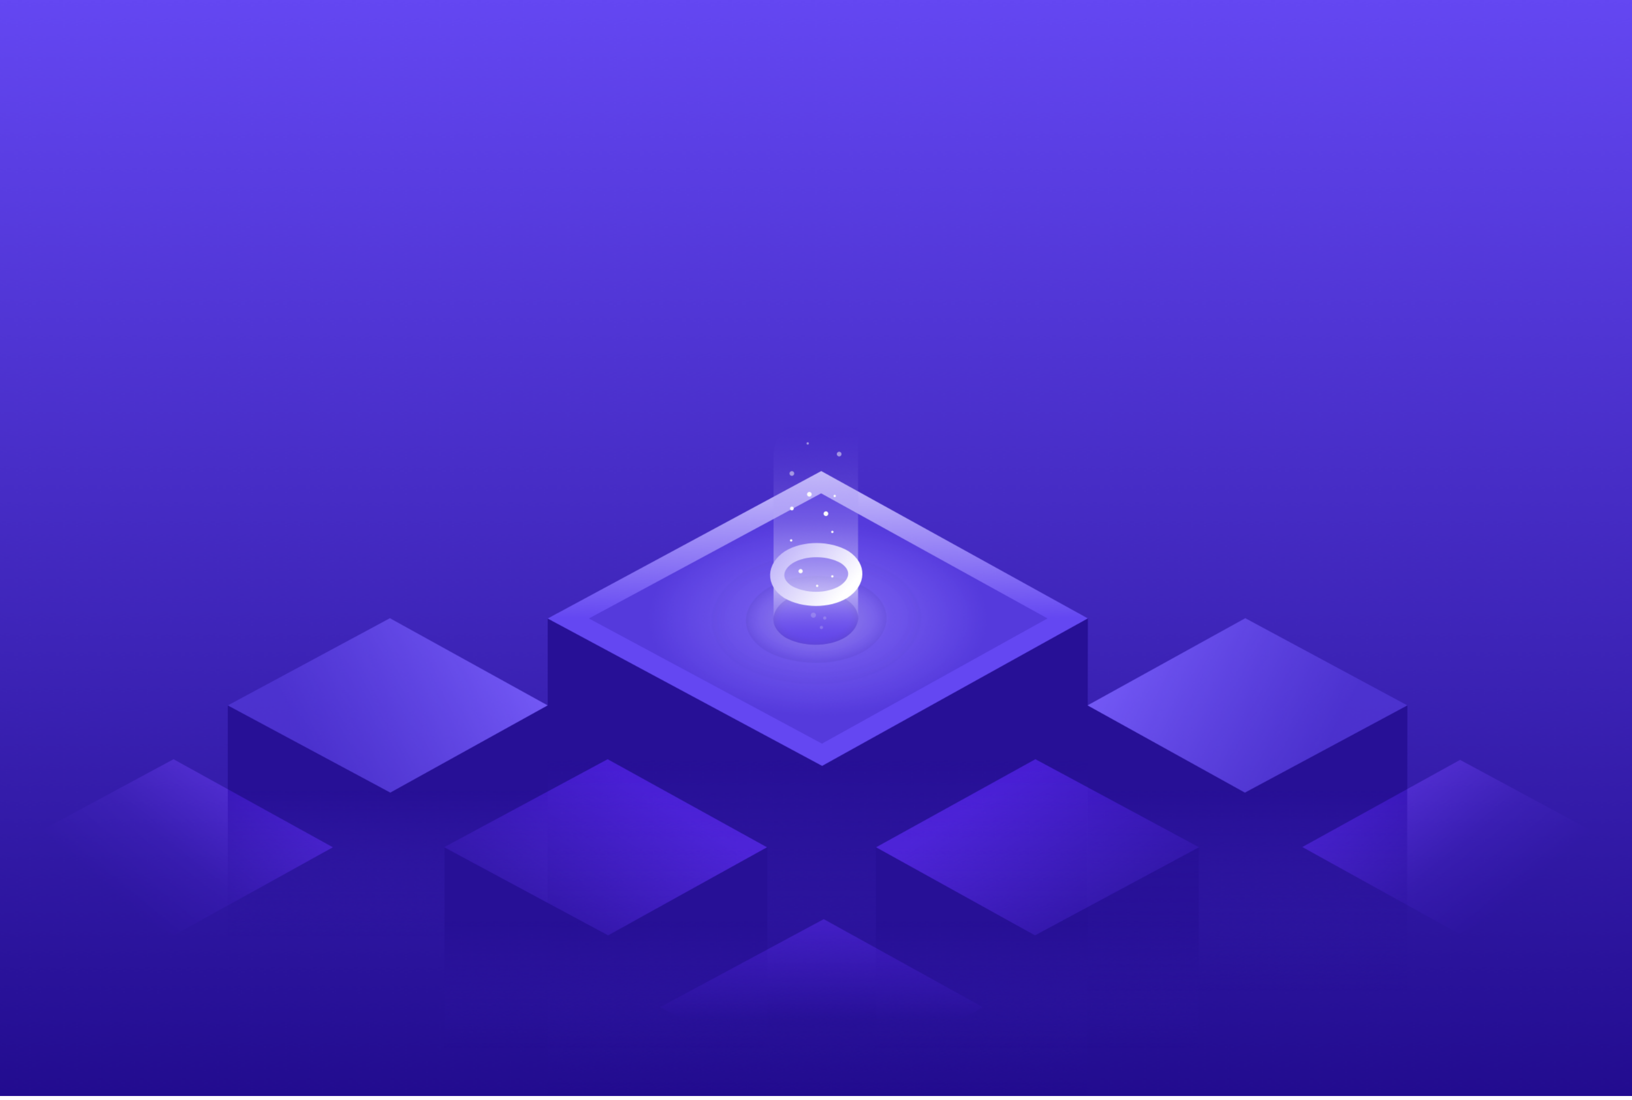

BREACH BAROMETER REPORT: YEAR IN REVIEW

5.6M Patient Records Breached in 2017, as Healthcare Struggles to Comprehensively and Proactively Detect Health Data Breaches

Protenus, Inc. in collaboration with DataBreaches.net

Introduction

Throughout 2017, patient medical information remained vulnerable to both internal and external threats. Hacking and ransomware attacks as well as malicious insiders continue to jeopardize the security of protected health information (PHI). Several insider incidents reported this year had gone undetected for several years before healthcare organizations even knew they had been breached. Although 2017 saw fewer massive health data breaches when compared to 2016, there was still an average of at least one health data breach per day throughout the entire year. Progress is being made, but there is still much that healthcare organizations must do in order to ensure that the patient data entrusted to them is properly secured.

This retrospective examines 2017 health data breaches with an eye towards lessons learned and a way forward for protecting patient privacy.

Overview of 2017 Findings

Our analysis is based on 477 health data breaches reported to HHS, the media, or some other source during 2017. We have numbers for 407 of those incidents, which affected 5,579,438 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 (450 in 2016 compared to 477 in 2017), but there was also a drastic decrease in the number of affected patient records. In 2016, 27,314,647 records were affected by health data breaches, over five times greater than the number of records affected in 2017, though it should be noted that there were a few massive hacking incidents that contributed to these very large numbers (figure 2).

The single largest breach reported in 2017 (figure 3) was the result of insider-wrongdoing. It involved two separate occasions in which a hospital employee inappropriately accessed the billing information of 697,800 patients on an encrypted USB and CD. The employee implied that they needed the data to perform their job, but the investigation found there was no work-related reason for the employee to access the information.

