PUBLIC REPORT OF THE COMMITTEE OF INQUIRY INTO THE CYBER ATTACK ON SINGAPORE HEALTH SERVICES PRIVATE LIMITED'S PATIENT DATABASE ON OR AROUND 27 JUNE 2018

10 JANUARY 2019

TABLE OF CONTENTS

Exchange of letters with the Minister

31 Dec 2018

Mr S Iswaran
Minister-in-Charge of Cybersecurity, and
Minister for Communications and Information
140 Hill Street
Old Hill Street Police Station
Singapore 179369

Dear Minister

SUBMISSION OF REPORT OF THE COMMITTEE OF INQUIRY INTO THE CYBER ATTACK ON SINGHEALTH

We were appointed on 24 July 2018 by you to inquire into the events and contributing factors leading to the cyber attack on SingHealth's patient database system on or around 27 June 2018, establish the response thereto and recommend measures to reduce the risk of such attacks. We are honoured to have been appointed and to have served in this Committee of Inquiry.

2 We submit the report as enclosed, which covers the assessment of the evidence, findings, attribution of the attack, and priority & additional recommendations. This report contains sensitive information, and is hence classified 'Top Secret'. The contents of the report are the unanimous view of all members of the Committee.

3 We thank the Secretariat from the Ministry for their unwavering support throughout the Inquiry, and for working closely and assiduously with the Committee in writing the report. We also thank the State Counsel team led by Solicitor-General Kwek Mean Luck, and the investigation team led by Mr Tay Cheong Beng Lawrence, comprising members from the Cyber Security Agency of Singapore (CSA) and the Criminal Investigation Department (CID). Finally, we thank CSA for its advice on technical matters pertaining to the Inquiry.

Yours faithfully,
The Committee of Inquiry

Richard Magnus	Lee Fook Sun	T.K. Udairam	Cham Hui Fong
Chairman	*Member*	*Member*	*Member*

Enc

MINISTER FOR COMMUNICATIONS
AND INFORMATION
SINGAPORE

31 December 2018

Mr Richard Magnus
Chairman, Committee of Inquiry into the Cyber Attack on SingHealth

Dear Richard

REPORT OF THE COMMITTEE OF INQUIRY INTO THE CYBER ATTACK ON SINGHEALTH

Thank you for the report of the Committee of Inquiry (COI) into the cyber attack on SingHealth's patient database system.

The COI report gives a thorough account of the events and contributory factors leading to the cyber attack. The Committee has also examined in great detail the responses to the incident, and submitted a comprehensive set of recommendations to better manage and secure the IT systems of SingHealth, as well as those of other public healthcare clusters and the public sector, against similar attacks.

The COI report is the result of an extensive fact-finding process and a rigorous inquiry over the past five months. I note that the COI held 22 days of hearings with 37 witnesses giving evidence. Many more weeks were spent deliberating and finalising the report.

On behalf of the Government, I would like to express my deepest gratitude to you and the members of COI, the COI Secretariat, the State Counsel and Investigation teams for your hard work and dedication.

The Government takes cybersecurity with utmost seriousness. We will carefully study the COI's detailed findings and recommendations, and issue a response in January 2019.

We will learn from this incident and take measures to further strengthen our public sector IT systems and uphold the trust of Singaporeans.

Yours Sincerely

S ISWARAN
Minister-in-charge of Cybersecurity

Executive Summary

A. INTRODUCTION

1. Between 23 August 2017 and 20 July 2018, a cyber attack (the "**Cyber Attack**") of unprecedented scale and sophistication was carried out on the patient database of Singapore Health Services Private Limited ("**SingHealth**"). The database was illegally accessed and the personal particulars of almost 1.5 million patients, including their names, NRIC numbers, addresses, genders, races, and dates of birth, were exfiltrated over the period of 27 June 2018 to 4 July 2018. Around 159,000 of these 1.5 million patients also had their outpatient dispensed medication records exfiltrated. The Prime Minister's personal and outpatient medication data was specifically targeted and repeatedly accessed.

2. The crown jewels of the SingHealth network are the patient electronic medical records contained in the SingHealth Sunrise Clinical Manager ("**SCM**") database. The SCM is an electronic medical records software solution, which allows healthcare staff to access real-time patient data. The SCM system can be seen as comprising front-end workstations, Citrix servers, and the SCM database. Users would access the SCM database *via* Citrix servers, which operate as an intermediary between front-end workstations and the SCM database. The Citrix servers played a critical role in the Cyber Attack.

3. At the time of the Cyber Attack, SingHealth was the owner of the SCM system. Integrated Health Information Systems Private Limited ("**IHiS**") was responsible for administering and operating the system, including implementing cybersecurity measures. IHiS was also responsible for security incident response and reporting.

B. THE EVENTS OF THE CYBER ATTACK AND INCIDENT RESPONSE BY IHIS AND SINGHEALTH

4. The Committee's Terms of Reference ("**TORs**") include (i) establishing the events and contributing factors leading to the Cyber Attack and the exfiltration of patient data ("**TOR #1**"), and (ii) establishing how IHiS and SingHealth responded to the Cyber Attack ("**TOR #2**"). The Committee's findings on these TORs are set out in Parts III-VI of the main report.

5. In the present section, the Committee will first provide a summary of the key events of the Cyber Attack and the incident response by IHiS and SingHealth. The Committee will then present five Key Findings in respect of TORs #1 and #2.

I. Summary of events

6. The attacker gained initial access to SingHealth's IT network around 23 August 2017, infecting front-end workstations, most likely through phishing attacks. The attacker then lay dormant for several months, before commencing lateral movement in the network between December 2017 and June 2018, compromising a number of endpoints and servers, including the Citrix servers located in SGH, which were connected to the SCM database. Along the way, the attacker also compromised a large number of user and administrator accounts, including domain administrator accounts.

7. Starting from May 2018, the attacker made use of compromised user workstations in the SingHealth IT network and suspected virtual machines to remotely connect to the SGH Citrix servers, and tried unsuccessfully to access the SCM database from the SGH Citrix servers.

