file sharing, and advanced persistent threats. Another resource is HHS's "Health Industry Cybersecurity Practices (HICP): Managing Threats and Protecting Patients" publication, which aims to provide voluntary cybersecurity practices to healthcare organizations of all types and sizes, ranging from local clinics to large hospital systems. The publication resulted from an industry-led effort to develop practical cybersecurity guidelines to cost-effectively reduce cybersecurity risks for the healthcare industry. The main document of the publication explores the five most relevant and current threats to the industry.[6]

[3] See Chapter 1 for an overview of the types of cybersecurity threats facing healthcare entities and Chapter 2 for a discussion of international cybersecurity threats.

[4] *See* Larry Wichman, *3 Key Ways to Improve Your Incident Response* (Oct. 8, 2018), https://www.optiv.com/blog/3-key-ways-improve-your-incident-response.

[5] U.S. Department of Health and Human Services, *Fact Sheet: Ransomware and HIPAA*, https://www.hhs.gov/sites/default/files/RansomwareFactSheet.pdf?language=es.

[6] U.S. Department of Health and Human Services, *HHS, in Partnership with Industry, Releases Voluntary Cybersecurity Practices for Health Industry* (Dec. 28, 2018), https://www.hhs.gov/about/news/2018/12/28/hhs-in-partnership-with-industry-releases-voluntary-cybersecurity-practices-for-the-health-industry.html.

B. Identify PHI: Data Mapping, Data Purging, and Reasonable Safeguards

Key to protecting an organization's data is knowing where it is and how it's being used. Said simply: "Know Your Data." What is the data and where is it located? What data is mission critical? Although answering these questions is generally important when evaluating cyber risks, it is especially critical in a healthcare environment. You should clarify data retention requirements and ensure that you aren't keeping unneeded data in the environment. Many cyber incidents have a greater impact than they would otherwise have had due to unnecessary data being left in the environment when it should have been deleted. To choose an appropriate data retention and backup strategy, entities should review data held in the context of operational, legal, and compliance requirements. Once the breadth of requirements is considered, it is important to implement the necessary processes and procedures to meet those requirements. It is also sensible to consider hardware/software (HW/SW) solutions to meet requirements. You must also ensure that systems and databases are updated, patched, and kept at the latest (stable) version. For critical systems, ensure that you have an update and patching test procedure in place to prevent system outages.

Many healthcare entities permit online credit card payments of patient invoices through their websites; this information is considered both electronic PHI subject to the HIPAA Security Rule and payment card information subject to the Payment Card Industry Data Security Standard.[7] Thus, it is critical to ensure that these systems are secure. One of the key risks to online credit card payment information is e-skimming, as highlighted by the FBI in a 2019 article.[8] E-skimming happens when a website has been infected by malicious code introduced by a bad actor, who typically gained access through phishing or a vendor's vulnerability that is connected to the business's server. The code enables the bad actor to capture the credit card information as the individual enters it into the website. The FBI recommends the following preventive measures to guard against e-skimming:

- Ensure that anti-virus (A/V) and/or Endpoint Detection & Response (EDR) software, as well as its definitions, is up to date across all systems—especially for mission-critical systems.

[7] For information on PCI DSS, *see* https://www.pcisecuritystandards.org.

[8] *See* Oregon FBI Tech Tuesday: Building a Digital Defense Against E-Skimming (October 22, 2019), https://www.fbi.gov/contact-us/field-offices/portland/news/press-releases/oregon-fbi-tech-tuesday-building-a-digital-defense-agaist-e-skimming.

- Change default log-in credentials on all systems and disable default "Administrator" accounts.
- Train and educate employees on e-mail hygiene and cybersecurity "Best Practices." Ensure that users understand that they are not to click on links or unexpected attachments in messages.
- Segregate and segment network systems to limit how easily cybercriminals can move ("pivot") from one system to another. This is especially useful in mitigating ransomware attacks as well.

C. Assessing Cyber Risks: External, Internal, and Physical Safeguards

One critical aspect of the HIPAA Security Rule's administrative safeguards is the requirement for a HIPAA security risk analysis and risk management review. This requirement requires "a covered entity or business associate to conduct accurate and thorough assessments of potential risks and vulnerabilities to the confidentiality, integrity, and availability of Electronic Public Health Information (EPHI) held by the entity."[9] The results of such analyses should then become the baseline (or "baseline" network/systems activity) for security processes within the entity. Failure to conduct a comprehensive risk assessment has been a common deficiency cited by OCR in recent enforcement actions.

To assist covered entities and business associates with conducting their security risk assessment(s), in March 2014, HHS released an online Security Risk Assessment (SRA) Tool. An updated version of the SRA Tool was released in October 2018 with the goal of making it easier to use and apply more broadly to risks of health information. The tool is designed for use by small to medium-sized health practices and business associates to help them identify risks and vulnerabilities to EPHI.[10] EPHI risks can be classified into External, Internal, and Physical security categories. These are defined as follows.

