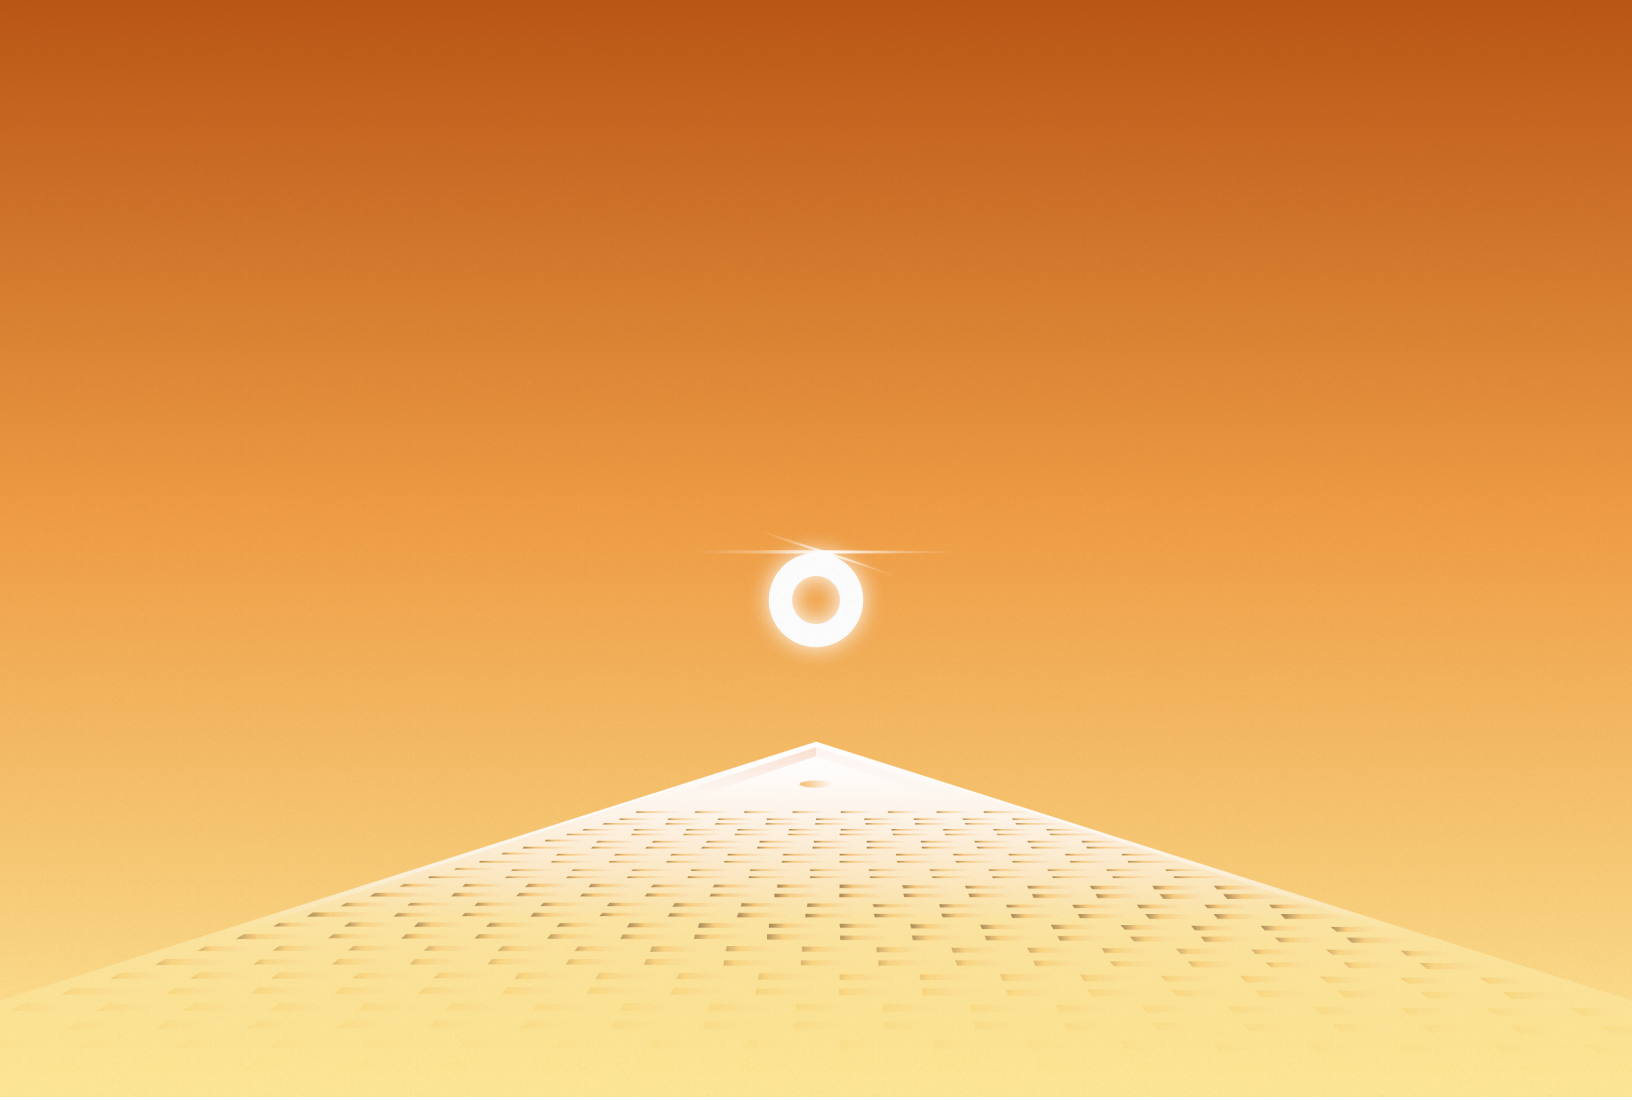

PROTENUS 2019 BREACH BAROMETER

15M+ Patient Records Breached in 2018 as Hacking Incidents Continue to Climb

Protenus, Inc. in collaboration with DataBreaches.net

Introduction

In 2018, the healthcare industry continued to be plagued by data breaches involving sensitive patient information, with public reports of hacking and phishing incidents reminding us how vulnerable patient data remains. Unfortunately, patient information can still be easily accessed and obtained by insiders and external actors alike. Even as the healthcare industry becomes increasingly aware of the importance of protecting patient data, the trend of at least one health data breach per day remains.

This retrospective examines 2018 health data breaches with an eye towards lessons learned and a way forward for protecting patient privacy.

Overview of 2018 findings

Our analysis is based on 503 health data breaches reported to HHS, the media, or some other source during 2018. We have details for 417 of those incidents, which affected 15,085,302 patient records. As shown in figure 1, comparing these numbers with those of last year, we see that there was a slight increase in the number of breaches reported (477 in 2017 compared to 503 in 2018), there was also an alarming increase in the number of affected patient records. In 2018, the total number of affected patient records almost tripled when compared to 2017 data (5,579,438 affected patient records in 2017) (figure 2). Also in 2018, there has been a substantial increase in the number of breached patients records each quarter throughout the year (figure 3).

The single largest breach reported in 2018 (figure 4) was the result of a hacking of a business associate. It involved a North Carolina-based health system vendor that had its patient information accessed by an unauthorized party. Hackers gained access to patient information over the course of a week, affecting 2.65M patient records. Compromised information included dates of birth, social security numbers, insurance policy information, dates of service, etc. The health system began notifying affected patients two months after detecting the incident. The delayed notification was the result of ongoing investigations by forensic investigators and the FBI.