8. IHiS' IT administrators first noticed unauthorised logins to the Citrix servers and failed attempts at accessing the SCM database on 11 June 2018. Similar malicious activities were detected on 12, 13, and 26 June 2018. Unknown to them, the attacker had obtained credentials to the SCM database on 26 June

2018. Starting from 27 June 2018, the attacker began querying the SCM database, stealing and exfiltrating patient records, and doing so undetected by IHiS.

9. On 4 July 2018, an IHiS administrator for the SCM system noticed suspicious queries being made on the SCM database. Working with other IT administrators, ongoing suspicious queries were terminated, and measures were put in place to prevent further queries to the SCM database. These measures proved to be successful, and the attacker could not make any further successful queries to the database after 4 July 2018.

10. Between 11 June and 9 July 2018, the persons who knew of and responded to the incident were limited to IHiS' line-staff and middle management from various IT administration teams, and the security team. On 9 July 2018, IHiS senior management were finally informed of the matter. On 10 July 2018, the matter was escalated to the Cyber Security Agency of Singapore ("**CSA**"), SingHealth's senior management, the Ministry of Health ("**MOH**"), and the Ministry of Health Holdings ("**MOHH**").

11. Starting from the night of 10 July 2018, IHiS and CSA carried out joint investigations and remediation. Several measures aimed at containing the existing threat, eliminating the attacker's footholds, and preventing recurrence of the attack were implemented. In view of further malicious activities on 19 July 2018, internet surfing separation was implemented for SingHealth on 20 July 2018. No further suspicious activity was detected after 20 July 2018.

12. After being notified of the Cyber Attack, SingHealth's senior management, in consultation with MOH, IHiS, CSA, and the Ministry of Communications and Information, began making plans for a public announcement, and for patient outreach and communications.

13. The public announcement was made on 20 July 2018, and patient outreach and communications commenced immediately thereafter. SMS messages were used as the primary mode of communication, in view of the need for quick dissemination of information on a large scale. Other modes of communication

included letters, telephone hotlines, and various online channels. In total, SingHealth intended to contact 2.16 million patients. At the time of the Inquiry, 2.9% of the patients could not be contacted despite SingHealth's efforts.

II. Key findings of the Committee

14. The Committee has made numerous findings in respect of TORs #1 and #2. From these findings, the Committee has identified five Key Findings.

Key Finding #1: IHiS staff did not have adequate levels of cybersecurity awareness, training, and resources to appreciate the security implications of their findings and to respond effectively to the attack

- A number of IHiS' IT administrators are commended by the Committee for their vigilance in noticing suspicious activity, such as unauthorised logins to the Citrix servers, suspicious attempts at logging in to the SCM database, presence of unauthorised software, and suspicious queries being run on the SCM database.
- However, these same IT administrators could not fully appreciate the security implications of their findings, and were unable to co-relate these findings with the tactics, techniques, and procedures ("**TTPs**") of an advanced cyber attacker.
- They were also not familiar with the relevant IT security policy documents and the need to escalate the matter to CSA. There was also no incident reporting framework in place for the IT administrators.
- Members of the Security Management Department, Computer Emergency Response Team, and senior members of IHiS' management were similarly unable to fully appreciate the security implications of the findings.

Key Finding #2: Certain IHiS staff holding key roles in IT security incident response and reporting failed to take appropriate, effective, or timely action, resulting in missed opportunities to prevent the stealing and exfiltrating of data in the attack

- The Security Incident Response Manager ("**SIRM**") and Cluster Information Security Officer ("**Cluster ISO**") for SingHealth, who were responsible for incident response and reporting, held mistaken understandings of what constituted a 'security incident', and when a security incident should be reported.
- The SIRM delayed reporting because he felt that additional pressure would be put on him and his team once the situation became known to management.
- The evidence also suggests that the reluctance to escalate the matter may have come from a belief that it would not reflect well in the eyes of the organisation if the matter turned out to be a false alarm.
- The Cluster ISO did not understand the significance of the information provided to him, and did not take any steps to better understand the information. Instead, he effectively abdicated to the SIRM the responsibility of deciding whether to escalate the incident.

Key Finding #3: There were a number of vulnerabilities, weaknesses, and misconfigurations in the SingHealth network and SCM system that contributed to the attacker's success in obtaining and exfiltrating the data, many of which could have been remedied before the attack

- A significant vulnerability was the network connectivity (referred to in these proceedings as an "open network connection") between the SGH Citrix servers and the SCM database, which the attacker exploited to make queries to the database. The network connectivity was maintained for the use of administrative tools and custom applications, but there was no necessity to do so.
- The SGH Citrix servers were not adequately secured against unauthorised access. Notably, the process requiring 2-factor authentication ("**2FA**") for administrator access was not enforced as the exclusive means of logging in as an administrator. This allowed the attacker to access the server through other routes that did not require 2FA.
- There was a coding vulnerability in the SCM application which was likely exploited by the attacker to obtain credentials for accessing the SCM database.
- There were a number of other vulnerabilities in the network which were identified in a penetration test in early 2017, and which may have been exploited by the attacker. These included weak administrator account passwords and the need to improve network segregation for administrative access to critical servers such as the domain controller and the Citrix servers. Unfortunately, the remediation process undertaken by IHiS was mismanaged and inadequate, and a number of vulnerabilities remained at the time of the Cyber Attack.

Key Finding #4: The attacker was a skilled and sophisticated actor bearing the characteristics of an Advanced Persistent Threat group

- The attacker had a clear goal in mind, namely the personal and outpatient medication data of the Prime Minister in the main, and also that of other patients.
- The attacker employed advanced TTPs, as seen from the suite of advanced, customised, and stealthy malware used, generally stealthy movements, and its ability to find and exploit various vulnerabilities in SingHealth's IT network and the SCM application.
- The attacker was persistent, having established multiple footholds and backdoors, carried out its attack over a period of over 10 months, and made multiple attempts at accessing the SCM database using various methods.
- The attacker was a well-resourced group, having an extensive command and control network, the capability to develop numerous customised tools, and a wide range of technical expertise.

Key Finding #5: While our cyber defences will never be impregnable, and it may be difficult to prevent an Advanced Persistent Threat from breaching the perimeter of the network, the success of the attacker in obtaining and exfiltrating the data was not inevitable

- A number of vulnerabilities, weaknesses, and misconfigurations could have been remedied before the attack. Doing so would have made it more difficult for the attacker to achieve its objectives.
- The attacker was stealthy but not silent, and signs of the attack were observed by IHiS' staff. Had IHiS' staff been able to recognise that an attack was ongoing and take appropriate action, the attacker could have been stopped before it achieved its objectives.