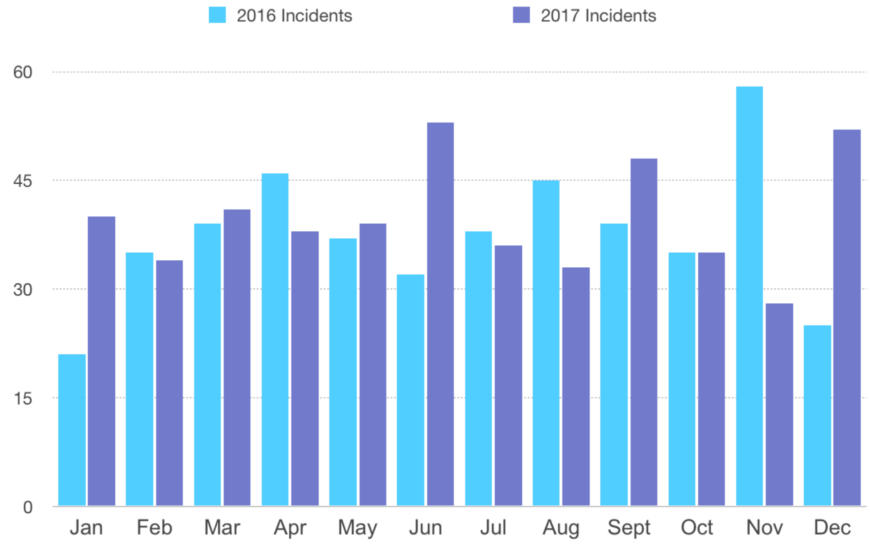

Figure 1. Total disclosed incidents, 2016 vs. 2017 health data breaches

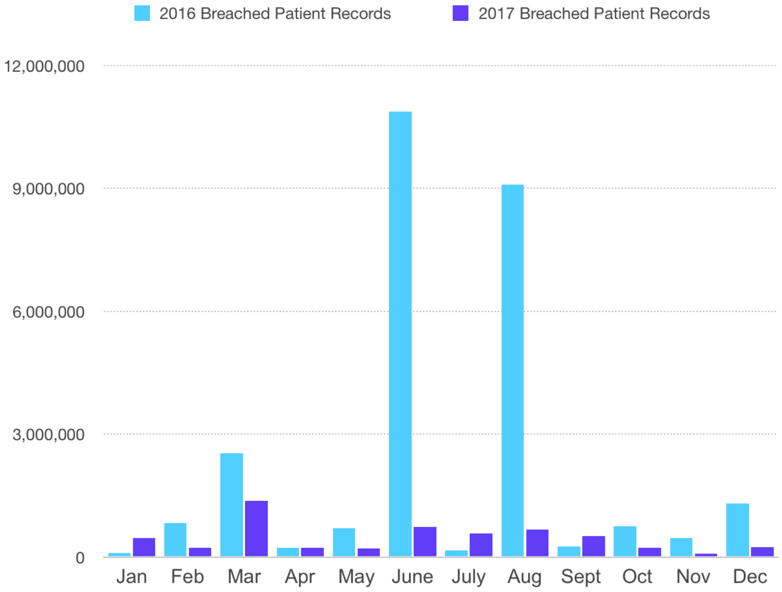

Figure 2. Total breached patient records, 2016 vs. 2017 health data breaches

2017 Largest Health Data Breaches	Organization type	Type of Breach	Number of affected patient records
January	Business Associate	Insider-error	220,000
February	Health Plan	Insider-error	100,000
March	Provider	Insider-wrongdoing	697,800
April	Provider	Hacking	93,323
May	Provider	Insider-error	75,000
June	Provider	Hacking	500,000
July	Provider	Hacking	300,000
August	Provider	Hacking	266,123
September	Provider	Hacking	128,000
October	Business Associate	Insider-error	150,000
November	Provider	Hacking	16,474
December	Provider	Insider-wrongdoing	29,579

Figure 3. Largest incidents, 2017 health data breaches

As figures 4 and 5 demonstrate, there was no linear trend in the number of breaches or number of affected patient records in 2017. June had the greatest number of breaches