



FOUR SECRETS OF US HIPAA COMPLIANCE

A Whitepaper By



Andy Green

Founder



TECHNOVERSE BLOG
Fighting the future one day at a time.



Overview

Passed in 1996, the Health Insurance Portability and Accountability Act (HIPAA) was landmark legislation covering the US healthcare industry. It is a complicated law that among other things established HIPAA's regulatory rules for data privacy and security. HIPAA's rules apply to "covered entities," the term used to describe doctors, hospitals, and insurers, as well as their "business associates," third-parties that perform additional data processing services. HIPAA requires that covered entities and their associates have in place and maintain:

"...reasonable and appropriate administrative, physical, and technical safeguards to ensure the integrity and confidentiality of the information, and to protect against any reasonably anticipated threats."

Based on this, the US Department of Health and Human Services (HHS), which is responsible for enforcing HIPAA, established the Security Rule—the main regulatory rules for securing protecting health information (PHI). Each of the aforementioned safeguards—administrative, physical, and technical—have a series of implementation details that are either required or in HIPAA-language addressable—i.e., where you're allowed to analyze whether the requirement is appropriate.

It is easy to get lost in the details of the HIPAA Security Rule, which can be found in 45 Code of Federal Regulation 160 and 164. Fortunately, there is an easier way to grapple with all of this. While the Security Rule is made up of many individual security controls—the actual measures and procedures that need to be implemented— they can be conveniently grouped into the following three areas:

1. **Identify (assets and risks):** These are the security controls for discovering where the PHI is located and then assessing the risks associated with the data. The risk assessment is based on the assets that have been identified, including current access rights, IT configurations, existing policies, and the actual threat environment.
2. **Protect:** The Security Rule includes many detailed controls for protecting PHI. They can be summarized as the policies and procedures for assigning access roles for employees and then implementing them as appropriate file and folder permissions to reflect these roles.
3. **Monitor and Respond:** This last group of controls relate to the actual detection and analysis of any anomalous events. Covered entities need to have in place systems to capture and correlate security events, analyze these events to determine if a threat—malware, ransomware—is in progress and then have a response plan for these threats.

This Identify-Protect-Monitor paradigm is simply a way to organize controls into broader classifications, thereby helping you to see "the big picture. In the US, the National Institute for Standards and Technology (NIST) has actually done this in the form of its Cybersecurity Framework, for many federal data security standards, including HIPAA, and other important private standards as well. Refer to the end of this white paper to see how Lepide products help support HIPAA compliance through the Identify-Protect-Monitor approach.

With this overall model in mind, let's explore four ideas to get you started with HIPAA compliance. Obviously, there's no one-shot magic solution — compliance is a really process that you continually engage with—but the following should make the process far more manageable.

1. Protected Health Information (PHI): What you really need to monitor!


What exactly is PHI? According to HIPAA, it's any health information that can be connected to an individual, or where there's a "reasonable basis to believe the information" can then be used to identify the individual. Interestingly, the HIPAA regulations don't explicitly define PHI as particular identifiers!


PHI clearly includes classic personally identifiable information (PII) such as name, address, phone number or some other unique identifier such as an insurance account number. The trickier part is what "reasonable basis" means. We're now in the realm of what's known as quasi-identifiers. These are a collection of identifiers that on their own can't link back to an individual, but taken together with very high probability can indeed identify a particular person.

Quasi-Identifiers Demystified

The easiest way to understand quasi-identifiers is through an example involving the trio of full birth date, zip code, and gender. This particular example comes by the way of a Harvard researcher who was able to show that a patient could be re-identified from public health records containing these three quasi-identifiers — even though the data was cleaned of traditional PII. It was a surprising result. The researcher discovered that to locate the person behind the data involved the use of a second data set. In this particular case, this data came from publicly available voting registration records. The zip code identifies the specific town where the person lives, and therefore the town's voter registration data. The voter registration records include birth date, gender, and of course the names of the voters. Typically the number of voter that match a specific birth date is very small and is often unique! This was an important discovery made by security researchers, and helped inform HHS's policies on securing PHI.