C. RECOMMENDATIONS BY THE COMMITTEE

15. The Committee's TORs also include recommending measures to (i) enhance the incident response plans for similar incidents ("**TOR #3**"); (ii) better protect SingHealth's patient database system against similar cyber attacks ("**TOR #4**"); and (iii) reduce the risk of such cyber attacks on public sector IT systems which contain large databases of personal data, including in the other public healthcare clusters ("**TOR #5**"). The Committee's recommendations on these TORs are set out in Part VII of the main report.

16. The Committee makes sixteen recommendations, comprising seven Priority Recommendations and nine Additional Recommendations, all of which have been explored and examined in great detail.

17. The seven Priority Recommendations include strategic and operational measures to uplift the cybersecurity posture of SingHealth and IHiS, and steps must be taken to implement these Priority Recommendations immediately. The nine Additional Recommendations relate to other specific concerns raised in the course of this Inquiry, including technical, organisational, training, and process-related issues. The measures, which are similarly aimed at uplifting the cybersecurity posture of SingHealth and IHiS, must be implemented or seriously considered.

18. All sixteen recommendations are made in respect of TORs #3 and #4, and apply equally to TOR #5. They range from basic cyber hygiene measures to more advanced measures which may be more relevant after a certain level of cybersecurity maturity has been attained by the organisation.

19. While some measures may seem axiomatic, the Cyber Attack has shown that these were not implemented effectively by IHiS at the time of the attack. For IHiS, SingHealth, and other organisations responsible for large databases of personal data, getting the fundamentals right is a necessary and vital step in building cybersecurity competencies and the ability to counter the real, present, and constantly evolving cybersecurity threats.

I. Priority Recommendations

Recommendation #1: An enhanced security structure and readiness must be adopted by IHiS and Public Health Institutions

- Cybersecurity must be viewed as a risk management issue, and not merely a technical issue. Decisions should be deliberated at the appropriate management level, to balance the trade-offs between security, operational requirements, and cost.
- IHiS must adopt a "defence-in-depth" approach.
- Gaps between policy and practice must be addressed.

Recommendation #2: The cyber stack must be reviewed to assess if it is adequate to defend and respond to advanced threats

- Identify gaps in the cyber stack by mapping layers of the IT stack against existing security technologies.
- Gaps in response technologies must be filled by acquiring endpoint and network forensics capabilities.
- The effectiveness of current endpoint security measures must be reviewed to fill the gaps exploited by the attacker.
- Network security must be enhanced to disrupt the 'Command and Control' and 'Actions on Objective' phases of the Cyber Kill Chain.
- Application security for email must be heightened.

Recommendation #3: Staff awareness on cybersecurity must be improved, to enhance capacity to prevent, detect, and respond to security incidents

- The level of cyber hygiene among users must continue to be improved.
- A Security Awareness Programme should be implemented to reduce organisational risk.
- IT staff must be equipped with sufficient knowledge to recognise the signs of a security incident in a real-world context**.**

Recommendation #4: Enhanced security checks must be performed, especially on CII systems

- Vulnerability assessments must be conducted regularly.
- Safety reviews, evaluation, and certification of vendor products must be carried out where feasible.
- Penetration testing must be conducted regularly.
- Red teaming should be carried out periodically.
- Threat hunting must be considered.

Recommendation #5: Privileged administrator accounts must be subject to tighter control and greater monitoring

- An inventory of administrative accounts should be created to facilitate rationalisation of such accounts.
- All administrators must use two-factor authentication when performing administrative tasks.
- Use of passphrases instead of passwords should be considered to reduce the risk of accounts being compromised.
- Password policies must be implemented and enforced across both domain and local accounts.
- Server local administrator accounts must be centrally managed across the IT network.
- Service accounts with high privileges must be managed and controlled.

Recommendation #6: Incident response processes must be improved for more effective response to cyber attacks

- To ensure that response plans are effective, they must be tested with regular frequency.
- Pre-defined modes of communication must be used during incident response.
- The correct balance must be struck between containment, remediation, and eradication, and the need to monitor an attacker and preserve critical evidence.
- Information and data necessary to investigate an incident must be readily available.
- An Advanced Security Operation Centre or Cyber Defence Centre should be established to improve the ability to detect and respond to intrusions.

Recommendation #7: Partnerships between industry and government to achieve a higher level of collective security

- Threat intelligence sharing should be enhanced.
- Partnerships with Internet Service Providers should be strengthened.
- Defence beyond borders – cross-border and cross-sector partnerships should be strengthened.
- Using a network to defend a network – applying behavioural analytics for collective defence.

II. Additional recommendations

Recommendation #8: IT security risk assessments and audit processes must be treated seriously and carried out regularly

- IT security risk assessments and audits are important for ascertaining gaps in an organisation's policies, processes, and procedures.
- IT security risk assessments must be conducted on CII and mission-critical systems annually and upon specified events.
- Audit action items must be remediated.

Recommendation #9: Enhanced safeguards must be put in place to protect electronic medical records

- A clear policy on measures to secure the confidentiality, integrity, and accountability of electronic medical records must be formulated.
- Databases containing patient data must be monitored in real-time for suspicious activity.
- End-user access to the electronic health records should be made more secure.
- Measures should be considered to secure data-at-rest.
- Controls must be put in place to better protect against the risk of data exfiltration.
- Access to sensitive data must be restricted at both the front-end and at the database-level.

Recommendation #10: Domain controllers must be better secured against attack

- The operating system for domain controllers must be more regularly updated to harden these servers against the risk of cyber attack.
- The attack surface for domain controllers should be reduced by limiting login access.
- Administrative access to domain controllers must require two-factor authentication.

Recommendation #11: A robust patch management process must be implemented to address security vulnerabilities

- A clear policy on patch management must be formulated and implemented.
- The patch management process must provide for oversight with the reporting of appropriate metrics.

Recommendation #12: A software upgrade policy with focus on security must be implemented to increase cyber resilience

- A detailed policy on software upgrading must be formulated and implemented.
- An appropriate governance structure must be put in place to ensure that the software upgrade policy is adhered to.