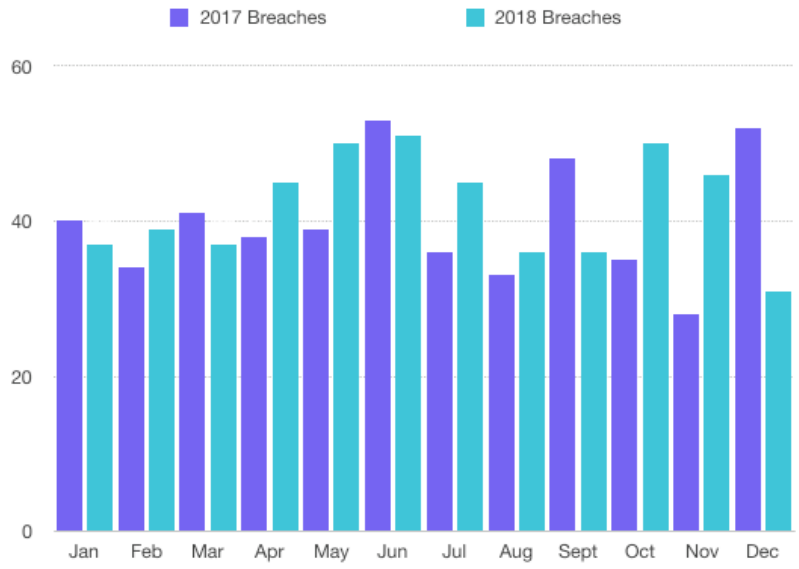

Figure 1. Total disclosed incidents, 2017 vs. 2018 health data breaches

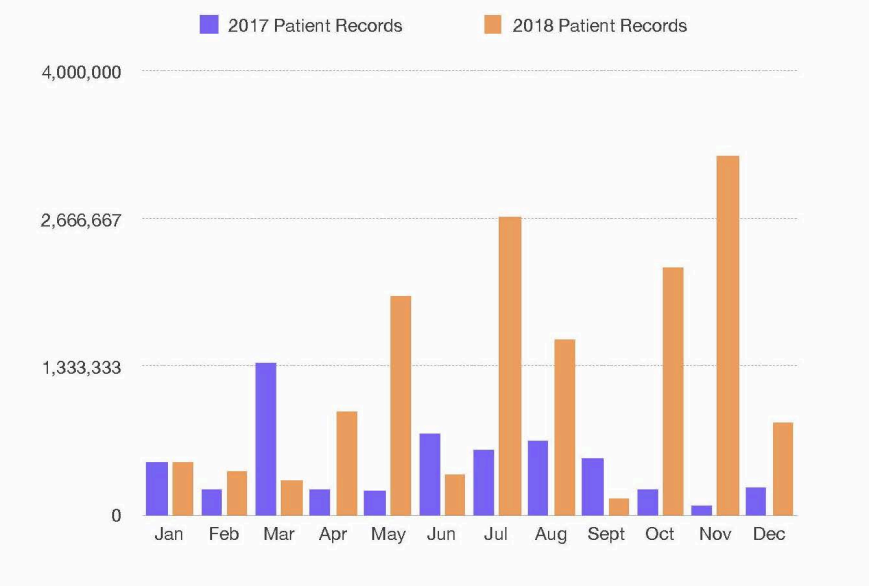

Figure 2. Total breached patient records, 2017 vs. 2018 health data breaches

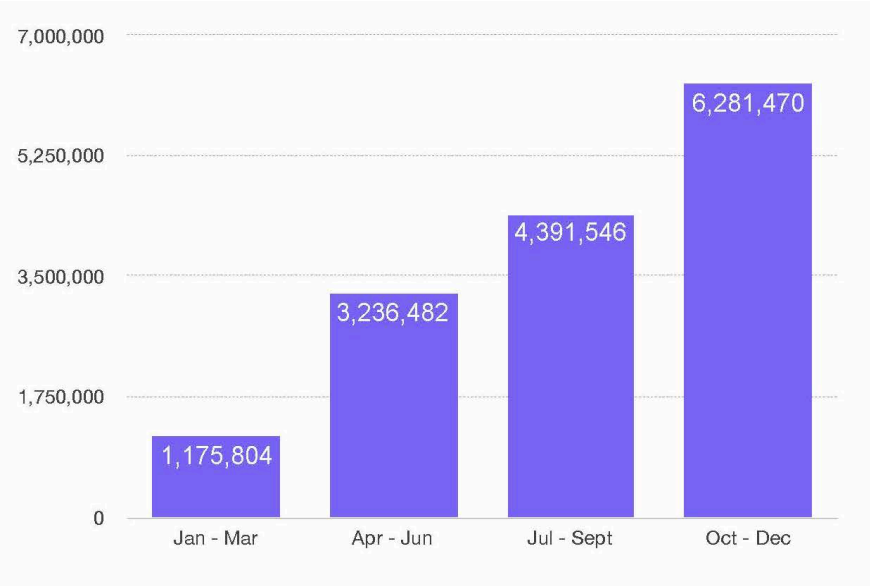

Figure 3. Affected patient records by quarter, 2018 health data breaches

2018 Largest Health Data Breaches	Organization Type	Type of Breach	Number of Affected Patient Records
January	Provider	Hacking	279,865
February	Provider	Hacking	135,000
March	Provider	I-E	63,551
April	Agency	Theft	582,174
May	Provider	Hacking	566,236
June	Business Associate	Hacking	276,057
July	Provider	Hacking	1,400,000
August	Business Associate	Hacking	502,416
September	Health Plan	I-W; BA	26,942
October	Health Plan	I-E	1,248,263
November	Business Associate	Hacking	2,652,537
December	Misc	Hacking	500,000

Figure 4. Largest incidents, 2018 health data breaches

As figures 1 and 2 demonstrate, there was no linear trend in the number of breaches or number of affected patient records in 2018. June had the greatest

number of disclosed breaches (51 incidents), followed closely by May and October, both with 50 incidents. November had, by far, the greatest number of affected patient records since this also included the largest breach incident of the entire year.

Insider-wrongdoing incident went undiscovered for 15 years

Healthcare has continued to suffer from insider incidents in 2018, with one insider snooping on patient records throughout their employment over the course of 15 years. As a result of the organization's investigation, this employee was terminated.

For the purpose of our analyses, we characterized insider incidents as either insider-error or insider-wrongdoing. The former included accidents and anything without malicious intent that could be considered "human error." Insider-wrongdoing included employee theft of information, snooping in patient files, and other cases where employees appeared to have knowingly violated the law.

Insiders were responsible for 28.09% of the total number of breaches this year, which is a slight decrease from the proportion in 2017 (37% of total incidents). There was information for 106 of those incidents, affecting 2,793,607 patient records (19% of total affected patient records).

Overall, while the number of insider-related incidents decreased when compared to 2017 data (176 insider-related incidents in 2017, down to 139 in 2018), there was a substantial increase in the number of affected patient records (figure 6). Protenus' proprietary data discovered that in 2018, on average, 3.86 healthcare employees breach patient privacy per every 1,000 employees.

There were 94 incidents that involved insider-error in 2018 and we have data for 76 of them. In contrast, 45 incidents involved insider-wrongdoing and we have information for 24 of these incidents. It is important to note that there are two incidents in which there was not enough information to classify them as either insider-wrongdoing or insider-error. Insider-error affected

2,056,138 patient records and insider-wrongdoing affected 386,469 records. Figure 6 highlights that significantly more patient records were breached by insider-error than by insiders with malicious intent.

While there were substantially fewer patient records breached by insider-wrongdoing, they are often more dangerous since employees with legitimate access to patient information can abuse their access with malicious intent, often undetected. In one case this past year, a medical assistant stole patient data by printing patient profiles and giving that sensitive information to others who used them to commit federal crimes. The medical assistant fraudulently collected more than $33,000 in unemployment benefits. This is just one example of the harm insider threats pose when they abuse their access to sensitive data while working for healthcare organizations. This entity may now face substantial post-breach costs that have been estimated to be close to $10M per breach. It is also important to remember that while the industry experiences a multitude of patient records affected from a single hacking incident, they are often quickly discovered due to the immediate disruption to hospital operations. Insider threats can remain undetected for long periods of time due to their legitimate access, as described in an example above.

