Information and Privacy Commissioner,
Ontario, Canada

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

PHIPA DECISION 249

Complaint HR22-00501

[A medical imaging clinic]

June 21, 2024

Summary: A medical imaging clinic 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* or *PHIPA*), following a ransomware attack against the clinic. The threat actor encrypted and exfiltrated files from the electronic medical records and file sharing servers and deleted the clinic's backups. The clinic shut off the servers immediately, and these remained off while the clinic engaged in discussions with the threat actor. The threat actor provided the clinic with a file tree indicating which files they had exfiltrated, and the clinic ultimately decided to pay the ransom. The clinic was then able to decrypt all information on the affected servers and recover all files.

The clinic's virtual private network kept logs of connections to its systems, but these logs had limited data storage. Because there was so much activity during the attack, the logs ran out of storage and the more recent data wrote over older events. The records from the earlier part of the attack were therefore lost before the clinic could review them. This limited the clinic's investigation, though they did find that a dormant account belonging to a former internal application developer had been active during the attack. They concluded that this account, which had significant administrative privileges, was used by the threat actor to first gain access to the system and then move to other servers.

The clinic posted notices regarding the cyberattack within its physical locations and online. These notices listed the categories of information in the file tree. Later notices also acknowledged that patient medical records were stored in the affected servers, but that there were no signs that these records had been accessed. The clinic clarified to the IPC that its forensic experts examined all files for indications of access by the threat actor. The areas for which there was evidence of access correlated with what the threat actor had set out in its file tree.

The clinic revised its guidance to include improved password security, limitation on privileges granted to accounts, deletion of dormant accounts, and improved patch management. It put in place additional security measures, including replacement of the virtual private network with a newer model with enhanced storage to keep logs indefinitely, and extended detection and response capabilities. The clinic now always keeps one backup offline. In light of the steps taken by the clinic to remediate the situation, 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, c. 3.

BACKGROUND:

The Attack

[1] On December 23, 2022, a medical imaging clinic (the clinic) notified the IPC that it had been the victim of a ransomware attack a week earlier, on December 16, 2022. During the attack, a known hacking group (the threat actor) encrypted and exfiltrated files from their systems. Based on the systems affected, the affected files could have included up to 550,000 patient records and 1,600,000 case files.

[2] The clinic discovered the attack due to unusual activity on its systems, and immediately engaged breach counsel and a team of cybersecurity experts to assist with the containment, investigation, and remediation of the breach. These agents found a ransom note left by the threat actor and engaged in communications with them.

[3] These experts determined that the attack had started five days before the clinic discovered it, on December 11, 2022, and that the initial attack vector was likely compromised virtual private network (VPN) credentials belonging to a deceased account holder. The account holder had been an internal application developer with elevated privileges. They found that years-dormant secure sockets layer (SSL) VPN credentials had been used at the time of the breach to gain access to the clinic's systems through this account. They found no other suspicious access from any other source. The clinic later deleted the compromised account.

[4] The clinic and its experts were not able to confirm their suspicions about the account used via VPN logs, as these were no longer available at the time of its investigation. The VPN appliance in use at the time only had enough storage to keep a limited number of logs. As new events occurred during the attack, these filled up the available storage, writing over the log entries from earlier in the attack. While normally log entries would be on the device for about a week, the high number of events and high traffic on the device during the attack resulted in the relevant VPN logs being overwritten by recent entries much sooner than was normal, before the clinic had an opportunity to review them.

[5] The threat actor used tools to gain access to administrative privileges and then

used these credentials to move to other servers. The threat actor was able to encrypt and exfiltrate records from both an EMR server and a file share server.

[6] The clinic decided to pay the ransom demanded by the threat actor in exchange for the decryption key and a promise to delete the exfiltrated files. The clinic states that it made this decision based on the sensitivity of the data, the proof that data had been exfiltrated, and the lack of reliable backups. After receiving the decryption key, the clinic was able to decrypt all information on the affected servers and recover all affected files.

Preliminary matters

[7] I find that, the clinic is a "health information custodian" under paragraph 4(i) of section 3(1) of the *Act* and that the information accessed by the threat actor included records containing "personal health information" within the meaning of section 4(1) of the *Act*. There is no dispute regarding these findings.

ISSUES:

[8] This decision addresses the following issues:

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

DISCUSSION:

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

[9] 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.