disclosed and March had, by far, the greatest number of patient records breached since this also included the largest breach incident of the entire year.

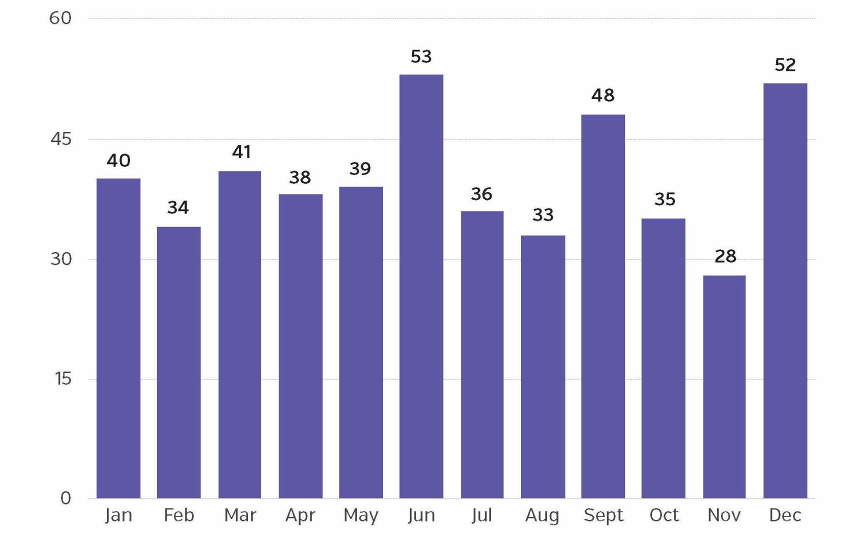

Figure 4. Total disclosed incidents, 2017 health data breaches

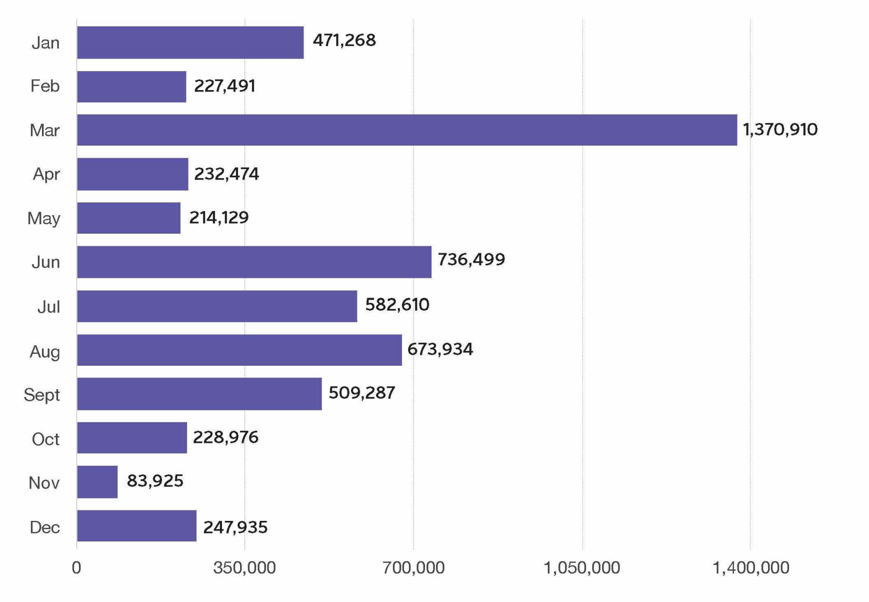

Figure 5. Breached patient records, 2017 health data breaches

Insiders Responsible for 176 Incidents in 2017

Unfortunately, insider incidents continue to plague the healthcare industry in 2017, with one incident remaining undiscovered for 14 years. Insiders were responsible for 37% of the total number of breaches this year (176 incidents), which is similar to 2016 findings. As highlighted in figure 6, we had information for 143 of those incidents, which affected 1,682,836 patient records (30% of total affected patient records). This year's insider-related incidents and patient records were lower then those in 2016, where 192 incidents were disclosed and 2,000,262 patient records were affected.