The *external* security risk categories are comprised of attackers or vendor issues. Attackers can be split into a variety of categories, such as state-sponsored attackers looking for private healthcare data specifically, general hackers looking to steal financial data, or cyber extortionists looking to deploy ransomware or steal sensitive or valuable data for ransom. Vendor issues can be viewed from different perspectives. Some vendors may hold large repositories of sensitive data, which represent a data breach risk and all of the financial

[9] HIPAA Security Rule § 164.308(a)(1)(ii).

[10] This tool is available at http://www.healthit.gov/providers-professionals/security-risk-assessment.

damage such a breach entails. Other vendors may represent more of a business disruption risk in the event they suffer a cyber attack. Lastly, some vendors may be directly connected from a technology perspective, which can represent both types of cyber risk. Vendors that represent a higher level of potential exposure should be carefully evaluated for their security controls to reduce the risk of a damaging cyber attack.

Another security risk category is *internal threats*, such as employees. These are more frequently unintentional than not. For example, employees are often susceptible to social engineering attacks, which can include generic phishing e-mails, spear phishing (targeted phishing), and malvertisement or malspam. Employees may share sensitive information with third parties in the course of their duties, or they may abuse their access to steal sensitive data such as PHI. It is also common to find employees who use personal e-mail, cloud storage, or USB devices to transfer PHI against company policy.

The last category, *physical*, particularly applies with theft or loss of hardware such as laptops, desktops, and mobile devices.

D. Developing an Incident Response Plan: Strategy and Planning

Developing an incident response plan (IRP) requires a diverse set of stakeholders. In fact, the modern cyber threat environment requires a wider variety of involvement than ever. Although legal (privacy in particular) and technical (forensic and incident response) expertise is still needed, the breadth of technical and business expertise has expanded as well. Stakeholders such as risk management and business resilience experts are essential to crafting a comprehensive plan.

1. *Preparation, Detection, Containment, Recovery, and Post-Incident Review*

- The preparation portion is really a risk assessment. Think about the weaknesses in the systems and how an attack on them could be detected and resolved. Utilize the data mapping process results to determine the operational risks.
- Build an incident response team with both internal and external resources and describe the roles of all team members.
- Determine all response requirements, including those outside of HIPAA (any other federal or state laws).
- Create a list of states in which patients reside in order to determine potential notification requirements. Ensure that you have the capability to keep the list up to date.

- Determine whether state laws have alternative notice provisions similar to HIPAA's in the event the entity does not have sufficient contact information to notify an affected individual. For example, can a website notice be utilized to satisfy the obligation to notify?
- Review OCR's guidance, such as:
  - Quick Response Cyber Attack Checklist and Infographic, which outlines the steps for a HIPAA-covered entity or business associate to take in response to a cyber-related security incident.[11]
  - Guidance on understanding and responding to the threat of ransomware.[12]

2. *Integrate with Business Continuity and Disaster Recovery (BCDR) Plans*

   - Many cybersecurity attacks, such as ransomware, will disable systems or access to information needed for the business. It is critical to integrate the IRP with the business continuity plan (BCP) and disaster recovery plan (DRP) to mitigate the impact of one of these attacks so that the business can continue to operate, and patient care is not affected.

3. *Cyber Insurance Policy Considerations*

Although relatively new to most incident response plans, cyber insurance interaction has become an essential component of the response to a cyber incident. The IRP should reference any cyber insurance policy that the company maintains, and should require timely notification to the insurer in accordance with policy requirements. Information about services provided via the coverage is useful to understand before a cyber incident occurs. It is also useful to determine if any other policies may include coverage for a cyber incident (e.g., cybercrime) and note this information in the IRP.

Any other limitations provided in the insurance policy should be discussed in the IRP, as failure to follow the terms of the policy could potentially result in reduced or negated coverage of an otherwise covered cybersecurity incident. See Chapter 5 for a more in-depth discussion on insurance coverage for cybersecurity incidents.

[11] U.S. Department of Health and Human Services, *Cyber Security Guidance Material*, https://www.hhs.gov/hipaa/for-professionals/security/guidance/cybersecurity/index.html.

[12] *Id.*

4. *Review and Update IRP*

In many organizations, there is a temptation to simply create a plan to meet checkbox requirements and consider the activity complete. However, the IRP should be tested and updated with regular frequency and at consistent intervals. An IRP sitting on the shelf gathering dust is not an effective IRP.

An IRP is a living document that will have to be reviewed and updated periodically. How often this occurs will depend on the organization: some entities will find quarterly is appropriate, whereas others may review biannually or annually. Senior leadership is ultimately responsible for the IRP and must ensure that it aligns with both current risks confronting the organization and any other mitigation efforts that may be needed when responding to a cybersecurity incident.