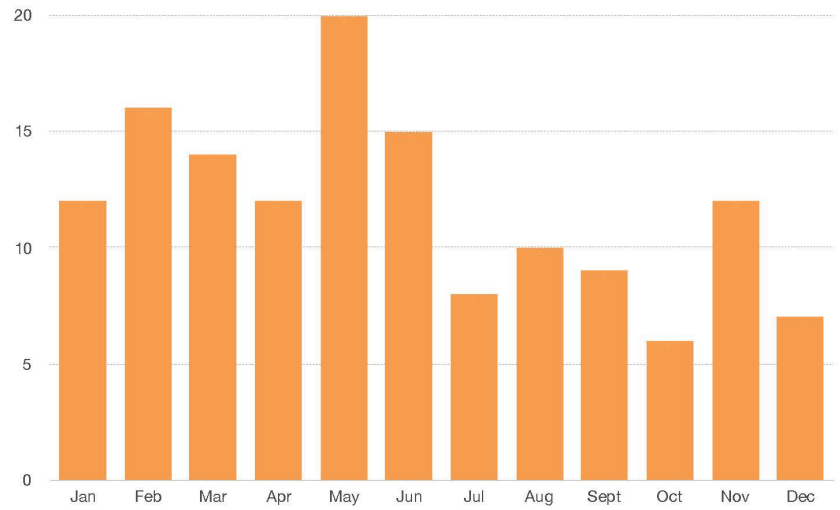

Figure 5. Insider-related incidents, 2018 health data breaches

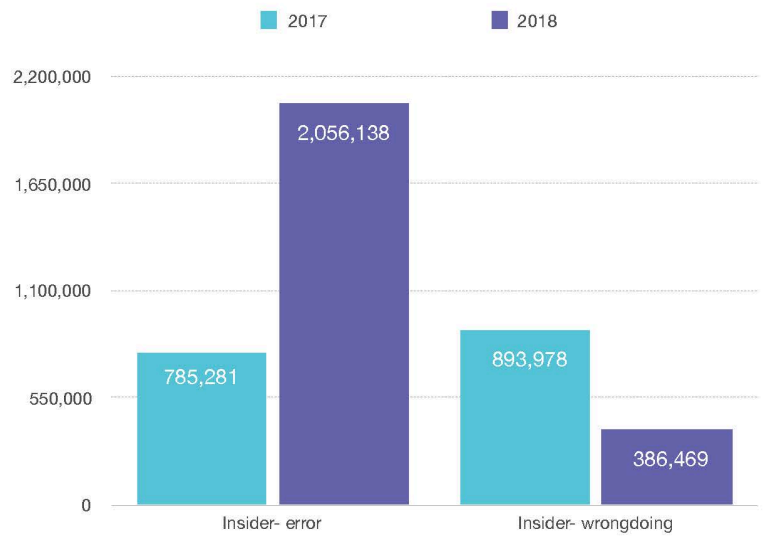

Figure 6. Patient records breached by insiders, 2017 vs. 2018 health data breaches

With the right tools, one compliance investigator can monitor an average of 7,500 employees

Due to the daily volume of accesses to health data, privacy, and security teams charged with reviewing and investigating breaches to patient privacy that occur within healthcare organizations are only able to investigate a small fraction of potential violations. Data available for 2018 shows just how difficult it is for investigators to keep up with the volume of insider threats.

The Protenus platform scans hundreds of millions of accesses to patient records every quarter to detect anomalous activity, using AI-powered analytics to review every access to patient data. The data shows that when leveraging AI-powered analytics, each hospital investigator's full-time equivalent is able to monitor the electronic access of an average of 7,500 active EHR users per month in 2018 (figure 7). This figure underscores the fact that manual audit processes, like ad-hoc or random audits, are insufficient to monitor such a large population, each of whom accesses numerous medical records per day.

Average Per Month, Per Quarter in 2018	
Number of cases per investigator	26
Active EHR users per investigator	7,500

Figure 7. Averages for privacy investigators, 2018 Protenus data

Family member snooping is the most common insider-related breach

According to Protenus data, healthcare insiders were most likely to snoop on their family members (67.38% of violations) when breaching privacy. Snooping on fellow co-workers (15.81% of violations) was the second most common insider-wrongdoing violation, followed by "other insider-wrongdoing" (6.95%) and VIP-related (6.66%) incidents. Snooping on a neighbor was the least common type of incident in 2018 (3.20%) (figure 8). It's important to note, the "other insider-wrongdoing" category included incidents like phishing attacks, insider credential sharing, downloading records for sale, identity theft, or other types of nefarious behaviors.

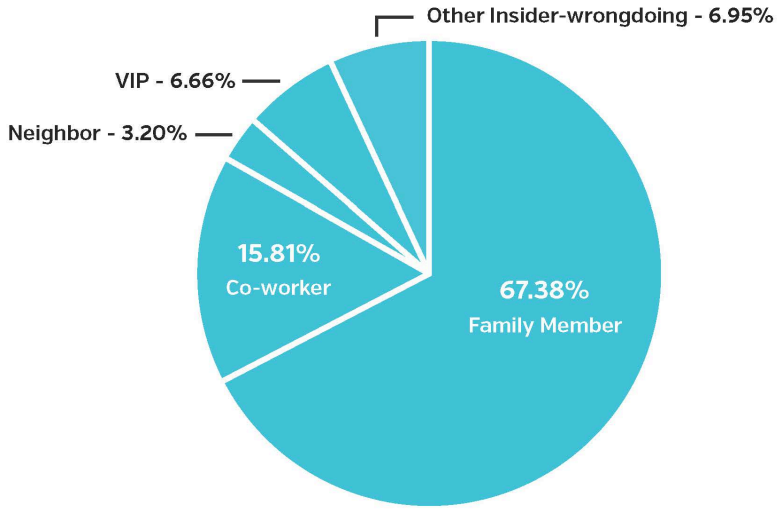

Figure 8. Insider incidents by category of violation, 2018 health data breaches

Insiders are more likely to breach privacy after first violation

If an individual healthcare employee breaches patient privacy once, there is a greater chance that they will do it again. In 2018, 51% of privacy violations were repeat offenders (figure 9). This evidence indicates health systems accumulate risk that compounds over time if proper reporting, education, and discipline actions do not occur.