To make it easier for healthcare covered entities to decide what constitutes PHI, HHS ultimately created a Safe Harbor list of 18 identifiers. In short: to be HIPAA compliant, covered entities need to apply HIPAA security controls to the following list:





HIPAA Identifier	Comment	HIPAA Identifier	Comment
Name	Standard PII	Social Security numbers	Standard PII
Geographic identifiers	Any geographic subdivision smaller than a state-level—including zip code or geo-location.	Internet Protocol (IP) addresses	IP addresses can be use similar to geo-identifiers.
Dates	All dates related to an individual: birth date, admission and discharge date.	Medical record numbers	Standard PII
Telephone Numbers	Standard PII	Biometric identifiers	Including finger and voice prints.
Vehicle identifiers	Standard PII	Health plan beneficiary number	Standard PII
Fax numbers		Full-face photographs	Particularly an issue with improved face identification software.
Device identifiers	Social media and other methods can be used to re-identify individuals.	Account Numbers	Standard PII
Email addresses		License Numbers	
URLs		Any other unique identifier	Example: "The vice-president of marketing at Lepide"

You'll notice that geographic identifiers below the level of a state region is considered PHI. This includes not only street addresses and zip codes, but also, for example, geo-location identifiers (from GPS devices).

Since the Safe Harbor list was introduced in 2010, there's been an information and social media explosion giving even greater opportunity for hackers and data thieves to take advantage of these quasi-identifiers.

Researchers noticed that local news sources often carry stories about residents undergoing serious medical care, and this can also be used as secondary data set, similar to voting records. So for example, using basic PHI, including admission dates and some geo information, it would be possible to re-identify a person by cross-referencing with the news or through similar information found on online forums!

Takeaway

HIPAA PHI is not just basic PII. To be HIPAA compliant, you'll need to identify and monitor the above complex list of Safe Harbor PHI. This requires efficient pattern matching and classification algorithms. A potential solution would have to scan terabytes of file system data, find files that contain HIPAA PHI from the Safe Harbor list, verify access rights of the files, and then continually monitor for abnormal events related to the PHI. In short: locating and monitoring PHI is a complex problem, requiring specialized software.

2. HIPAA's Minimum Access Principle


As with other data security and privacy standards, there's often a "hidden" philosophy underlying the controls. In the case of HIPAA's Security Rule, there is the minimum necessary principle. To help protect and secure PHI, covered entities are required to have in place practices and safeguards that "limit unnecessary or inappropriate access to and disclosure of protected health information."

In practice, this means that HIPAA requires covered entities to implement a least-privileged access model, which is a long-standing computer science concept. The idea is to limit who gets to view and update PHI within an organization's file systems to only those who need access as part of their job function. In fact this or more or less spelled out in the Security Rule's access control requirement (45 CFR §164.312).

Least Privileged Access and Risk Reduction Procedures

Why does HIPAA require this? Least privileged access controls help limit potential *risk exposure*. If a hacker manages to steal credentials of an employee in a healthcare organization, say through hash dumping or a Pass-the-Hash style attack, the chances that a particular user has access to sensitive PHI is greatly reduced! In short: fewer employees who can view PHI reduces the case of a hacker getting "lucky" and hitting the jackpot: millions of records containing social security numbers, account numbers, and more.

A least privileged access programs would typically have a few major parts. The first is accomplished during the Identify phase. In addition to scanning for PHI information, you're also collecting file permissions or ACLs, and relevant Active Directory groups. The more difficult work is to then decide whether these groups have the correct employees, and that the ACLs reflect minimum permissions: it



would involve discussions with the appropriate managers who would best know the actual job roles of the staff.