5. *Training*

Essential to the success of an IRP planning activity is employee change management, especially training. Employees need to know how the process works and understand what to do in the event of an incident. Employees need to be prepared on how to respond by receiving in-depth instruction as to what is in the policy and what their role is in the response process.

6. *Practice/Tabletop Exercises*

Training is critical, but so is practicing implementation of the IRP. Tabletop exercises assist employees with really understanding their roles in responding to an incident and gives them hands-on experience that can show a business any gaps in the IRP that may need to be addressed *before* an incident occurs. An IRP should not be considered complete until it has been tested and practiced.

II. Incident Response Best Practices

There are a number of best practices to consider while developing plans to respond to an incident.

The first best practice is to actually implement the IRP (as well as business continuity and disaster recovery plans, as applicable). If or when a breach occurs, follow the IRP. The IRP should contain step-by-step instructions on how to properly document, investigate, and take action during the incident. In the event that the attack has disabled systems or accessibility of data needed to operate, implement the continuity plan to keep the business functioning.

The next best practice is to determine the resources required for adequate response. Although an incident will require sufficient technical expertise to respond, there are also often requirements for additional legal (especially privacy breach coaches in the early stages), public relations, and even call-center personnel in the event the internal staff cannot handle the volume of external queries after an incident. These internal and external resources must be coordinated and scheduled, and information shared among them. Although most cyber incidents do not require a call center to handle the volume of incoming queries, it is helpful to have a plan for augmenting these resources in the event of an incident that exceeds the internal infrastructure available within the entity.

It is also best practice to plan for the response to a breach by a business associate. In order to do so, it is essential to review the business associate's capability and obligations in regard to a cyber incident. First, look at the business associate agreement for obligations that were negotiated between the parties and determine which party is responsible for reporting. Also, determine whether the business associate has a sufficient mitigation plan in the event of an incident.

One area of incident response best practice that is rapidly evolving is the use of cyber insurance in responding to an incident. Due to the changes in cyber insurance policies and claims services, it is essential to involve insurance in the response. In the event of a significant cyber breach, it is advisable to notify the cyber insurer as soon as possible if and when a breach occurs. The policy will spell out notification requirements, and those should be understood and followed. In the event that a privacy breach coach is assigned or engaged, it is advisable to proactively communicate with counsel. Many cyber insurers will have a panel of providers for such key services, and working with these firms before an incident will help ensure a smooth engagement if and when an incident occurs.

If a cyber incident occurs, actions should be taken as needed to minimize the impact of the incident even before coverage is affirmed. As mentioned, insurance can also give support with the provision of incident response services, such as legal counsel (for privacy or defense), digital forensics, incident response, and public relations.

Another critical best practice to consider is how to respond to regulatory and litigation requirements. In a heavily regulated industry such as healthcare, it is essential to understand and plan for those impacts on healthcare breach response.

At the federal level, HHS is required by the Breach Notification Rule to investigate each reported breach of unsecured PHI. Depending on the type of breach, this may start with a phone call to scope the entity's response. Also depending on the incident, this may lead to follow-up requests for additional information and copies of HIPAA policies, risk assessment, breach notification documentation, etc., to prove HIPAA compliance.

At the state level, state-level attorneys general may investigate as well. This is common when there is a data breach of the personal data (including PHI) of the citizens of a specific state. Potential litigation is dependent on applicable state laws, the type of incident, and whether harm must be shown. It is important to show that a reasonable IRP/process was in place and followed. See Chapter 3 for an in-depth discussion on how potential litigation defines issues.

III. Notification: Overview of Federal and State Reporting Requirements

Reporting obligations vary depending on the laws to which an entity in the healthcare industry is subject. In general, entities will be aware of and complying with HIPAA's Breach Notification Rule; however, other federal and state laws could also apply depending on the type of entity and geographic location of the services. This section gives a broad overview of the HIPAA Breach Notification Rule and state data breach notification laws.

A. HIPAA Breach of Unsecured PHI

The HIPAA Breach Notification Rule outlines the requirements for covered entities and business associates to follow when a breach of unsecured PHI occurs.[13] When investigating a potential breach of unsecured PHI, conducting and documenting a thorough assessment of the incident and confirming that the incident falls within the definition of a breach is critical. A breach under the HIPAA Breach Notification Rule is defined as the acquisition, access, use, or disclosure of unsecure PHI that is impermissible under the Privacy Rule and that compromises the security or privacy of the PHI.[14] A covered entity or business associate must conduct the following four-prong inquiry to determine if a breach has occurred.