Resources provided to healthcare organizations are pivotal in reducing the number of breach incidents that occur. Educating staff on EHR policy and procedures has been shown to reduce the frequency of repeat offenders within the organization.

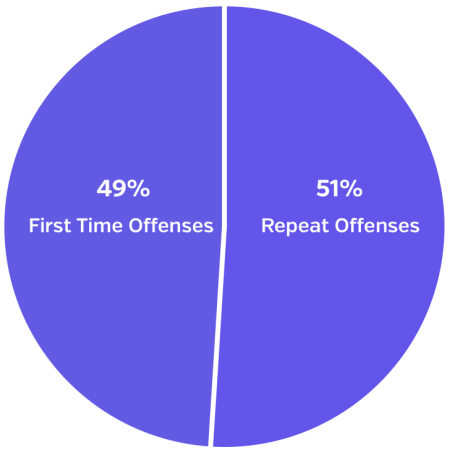

Figure 9. First time vs. repeat offenses to patient privacy, 2018 health data breaches

Hacking incidents have continued to climb since 2016

The healthcare industry experienced yet another increase in hacking incidents, a trend first noted in the 2016 Breach Barometer. As figure 10 illustrates, hacking incidents were constant throughout the year with a total of 222 incidents in 2018 (44.22% of all 2018 breaches) (figure 11). We have data on 180 of those incidents, which affected 11,335,514 patient records (figure 12). For comparison, in 2017, there were 178 hacking incidents, which affected 3,436,742 patient records.

In one hacking incident this past year, hackers used phishing techniques (e.g., emails disguised to look like official emails to employees) to gain access to hospital systems. The hackers successfully gained access to sensitive patient information such as diagnoses, types of care, and possibly bank account numbers. This serves as a critical reminder for healthcare organizations to provide frequent training and education on how to spot phishing emails and what to do if they receive one.

Besides hacking and insider incidents, there were also 61 breaches due to theft. We have data for 47 incidents, which affected 771,656 records. 11 incidents involved missing or lost records, these incidents affected 23,559 patients records.

Finally, there were 67 incidents in which not enough information was available to categorize them. We have numbers for 66 such incidents, affecting 147,216 records.