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.

There were 102 incidents that involved insider-error in 2017 and we have data for 86 of them. In contrast, 70 incidents involved insider-wrongdoing and we have information for 57 of these incidents. It should be noted that there are four incidents in which there was not enough information to classify the incident as either insider-wrongdoing or insider-error. Insider-error affected 785,281 patient records and insider-wrongdoing affected 893,978 records. Figure 7 highlights that more patient records were breached by insiders with malicious intent than by insider-error even though there were fewer insider-wrongdoing incidents.

One particular incident of insider-wrongdoing serves as a reminder of just how dangerous insider threats can be, since employees' legitimate access to patient information can be detrimental when their access is abused with malicious intent. In this case, a hospital employee was snooping on patient information for 14 years before the breach was discovered. The breach affected 1,100 patient records and remained undetected until one of the patients called in with a complaint. This is an unfortunate example of how detrimental insider threats can be for a healthcare organization. This entity will now face a multitude of costs associated with a breach in addition to already taking additional measures to further secure their patients' sensitive medical information. It's important to note that while hacking incidents are often quickly discovered because of the immediate disruption they have on an organization's day-to-day operations, insider threats can remain undiscovered for long periods of time.

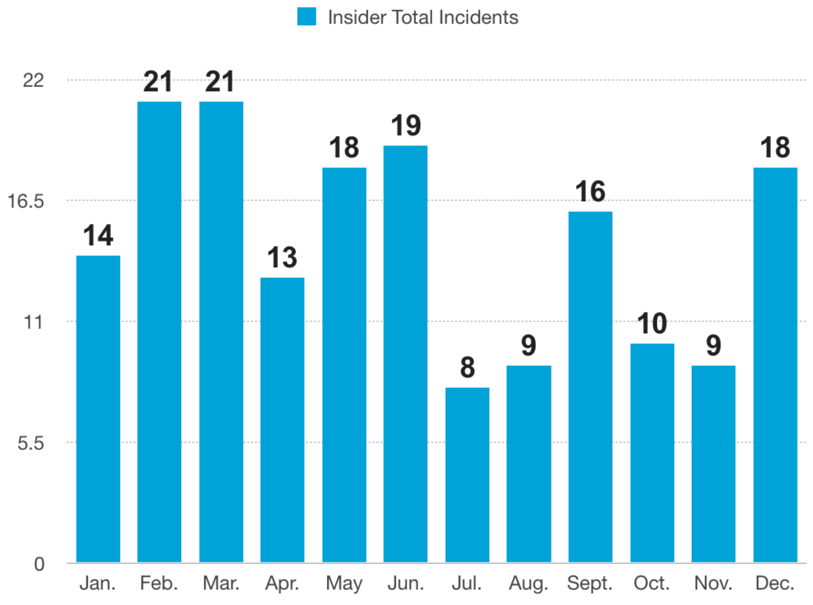

Figure 6. Insider-related incidents, 2017 health data breaches

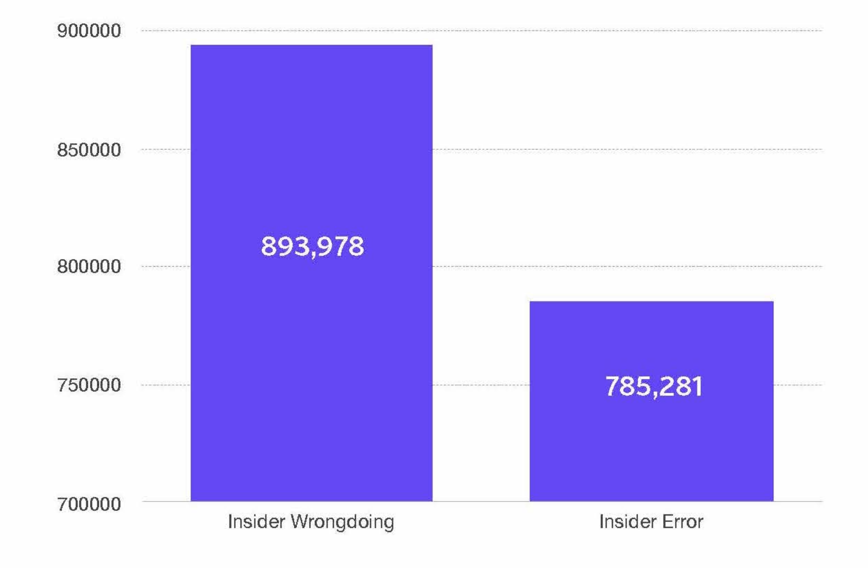

Figure 7. Patient records breached by insiders, 2017 health data breaches

Hacking incidents involving ransomware and malware seemingly double from 2016 to 2017