1. Does the potential "breach" involve unsecured PHI? PHI is unsecured if it is not rendered unusable, unreadable, or indecipherable to unauthorized individuals through the use of a technology or methodology specified in guidance published by HHS.[15]

[13] 45 C.F.R. pt. 164, subpart D.

[14] *Id.* § 164.402.

[15] *Id.* § 164.402; *see* Department of Health and Human Services, *Guidance Specifying the Technologies and Methodologies that Render Protected Health Information Unusable, Unreadable, or Indecipherable to Unauthorized Individuals for the Purposes of the Breach Notification Requirements under the HITECH Act* (Apr. 17, 2009), http://www.hhs.gov/ocr/privacy/hipaa/administrative/breachnotificationrule/brguidance.html.

2. Has there been an impermissible acquisition, access, use, or disclosure? Did the alleged impermissible acquisition, access, use, or disclosure violate the HIPAA Privacy Rule?

3. Is the probability low that PHI was compromised? An impermissible acquisition, access, use, or disclosure of PHI is presumed to be a breach unless there is a low probability that the PHI has been compromised based on a risk assessment of at least the following factors: (1) the nature and extent of the PHI involved, including the types of identifiers and the likelihood of reidentification; (2) the unauthorized person who used the PHI or to whom the disclosure was made; (3) whether the PHI was actually acquired or viewed; and (4) the extent to which the risk to the PHI has been mitigated.[16]

4. Does an exception apply? There are three exceptions to the definition of *breach*. Two of these exceptions generally capture benign incidents of unintentional acquisition, access, use, or disclosure of PHI by or to a workforce member or person acting under the authority of a covered entity or business associate. To meet these exceptions, the PHI cannot be further used or disclosed in a manner not permitted by the Privacy Rule. The final exception applies if the covered entity or business associate has a good-faith belief that the unauthorized person to whom the impermissible disclosure was made would not have been able to retain the information.[17]

Upon discovering a breach of unsecured PHI, covered entities must notify affected individuals, HHS, and, if more than 500 residents of a state or jurisdiction are affected, the media. These notifications must be reported without unreasonable delay but no later than sixty days from discovery, unless otherwise directed by law enforcement.[18] If the breach is discovered by a business associate, the business associate must notify the covered entity of such breach without unreasonable delay, but no later than sixty days following discovery, unless otherwise directed by law enforcement. To the extent possible, the business associate should provide the covered entity with the identification of each individual affected by the breach, as well as any other available information required to be provided by the covered entity in its notification to affected individuals.[19] In addition to the Breach Notification Rule, business associates

[16] 45 C.F.R. § 164.402(2).

[17] *Id.* § 164.402(1).

[18] 45 C.F.R. §§ 164.404, 164.406, 164.408.

[19] *Id.* § 164.410.

may be subject to contractual requirements relating to a breach as outlined in the business associate agreement with the affected covered entity.[20]

B. Breach of Personally Identifiable Information under State Laws

All fifty states, the District of Columbia, Guam, Puerto Rico, and the Virgin Islands have enacted some form of a data breach law. These laws vary greatly in how the term *personal information* or *individually identifiable information* is defined, whether a harm threshold applies before requiring notification, which entities/individuals must be notified, and whether any exceptions apply (such as compliance with HIPAA's obligations). Most states provide a harm threshold before triggering breach notifications.[21] Currently, three states plus Puerto Rico require notifications upon unauthorized "access" to electronic information, not just upon unauthorized acquisition.[22] Many states require notification to the state's attorney general or another agency.[23] For example, California requires multiple notifications for a health insurance data breach, including to the attorney general (if more than 500 California residents are notified of the breach), the California Department of Health Services (triggered by access to health information of a licensee), and the California Insurance Department's insurance commissioner (copies of information provided to the attorney general must be sent to the insurance commissioner as well).

C. Notifications to Individuals, Governmental Agencies, Credit Reporting Agencies, and Media

The Breach Notification Rule includes specific direction regarding what information must be reported to individuals, HHS, and the media.

Notification to an individual should be written in plain language and include, to the extent possible, the following: (1) a brief description of what happened, including the date of the breach and the date of the discovery of the

[20] *Id*. §§ 164.404, 164.406, 164.408.

[21] For example, Alabama's data breach notification requirement triggers based on the unauthorized acquisition being reasonably likely to cause substantial harm to the individuals to whom the information relates. Alabama S.B. 318.

[22] Connecticut (Conn. Gen. Stat. § 36a-701b), Florida (Fla. Stat. § 501.171), New Jersey (N.J. Stat. § 56:8-163), and Puerto Rico (10 L.P.R.A. § 4051 *et seq*.).