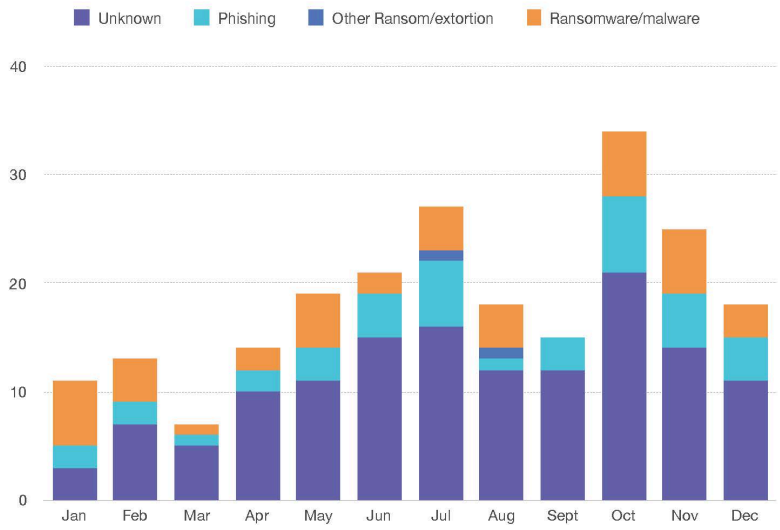

Figure 10. Total hacking incidents, 2018 health data breaches

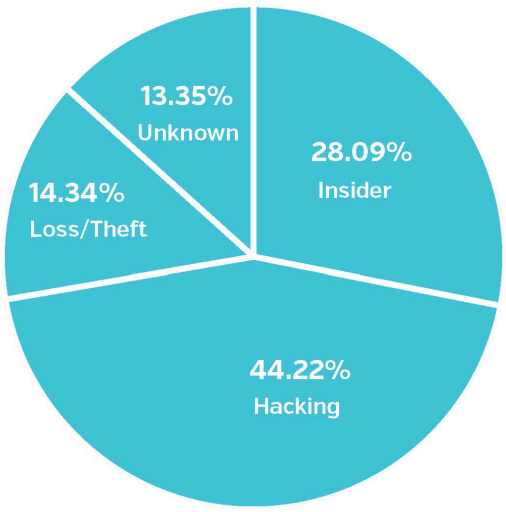

Figure 11. Type of Incidents, 2018 health data breaches

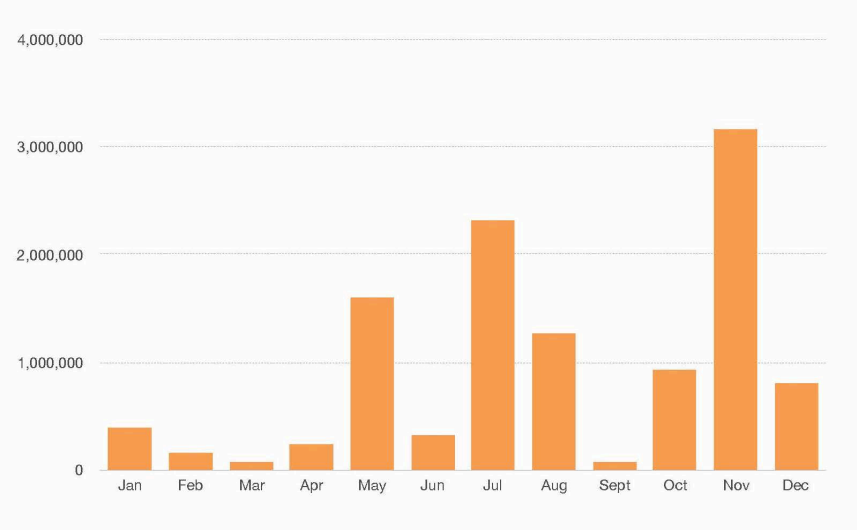

Figure 12. Patient records breached by hacking, 2018 health data breaches

BA-related incidents affect 5.3M patient records

Of the 503 reported incidents in 2018, 353 involved healthcare providers (70% of all reporting entities), 62 involved health plans (12%), and 39 (8%) involved some other type of entity (figure 13).

For the purpose of this report, Business Associates (BA) are defined as third-party vendors that are contracted by health systems to conduct business or provide services on behalf of the healthcare organization.

Of note, there were 49 incidents disclosed by business associates or third parties (10% of total incidents), and at least 102 breaches disclosed by other entities (20% of total incidents) involved a business associate or third party (figure 14). We had information for 84 of these incidents, and they affected 5,328,525 records. Even with the large increase in affected patient records from BA-involved incidents, it should be noted that there could be more incidents involving third parties, but there was not always enough information to make that determination.

Finally, even though most healthcare organizations have already switched over to digitized patient records, 89 incidents involved paper records. These incidents affected 586,728 patient records. It is possible that there are more breaches involving paper records, but again, some reports lacked sufficient detail to make that determination.