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

[11] Health information custodians, when confronted with a breach of personal health information, should take appropriate steps in response. These include identifying of the

scope of the breach, containing the personal health information involved, notifying those affected, and investigating and remediating 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).[1]

Scope

[12] The clinic reported that two servers were affected: an electronic medical record (EMR) server and a file share server. These contained up to 550,000 patient records and 1.6M case files, respectively. Some of these records were duplicates, so the clinic was not able to determine the number of affected individuals.

[13] The threat actor provided the clinic with "Proof of Data Exfiltration" in the form of a file tree. A file tree (or directory tree) is a representation of the folders and files within a computer, server, or directory, showing how these relate to each other. In this context the file tree showed that, at the very least, the threat actor had sufficient access to view the structure and contents of the files and folders set out in that directory tree.

[14] Based on this file tree, the clinic reported that the following types of information were affected:

Patient Information:

- contact information (home and mobile phone numbers, addresses)
- health card information (last three digits of OHIP number and version code)
- date of birth
- sex

Employee Information:

- name
- contact information (phone numbers and email addresses)
- tax information

[15] In addition to the patient and employee information, the impacted information also included information from the clinic's billing system, such as billing codes, the hospital for which the billing was being done, the amount associated with each billing code, and business addresses of referring physicians. The clinic confirmed that the billing information accessed was related to its health care partners, rather than patients, and

[1] https://www.ipc.on.ca/resource/responding-to-a-health-privacy-breach-guidelines-for-the-health-sector/

did not contain any related personal information or personal health information.

[16] The clinic initially stated that the records described in the file tree did not contain any patient medical records. However, the clinic later reported to the IPC that the 550,000 patient records on the EMR server were medical records, containing more extensive personal health information than that listed by the threat actor. Similarly, the 1.6M case files stored on the file sharing server contained both patients' personal health information and employees' personal information, including more data elements than those noted. This meant that there was an apparent discrepancy between the information the clinic determined the threat actor accessed (the records reflected in the file tree) and the records that the threat actor could have possibly accessed (the entirety of the records on the servers).

[17] The IPC inquired about this discrepancy in the scope of the breach, and the clinic's reasons for determining that the threat actor had only accessed the types of information described in the file tree. The clinic stated that its forensic experts determined that there was no forensic proof that the threat actor accessed more information than the types of information listed in the file tree.

[18] In making this determination, these experts reviewed the clinic's systems to identify any indications of malicious activity, looking for signs of data access and possible data exfiltration. This included a review of the available server activity logs for forensic artifacts. They also looked for LNK[2] files, which show if a target file was moved or deleted. They also examined shellbags, computer registry entries that help track views, sizes, and positions of folder windows. Taken together, these can provide evidence of what areas threat actors accessed during an intrusion.

[19] After completing the review, the experts came up with a list of files that showed signs of having been potentially accessed by the threat actor. They compared these findings with the files that the threat actor stated they had accessed, as set out in the file tree. These overlapped, as the records that showed signs of potential access were found only among the records set out in the file tree. The clinic noted that the experts' search encompassed records beyond just those included in the file tree. Had the threat actor accessed other records, there would have been signs of accesses in files outside the records included in the file tree. That was not the case here, so the clinic concluded that the accessed files were only those contained in the file tree provided by the threat actor.

Containment

[20] To contain the incident, the clinic blocked internet access to the affected EMR and file sharing servers and severed VPN connections. The clinic shut down its critical servers from December 16, 2022 to December 28, 2022, resulting in its closure from December 16, 2022 to January 3, 2023. The clinic also performed malware scanning on all endpoints

[2] LNK files are Windows shortcut files that are either automatically created whenever a user opens their files, or manually created by users.

on December 16, 2022.

[21] The clinic's forensic team negotiated with the threat actor, and the clinic ultimately paid the ransom on December 21, 2022. They then received a decryption key and the threat actor confirmed that it had deleted all exfiltrated data and would not use it further.

[22] The clinic successfully decrypted all information on the affected servers and recovered all the affected files. They also scanned their data stored onsite at different healthcare partners' locations and found the data to be clean. The clinic also found that the offsite data was not connected to the clinic's site during the incident period.

[23] Following the incident, the forensics team conducted a dark web monitoring exercise for a 3-month period from March 1 to May 31, 2023. This included searches of the threat actor's known affiliates' sites and dedicated leak sites of other ransomware groups, as well as forums and chat rooms. As of March 31, 2023, the team was not able to find any evidence that the data had been marketed, resold, or repackaged on the dark web. During the IPC's investigation, the clinic committed to continue dark web monitoring for a further two years.

