Information and Privacy Commissioner,
Ontario, Canada

Commissaire à l'information et à la protection de la vie privée,
Ontario, Canada

PHIPA DECISION 210

Complaint HR22-00036

[A Public Hospital]

June 16, 2023

Summary: A public hospital notified the office of the Information and Privacy Commissioner of Ontario (the IPC) of a breach under the *Personal Health Information Protection Act* (the *Act*), as there had been a cyberattack against the hospital. After the hospital self-reported the breach, the IPC opened a file relating to this breach, and subsequently received four complaints from affected individuals. During the cyberattack, the threat actor accessed numerous hospital systems, via a password-spraying attack that compromised an account with privileged access. The hospital took immediate steps to disable the affected accounts and fix the firewall issue that had allowed for the access to occur. The hospital found that the threat actor had exfiltrated large amounts of information, but was not able to determine the exact data that had been taken. The hospital did determine the types of personal health information that may have been accessed, and estimated the number of patients who may have been affected. The hospital provided public notice of the breach, and has agreed to continue to monitor the dark web for two years for any activity relating to this breach.

The hospital provided the IPC with numerous guidelines in place addressing information security, all of which were revised following the cyberattack. These included guidance on strength of passwords, limitation on privileges granted to accounts, and firewall protections. The hospital also provided the IPC with a breach protocol specific to cybersecurity incidents, which was put in place following the incident. In light of the steps taken by the hospital to remediate the situation, including the guidance now in place, I have concluded that it is not necessary to pursue a review of this matter under Part VI of the *Act*.

Statutes Considered: *Personal Health Information Protection Act*, 2004, S.O. 2004, section 12(1) and 12(2).

BACKGROUND:

[1] On February 1, 2022, a public hospital (the hospital) notified the Information and Privacy Commissioner/Ontario (the IPC) of a breach under the *Personal Health Information Protection Act* (the *Act* or *PHIPA*), after having been subject to a cyberattack. The IPC opened Complaint HR22-00036 to address this matter.

[2] In early 2022, the hospital noticed unusual activity in its systems and discovered a cyberattack that had begun four days earlier. During this attack, a third party (the threat actor) accessed and transferred personal health information out of the hospital's servers. The hospital later determined that the threat actor had found a hospital server connected to the internet, and performed a password-spraying attack[1] to gain access to it.

[3] This attack compromised fourteen hospital accounts, including one legacy privileged service account. The threat actor used this privileged account to access data across dozens of servers.

[4] During its investigation, the hospital found that the server was normally protected by a firewall. However, a change to the firewall was made and inadvertently not changed back which led to the hospital server being exposed to the internet.

[5] Upon learning of the breach, the hospital severed its servers from the internet and third-party networks, and isolated all systems that showed signs that they may have been compromised. The hospital disabled the privileged account the threat actor used, as well as other compromised accounts, and forced password resets for all accounts in its active directory.

[6] The hospital notified the IPC of this breach within a week of learning of it. After a dialogue with the IPC, the hospital also published a notice, alerting past and present patients to the fact that some of their personal information may have been affected. This notice was posted several months after the breach itself occurred.

[7] During this investigation, the hospital asked that some of the specific information it provided to the IPC not be included in the *PHIPA* Decision, citing concerns that to do so would result in cyber-security risks to the hospital, rendering it more vulnerable to future breaches. Wherever possible, I have left out references to those specifics.

---

[1] See Canadian Centre for Cyber Security definition at https://cyber.gc.ca/en/guidance/strategies-protecting-web-application-systems-against-credential-stuffing-attacks. During a password-spraying attack, the attacker brute forces logins to a system, using a list of many usernames with a small set of common passwords. The threat actor may deduce the usernames from publicly-available information or obtain them via social engineering.

Preliminary matters

[8] There is no dispute, and I find that, the hospital is a "health information custodian" under paragraph 4(i) of section 3(1) of the *Act*.

[9] The hospital provided a list of the data that may have been present on the systems to which the threat actor had access.[2] The hospital stated that the types of information at issue were personal health information pursuant to section 4(1) of the *Act*.[3] I agree, and find that the data on the affected hospital systems included personal health information, as defined under section 4(1) of the *Act*.