Recommendation #13: An internet access strategy that minimises exposure to external threats should be implemented

- The internet access strategy should be considered afresh, in the light of the Cyber Attack.
- In formulating its strategy, the healthcare sector should take into account the benefits and drawbacks of internet surfing separation and internet isolation technology, and put in place mitigating controls to address the residual risks.

Recommendation #14: Incident response plans must more clearly state when and how a security incident is to be reported

- An incident response plan for IHiS staff must be formulated for security incidents relating to Cluster systems and assets.
- The incident response plan must clearly state that an attempt to compromise a system is a reportable security incident.
- The incident response plan must include wide-ranging examples of security incidents, and the corresponding indicators of attack.

Recommendation #15: Competence of computer security incident response personnel must be significantly improved

- The Computer Emergency Response Team must be well trained to more effectively respond to security incidents.
- The Computer Emergency Response Team must be better equipped with the necessary hardware and software.
- A competent and qualified Security Incident Response Manager who understands and can execute the required roles and responsibilities must be appointed.

Recommendation #16: A post-breach independent forensic review of the network, all endpoints, and the SCM system should be considered

- IHiS should consider working with experts to ensure that no traces of the attacker are left behind.

III. Implementation of recommendations

20. IHiS and SingHealth should give priority to implementing the recommendations. Adequate resources and attention must be devoted to their implementation, and there must be appropriate oversight and verification of their implementation. Most importantly, implementation of the recommendations requires effective and agile leadership from senior management, and necessary adjustments to organisational culture, mindset, and structure.

21. These imperatives apply equally to all organisations responsible for large databases of personal data. We must recognise that cybersecurity threats are here to stay, and will increase in sophistication, intensity, and scale. Collectively, these organisations must do their part in protecting Singapore's cyberspace, and must be resolute in implementing these recommendations.

Part I – Introduction

1 APPOINTMENT AND TERMS OF REFERENCE OF THE COMMITTEE OF INQUIRY

1.1 Introduction

1. From around 23 August 2017 to 20 July 2018, a cyber attack (the "**Cyber Attack**") of unprecedented scale and sophistication was carried out on the patient database of Singapore Health Services Private Limited ("**SingHealth**"). The database was illegally accessed and the personal particulars of almost 1.5 million patients, including their names, NRIC numbers, addresses, genders, races, and dates of birth, were exfiltrated over the period of 27 June 2018 to 4 July 2018. 159,000 of these 1.5 million patients also had their outpatient dispensed medication records exfiltrated. The Prime Minister's personal and outpatient medication data was specifically targeted and repeatedly accessed.

1.2 Appointment and members of the Committee

2. Given the extraordinary nature of the incident, the Minister-in-Charge of Cybersecurity, Mr S Iswaran, appointed a committee of inquiry (the "**Committee**") under Section 9(1)(b) of the Inquiries Act (Cap. 139A, 2008 Revised Edition) (the "**Inquiries Act**") on 24 July 2018 to inquire into the events and contributing factors leading to the cyber attack on SingHealth's patient database system.

3. The Committee comprises four members, appointed by Minister Iswaran for their expertise in legal, technical, healthcare, and management fields. The Chairman of the Committee, Mr Richard Magnus, is a retired Senior (subsequently termed Chief) District Judge and is currently a member of the Public Service Commission. The other members are Mr Lee Fook Sun, Chairman of Ensign InfoSecurity Private Limited; Mr T K Udairam, Group Chief Operating

Officer of Sheares Healthcare Management Private Limited; and Ms Cham Hui Fong, Assistant Secretary-General of the National Trades Union Congress.

1.3 Terms of reference

4. The Committee's terms of reference ("**TORs**") are to:

    1. Establish the events and contributing factors leading to the cybersecurity attack on SingHealth's patient database system on or around 27 June 2018, and the subsequent exfiltration of patient data therefrom;

    2. Establish how the Integrated Health Information Systems Private Limited ("**IHiS**") and SingHealth responded to the cybersecurity attack;

    3. Recommend measures to enhance the incident response plans for similar incidents;

    4. Recommend measures to better protect SingHealth's patient database system against similar cybersecurity attacks;

    5. In light of the cybersecurity attack and the findings above, recommend measures to reduce the risk of such cybersecurity attacks on public sector IT systems which contain large databases of personal data, including in the other public healthcare clusters;

    6. Conduct itself in accordance with the provisions of the Inquiries Act, with the discretion to determine which, if any, part(s) of the inquiry shall be held in public, and consider the evidence put before the Committee as led by the Attorney-General or his designates; and

7. Make and submit a report of its proceedings, findings and recommendations to the Minister-in-Charge of Cybersecurity by 31 Dec 2018.

2 ASSISTANCE TO THE COMMITTEE

2.1 Appointment of the Attorney-General to lead evidence

5. On 24 July 2018, pursuant to section 9(2)(d) of the Inquiries Act, Minister Iswaran appointed the Attorney-General or his designates to lead evidence in the inquiry. The Attorney-General subsequently directed, in accordance with paragraph 11(1) of the Schedule to the Inquiries Act, that the Solicitor-General Mr Kwek Mean Luck lead evidence in the inquiry. The Solicitor-General was assisted by Senior State Counsel Mr G Kannan, Senior State Counsel Ms Kristy Tan, Deputy Senior State Counsel Ms Sarah Shi, Deputy Senior State Counsel Mr Sivakumar Ramasamy, State Counsel Ms Sheryl Janet George, and State Counsel Mr Alexander Woon.

2.2 Appointment of investigators

6. On 31 July 2018, the Committee requested, pursuant to paragraph 9(1) of the Schedule to the Inquiries Act, that the Public Prosecutor cause any matter relevant to the inquiry to be investigated.

7. On 2 August 2018, the Public Prosecutor appointed, under paragraph 9(2) of the Schedule to the Inquiries Act, Mr Tay Cheong Beng Lawrence, Director of the Regulations Division of the Cyber Security Agency of Singapore ("**CSA**"), and officers from CSA and the Criminal Investigation Department ("**CID**") of the Singapore Police Force ("**SPF**") supporting him, to investigate any matter relevant to the inquiry. The investigation team comprised a total of six officers from CSA and CID, with their respective skills and expertise.