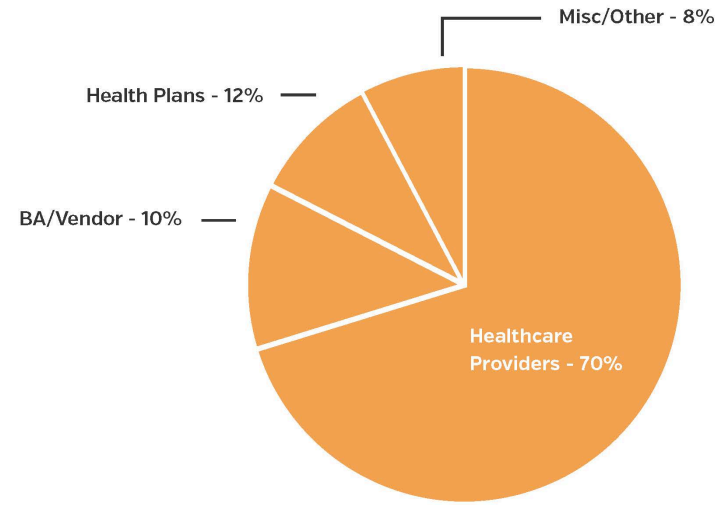

Figure 13. Types of entities reporting, 2018 health data breaches

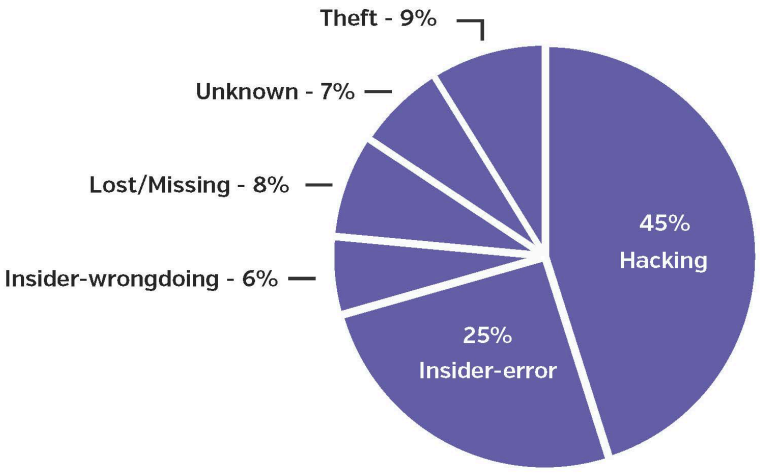

Figure 14. Business Associate/Third Party involvement, 2018 health data breaches

Several insider incidents took over four years to discover

As illustrated in figure 15, of the 141 health data breaches for which we have data, it took an average of 255 days for an healthcare organization to discover

that it had suffered a breach. This represents an improvement from 2017, when it took an average of 308 days for breach detection. The median discovery time in 2018 was 28 days. There were a wide variety of time frames for discovery, with the shortest discovery time being one day and the longest being 5,605 days (15.36 years).

Of the 227 health data breaches for which we have data, it took an average of 73 days for organizations to report a breach to HHS, the media, or other sources after it was discovered (figure 16). These averages seem to be holding steady as this is the same average the industry experienced in 2017. The median disclosure time was 59 days, just squeaking in under the HHS required 60-day reporting window. We hope to report in 2019 that detection and reporting continue to improve as healthcare organizations continue to leverage healthcare privacy best practices.

It's important to note that the data set for this analysis varies greatly from month to month, and data wasn't available for every incident that occurred in 2018. As a result, the smaller data set may not provide a complete picture of reporting times throughout the year.