[10] The hospital reported this matter to the IPC as a breach of the *Act*. I agree, and find that the incident involved access to individuals' personal health information that was unauthorized under the *Act*.

ISSUES:

[11] In this decision, the following issues will be discussed:

1. Did the hospital take reasonable steps to protect personal health information?
2. Is a review warranted under Part VI of the *Act*?

DISCUSSION:

Issue 1: Did the hospital take reasonable steps to protect personal health information?

[12] Section 12(1) of the *Act* sets out the security obligations of health information custodians as follows:

> A health information custodian shall take steps that are reasonable in the circumstances to ensure that personal health information in the custodian's custody or control is protected against theft, loss and unauthorized use or disclosure and to ensure that the records containing the information are protected against unauthorized copying, modification or disposal.

[13] Section 12(2) requires that in the event that personal health information in their

---

[2] This list is set out in full later in this Decision, at paragraph 17.

[3] The systems also may have contained employee names and numbers, which the hospital states are not personal health information under the *Act*. The inclusion of that data did not affect the hospital's response to the breach of personal health information or my investigation of the hospital's response, and as such, will not be addressed further.

custody or control is stolen or lost, or used or disclosed without authority, the health information custodian must notify the individual at the first opportunity.

[14] Health information custodians, when confronted with a breach of personal health information, should take appropriate steps in response. These include identification of the scope of the breach, containment of the personal health information involved, notification of those affected, and investigation and remediation of the breach. The IPC guidance on these steps is set out in more detail in *Responding to a Health Privacy Breach: Guidelines for the Health Sector* (the PHIPA Breach Guidelines)[4]

Scope and Containment

[15] The hospital determined that during the attack, the threat actor was able to access several dozen hospital systems. Of these, the hospital inferred that the threat actor had browsed certain folders on about half of the systems that had been accessed and exfiltrated some of the data. By looking at its firewall logs, the hospital was able to determine the quantity of data leaving the network to the internet and which systems the data came from. The hospital concluded that the threat actor had transferred data from approximately one in every five systems that they were able to access. However, the firewall logs could not be used to determine the exact scope of the breach, as they do not detail what data was exfiltrated, only the quantity of it.

[16] The hospital conducted further forensic searches in the form of a "targeted, preliminary eDiscovery" on a representative portion of the total affected systems. They did so to better understand the volume, distribution, and composition of the data resident on the systems, which was needed to provide notification. The eDiscovery software searched for data matching particular terms related to personal information or personal health information. This produced a list of 2.8 million responsive documents, though due to the nature of the search, the hospital assumed that this list included duplicate documents. The hospital used the search results to estimate the upper range of the potential number of individuals who may have been impacted, which they estimated to be over one million.

[17] The search also identified the type and quantity of data present on the affected systems at the time of the incident. The hospital found that the categories of personal health information contained on the servers that the threat actor accessed may have included:

- Patient ID (MRN)
- Accession number (further unique patient identifier)
- Patient name (first, last, middle)

[4] Available online at https://www.ipc.on.ca/wp-content/uploads/2018/10/health-privacy-breach-guidelines.pdf

- Patient gender
- Patient date of birth
- Patient marital status
- Patient home address
- Patient postal code
- Patient residential area
- Patient phone number
- Patient email address
- Presenting complaint (i.e., for first responder)
- OHIP number and version
- Insurance policy number
- Provider name
- Provider HCP number (i.e., CPSO)
- Procedure description/ordered performed
- Order
- Result
- Attending and/or ordering physician name, number
- Medical/clinical/diagnosis information/findings/reports
- Lab reports/results
- COVID immunizations

[18] The hospital took immediate steps to contain the breach when it learned of it. It severed its servers from the internet and third-party networks, and isolated any systems showing signs of compromise. The hospital disabled all compromised accounts, including the one used by the threat actor, and forced password resets for all accounts in the hospital's active directory.