2.3 Legal representatives for MOH, MOHH, SingHealth, and IHiS

8. The Ministry of Health ("**MOH**") was represented by State Counsel from the MOH Legal Office, Director Ms Chua Ying-Hong, assisted by Senior Assistant Director Mr Terence Ang. MOH Holdings Private Limited ("**MOHH**") was represented by Senior Counsel Andre Maniam from WongPartnership LLP, assisted by Ms Lim Wei Lee and Mr Russell Pereira. SingHealth was represented by Senior Counsel Dr Stanley Lai from Allen & Gledhill LLP, assisted by Ms Leong Yi-Ming, Mr Afzal Ali, and Mr Joshua Hiew. IHiS was represented by Senior Counsel Mr Philip Jeyaretnam of Dentons Rodyk & Davidson LLP, assisted by Mr Gilbert Leong, Mr Amogh Chakravarti, Mr Desmond Chew, Mr Francis Wu, Mr Joshua Woo, and Ms Joy Yee.

2.4 The Secretariat

9. Minister Iswaran also appointed, under section 12 of the Inquiries Act, Mr Thng E-Shen and Ms Melanie Huang, respectively Director and Deputy Director of the Cyber Security and Resilience Division of the Ministry of Communications and Information, as Secretaries of the Committee. They were assisted by Ms Daphne Chang, Mr Winston Chai, Mr Goh Chian Hao, Ms Alice Yeo, Mr Lim Zhen Xiong, and Mr Ng Song Yeong.

3 ACTIONS TAKEN BY THE COMMITTEE BEFORE THE HEARINGS

10. The Committee was convened on 24 July 2018, and the Members held their first administrative meeting on 25 July 2018. The Chairman of the Committee also met the Solicitor-General on 25 July 2018, and they discussed the appointment of investigators to assist the Committee with investigating matters relevant to the inquiry.

11. On 15 August 2018, the Committee made a site visit to the SingHealth Academia[1] building (the "**Academia**"), located at 20 College Road, Singapore. There, the Committee was given an overview of (i) the SingHealth Electronic Medical Records System (the "**EMR System**"), and (ii) SingHealth's IT network and security measures as they existed before the Cyber Attack.

12. The Committee also held three Pre-Inquiry Conferences ("**PIC**") on 7, 15, and 28 August 2018. The purpose of the PICs was for the Solicitor-General to update the Committee on the progress of investigations, and for interested parties to take directions from the Committee on administrative and procedural matters before the commencement of the hearings.

4 CONDUCT OF THE INQUIRY

13. Hearings commenced on 28 August 2018, and were held over a total of 22 days spread over four tranches:

(a) 28 August 2018;

(b) Between 21 September and 5 October 2018;[2]

(c) Between 25 October and 14 November 2018[3]; and

(d) 30 November 2018.

---

[1] The Academia is part of the SingHealth Academy. It is an education and research centre which has been designed to facilitate interconnectivity among clinical scientist, researchers, pathologists, educators and medical students.

[2] In this period, hearings were conducted on 21 and 24-28 September 2018, and on 1, 2, and 5 October 2018.

[3] In this period, hearings were conducted on 25, 26, 29, and 31 October 2018, and on 1, 2, 5, 9, and 12-14 November 2018.

14. Another three PICs were also held in the course of the hearings, on 21 September 2018, 25 October 2018, and 9 November 2018.

15. Hearings were generally open to members of the public and the media. A gag order[4] was granted in respect of the following categories of evidence, which were heard in private, and with interested parties' counsel present:

(a) Details of SingHealth's network architecture;

(b) Details of the means by which SingHealth's systems were accessed or compromised;

(c) Details of the technical vulnerabilities that were exploited by the attacker;

(d) Information about CSA's forensic processes and capabilities; and

(e) Any information that leads to, or is likely to lead to, the disclosure of the foregoing.

16. In addition to the above, evidence concerning matters of national security were presented by the witnesses with only the Committee and the Solicitor-General's team present.

17. In total, the Committee heard testimony from 37 witnesses, comprising 34 witnesses of fact and three expert witnesses. The Committee also received a written report from one other expert. The testimony of the witnesses was adduced by way of Conditioned Statements ("**CS**") or reports, supplemented by oral evidence.

[4] The order states that "*no person shall publish, broadcast, or disseminate any of the following categories of evidence, or do any other act which is likely to lead to such publication, broadcast, or dissemination of such evidence*".

18. Documentary evidence was adduced by the Solicitor-General and the interested parties. The documentary evidence was compiled into a bundle of exhibits, comprising documents ("**D**"), emails ("**E**"), logs ("**L**") and miscellaneous items ("**M**"). In total, the bundle of exhibits comprised 166 exhibits marked "**D**", 95 exhibits marked "**E**", eight exhibits marked "**L**", and 21 exhibits marked "**M**". The exhibits adduced by the interested parties were marked as "**S**" for SingHealth (totalling four exhibits), "**I**" for IHiS (totalling seven exhibits), and "**H**" for MOH (totalling two exhibits). No exhibits were tendered by MOHH.

19. The general public was also invited to submit written representations on any matter falling within TORs #3, #4, and #5. The Committee received a total of 26 written representations from various individuals and organisations. The Committee studied these representations with care and found them to be useful. The Committee expresses its thanks to all members of the public for their participation and assistance.

20. On 31 December 2018, the Committee submitted its unanimous report to Minister Iswaran. As the report contains sensitive information, it is classified 'Top Secret'. The Minister has directed that a version of the report be made available to the public. The present public report contains all of the material findings and recommendations by the Committee, save for redactions made in accordance with the parameters of the gag order (as stated in paragraph 15 above), and editorial amendments to maintain the clarity of the document.

Part II – Background information relevant to the Inquiry

TABLE OF CONTENTS – PART II

5 INTRODUCTION TO THIS PART

21. The sections in this Part contain background information on various organisational and technical matters – including structures, policies, roles and responsibilities, practices, and systems – that are relevant to the Cyber Attack and the Committee's findings and recommendations. Unless otherwise mentioned, the information in this Part was correct and in effect at the time of the Cyber Attack.