Finally, you'd have to implement any changes in access controls that result from these discussion to create a solid baseline. Even more importantly, you'll need to establish a long-term process to control access requests. In the language of HIPAA, it's the "technical policies and procedures" to allow access only to those employees who truly need it.

Ultimately, you're trying to avoid ad hoc solutions that lead back again to the problem you were trying to address in the first place: reducing the risk of hackers stealing the credentials of an overly privileged user.

Takeaway

Hackers often achieve their goals of stealing sensitive or monetizable data less through their own cleverness but rather as a result of overly broad permissions given to employees who don't necessarily need access. With these generous permission rights, attackers then have a far great probability of accessing and copying PHI. The goal of HIPAA's minimum access principle is to reduce the risk—not necessarily to eliminate it—of hackers finding the few employees who have legitimate PHI access rights.

3. HIPAA Breach Reporting Rules

In 2009, the HITECH Act updated parts of the original HIPAA law, and specifically added a breach reporting rule that was not in the original law. This new reporting rule (45 CFR § 164.400-414) asks covered entities to notify affected individuals following the "discovery of a breach of unsecured protected health information (PHI)."

Unsecured PHI

Let's unpack this definition. According to the HIPAA definition, unsecured PHI effectively means unencrypted data. There is HIPAA guidance on what constitutes a valid encryption algorithm —keep in mind that HIPAA is technology agnostic—but as a practical matter, encryption algorithms available on Windows operating system would meet HIPAA's rules.

For example, suppose a laptop with a large encrypted file containing patient health data is lost. Do you have to report it? The answer is no. And the same would be true if a hacker accessed and copied an encrypted file on a healthcare organization's servers.

Unauthorized Employee Access

HIPAA further expands on the technical aspects of a breach. A breach is considered "the unauthorized acquisition, access, use, or disclosure" of PHI. There are more that few subtleties for covered entities dealing with a potential breach.

Unauthorized access can include employees, not just outside external attackers, viewing or copying PHI. HIPAA, though, does make allowances for unintentional employee access if it was “made in good faith” and does not result in “further use or disclosure.”

HHS, by the way, has a long history of enforcing these rules with respect to employees snooping on health records, particularly regarding celebrities. In a well-known case in California in 2012, employees were viewing medical records of certain Hollywood actors, who then filed a complaint, and the hospital was fined over \$800,000: the investigation by HHS found that the hospital workers did not have “permissible reason” and the access was not unintentional.

Rules for Breach Reporting

Let’s say a hospital has been the victim of an attack by a hacker in which unsecured PHI has been copied. Does it then have to automatically report the breach? This was a confusing matter until somewhat recently.

When the initial breach rules were established by the regulators in 2009, they set up a “risk of harm” standard. It gave health organizations leeway to decide whether the disclosure or impermissible use or disclosure posed significant “financial, reputational, or other harm.” In other words, it would be possible to argue under the initial ruling that a breach of, for example, email addresses or perhaps URLs searched by patients on a hospital’s website would not cause significant harm.

However, this approach was criticized by consumer groups, and HHS backtracked in 2013 in their final ruling — the Omnibus Rule, which made important tweaks to HIPAA—to completely drop the harm standard. In short: if a hacker gets access to any PHI, a covered entity would have to report it!


The following information would have to be sent to the affected individuals:

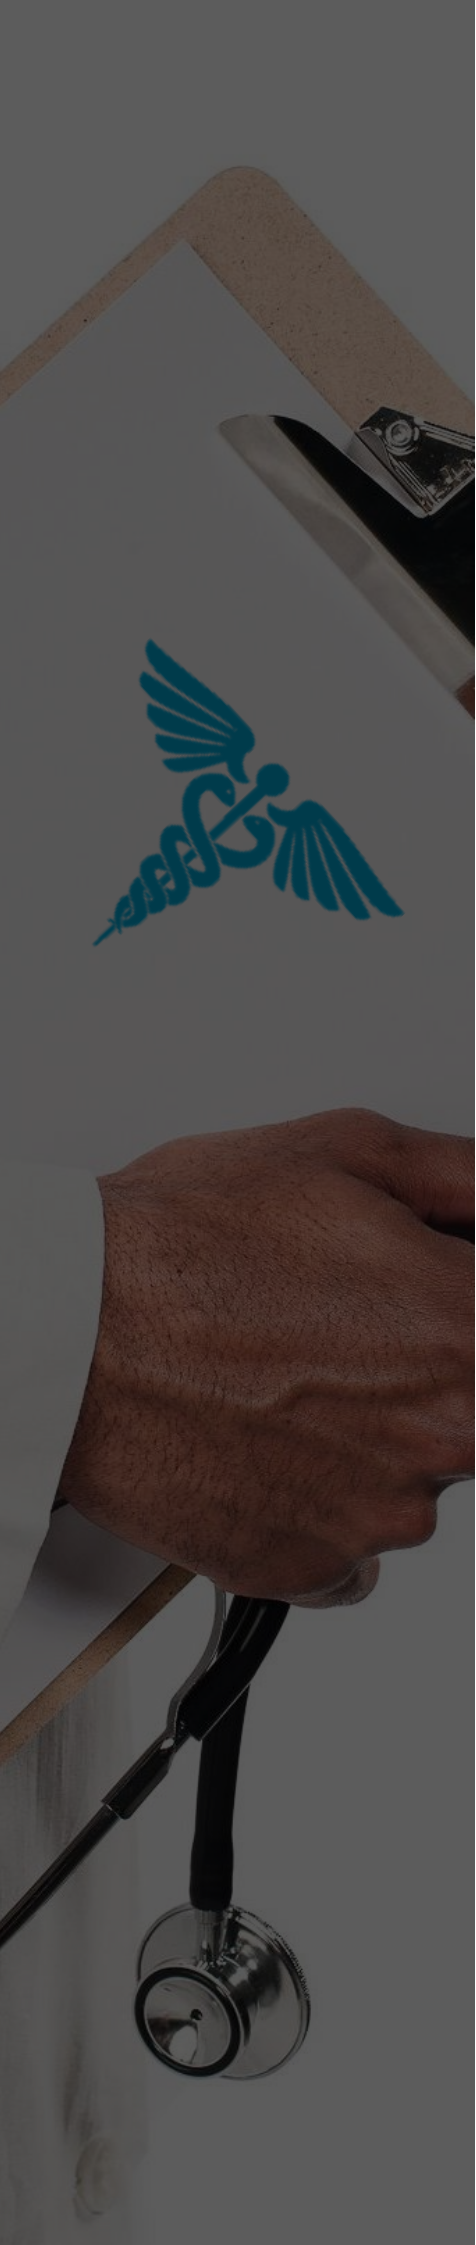
1. A brief description of the breach,
2. A description of the types of information that were involved in the breach,
3. The steps affected individuals should take to protect themselves from potential harm,
4. A brief description of what the covered entity is doing to investigate the breach, mitigate the harm, and prevent further breaches, as well as contact information for the covered entity.

They’d have to be notified —via email or postal mail—“without unreasonable delay” and no later than 60 days after discovery. In the case where more than 500 individuals are affected, the covered entity is also required to report the breach directly to HHS, where it can be viewed on their website—the so called wall of shame.

Reporting Security Incidents

HIPAA has another designation for more general types of cyber-attacks that may not necessarily be related to PHI. Under HIPAA, a security incident (45 CFR § 164.304) is “the attempted or successful unauthorized access, use, disclosure, modification, or destruction of information or interference with system operations in an information system.”





As with the breach definition, there are subtle aspects to these types of incidents. A HIPAA security incident covers access to any information, not just PHI. Suppose a hacker is snooping around some folders, which don't contain any PHI in the underlying files, and then leaves the system. HIPAA would indeed consider that a security incident, but not a breach. Similarly, if a DDoS attack affected the availability of any data, then based on the second part of the definition, it too would be considered an incident.