[19] The hospital was not able to contain the data that the threat actor had already transferred out before the hospital found out about the breach. However, it did make

efforts to limit any further spread of this data, by monitoring the dark web for signs of any data that may have been obtained from this breach. This included monitoring known threat actor channels, forums, and marketplaces for: evidence of the intrusion; attempted sale or transfer of the exfiltrated data; any discussion or message board "chatter" about the hospital or the cybersecurity incident itself; and any indications of ongoing of future attacks against the hospital.

[20] The hospital states that it has not yet found any evidence of hospital data on the dark web, but is continuing to monitor this. In its response to the IPC, it noted that dark web monitoring services usually remain in place for 6-12 months after a cybersecurity incident.

[21] In the present case, the threat actor was able to exfiltrate a very large amount of personal health information, presumably causing concern to the patients who may have been affected. While the hospital cannot contain the breach by retrieving the information already exfiltrated, monitoring the dark web allows the hospital to discover if that information has been sold or otherwise used for bad purposes. As such, continuing this monitoring provides some level of assurance to the affected individuals that the hospital is taking reasonable steps available to it to address the breach and mitigate risk of harm. Although the hospital initially noted that, in its view, dark web monitoring services usually remain in place for 6-12 months after a cybersecurity incident, following discussions with the IPC, the hospital agreed to continue to provide dark web monitoring for a period of two years following the date of this decision, which I believe is more in keeping with comparable cases.[5]

Notification

[22] As set out above, the hospital found that a large number of patients may have been affected by the breach, but was not able to determine exactly what information the threat actor accessed or transferred out. The investigation was however able to determine the categories of information on the servers that the threat actor had access to, and may have exfiltrated.

[23] Given the number of patients who may have been affected, and the hospital's inability to say with certainty which individuals were affected, the hospital determined it would not be possible for it to notify the affected individuals directly. Instead, they provided notice by posting a Personal Information Public Notice (the Notice) on the hospital's website. This Notice included an outline of the incident, the types of data that may have been accessed during the incident, and the efforts the hospital made to contain and address the breach. The Notice also stated that the hospital had notified the IPC, and provided a link to file a privacy complaint with the IPC for those who wished to avail themselves of that right. The IPC received four complaints from individuals in relation to this breach.

---

[5] See *Re Nova Scotia Health Authority*, 2020 NSOIPC 2.

[24] The Notice stated that current and former hospital patients could obtain two years of online credit monitoring, without cost to themselves, and provided a number to call to access that service. Also included was a telephone number to contact at the hospital for more information and an apology for the incident.

[25] The hospital posted the Notice four months after the incident occurred, and also issued a news release at that time, resulting in coverage of the breach in many Ontario media outlets.

[26] Section 12(2) of the *Act* mandates notification in the case of the theft, loss, or other unauthorized disclosure of personal information. It requires that health information custodians notify the affected individual at the first reasonable opportunity and that the notice include a statement that the individual is entitled to make a complaint to the IPC regarding that breach. While there was some time between the breach and the posting of the public notice, the hospital was in contact with the IPC throughout, during which it provided the reasons for this timing. The hospital board needed to approve the notice the hospital would be providing, and the delay was due to a postponement of a board meeting. In my view, this is not a valid reason for delaying notice to affected individuals. Given the magnitude and seriousness of this breach, a board meeting should have been specially convened to approve the public notice much sooner.

[27] It is generally better to provide direct notification to individuals who may have been affected by a privacy breach. Correspondence alerts them to their possible involvement in a way that is more likely to draw their attention than a general announcement. However, in this case the hospital knew the servers affected, and the individuals who had personal health information present on those servers, but could not say whose personal information was contained in the data that was exfiltrated. In addition, the hospital estimated the number of individuals who may have been affected at over a million.

[28] Given the very large number of individuals estimated to be affected, and the hospital's inability to say who among them had their personal information exfiltrated, it was reasonable for the hospital to determine that direct notification was not feasible. Direct notification remains the standard notice that health information custodians should provide, and the resort to public notice in this case was dictated by the scale and specific circumstances of this breach.