6 ROLES OF MOH, MOHH, SINGHEALTH AND IHIS IN IT ADMINISTRATION FOR THE PUBLIC HEALTHCARE SECTOR

22. This section elaborates on the relationship between the parties who own, manage, and/or have oversight over the SingHealth IT network and the Sunrise Clinical Manager ("**SCM**") System, which were key systems compromised in the Cyber Attack.

6.1 The Ministry of Health (MOH) – Healthcare sector regulator

23. MOH oversees the public healthcare system in Singapore, where IT plays an important role for patient care and services, and is also key to enabling longer-term healthcare sustainability. To support this work, MOH has a Chief Information Officer ("**MOH CIO**") and a Chief Information Security Officer ("**MOH CISO**"):

(a) The MOH CIO is responsible for national healthcare plans such as the National Electronic Health Records ("**NEHR**") as well as systems in MOH Headquarters. The MOH CIO is Bruce Liang ("**Bruce**"), who is concurrently the Chief Executive Officer ("**CEO**") of IHiS.

(b) The MOH CISO was a new position created in January 2015. The MOH CISO's role is to make sure all Information and

Communications Technology ("**ICT**") security measures for MOH and its related agencies are implemented according to Government Instruction Manual 8 ("**IM 8**")**.** The MOH CISO is Chua Kim Chuan ("**Kim Chuan**"), who is concurrently the Director of Cyber Security Governance ("**CSG**") in IHiS. Kim Chuan reports to Bruce in both MOH and IHiS capacities.

6.2 MOH Holdings (MOHH) – Operating arm of MOH, in charge of infrastructure for public healthcare

24. The public healthcare system comprises corporatised institutions owned by their holding company MOHH, which in turn is owned by the Government through Minister for Finance (Incorporated)[5]. The corporatised institutions comprise, *inter alia:*

(a) Since January 2018, three Clusters of public healthcare institutions ("**PHIs**"; collectively, the "**Clusters**"), (i) Singapore Health Services Pte Ltd ("**SingHealth**") (into which the Eastern Health Alliance Cluster had merged); (ii) National University Health System Pte Ltd ("**NUHS**") (into which the Jurong Health Services Cluster had merged); and (iii) National Healthcare Group Ptd Ltd ("**NHG**") (into which the Alexandra Health System Cluster had merged);

(b) Integrated Health Information Systems Pte Ltd ("**IHiS**"); and

(c) Agency for Integrated Care ("**AIC**").

[5] Minister for Finance (Incorporated) is a statutory body constituted by the Minister for Finance (Incorporation) Act (Cap. 183, 2014 Revised Edition).

25. An organisational chart of MOHH and its subsidiaries is set out below.

Figure 1: Organisational chart of MOHH and its subsidiaries

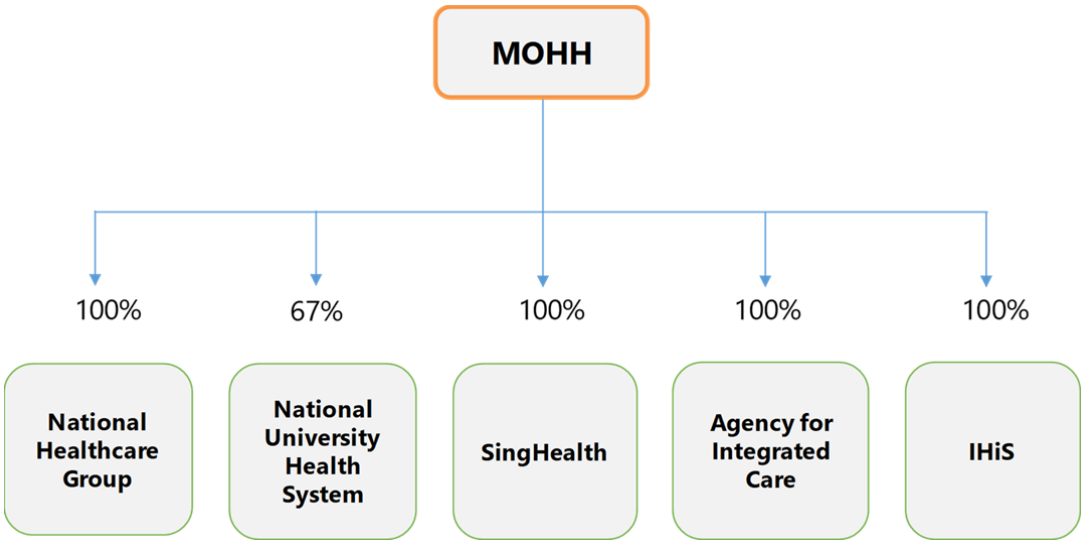

26. From around 2007, MOHH took on more operational functions, such as the central employment and management of junior doctors, and the development of public healthcare infrastructure.

27. Until the merger of MOHH's Information Systems Division ("**MOHH ISD**") into IHiS on 1 November 2016 (see paragraph 36 (pg 14) below) , MOHH ISD was responsible for upstream IT master planning for the public healthcare sector and the development of national healthcare IT.

28. Prior to the above-mentioned merger, MOHH was also responsible for the internal audit function for its subsidiaries, a task performed by the Group Internal Audit ("**GIA**") department. MOHH continues to perform this function today.

29. MOHH has a Board of Directors that is appointed by MOH to guide MOHH's strategic efforts to ensure that its direction is in line with national healthcare policies and goals, and chaired by Permanent Secretary (Health).

6.3 SingHealth – Healthcare Cluster and legal owner of SCM system

30. SingHealth is the largest of the three healthcare Clusters in the public healthcare sector[6]. SingHealth comprises Singapore General Hospital ("**SGH**"), Changi General Hospital ("**CGH**"), Sengkang General Hospital, KK Women's and Children's Hospital, National Cancer Centre ("**NCC**"), National Dental Centre Singapore, National Heart Centre Singapore, National Neuroscience Institute, Singapore National Eye Centre, SingHealth Community Hospitals and SingHealth Polyclinics.

31. Since 2012, Group Chief Executive Officer ("**GCEO**") of SingHealth has been Prof. Lim Swee Lian Ivy ("**Prof. Ivy**").