Unlike a breach, a security incident does not need to be reported to an outside agency. However, according to the Security Rule, the incident would still have to be documented (45 CFR § 164.308(a)(6)) internally, and if necessary mitigations put in place.

Data Breaches and Ransomware

In recent years, healthcare organizations in the US have seen a rise in ransomware attacks, in which hospital data, including PHI, has been encrypted and then held “hostage” until a payment is made. Generally, in a ransomware attack the data is not copied by the attackers to a remote server. Ransomware is clearly a security incident. But may not at first seem to be an actual breach—the data is encrypted and locked in place and not copied or “disclosed” to an outsider.

However, you'll note in the definition of a HIPAA breach that unauthorized access, modification, or destruction of PHI is considered a breach. At a minimum under a ransomware attack affecting PHI, the data has been accessed! It's also been modified — encrypted. Therefore according to HIPAA, a ransomware attack would be a breach.

Your responsibilities don't end with notifying affected individuals and potentially HHS. HIPAA also requires that you have a recovery plan (45 CFR § 164.310) to respond and that you have backups (45 CFR § 164.310) to restore the affected data. As with all incidents, you're required to investigate and then improve security procedures to reduce your risk exposure for the next event. HHS has published a fact sheet on the HIPAA implications of ransomware that is worth reviewing.

Takeaways

A PHI breach effectively becomes a test of HIPAA compliance. On its own, a breach is not a violation. However, how you respond to the breach and the security shortfalls that led to the breach can very well be a violation. An HHS investigation can be triggered by not reporting the breach within the 60-day limit, or offering incomplete details on the PHI that was compromised, the number of records affected, or the response plan. Of course, if you've done the work required by HIPAA, particularly the risk assessment, then an adequate breach detection and response program should follow as a matter of course.

4. Surviving a HIPAA Audit

There are a few ways an organization can become the subject of a HIPAA audit. Let's first cover audits conducted by HHS to measure overall industry compliance, and then those initiated by a complaint or as a result of a reported breach.

Under the HITECH Act, HHS is required to periodically audit covered entities to gauge overall compliance with HIPAA rules. Beginning in 2016, HHS ran a program to audit a sample of 160 covered entities. These were off-site "desktop" audits in which HHS was looking for documentation of policies and procedures and also, more specifically, records regarding breach and incident reporting.

The results of the [survey](#) were not encouraging with far too many falling below the minimum standards. Some of the more significant shortfalls involved the basic controls of the Security rule — minimal or negligible compliance related to risk assessment and breach reporting.


Though there have been calls for more random audits, the most common path to an HHS investigation comes as the result of a complaint filed by individuals whose rights were violated, or a data breach that's been reported directly to HHS. The investigators will initially be looking for policies and procedures, and more specifically, the results of the organization's security risk assessment, and then evidence of a response plan.