[29] Taking into account the considerations noted above, I am satisfied that the hospital has provided the notification required by section 12(2) of the *Act* although it should have done so much sooner.

Investigation and Remediation

[30] The IPC's PHIPA Breach Guidelines state that the investigation and remediation

of a breach should include both a review of the circumstances surrounding the breach and a review of the adequacy of existing policies and procedures protecting personal health information. The latter also includes training on these policies and procedures.

Factors contributing to the Cybersecurity Incident

[31] Based on the hospital's response, there were two key elements that led to the threat actor accessing the personal health information at issue in this breach: a change to a firewall, which allowed the threat actor to gain initial entry, and their subsequent use of a heavily privileged legacy account. This allowed the threat actor, once inside, to access large amounts of data across several systems.

[32] The hospital states that prior to the firewall change, systems accessible from the internet were protected by geo-blocking rules, port and protocol filtering, and source IP filtering in some cases. They provided the reasons for the firewall change, and noted that, due to hospital staff oversight, the firewall was not changed back immediately. This error led to some of the hospital's infrastructure being exposed to the internet for one week, rather than being behind the usual protections in place, allowing the threat actor to access systems that were normally protected.

[33] The rule for this firewall was changed back when the hospital discovered the breach. The hospital has since reviewed all firewall rules and configurations to identify and remediate high-risk configurations.

[34] The hospital provided the IPC with its policy regarding network security, which implements a firewall management process. Under this process, any changes to firewalls need to be approved. It also mandates that the default setting is to block traffic, rather than allow it. The firewall rules for high-risk firewalls (such as those connected to the internet) are reviewed every six months. This policy also sets out the protection requirements for all untrusted network perimeters, such as those connected to the internet. In addition, the hospital has adopted a review process under which all firewall changes are peer-reviewed, and firewall logs are reviewed by a security team.

[35] The other key factor contributing to the scale of the breach was the discovery and use of a privileged account by the threat actor. While fourteen accounts were accessed by the threat actor in the initial attack, the hospital found that once the threat actor was inside the hospital's systems, the attack was carried off using only a single account. The hospital described this account as a legacy account that had been a member of a highly privileged group. Via this account, the threat actor utilized several tools to strengthen their position and move between hospital systems. After the breach was discovered, the hospital disabled the fourteen compromised accounts, including the privileged account.

[36] Prior to the incident, the hospital did have access management controls in place, providing separate privileged accounts for system administrators. They also followed

the principle of least privilege when assigning access to users. Nevertheless, the legacy account with its privilege continued to exist, such that the threat actor was able to use it to his advantage during the cyber attack.

[37] Restrictions on account privileges are set out in the hospital's Access Control Standard, which was revised following the incident. This standard continues to apply the principle of least privilege, stating that the hospital must minimize the assignment and use of privileged accounts. It specifies that privileged accounts must not be used for purposes other than operational system support or administration, and that users with privileged accounts must conduct routine business operations using a non-privileged account. Local administrative access rights are not to be enabled by default, and are to be granted on a temporary basis, and then only when that level of access is necessary to do the user's job effectively. The Access Control Standard also requires that the hospital review the privileges for all privileged user accounts on a quarterly basis.

[38] The hospital referred to the privileged account that the threat actor used as a legacy account, implying that its privileges had not been reviewed in some time. The current standard limits which accounts may be granted privileges, and requires periodic review of all privileged accounts. These restrictions help to minimize the number of privileged accounts within the hospital's system, which reduces the likelihood that a breach of this scale will be repeated. Under the current guidance, any similar legacy privileged accounts should have been reviewed no later than the end of May 2022 (which is three months after the Access Control Standard was revised) and had any unnecessary privileges revoked. The hospital confirmed that all legacy accounts in the same highly privileged group of accounts have either been deleted or had their privileges reduced.

[39] Beyond the firewall and account privilege issues, the method that the threat actor successfully used also raises concerns about password security – specifically, concerns about accounts having weak passwords that are easy to guess. As noted above, password-spraying attacks are conducted by a threat actor using the same password or set of passwords on many accounts. Accounts with default or easily discerned passwords are most susceptible to password-spraying attacks. The account at issue has since been disabled, but concerns over weak passwords remained for other active accounts, especially privileged ones.