Notification

[24] The clinic chose to provide notice of the breach via a "pop up" on its website (the Notice), beginning on December 23, 2022. A physical copy of the same Notice was also posted at the entrance and information desks at the clinic's physical locations. The clinic provided indirect notice because of the number of individuals potentially affected – up to 550,000 patient records, though some were duplicates – and the clinic was not able to narrow down who among them had their personal information potentially accessed or exfiltrated. Based on this, the clinic determined that direct notice to the individuals affected was not feasible.

[25] The Notice describes the incident, the measures taken in response, including notifying the IPC, and sets out the steps that individuals may want to take to protect themselves. It also provides a contact email address and informs individuals that they may file a complaint with the IPC.

[26] The Notice describes the data affected as follows:

> Based on our initial investigative efforts, the data accessed during this attack includes past and present patient data contained on our servers. The data potentially accessed may include:
>
> - Patient contact information (name, sex, date of birth, home address, phone number)
>
> - Patient health card information (Last three digits of OHIP number and version)

- Billing and procedural code information related to procedures our practitioners have done at affiliated hospitals

- Business addresses of referring physicians

[27] The Notice later states "[we] are unable to provide individuals with specific information about the data accessed."

[28] During the investigation, the clinic posted updated notices within its clinics in August 2023, which include the following statement:

> [W]hile patient medical records were stored within the affected servers, the results of the investigation did not show any signs of unauthorized access or exfiltration of these records"

[29] The clinic later added this statement to the Notice posted online.

[30] The servers accessible to the threat actor contained information beyond the information listed in the Notice. However, as discussed above, the clinic's forensic experts reviewed the indicators of intrusion on the server logs and found that what had likely been accessed correlated with the file tree provided by the threat actor. The clinic states that "the notices depict as accurately as possible the type of information that may have been potentially compromised." I am satisfied that the information contained in the Notice conveys the clinic's experts' assessment of the information likely accessed by the threat actor, while acknowledging that other records were present on the servers.

[31] In addition to the notification to the public, the clinic sent notification letters to over 14,000 referring physicians and to the clinic's employees. The clinic also sent letters to its healthcare partners notifying them of the breach and noting that the information accessed by the threat actor may have included the partner organization's billing codes.

[32] 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 publicly posted notice. However, in this case, there may have been over half a million affected individuals, based on the number of patient records, and the clinic was not able to determine whose information the threat actor accessed. Given this, it was reasonable that the clinic deemed direct notification to patients not feasible and proceeded by way of public notice instead.

[33] I do note that direct notification remains the standard notice that health information custodians should provide whenever possible. The clinic's letters to referring physicians, employees, and healthcare partners demonstrate its willingness to provide direct notification when feasible to do so.

[34] I am satisfied that the clinic has provided the notification required by section 12(2) of the *Act*.

Investigation and Remediation

[35] The IPC's *PHIPA Breach Guidelines* state that the investigation and remediation of a breach should include a review of the circumstances surrounding the breach and the adequacy of existing policies and procedures protecting personal health information. The latter also includes training on these policies and procedures.

Investigation of the Attack

[36] As noted above, the clinic's experts determined that the likely initial attack vector was the threat actor's use of VPN credentials to log into a dormant account. This could not be confirmed with logs, due to high degree of activity quickly pushing the relevant records out before they could be reviewed. However, the fact that an account belonging to a deceased former internal application developer, which had not been used for years, was active shortly before the attack does point to the likely cause.

[37] The clinic was able to determine the tools the threat actor deployed to gain access to two critical servers, which the threat actor then encrypted and exfiltrated data from. The clinic decrypted and regained full access to these servers, after it paid the ransom and the threat actor provided the decryption key.

[38] The clinic paid the ransom in part because it was not able to restore the relevant systems using backups. As part of the attack, the threat actor deleted both of the clinic's backups. The threat actor was able to access the backups via a Secure Shell Protocol (SSH) Command and then deleted the backups from those devices. The clinic states that the threat actor was not able to access the data on the backups, only delete it.

[39] At the time of the incident, the clinic stored its backups on two network-attached storage devices. Each of the devices stored a self-encrypted copy of the backup and were always kept online. The clinic scheduled daily backups for any modified data, with full backups completed every two months. The backups were tested once every two weeks, while the backup program was tested weekly to ensure that the backups were successful.

Remediation Efforts

[40] The clinic set out security measures it had in place prior to the incident and continues to have in place. The domain authentication infrastructure provides some protection from email-based attacks and domain users require unique passwords to access the systems. The systems are equipped with anti-virus protection. The network sits behind a firewall with site-to-site VPN connection between the main site and the satellite locations. External users must use SSLVPN to connect to the network using their own unique credentials. Their access is limited to only what they need.