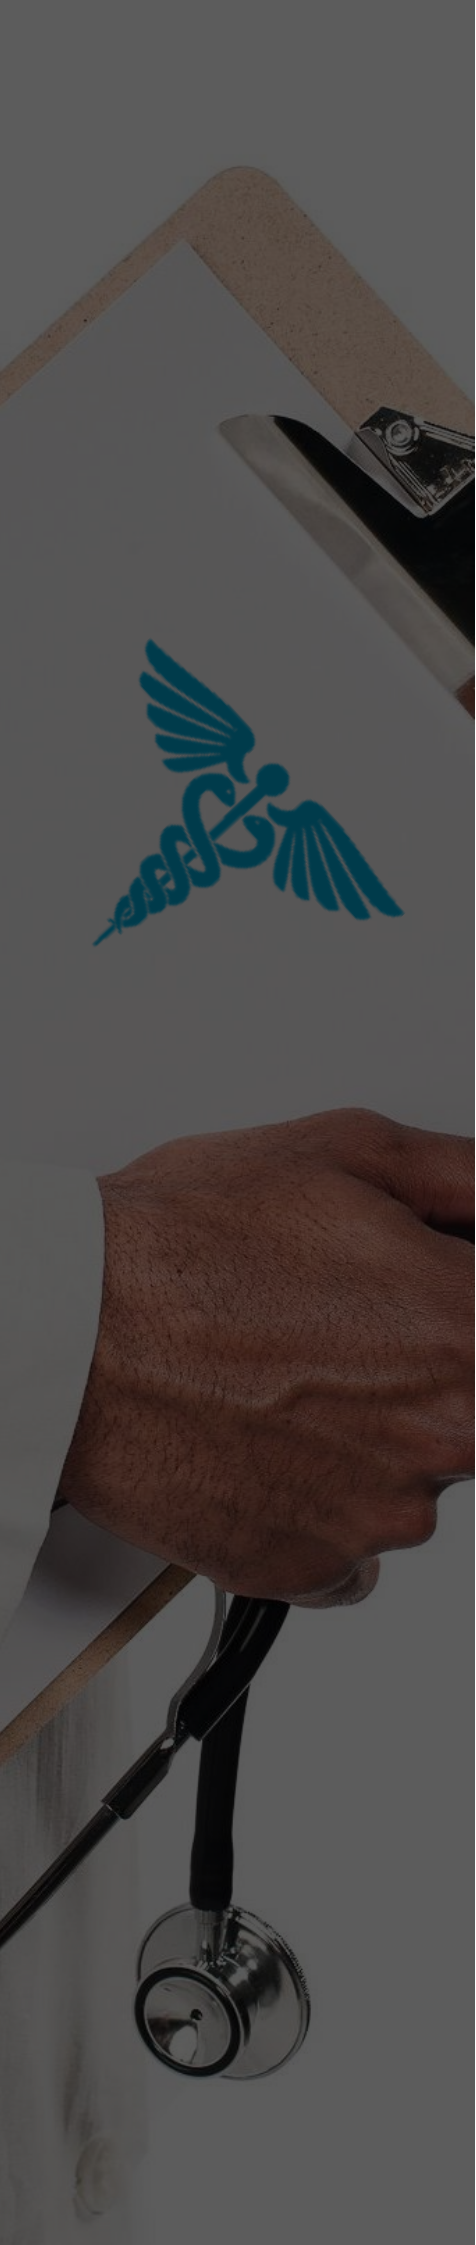
If the investigators find gaps, then further investigations, including onsite visits, can ensue, ultimately leading to potential fines.

What are some of the more common violations? HHS publishes the [results of their investigations](#) on their website. Perhaps not surprisingly, they are similar to the shortfalls discovered in the random HHS audits: failure to have performed a risk assessment, and failure to have controls in place to limit access to PHI.

To pass an investigation, HHS requires having all the required policies and procedures documents, and then making staff available for more specific questions. During a HIPAA investigation, you can expect some of the following requests:

1. Show documentation that you have a risk assessment plan—finding and categorizing PHI, analyzing and limiting access rights, and evaluating and responding to the current threat environment.
2. Show documentation that you have the security implementations based on the risk assessment.
3. Show the audit log of user activity as it relates to the incident including, file access activity, and user permissions.
4. What is your disaster recover plan?
5. What software do you use to limit PHI access?
6. Show a record of recent security incidents and your responses.
7. Describe employee security and privacy training.





Typically, after the paper work is filed, covered entities enter into a voluntary resolution agreement to remediate any issues without any fines. However, if the documentation is inadequate, further investigation and enforcement actions, including fines, can be expected.