Continuing a trend that began in 2016, the healthcare industry in 2017 was rocked by a barrage of ransomware and malware attacks. Many healthcare organizations recognize the threat these external attacks represent. A recent survey found that the majority of information technology and security officials believe that giving employees proper training when it comes to cybersecurity was second only to preventing malware/ransomware attacks when it came to implementing a cybersecurity strategy. This reinforces that healthcare leaders recognize the importance of preventing and mitigating hacking incidents, specifically those that have involved ransomware or malware.

As figure 8 illustrates, hacking incidents were constant throughout the year with a total of 178 incidents in 2017 (37% of all 2017 breaches). We have data on 144 of those incidents, which affected 3,436,742 patient records (figure 9). In 2016, there were 120 hacking incidents - those incidents accounted for 87% of all affected records (23,695,069 patient records). As a result, although there were 58 more hacking incidents in 2017, there was a significant decrease in the number of records that those incidents affected. This can be attributed to the lack of the massive hacking incidents like those we reported in 2016.

Of note, healthcare organizations in 2017 reported many more incidents of ransomware and malware. There were only 30 incidents reported in 2016, whereas in 2017, 64 incidents were reported that specifically mentioned ransomware or malware. It is entirely possible, however, that this increase in ransomware attacks is simply due to the fact that more organizations are better about reporting ransomware and have taken OCR's guidance on what to do when an organization has experienced a ransomware attack. 18 of the 178 hacking incidents mentioned the use of other types of ransomware or extortion methods, and 31 incidents involved phishing attacks.

Besides hacking and insider incidents, there were also 58 breaches due to theft. We have data for 53 incidents, which affected 217,942 records. 18

incidents involved missing or lost records, these incidents affected 20,019 patients records.

Finally, there were 47 incidents in which not enough information was available to categorize them. We have numbers for 46 such incidents, affecting 221,899 records.

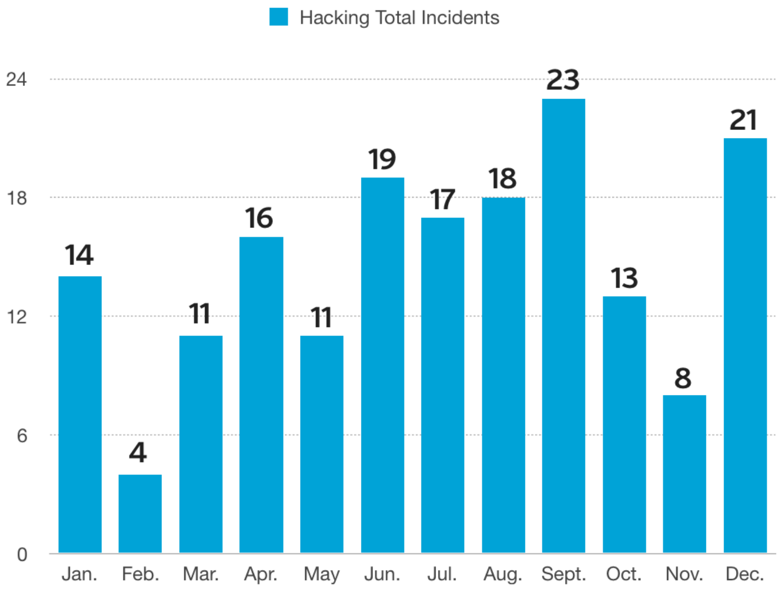

Figure 8. Total hacking incidents, 2017 health data breaches

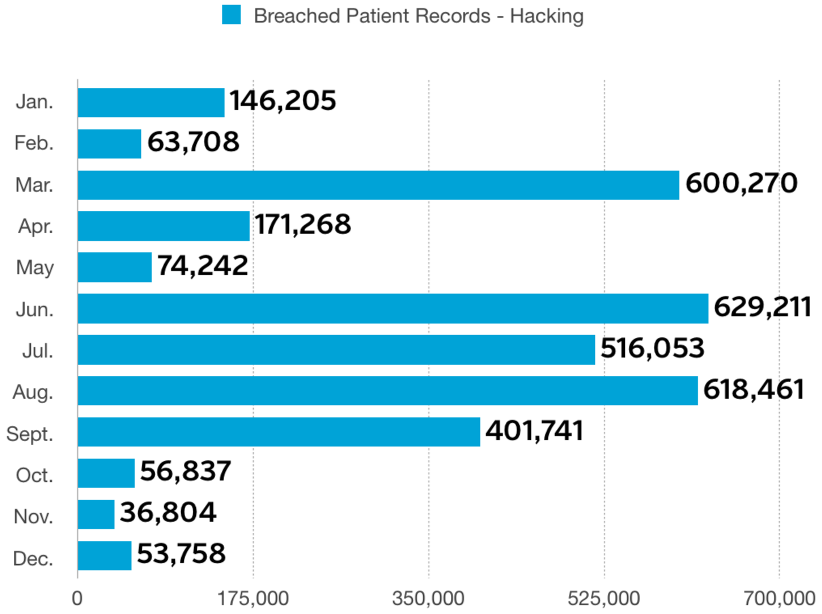

Figure 9. Patient records breached by hacking, 2017 health data breaches

Will healthcare proactively detect health data breaches in 2018?