[23] E.g., N.Y. Gen. Bus. Law § 899-AA (2005) (must notify the Attorney General, Consumer Protection Board, and the state Office of Cyber Security and Critical Infrastructure); and Texas Department of Insurance Commissioner's Bulletin #B-0022-16A (requiring a domestic insurer or HMO to notify its assigned financial analyst at the Texas Department of Insurance if the insurer or HMO experiences or discovers an unauthorized acquisition, release, or use of personal information or sensitive company information).

breach, if known; (2) a description of the types of unsecured PHI that were involved in the breach; (3) any steps individuals should take to protect themselves from potential harm resulting from the breach; (4) a brief description of what is being done to investigate the breach, to mitigate harm to individuals, and to protect against any further breaches; and (5) contact procedures for individuals to ask questions or learn additional information, which shall include a toll-free telephone number, an e-mail address, website, or postal address.

If a breach of unsecured PHI involves more than 500 residents of a state or jurisdiction, notification shall be provided to prominent media outlet(s) serving the state or jurisdiction (such as a newspaper). This notice should be written in plain language and include, to the extent possible, the same information provided to individuals.

Every breach of unsecured PHI must be reported to the Secretary of Health and Human Services. If 500 or more individuals are affected, written notification must be provided to the Secretary, in the manner specified on the HHS website.[24] If fewer than 500 individuals are affected, the covered entity shall maintain a log of such breaches and shall submit notification to HHS within sixty days of the close of each calendar year. Notification must be provided in the manner specified on the HHS website.[25]

D. Breach Notifications: State Requirements

Similar to HIPAA's notification requirements, most state laws require notification to the individuals affected, and sometimes to a governmental agency (state's attorney general or other agency). Unlike HIPAA, there are currently no state laws that require notification to the media as a result of a data breach. Typically, the state laws provide for notification to be given in the most expedient time and manner possible and without unreasonable delay, consistent with the legitimate needs of law enforcement. However, more and more state laws now provide a specific deadline for notification. Many of such states utilize forty-five days after discovery of the breach, but some utilize a shorter timeframe.[26] It is important to analyze all state laws that may apply and to chart out the timelines and requirements for notifications in advance of a reportable event.

[24] *See also* HHS, Submitting Notice of a Breach to the Secretary (Jan. 5, 2015), http://www.hhs.gov/ocr/privacy/hipaa/administrative/breachnotificationrule/brinstruction.html.

[25] *See also id.*

[26] *See* Or. Rev. Stat. § 646A.600 *et seq.* (forty-five days); Col. Rev. Stat. tit. 6, art. 1, § 6-1-716 (thirty days).

E. Special Cases: Minors, Deceased, Missing Contact Info (Substitute Notice)

In certain circumstances, the requirement to notify the individual about the breach is modified. For instance, notice can be provided to a parent or personal representative of a minor or individual who otherwise lacks legal capacity due to a physical or mental condition.[27] In addition, in the case of a deceased patient, notification is made to the next of kin or personal representative if the covered entity has the address of such individual.[28] Note that if the covered entity does not have the contact information, no substitute notice (described in this subsection) is required.[29]

There are also circumstances in which alternative notice is permitted (referred to as "substitute notice"). When a covered entity does not have current contact information or the information is insufficient to send written notification to the patient, the covered entity can provide substitute notice that is reasonably calculated to reach the individual.[30] If fewer than ten individuals are affected, alternative written notice, a phone call, or other means can be used.[31] If ten or more individuals are affected, the covered entity must conspicuously post the notice on its homepage for its website for ninety days, or provide a conspicuous notice in major print or broadcast media in the geographic areas where the affected individuals are likely to live.[32] The substitute notice must include a toll-free phone number, valid for ninety days, that affected individuals can call to learn about whether their PHI was included in the breach.[33]

F. Best Practices for Breach Notifications

A number of steps can be taken to prepare for a breach notification, which will simplify the response in the event a notification must be made. The first is to maintain comprehensive system and network logs when possible and reasonable. This will improve the organization's ability to reconstruct the incident and to better determine the affected population to be notified. Keeping patient contact information up to date will simplify the process of notifying

[27] Preamble to 2013 amendments to Breach Notification Rule, citing to 45 C.F.R. § 164.502(g).

[28] 45 C.F.R. § 164.404(d)(1)(ii).

[29] 45 C.F.R. § 164.404(d)(2).

[30] 45 C.F.R. § 164.404(d)(2).

[31] 45 C.F.R. § 164.404(d)(2)(i).

[32] 45 C.F.R. § 164.404(d)(2)(ii)(A).

[33] 45 C.F.R. § 164.404(d)(2)(ii)(B).

individuals of a breach. Preparing form notification letters that can be customized when an incident occurs will expedite notification and help ensure that the correct information is communicated to an affected population.

A critical step is to establish all of the key vendor relationships in advance and integrate their capabilities into the response process. This includes print, mail notifications, and vendors required to provide likely appropriate substitute notice. It is also important to be ready to engage call-center services in the event that the incident requires a large number of individuals to be notified. Call centers can assist with handling the influx of phone calls from patients and other affected individuals with questions regarding the breach.