[41] In the immediate aftermath of the incident, the clinic reset the passwords for all privileged, VPN, and user accounts, applying a strong password policy to the reset. It compiled a list of known malicious IP addresses using public threat intelligence reports

and blocked these addresses. It also blocked IP addresses from countries where threat actors are most active.

[42] The clinic also put in place longer term remediation measures following the incident. The Password Policy was revised in August 2023, and now requires users to create passwords with a minimum length and mix of characters. Under that policy, the passwords must be changed at least every 90 days, and cannot be re-used within a twelve-month period. Users are also advised to use a secure password manager and multi-factor authentication whenever possible. This policy was communicated to staff via an internal memo, which noted that non-compliance may result in account suspension or other measures.

[43] The clinic revised its System of Least Privilege Access Policy, which now states that users are provided the "bare minimum access to all servers, limiting access to what is an absolute necessity to complete the intended tasks." Since this implementation, only two administrative staff now have domain access privileges for their daily work.

[44] The clinic now has a Policy for Identification and Removal of Dormant User Accounts. Dormant accounts are defined as accounts that have had no logins for 90 days. The IT Security team is required to conduct quarterly reviews to identify dormant or unnecessary accounts, and account owners are notified at least fifteen days in advance of the accounts being removed due to dormancy. If a staff member ends work at the clinic, for reasons including death, termination, etc., the human resources administrator is required to notify the IT administrators to remove the user's access immediately.

[45] In addition, the clinic took other steps to minimize the chance of future intrusions. Emails now have spam filtering and quarantine measures in place. Only accounts that have enabled two-factor authentication have external access. The VPN was replaced with a newer system with up-to-date firmware. Logs are now enabled for firewall and VPN activity. These logs are uploaded daily and kept indefinitely. This should prevent the investigation hurdle the clinic faced in this case, of not being able to determine with certainty how the intrusion occurred.

[46] The clinic now has a Patch Management Policy in place, which requires their IT Security team to monitor vendor sources and security advisories to identify relevant security and operating system patches. When identified, critical patches will be immediately deployed, while routine patches will be put in place during regular maintenance times.

[47] Extended detection and response (XDR) capability is now in place and is configured to monitor and send alerts via email to managed provider staff if malicious or suspicious activity occurs. The personal information and personal health information that the clinic holds are now encrypted, both at rest and in transit; backup systems are also encrypted.

[48] The clinic has also changed its approach to backups in response to the incident.

The devices are now set up to ensure that one is always offline, so that at least one viable copy of the backup would remain unaffected and uncompromised if another cyberattack occurs. There is now more limited access to the devices from within the network and the clinic has enabled two-factor authentication for the devices. Finally, the SSH Command access on those devices has since been disabled to prevent any unauthorized access or destruction of backups.

Security Guidance Documents

[49] As part of the investigation, the clinic also provided the IPC with a copy of its Privacy Policy and its Workplace Technology Security Policy (the Security Policy). The Security Policy was revised following the attack.

[50] The Privacy Policy includes a section on information security that outlines the precautions taken to ensure that patient information is kept safe, including both physical and technological measures. This policy notes that it uses passwords, encryption, and firewalls among other technological safeguards. It also states that the clinic uses security clearances and limits access to information on a "need to know" basis, as well as noting that its security measures are reviewed and updated on a regular basis.

[51] The Privacy Policy largely covers the same information as the other policies discussed in more detail above. However, it does not provide much in the way of specific and actionable guidance for those within the clinic who are responsible for the security of patients' personal health information. Rather, the clinic's Privacy Policy is addressed to patients and is written in plain language. As such, it provides a useful description of patients' access and privacy rights, and the general security measures that the clinic undertakes to protect them.

[52] The Security Policy, on the other hand, is directed at employees, and lists the following as one of its objectives:

> The primary objective of this Policy is to ensure cybersecurity of company communications, to foster secure transactions between all stakeholders, maximize organizational awareness on how to prevent and mitigate risks, establish company protocols to manage risk, and to ensure compliance with legal and regulatory requirements.

[53] It goes on to note that all users must participate to "prevent and mitigate risk of cyberattacks, vulnerabilities, information leakage, and system and network compromises."

[54] The Security Policy's "User Access, Passwords and Control" section sets out safety measures that users must comply with, many of which are set out in more detail in the individual policies as set out above. In addition, the clinic is to use segregated networks to maintain the safety of the information on its servers and install network firewalls "where applicable." Access and traffic to the network will be monitored, as well as

company dataflow. The clinic will employ malware prevention software detection tools, as well as implement DNS filtering to keep users from accessing malicious webpages and web applications. New information systems are required to include user authentication and privileged access to limited users.