In 2016, we predicted that it would be the 'Year of Insider Breach Awareness'. Although we did not see a substantial shift in insider-related breaches, overall awareness has increased with better reporting and guidance from HHS and OCR on what to do in the aftermath of a breach. It's important to note that hacking and insider incidents were nearly equal in terms of how often they occurred throughout the year (figure 10). To see continued improvement in detection and reporting in 2018, healthcare leaders will need to build upon the progress made this past year by comprehensively auditing every access to the EHR to ensure threats to patient privacy are proactively detected and mitigated.

Overall, in 2017, there was a significant decrease in the total number of records breached but experts are unsure if this is an indicator of breach prevention or if malicious actors are taking a breath before a resurgence of attacks in 2018. One thing is for certain: the healthcare industry needs to continue to work hard to proactively detect and mitigate these breaches in order to reduce the overall devastation these incidents leave in their wake.

While there were fewer health data breaches that affected massive amounts of patient records in 2017, the trend of at least one breach per day that began in 2016 is expected to continue into 2018. In fact, we could see an increase in the number of incidents reported to HHS next year, but this would most likely be the result of the industry getting better at breach detection, rather than there actually being more incidents. As healthcare 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 continue to see a decrease in the number of records affected by health data breaches in 2018. However, it could also mean a short-term increase in incidents detected as we improve detection rates on our way to changing culture for the better in the long-term.

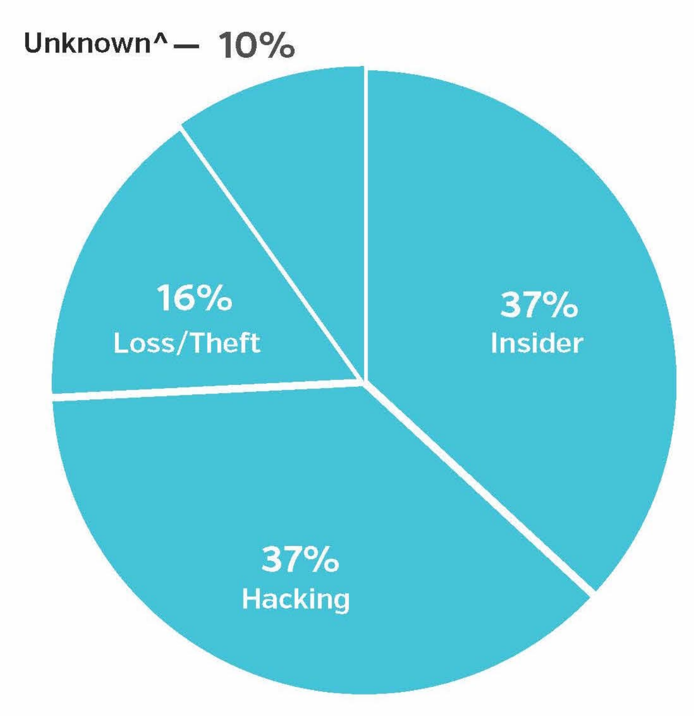

Figure 10. Type of incidents, 2017 health data breaches

Types of Entities Reporting

Of the 477 reported incidents in 2017, 379 involved healthcare providers (80% of all reporting entities), 56 involved health plans (12%), and 19 (4%) involved some other type of covered entity such as schools or law firms (figure 11).

Significantly, while there were 23 incidents reported by business associates or third parties (5% of total incidents), at least 66 breaches reported by other entities (14% of total incidents) involved a business associate or third party (figure 12). We had information for 53 of these incidents, and they affected 647,198 records. While both the number of breaches involving a BA and the number of records affected by these breaches are lower than in 2016, 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, 78 incidents involved paper records. It is possible that there are more breaches involving paper records, but again, some reports lacked sufficient detail to make that determination.