In fact, the 2009 HITECH Act updated HIPAA with steeper financial penalties. There are now four tiers of violations with different maximum fines:

- Tier 1: Violations in which the covered entity was unaware and could not have realistically avoided (maximum: \$250,000 per year).
- Tier 2: Violations that the covered entity should have been aware of but could not have avoided even with a reasonable compliance effort (maximum: \$100,000 per year).
- Tier 3: A violation suffered as a direct result of “willful neglect” of HIPAA, in cases where an attempt has been made to correct the violation. (maximum: \$250,000 per year).
- Tier 4: A violation constituting willful neglect, where no attempt has been made to correct the violation. (maximum: \$1.5 million per year).

In addition to the fines, violators can be placed under a corrective action plan that would involve remediating and implementing security policies and procedures, and then reporting back to HHS the progress on meeting the plans.

Takeaways

The lesson from HHS audits is the same as what we discussed at the beginning of this paper: conducting a risk assessment and having the documentation is key. Obviously, HIPAA compliance is more than just documentation of policies—you also have to prove you have carried out the implementation. However, if your paperwork is below HIPAA’s minimum standards, it’s a red flag for HHS investigators.

Conclusion and Lepide Compliance Chart

By itself, HIPAA is not a data security compliance standard, such as ISO 27000. Healthcare organizations are free to continue to work with whatever security standards they are currently using. As was mentioned earlier, NIST has provided convenient mappings for key data security standards back into the HIPAA requirements.

What HIPAA has effectively done is turn standard IT security practices into a law for US healthcare covered entities. If you don’t follow the law, HHS can enforce fines and remediation plans. Or perhaps worse, the hackers discover your organization’s weaknesses first, and you’re forever enshrined in HHS’s wall of shame!

But there’s no need to do it alone. The Lepide Data Security Platform is here to help you meet the core HIPAA requirements. Please review the following table that shows how we support HIPAA through our platform and contact our sales staff with any questions.

	HIPAA Requirements	Description	Lepide Data Security Platform
Identify	Security Rule: 45 CFR 164.308(a)(1)(ii)(A)	Conduct an accurate and thorough assessment of the potential risks and vulnerabilities to the confidentiality, integrity, and availability of PHI	Spot risks to PHI with pre-defined threat models, real time alerting, anomaly spotting and risk assessment reports.
	45 CFR 164.308(a)(2), (3)	Establish cybersecurity roles and responsibilities for the entire workforce and third- party stakeholders	N/A
	45 CFR 164.308(a)(1), CFR164.312(a)(1)	Asset vulnerabilities are identified and documented, assess access controls	Identify which users have access to PHI and report on excessive permissions. Govern access from within the solution.
	45 CFR 164.308(b)	Identify threat activity (audit logs)	Detailed audit reports for all changes/ interactions with PHI. Automated anomaly spotting helps to detect threat activity and threat models enable you to execute responses in real time.
Protect	164.308(a)(4),64.312(a)(1), 164.312(a)(2)(i),	<p>Polices and procedures for access authorization.</p> <p>Permissions are managed, incorporating the principles of least privilege and separation of duties</p>	Lepide allows you to identify which users have access to sensitive data and enables admins to govern access from within the solution. Specific reports for users with excessive permissions are also provided.
Monitor	164.312(b)	Record and examine user activity accessing PHI	Detailed audit reports for user activity accessing PHI.
	164.308(a)(6)	Anomalous event data is aggregated and correlated	Behavioral analysis automatically spots anomalies based on "normal" user behavior and generates real time alerts for admins and a detailed audit trail.
	164.308(a)(6),164.308(a)(7)	Response plan is executed during or after an event	Lepide can automatically respond to detected threats through threat models, enabling admins to isolate and shut down a threat in real time.


ABOUT LEPIDE

Lepide are the fastest growing provider of data-centric audit and protection solutions to enterprises all over the world. The award-winning Lepide Data Security Platform enables you to put your data at the heart of your security strategy; mitigating the risks of data breaches and helping to meet compliance requirements.

Protecting the Data of Thousands of Organizations Worldwide



HOGUE · FENTON



FUJITSU

NHS

Deloitte.