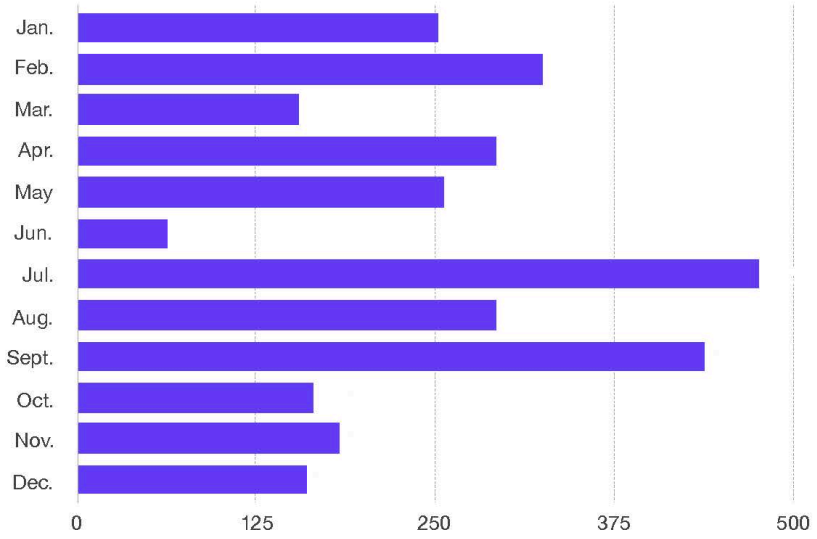

Figure 15. Average number of days from breach to discovery, 2018 health data breaches

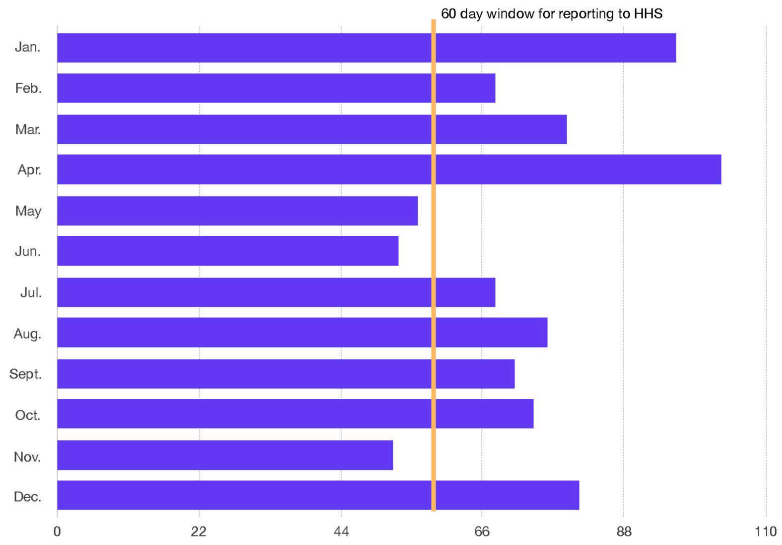

Figure 16. Average number of days from discovery to reporting to HHS, 2018 health data breaches

In general, healthcare entities are able to detect hacking incidents quicker than insider incidents. In many cases, hacking incidents have been discovered in one day, as noted above, where insider incidents can take place for years before discovery. While hacking incidents may be discovered quickly, they also tend to have longer gaps between the discovery of the breach and reporting it.

Insider incidents were associated with the longest gaps between the breach occurring and it being detected. This can be the case because insiders have legitimate access to the EHR, making it easier for inappropriate accesses to fall under the radar. As we discussed above, the longest breach reported this year continued for 15 years before it was discovered. And this incident is not alone. There were seven other health data breaches for which we had data that took four year or more to detect.

State Frequency

48 states (96%) are represented in the 503 incidents for which we had location data, in addition to Puerto Rico. One incident did not have enough information to determine its location, and two states did not have any reported breaches: Delaware and South Dakota. California had the most reported incidents with 63, followed by Texas with 38, and Florida with 31. Please note that numbers for some states are inflated because the analysis uses the state where the BA/vendor is located, not where the client is located.

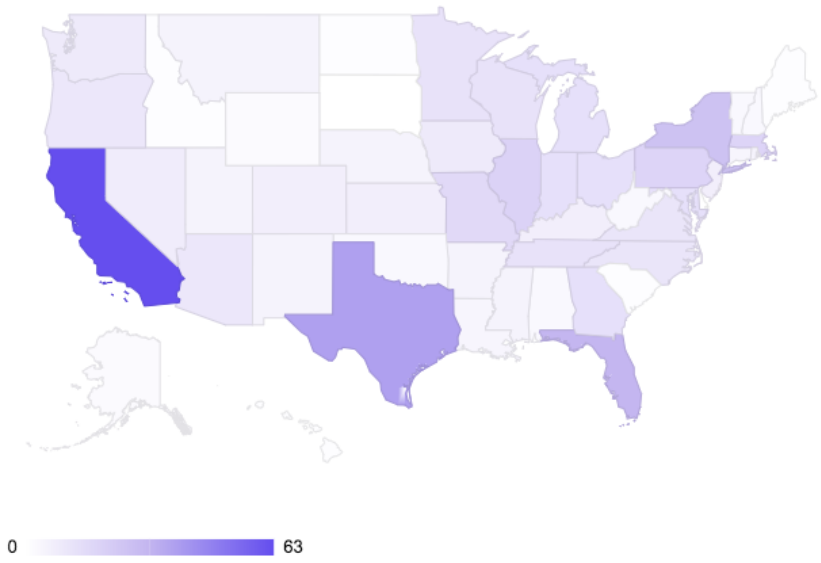

Figure 17. Number of incidents by state, 2018 health data breaches

Conclusion

Healthcare continues to be highly targeted by hackers and other malicious attackers, with the trend of at least one health data breach per day continuing throughout the year. There was a notable increase in breach incidents and nearly a tripling of the amount of affected patient records. In last year's report, experts noted that the decrease in breached records might have been

due to malicious actors taking a break before a resurgence in 2018. With the increase in numbers and the return of thedarkoverlord (TDO), a notable hacking group, it appears this break is now over. It will be imperative for the healthcare industry to continue to innovate and to proactively detect and mitigate these breaches, reducing the devastation these incidents can cause.