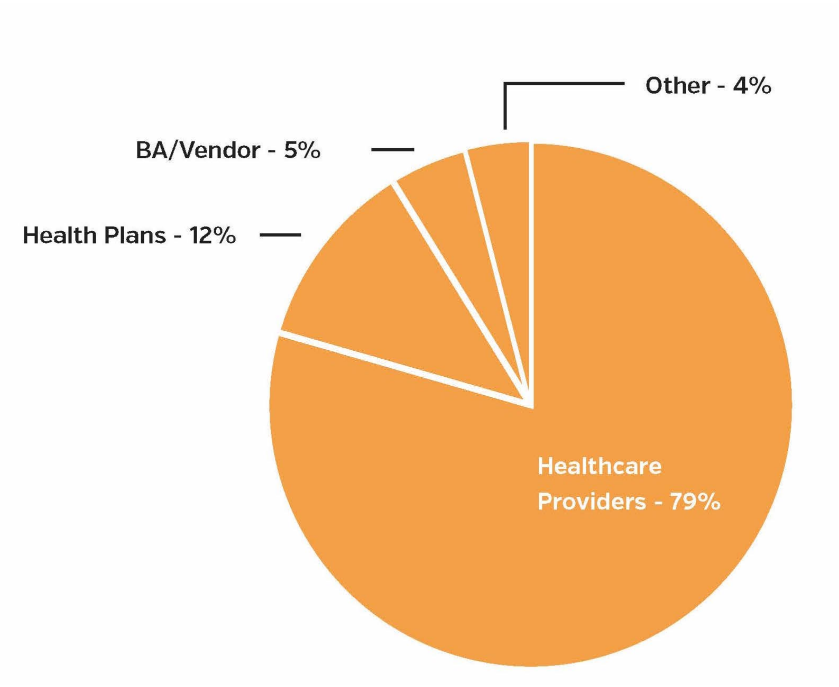

Figure 11. Types of entities disclosing, 2017 health data breaches

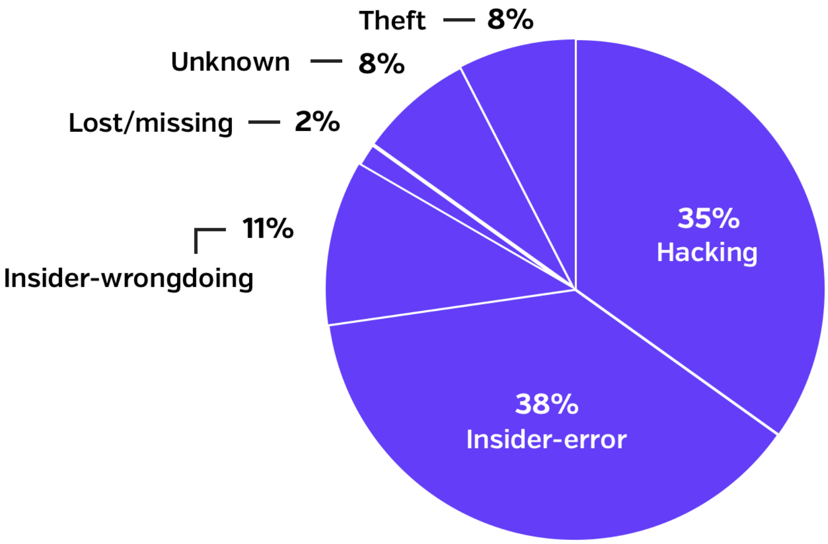

Figure 12. Business associate/Third-party involvement, 2017 health data breaches

Healthcare Entities Suffer Setback in Average Time Taken for Breach Detection

As illustrated in figure 13, of the 144 health data breaches for which we have data, it took an average of 308 days for an organization to discover that it had suffered a breach. This represents a significant setback from last year, when it took an average of 233 days for breach detection. This setback is partially due to the number of breaches reported in 2017 that had occurred for several years, some over a decade, before they were discovered.

Of the 220 health data breaches for which we have data, it took an average of 73 days for organizations to report a breach to HHS after it was discovered (figure 14). This seems to be a vast improvement from 2016, when it took an average of 344 days to report to HHS. Even still, health data breaches need to be reported to HHS within their required 60-day window, or civil monetary penalties could be levied. While this improvement is a great sign, we hope to report in 2018 that the yearly average fell within that 60-day window.

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 2017. As a result, the smaller data set may not provide a complete picture of reporting times throughout the year.

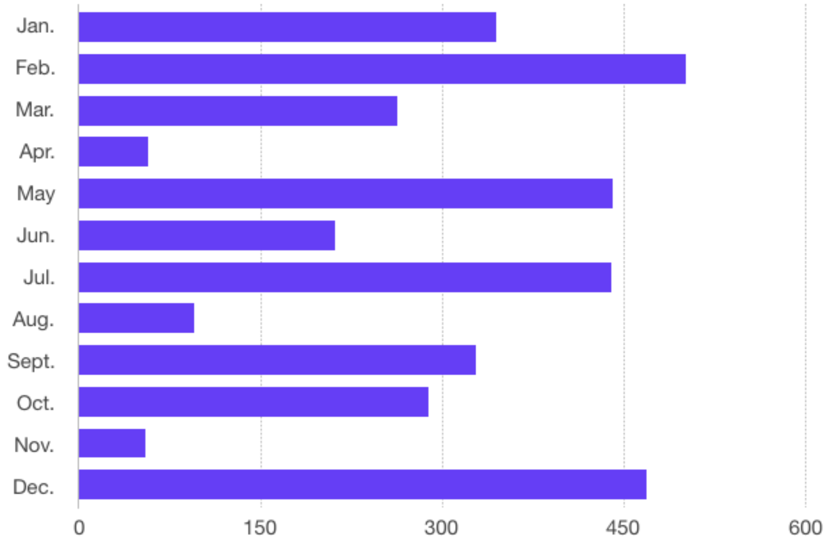

Figure 13. Avg. number of days from breach to discovery, 2017 health data breaches