Lastly, identify notification roles and responsibilities to ensure accountability and continuity.

IV. Cyber Extortion in Healthcare

A. Introduction to Cyber Extortion in Healthcare

When malware (henceforth used to include ransomware) affects a healthcare entity, the stakes immediately become high. An infected and/or compromised healthcare entity can be extorted for resources (including money and information) at its own expense, and to the expense of its patients' privacy, data, and even health. The level of risk and negative impact from cyber extortion is further increased if a healthcare entity lacks proper network/systems security, adequate training and response plans, and/or dedicated information technology (IT) and IT security staff.

Currently, the 2019 guidance from the United States Federal Bureau of Investigation (FBI) states that the FBI "does not advocate paying a ransom, in part because it [paying the ransom] does not guarantee an organization will regain access to their data."[34] Consider the healthcare organization that has been infected/encrypted by ransomware, and that has neither a dedicated cybersecurity staff nor a preexisting response plan. Such an organization may be more likely to pay the ransom, yet less likely to successfully recover from the attack, as paying the ransom does not necessarily mean that the infected systems and users will automatically recover from the attack. To this end, healthcare entities should seek to use "an ounce of prevention" rather than "a pound of cure," and should exercise caution and due diligence when considering, implementing, and reviewing their response plans.

[34] IC3.gov, Alert # I-100219-PSA.

B. Protecting Against Extortion: HIPAA and Breach Notification Requirements

The critically sensitive nature of personal health information means that PHI is highly sought-after and highly valuable to attackers: PHI can be used in conjunction with or in support of cyber extortion, and may be leveraged to initiate or to further social-engineering attacks. One such example of cyber extortion via social engineering is "spear phishing," where an attacker first conducts reconnaissance to gather information about the target, and then attempts to leverage personally identifiable information (PII) and/or PHI to manipulate the target into disclosing sensitive information. Within the context of healthcare, spear phishing could be used against a variety of targets—from the administrative and billing departments to the executive level— with a variety of effects; in such an example, the level of negative impact to the organization depends upon the intended spear phishing targets' responses to the attempted extortion.

As noted earlier, HIPAA requires that HHS be notified of a breach of unsecured PHI, which will result in HHS investigating the entity to determine compliance with HIPAA (see § III.A in this chapter). The larger the attack, the larger the penalties; the Anthem breach resulted in $16 million in penalties as a result of a spear-phishing attack. In that attack, the intent was to acquire the PHI for identity theft purposes, so the magnitude of the vulnerability in the system was evident. However, bad actors may find it more valuable to hold PHI for ransom instead of selling it on the dark web.

C. Threat and Incident Response in Healthcare: Best Practices

To develop a thorough and comprehensive response strategy, organizations must first understand what it is they are trying to secure: Which systems and networks contain PHI/PII, and "where" are they located? This is a broad question, and as the size of the healthcare organization increases, so too does the difficulty of identifying the answers. These basic questions hint at the larger security best practices of the healthcare entity. Does the organization maintain some type of "asset lists" for systems/networks containing PHI/PII (lists of known-good or known-bad devices or software, per #1 and #2 of the SANS Top 20 CSC)?

Organizations should proactively solicit input not only from security staff, but also from stakeholders, in this instance the owners of systems or services that monitor or store PHI/PII. These incident response best practices and "response strategies" should be tailored for and adopted by the individual healthcare entity, as no policy will be fully comprehensive across healthcare entities. This means that a significant amount of work and involvement is necessary for healthcare entities to identify and craft their own response strategies. In other words, it is not until these "asset lists" and general prerequisites have been satisfied that the response strategy can be formalized. The examples in this chapter should serve as a jumping-off point to assist organizations in developing their own, organization-specific response strategies (see Appendix 3-1).

Setting aside the particulars of an organization-specific response strategy, a key task within any response is understanding which variant(s) of ransomware is involved in the attack and its characteristics, both in terms of the attackers and the malware itself.

To begin, does this ransomware gang provide one key per PC, or a universal decryptor? One key per PC may be less costly per device but ultimately more costly and more time-consuming to decrypt the environment. A universal decryptor may appear more costly but can ultimately save more time, hours, and in the end money, as it will decrypt multiple systems simultaneously.

Another important factor is whether or not the ransomware variant encrypts critical forensic artifacts and clouds visibility while reconstructing the attack. In many cases, it may actually leave the artifacts intact, which will simplify the investigation somewhat.