[40] To address the immediate threat, the hospital reset passwords for all user accounts in the domain. The Access Control Standard also has password requirements in place with minimum standards for character variance and length, and prohibits categories of passwords that could be most readily guessed. Subsequent verification testing found that concerns over shared, weak, or trivial passwords were not identified again.

[41] In addition to the steps described above, the hospital also took the following

steps as part of its remediation plan to address the cybersecurity incident:

- Verified coverage of existing multi-factor authentication;
- Implemented a leading Endpoint Detection & Response solution across the hospital;
- Reviewed the active directory to identify and remediate high-risk configurations;
- Reviewed vulnerability reports to identify and remediate high-risk vulnerabilities;
- Implemented additional security hardening and audit logging configurations on key systems;
- Rebuilt or restored from backup impacted systems where possible, and performed an exception process for other systems. This process included a targeted antivirus scan and indicator of compromise checks, manual removal of any remnants from the incident, additional security hardening controls, targeted threat hunting, and ensuring all patches are up to date; and
- Enhanced Security Information Event Management and Endpoint Detection & Response monitoring capabilities through a Managed Security Service Provider.

[42] The hospital stated that since the incident, its cybersecurity team, in consultation with independent advisors, has developed a security roadmap to help ensure the progress of its ongoing privacy and cybersecurity objectives. The hospital provided the IPC with a list of initiatives that are part of this project, the specifics of which I am not including in this Report due to the hospital's security concerns.

Security Guidance Documents

[43] At the time of the incident, the hospital was in the process of drafting guidance on how to respond to cybersecurity attacks. The hospital's Cyber Incident Response Process (the Response Process) came into force two months following the breach.

[44] The Response Process sets out its purpose as "to define a set of process steps which provides direction on how to respond to cyber incidents which threaten [the hospital's] information security and/or privacy." It notes that security incidents can be detected or reported by a number of sources (such as employees, service providers, or business partners) and states that once reported, a dedicated team must be formed to address them. It categorizes the types of incidents that may occur, together with examples of them, and defines different levels of cyber incident severity.

[45] The Response Process includes specific steps the hospital should take to contain the breach under the direction of the dedicated team, in order to ensure that the containment does not impact any forensic investigation that may follow. It sets out

concrete measures that the dedicated team should consider taking to assist in containment, which will depend on the situation and the nature of the attack. Due to the hospital's security concerns, I will not detail these steps in this Decision, but note that they provide clear and useful guidance.

[46] The next step in the Response Process requires determining the scope of the breach, how it happened, and who needs to be notified. This includes identifying the source of the breach, the timeline involved, whether the intruder exported or deleted any information, and whether the intruder left any remnants that may enable them to gain access in future. The hospital must review the network to identify all compromised or affected systems, and the Response Process sets out where to look and how to do so. It also requires that systems continue to be monitored for any signs of continued intruder access.

[47] As part of this stage, the hospital must determine the classification of the information at issue – is it personal information, personal health information, or corporate information? If the information at issue is either of the two former categories, the dedicated team must immediately notify the hospital's privacy office, so that they can assess the situation and engage with the IPC on the breach. The dedicated team initiates the breach notification process, which is then carried out by the hospital's privacy or communication teams. Finally, the dedicated team must then determine if the incident requires computer forensic investigation from an outside specialist. If it does, the Response Process has requirements in place that the chosen forensic investigator must adhere to in conducting and documenting their investigation.

[48] The Response Process also requires that its information security department provide an incident report to the hospital's executive team and board. There is a review process, under which stakeholders should review the lessons learned from the incident. If the incident report includes recommendations, management should be provided with updates when the recommended activities are completed.