[55] The Security Policy also includes a section setting out the steps to take in the event of a breach of information. It lists the individuals who are to respond, and the immediate steps to take, including: password changes; exploring account recovery options; scanning hardware for threat or suspicious activity detection; removal of sensitive data; and conducting security audits. Designated personnel will conduct an investigation to determine the source of the compromise or interception, devise a solution, and notify employees of the issue.

Training

[56] The clinic states that before the incident, privacy issues were discussed during staff meetings, but it did not provide any formal cybersecurity and privacy training to its staff. Since that time, the clinic retained a third-party vendor to provide specialized cybersecurity training to its staff. The clinic states that this training is intended to ensure that staff are well-prepared, informed, and up to date on cybersecurity matters. The training sessions have included the following topics:

- Introduction to information security
- Password best practices
- Social engineering awareness
- Web phishing prevention
- Incident reporting procedures
- Malware detection and mitigation.

[57] All employees have completed the first series of training sessions, which took place in September and October 2023. Further training sessions are scheduled for June 2024. Cybersecurity training is included in new employees' onboarding process, and all employees perform cybersecurity training every quarter. Employees are also reminded of cybersecurity best practices at clinical meetings.

Analysis

[58] The breach was likely caused by a threat actor gaining access to a long-dormant account with significant administrative privileges. The high level of activity during the attack caused logs to be overwritten before they could be reviewed. This meant that the clinic could not determine exactly how the intrusion occurred, or what tactics were used

to gain access to the account credentials. It was able to find evidence of intrusion on portions of the servers; this correlated to the areas that the hacker claimed to have accessed. The clinic stated that it decided to pay the ransom demand largely because the hacker deleted the backups that it had in place.

[59] The clinic has since put in place policies directed at preventing similar situations occurring in future. Password strength and complexity requirements are in place, and the clinic now monitors for and deletes dormant accounts. Implementation of the least privilege access principle has limited regular administrative privileges to only two accounts. The clinic now has extended detection and response measures in place and segregates its networks. Firewalls are now in place and the clinic has provided guidance to ensure that security patches are kept current. The combination of these measures means that threat actors should face much more difficulty in obtaining credentials, and if they do gain access, should have less opportunity to access information or otherwise hijack the clinic's systems.

[60] Regarding the threat actor's deletion of the backups, the clinic now has a system in place where one backup is always kept offline. If an intrusion does occur, the clinic should be able to use the offline backup to restart its operations. Similarly, the clinic now uploads its VPN and firewall logs daily and stores them indefinitely. Should another attack occur, the clinic should be able to better investigate the matter with these logs in place.

[61] Under section 12(2) of the *Act*, a health information custodian must notify the individual if personal health information in its custody or control is stolen or lost, or used or disclosed without authority. A health information custodian that minimizes the true extent of the information accessed by a hacker is not complying with its obligations under the *Act*.

[62] The initial version of the clinic's Notice set out the types of personal health information accessed by the threat actor as including the individual's contact information, health card information, date of birth, and sex. The clinic stated that this was the information outlined in the file tree provided by the hacker as part of their ransom demand, to demonstrate their access to the clinic's systems. However, the clinic later stated that the threat actor also had access to both servers, which included patient records and medical records.

[63] I asked the clinic to address this discrepancy and explain why their initial notice to patients only mentioned the information outlined in the file tree. In response, the clinic detailed the steps its forensics team took to look for independent evidence showing where the hacker had been on the clinic's servers and confirmed that their findings corresponded with the file tree provided. The clinic nonetheless amended the notice, acknowledging that there were medical records stored on the affected servers but that their investigation did not show signs that these had been accessed by the threat actor. Based on the additional information provided, I am satisfied that the clinic has taken sufficient efforts to determine the scope of the breach and provide the appropriate notice. I am also

satisfied that it responded adequately to the breach that occurred, especially in light of the remediation steps that the clinic has taken to address this matter.

Issue 2: Is a review warranted under Part VI of the *Act*?

[64] 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.

[65] 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 in the circumstances, given the notice provided and the remedial measures that have already been implemented to minimize the risks of such a breach reoccurring in the future.

DECISION:

For the foregoing reasons, no review of this matter will be conducted under Part VI of the *Act*.

Original signed by: June 21, 2024

Jennifer Olijnyk
Investigator