Other characteristics of the ransomware and attacker that are useful to understand are those of the attack methodology. An attack example, as described in Appendix 3-1 using the MITRE ATT&CK Framework, can be extremely useful in illuminating a proper response strategy. Does this threat actor rely on remote tools and group policy to orchestrate the attack from a single machine, or does the actor rely on lateral movement and manual execution? Lastly, does this ransomware variant partner with secondary or follow-on malware that will require extensive threat hunting and advanced software to eradicate? For example, persistent malware (such as the Emotet banking trojan) combined with a network propagation exploit (such as MS17-010 "EternalBlue") significantly increases the difficulty of containing and remediating infected systems and networks.

D. Remediation: "Recover and Rebuild" vs. "Pay Ransom and Decrypt"

Healthcare entities with comparatively higher resources and better-established cybersecurity procedures (dedicated IT staff, system and configuration backups, system and network logging) may be able to recover from a ransomware attack with or without paying the ransom. This approach typically requires a higher level of preparation and may only be suitable for larger healthcare entities or healthcare entities with access to financial capital.

Entities that do not have a dedicated IT staff are more likely to have improperly secured systems and network security. Examples of insecure systems and network security include outdated, unpatched, and otherwise vulnerable operating systems (OSs), applications, and databases. Healthcare entities lacking a formalized cybersecurity strategy are less likely to recover from a ransomware attack without paying a ransom.

The first approach we will discuss is the "Recover & Rebuild" method. In this example, let's say a healthcare entity (or company) has robust backups dating back to just before the attack occurred, and the healthcare entity opts to restore from backups instead of paying the ransom. On the bright side, it avoids paying a potentially crippling ransom demand and dealing with volatile threat actors whose decryption software may be mediocre and slow. On the not-so-bright side, restoring systems from backups (even recent backups) destroys the logs and artifacts from the time of attack, thereby leaving investigators with no means to disprove the access or acquisition of sensitive patient data. Without the ability to disprove such activity, the company may have no choice but to notify regulators, employees, and patients. A formalized response strategy should include, but is not limited to, the following considerations.

The first is to think about how you will manage expectations. Regardless of whether you restore from backups, or pay the ransom, you will not be back up and running overnight. Depending on the type of ransomware, it could take up to a week to restore critical systems and another one to two weeks for non-critical systems. Setting unrealistic deadlines will only add undue stress on all parties involved.

Next, it is essential to have a strategy for device preservation. For physical servers and workstations, replace the affected hard drives and securely store them for the investigators. For virtual machines, preserve snapshots of each one and archive them on an encrypted storage device. It may sound elementary, but it would be wise to keep track of the closest computer hardware stores in case you need to buy storage media in bulk and cannot afford to wait for next-day shipping. In addition, have contingency strategies for machines whose drives cannot be dismounted or removed from circulation. This strategy may include collecting a forensic image (typically will take three to four hours), migrating event logs and registry data to a secure storage device, etc. One of the chapter authors once had as a client a hotel chain located on a remote Caribbean island; all six of its hotels were hit with ransomware on a holiday weekend. They needed to be up and running in forty-eight hours, but the only computer store on the island was closed. Fortunately, the IT director had six TB of empty storage on hand and we were able to walk them through preserving images of the six servers that stored PII so they could be wiped and rebuilt in a timely fashion.

Lastly, assess the availability of boots-on-the-ground resources. The more hands you have, the faster machines can be rebuilt. If the infection has affected multiple sites, have a device count and priority assessment ready for each location so onsite resources can be triaged appropriately.

The alternate approach to the above is the "Pay Ransom & Decrypt" method. Backups are not always a viable solution. If a healthcare entity is recovering from a ransomware attack and is deciding whether to proceed by paying the ransom to restore the data, incident response staff are key in providing subject matter expertise to the leadership team(s). Dialogue between leadership and incident response staff is essential for properly weighing the pros and cons of each approach. Questions to be asked in a "Pay Ransom & Decrypt" approach should include:

1. What are the chances that we pay, and the threat actor doesn't give us anything?
2. What if we pay, and the threat actor asks for more money?
3. How long will the decryption process take?
4. Will the decryption software bring everything back 100%?
5. Is there a free way to decrypt our data?

To safeguard against unauthorized disclosure of data, healthcare organizations may wish to apply a data classification standard of "Highly Restricted" to PHI. This means that the data carries a higher degree of risk, and therefore should receive an increased level of awareness and protection. In the simplest terms, healthcare entities must identify, secure, and monitor systems and networks containing PHI/PII for both compliance requirements and operational security.

Appendix 3-1

Using Example Attacks to Define a Response Strategy

Using the MITRE ATT&CK Framework, an attack generally progresses in the following steps:

1. Initial Access
2. Execution
3. Persistence
4. Privilege Escalation
5. Defense Evasion
6. Credential Access
7. Discovery
8. Lateral Movement
9. Collection
10. Command & Control
11. Exfiltration
12. Impact

Example Attack Progression / Example Attack Framework

The following is an example of how the MITRE ATT&CK framework looks during the course of a possible ransomware attack.