It's also important to note that the trend of at least one breach per day that began in 2016 is expected to continue in 2019. In fact, we may continue to see an increase in the number of incidents reported to HHS next year. The industry is getting better at breach detection by using advanced analytics to reduce overall risk to their organization, but phishing techniques are of concern and seem to be increasingly popular with hackers. Hospital employee education and training to detect and not fall victim to these attacks will be imperative to get ahead of the hacking incidents currently plaguing healthcare.

As organizations gain the ability to monitor every access to the EHR and detect suspicious behavior as soon as it occurs, this will hopefully mean that the industry will see a decrease in the number of records affected by health data breaches in 2019. For this to happen, it is vital for health systems to make health data security a top priority, gaining insight into how patient data moves through the organization and gaining the ability to differentiate between appropriate and inappropriate access to patient information. The healthcare industry needs solutions that are tailor-made to meet the unique challenges and requirements these entities face in enforcing best practices within their organizations. Armed with the latest information and utilizing the latest advances in technology, the healthcare industry can gain unprecedented visibility into EHR access which will ultimately make their institutions more secure and ensure patient trust.

About Protenus, Inc.

Protenus is a healthcare compliance analytics platform that uses artificial intelligence to detect inappropriate activity in hospital EHR systems. The Protenus platform uniquely understands the clinical behavior and context of each person accessing patient data to determine the appropriateness of each action, elevating only true threats to privacy, security and compliance teams. Protenus and its partner health systems are fundamentally improving the way hospitals protect their patient data—further ensuring trust in healthcare.

About DataBreaches.net

DataBreaches.net is a web site devoted to reporting on data security breaches, their impact, and legislative developments relevant to protecting consumer and patient information. In addition to providing news aggregation from global sources, the site also features original investigative reporting and commentary by the site's owner, a healthcare professional and privacy advocate who writes pseudonymously as "Dissent."

Methodology

The purpose of this section is to explain decisions that were used to guide our analyses.

Sources

Incidents included in the analyses for this report were compiled for Protenus by DataBreaches.net, and include:

- Incidents reported to HHS between January 1, 2018 – December 31, 2018 that appear on their public breach tool. Incidents reported to HHS before December 31 that were not added to the breach tool in time have not been included.

- Incidents that were reported to other federal or state regulators such as SEC filings or state-mandated notification to state attorneys general or consumer protection agencies;

- Publicly disclosed incidents involving U.S. organizations or entities that are not HIPAA-covered entities but that involved what would be considered protected health information under HIPAA;

- Incidents based on research by DataBreaches.net that may not have been reported to federal or state regulators.

Coding

In addition to going beyond HHS's public breach tool to find breach incidents, this report also uses significantly different coding and analysis than HHS's public breach tool, permitting analyses that are not readily conducted based on HHS's tool, as follows:

- HHS's "unauthorized access/disclosure" category was abandoned in favor of a more refined analysis that allowed us to do a deeper dive into the rate and scope of insider/human error breaches vs. insider/intentional wrongdoing breaches.

- HHS's "Hacking/IT incident" led to further analysis of incidents reported in that category to determine if there was actually an external attack or if – as was the case in a number of incidents – entities were reporting being "hacked" when it might be more accurate to describe the incident as an unintended exposure of PHI on public FTP servers that researchers or others then accessed. In those cases, regardless of how the entity submitted the incident to HHS, our analysis coded those incidents as "insider-error," just as failures to restore firewalls after an upgrade that resulted in data acquisition were coded as "insider-error."

Calculating Time to Reporting

The inclusion of numerous third-party incidents resulted in the decision that for purposes of determining time intervals for "date of breach to date of discovery" and "date of discovery to date of public report," we would define the "discovery date" as the date that the third party first discovered the breach, and not the date that they first informed the covered entity about it.

In calculating time intervals between date of breach and date of public report, we defined the date of public report as the date that the entity first reported the incident to HHS or a regulator, or the date that there was a media report or something like a Twitter announcement that made the public aware that there had been an incident.

In some cases, we did not have exact dates, but only knew the month or year the breach first occurred. In calculating the interval between the breach to discovery and between the breach and reporting:

- If data was only available for the month or year of the breach, the first day of the year or month was used for calculation purposes.

- The date a BA/vendor first discovered the breach was used as the discovery date and not the date the covered entity first learned of the breach.

State Data

For state frequency data, if a Business Associate or vendor was responsible for the breach, we assigned the breach to the state where the BA or vendor is headquartered or located, if the third party's identity was known. In cases where the third party's location could not be determined, the incident was assigned to the covered entity's state.

Any inquiries about the data collection or analyses should be directed to kira@protenus.com.

Disclaimer

This report is made available for educational purposes only and "as-is." Although we have tried to provide accurate information, as new information or details become available, any findings or opinions in this paper may change. Despite our diligent efforts, we remain convinced that the breaches we find out about publicly are only the tip of a very, very large iceberg, and any patterns we see in publicly disclosed breaches may not mirror what goes on beneath the tip.