[49] In addition to the Response Process, the hospital has a number of other guidance documents in place that address information security, all of which were revised following the cybersecurity incident. These include:

- Acceptable Use Policy
- Access Control Standard
- Application Security Standard
- Asset Management Standard
- Business Continuity & Disaster Recovery Standard
- Cryptography and Key Management Standard

- Incident Response Standard
- Information Classification Standard
- Information Security Policy
- Information Security Risk Management Standard
- Logging, Auditing, and Monitoring Standard
- Mobile Device Management Standard
- Network Security & Remote Access Standard
- Operation Security Standard
- Physical Security Standard
- Third Party Risk Management Standard

[50] The purpose of the Access Control Standard is to control, manage, and limit access to hospital information to mitigate information security risks, and its relevant provisions are outlined in paragraphs 37 and 40.

[51] The Incident Response Standard requires the hospital to have a security incident management process in place, which includes detection, containment, investigation, eradication, and follow-up, and which addresses different types of security incidents. It also refers to the hospital's Cybersecurity Incident Response Process (as described above) for specifics on such incidents, and mandates that the hospital perform a post incident analysis for all security incidents.

[52] In addition, the Network Security & Remote Access Standard sets out that the hospital network must be segmented into zones via firewalls, according to trust level. It states that personal health information is to be located in a high trust zone and protected by network firewalls and intrusion prevention, where technically feasible and based on risk. It requires that the hospital maintain and implement a firewall rule management process, which includes required protective baseline rules, and reviews of the rule lists for high-risk network zone firewalls every six months.

Security Training

[53] The hospital stated that it provides security awareness training to all users annually, via online interactive training courses. Since the cybersecurity incident, it has increased the number of courses users are required to take, which include:

- Introduction to Phishing Emails and Website;

- Healthcare Data Overview;
- Healthcare Data Protection;
- Healthcare Privacy Violation;
- Multi-Factor Authentication;
- Internet Basics: Internet and Cloud Security;
- Security Basics: Compliance; and
- Security Basics: Passwords and Authentication.

[54] The hospital also performs periodic simulated phishing attacks to ensure that users can identify a phishing email, and proactively alert the hospital after identifying it.

Analysis

[55] The nature of this cybersecurity attack raised concerns regarding the privileges attributed to hospital accounts, the protections the hospital placed on systems containing personal health information, and the strength of passwords used to access accounts.

[56] The hospital has demonstrated to the IPC how it now limits the number of privileged accounts to only those that require additional privileges, and minimizes the privileges granted to those account holders. The hospital has also explained the oversight that led to the system being connected to the internet without firewall protections, and provided their guidance on firewall requirements. They have implemented a new process under which all changes to the firewall are now peer-reviewed. Finally, the hospital has demonstrated that it now has increased its strength requirements for passwords, and has conducted testing to ensure that there were no ongoing vulnerabilities due to weak passwords. The hospital also provided guidance documents addressing various aspects of information security and outlined the annual security awareness training that its staff is required to take.

[57] In addition, the hospital now has a response process in place specific to cybersecurity incidents, which should allow it to respond promptly and rigorously to any such incidents that may occur in future. I recommended that the hospital conduct a breach simulation or table top exercise at least once a year in order to ensure that roles and responsibilities are clearly understood throughout the organization and well-rehearsed in advance of another real breach occurring. The hospital confirmed that it would be adopting this recommendation, and would be including such exercises in their cybersecurity roadmap.

[58] I am satisfied that the hospital has responded adequately to the breach that

occurred. In light of the remediation steps the hospital has already taken to address this matter, I conclude that pursuing a review under Part VI of the *Act* is not warranted.

NO REVIEW:

[59] Section 58(1) of the *Act* sets out the Commissioner's discretionary authority to conduct a review as follows:

> The Commissioner may, on his or her own initiative, conduct a review of any matter if the Commissioner has reasonable grounds to believe that a person has contravened or is about to contravene a provision of this Act or its regulations and that the subject-matter of the review relates to the contravention.

[60] In accordance with my delegated authority to determine whether to conduct a review under section 58(1) of the *Act*, and for the reasons set out above, I find that pursuing such a review is not warranted.

Original Signed by: June 16, 2023

Jennifer Olijnyk
Investigator