1. **Reconnaissance:** This initial stage is where attackers use freely available ("open source intelligence") information to begin to identify their targets. In a traditional crime scenario, this is analogous to driving through a neighborhood looking for a property to burglarize.
    - Searching shodan.io for Fortune 500 companies' IP addresses with port 3389 (RDP) open
    - "Googling" for e-mail addresses for VIP and other "high-value targets" (HVTs)
2. **Scanning:** This stage is where the attackers run active scans against targets to gain a more in-depth understanding of the environment. This would be similar to a burglar checking if a house was locked by turning the door handle. Similarly, it is the first stage where you might be detected.
    - Searching for open ports, services, vulnerabilities, etc. associated with an RDP IP address gained in step 1
    - Searching for trusting people and establishing relationships with potential targets
3. **Gaining and Maintaining Access:** This is the stage where the attacker is finally into the network. The burglar has entered the home.
    - Exploiting remote desktop protocol (RDP) vulnerabilities for IP address gained in step 1
    - Sending infected documents (Emotet, Trickbot) to high-value corporate targets

- Establishing backdoors after exploiting the "BlueKeep" RDP vulnerability
- Credential scraping (capturing credentials from admins logging onto the endpoints) for HVTs

4. **Clearing Tracks:** This is where the attacker is looking to minimize or remove any evidence of being on the network. The burglars are making sure they didn't leave any fingerprints or other identifying items.
    - During ransomware encryption, both the system and logs are generally encrypted
    - This reflects additional risk if the system is not sending to syslog and/or if also on an active directory (AD) domain
5. **Encrypting Data:** This is the stage where the attacker activates the malware and holds the data for ransom. This would be similar to a burglar taking a precious heirloom.
    - Once access is gained and contingency plans are in place, attacker starts ransomware encryption
    - Attacker will start encrypting your data and, in some cases, the entire system, leaving it not bootable
    - Alternately, attackers may leave the system with just the important files encrypted, leaving the system somewhat usable
6. **Ransom Demand and Recovery:** Fairly straightforwardly, this is where the demand for money in exchange for data recovery occurs. This would be similar to asking for a ransom to return an item with much more value to a person than the attacker, such as the family heirloom referenced earlier.
    - After ransomware encryption, ransom notes are left in the system
    - Ransom notes are usually in html format, have an e-mail contact, and contain the Bitcoin wallet ID
    - If the user pays ransom, attacker *may* send recovery keys in exchange for Bitcoin (money)
    - Recovery keys are used to decrypt ransomware-encrypted files
7. **Re-encryption/Reinfection:**
    - If the ransom is paid and the restoration completed without fixing the underlying security issue, the attackers may target the organization again within one to three months when they notice the issues during stage 1.

After defining, describing, and illustrating an example attack (using the MITRE ATT&CK Framework just described), a healthcare entity can begin to review and formulate an attack prevention strategy. Preventing malware and ransomware from affecting a healthcare organization begins by defining core organizational processes, because these processes are key elements in an effective response strategy. First hand accounts of situations wherein healthcare entities attempted to use security tools alone typically reflected an organizational breakdown at some level, due to the lack of these key, foundational organizational processes, and a lack of understanding from the security and monitoring staffs.

As described here, healthcare entities often lack a comprehensive, well-understood, and thoroughly tested attack response strategy. Organizational processes are essential foundations to preventing and defending against attack(s) as a whole and will provide a solid starting point for organizations to begin honestly reviewing, "What went right," "What went wrong," and "What can be done differently next time?"

Let's assume the organization has done everything right regarding the recon and scanning phases and the attacker still gained access through a zero-day (unknown) vulnerability. If the malware is on the system, the next step is to grab any credentials stored on that system, including the ones that logged in as a service account that is a part of the domain administrator's group. At that point your entire domain is owned by the attackers. One way to mitigate this attack path is by implementation of NIST 800.53(r4) technical controls around separation of duties, access enforcement, least privilege, and permitted actions without identification or authentication. (AC-5, AC-3, AC-6 and AC-14).

A lack of account segregation, or for medical administrators, "cross-contamination," is what allows attackers to gain the level of access that can cripple systems and take the network(s) offline. If the accounts never hit the system, they can't be scrapped. Create admin accounts as shown in the following table. For more detailed explanation, please refer to https://attack.mitre.org/techniques/T1003.

Account / Access	Domain Controllers	Servers	Workstations
Enterprise Admin	Yes	No	No
Domain Admin	Yes	No	No
Server Admin	No	Yes	No
Workstation Admin	No	No	Yes
Domain User	No	No	Yes

Security Technical Implementation Guides (STIGs) can be found at https://nvd.nist.gov/ncp/repository.