32. SingHealth is the legal owner of the SCM system. As the SCM system is also a Critical Information Infrastructure ("**CII**"), SingHealth is the CII Owner ("**CIIO**").

33. The SCM system provides real-time patient data to physicians, nurses, and other clinicians to facilitate delivery of medical services.

6.4 IHiS – Healthcare Sector Lead and central IT agency for the public healthcare system

34. IHiS is the central IT agency for the public healthcare system, and serves all the IT needs of the public healthcare Clusters, including SingHealth. IHiS is accountable to MOH for matters such as IT policy, governance, planning, and implementing IT projects; and serving the IT needs of the Clusters.

---

[6] The other two Clusters are the National Healthcare Group ("**NHG**") and the National University Health System ("**NUHS**").

6.4.1 Consolidation of public healthcare system's IT function under IHiS

35. IHiS was formed as a subsidiary of MOHH in July 2008. Prior to the formation of IHiS, the various Clusters managed their IT systems separately – different healthcare Clusters had their own IT departments and the Clusters would be responsible for their own IT security. In 2008, Cluster IT resources and capabilities (including Cluster IT staff) were consolidated under IHiS to better align public healthcare IT, promote interoperability between systems, and give Clusters access to additional IT expertise that they would not have with standalone IT units.

36. There was a further consolidation of IT resources in November 2016, when MOH decided to merge the MOHH ISD into IHiS. With this merger, national healthcare systems which were originally managed by MOHH ISD came under IHiS' management as well. This merger reduced uncertainty in accountability and mandate between MOHH ISD and IHiS' teams, avoided misalignment between policy and implementation, and reduced transaction and coordination costs. As a result of the merger, MOHH no longer has any IT staff.

37. Following the above-mentioned merger:

(a) IHiS became the central IT agency for the public healthcare system. It was accountable to MOH for matters such as IT policy, governance, planning, and implementing IT projects; and serving the IT needs of the Clusters;

(b) MOH appointed Bruce as CEO of IHiS, in addition to holding his concurrent appointment as MOH CIO;

(c) MOHH ISSD was restructured to become a new Cyber Security Governance ("**CSG**") division in IHiS. Kim Chuan, then-Director of MOHH Identity & Security Services Department ("**ISSD**"; a department within MOHH ISD), became Director of CSG. Kim

Chuan continued to concurrently hold his appointment as MOH CISO; and

(d) IHiS took over from MOHH the responsibility as Sector Lead[7] for the healthcare sector, and CSG took over from MOHH ISSD the day-to-day operational activities of the Sector Lead. Kim Chuan became the Sector Lead point-of-contact within IHiS, an appointment that he had held when MOHH was the Sector Lead.

38. More details on CSG's role may be found at section 9.2 (pg 29) below.

39. Operationally, the employment of IT personnel is centralised at IHiS, though IT personnel employed by IHiS are deployed back to the Clusters to deliver their IT projects and maintain their IT systems. The key IT personnel deployed to the Clusters include the Cluster Group Chief Information Officers ("**GCIOs**") and Cluster Information Security Officers ("**Cluster ISOs**"). The Cluster GCIOs and Cluster ISOs are charged with ensuring that Cluster IT initiatives are aligned to the broader objectives, strategies and policies for the public healthcare sector.

40. To balance the Clusters' need for some autonomy and flexibility in implementation, each Cluster continues to control its own IT budget, for which it remains responsible and accountable. Each Cluster GCIO prepares his Cluster's IT workplan and budget for his Board IT Committee's approval, while keeping in mind public healthcare sector-wide cybersecurity policies and strategies. To date, SingHealth has approved all, and not moderated down, any budget request relating to cybersecurity.

[7] The Sector Lead is the organisation responsible for coordination and information dissemination regarding the protection of CII in the respective sectors. CSA works together with and assists Sector Leads to fulfil their roles and responsibilities for CII protection.

6.4.2 Scope and scale of IHiS' IT operations

41. Overall, IHiS manages a sizeable IT environment across the public healthcare system. The scale of IHiS' operations and security measures is summarised in Figure 2 below. Notably, IHiS manages a total of 60,852 endpoints, 6,232 servers, monitors three terabytes of internet traffic per day, and there are 500 firewall denied attempts per second.