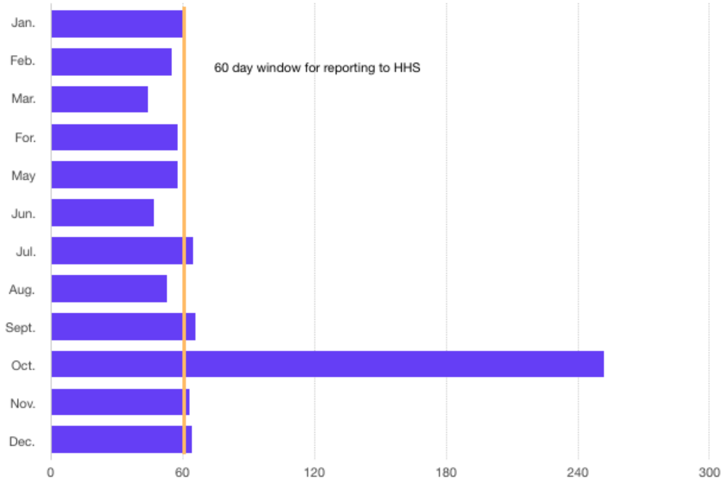

Figure 14. Avg. number of days from discovery to reporting, 2017 health data breaches

In general, healthcare entities are able to detect hacking incidents quicker than insider incidents, but hackings tended to have longer gaps between the discovery of the breach and reporting it. This may be due, in part, to law enforcement officials asking organizations not to disclose the breach publicly as they can continue their investigation.

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 14 years before it was discovered. And this incident is not alone. There were five other health data breaches for which we had data that took three or more years to detect.

In 2018, we estimate the gap will continue to close and the average time for reporting will fall within the mandated 60-day window as healthcare leaders continue to be responsible for reporting their incidents in a timely manner. Reporting will also become easier as healthcare organizations perform comprehensive reviews on how their data is accessed and used. Armed with this information and utilizing the available tools that are tailored for healthcare will enable organizations to immediately detect threats to their organization, take appropriate steps to mitigate their risk, and enforce security policies within the organization.

State Frequency: What's Going on in Hawaii, Idaho and New Mexico?

47 states (94%) are represented in the 477 incidents for which we had location data, in addition to Puerto Rico and the District of Columbia. Two incidents did not have enough information to determine its location, and three states did not have any reported breaches: Hawaii, Idaho, and New Mexico. California had the most reported incidents with 57, followed by Texas with 40, 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.

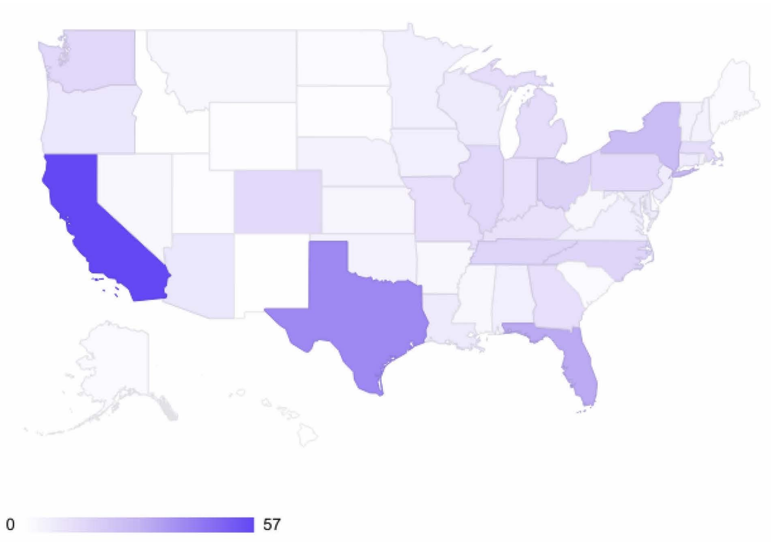

Figure 15. Number of incidents by state, 2017 health data breaches

Conclusion

As we reflect on 2017 and move into 2018, we can expect the healthcare industry to be, yet again, the most targeted industry by hackers and other malicious attackers, with the trend of at least one data breach a day continuing throughout the year. This means it is absolutely vital that healthcare organizations make data security a top priority. In particular, comprehensively understanding the clinical context of a user's behavior allows organizations to easily 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.

In 2018, healthcare has the opportunity to build upon the great strides that have been made to proactively combat the data breaches that are still plaguing the industry. 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 health data analytics platform that uses the latest big data techniques and Protenus-led advances in data science, machine learning, visualization, and software engineering 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, 2017 – December 31, 2017 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.

As a result of our broader approach to investigating breaches that put protected health information at risk, this report includes the 298 incidents on HHS's breach tool plus an additional 179 incidents, for a total of 477 in our sample.

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 "inside – 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.