Figure 2: Scale of operations & security across public healthcare

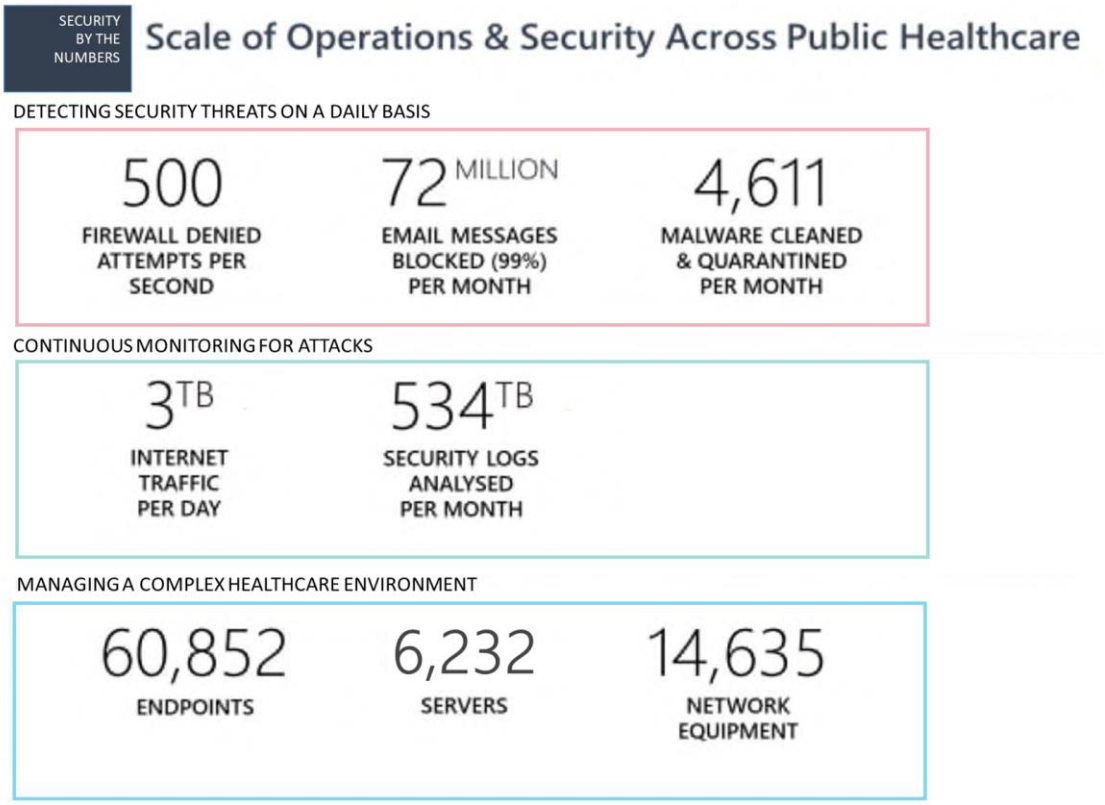

7 THE SUNRISE CLINICAL MANAGER SYSTEM

7.1 Overview of the SCM system

42. The crown jewels[8] of the SingHealth network are patient electronic medical records, contained in the SingHealth Sunrise Clinical Manager ("**SCM**") database. The SingHealth SCM database was the target of the Cyber Attack.

43. SingHealth uses SCM, an electronic medical records software solution from Allscripts Healthcare Solutions, Inc ("**Allscripts**"), a U.S. company whose products and solutions are used by healthcare institutions around the world. The SCM integrates inpatient, emergency and ambulatory care through a single enterprise-wide electronic medical record. This enables physicians, nurses, and other clinicians to access real-time patient data.

44. The SCM system is vital to SingHealth's operations and is extensively used in day-to-day care delivery.

45. The SCM system was identified as a CII system within the healthcare sector and also as a mission-critical system for SingHealth.

7.2 The SingHealth SCM database

46. The SingHealth SCM database that was illegally accessed contains the following information:

(a) Patient demographic data;

[8] 'Crown jewels' refers to an organisation's mission-critical information assets – which could include content, customer data, product designs or other business-critical intellectual property, which would cause major business impact if compromised.
(Source: Mark Lobel, "Cybersecurity: keeping the 'crown jewels' safe online is everyone's business", Feb 2015 <https://pwc.blogs.com/ceoinsights/2015/02/cybersecurity-keeping-the-crown-jewels-safe-online-is-everyones-business-.html.>)

(b) Clinical episode information (*e.g.* A&E, inpatient, outpatient);

(c) Orders (*e.g.* laboratory, radiology, cardiology, medication, nursing);

(d) Results (*e.g.* of diagnostic tests and orders);

(e) Clinical documentation (*e.g.* from doctors, nurses, rehabilitation);

(f) Vital signs (*e.g.* blood pressure, pulse);

(g) Medical alerts and allergies;

(h) Diagnosis and health issues;

(i) Vaccination details;

(j) Discharge summaries;

(k) Medical certificates; and

(l) Outpatient medication dispensed (with associated patient demographics).

47. As at July 2018, the SCM database contained patient data belonging to over 5.01 million unique patients. The data that was illegally accessed belonged to almost 1.5 million unique patients.

7.3 User access to SCM and the SCM database

48. A SingHealth user would access the SCM system through a virtualised version of the SCM client application hosted on the Citrix servers located in the H-Cloud data centre ("**HDC**"). The Citrix servers operate as an intermediary between user workstations and the SCM security server. Citrix servers allow for application virtualisation as opposed to installing applications locally on client

workstations. This means that only screen images of the SCM application are viewed by users on the client workstation. There is no transactional data that flows directly between user workstations and the SCM servers – the only thing that is passed from the Citrix Receiver installed on the workstation, to the Citrix servers, are the users' keystrokes and mouse clicks.

49. When a user launches the virtual SCM client application, the user is required to enter user credentials to log in to SCM. The user credentials are then sent through a Citrix server to the SCM security server for authentication. Upon successful authentication, the user will be logged into the SCM system and can access the SCM database with permissions based on the role that the user is associated with.

50. A simplified illustration of the user authentication process is as follows:

Figure 3:SingHealth user authentication process to access the SCM Database

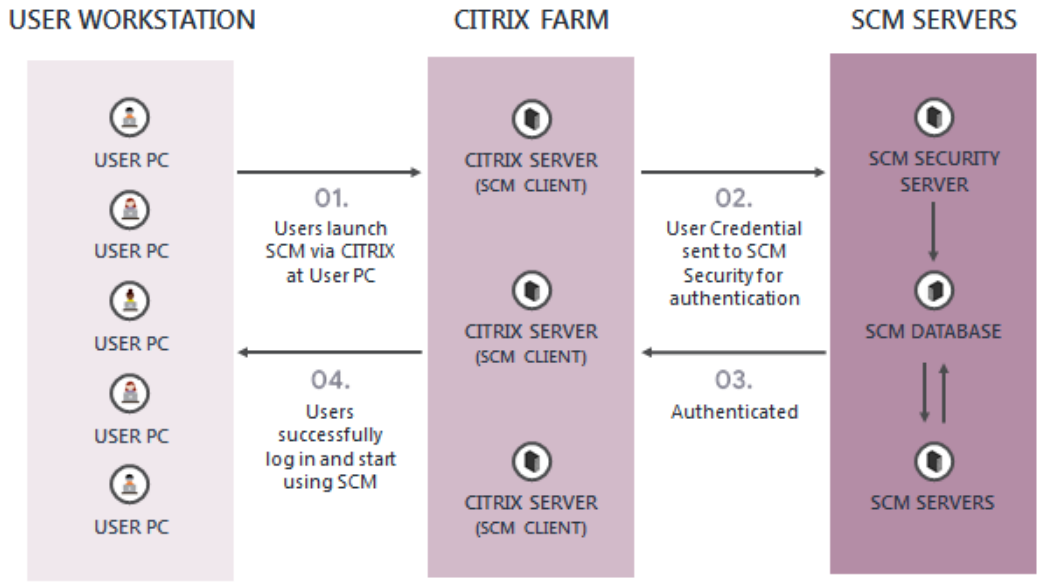

51. The SCM allows for the creation of roles in the SCM system (*e.g.* 'doctor role', 'nurse role'). Permissions can be set for each role, allowing that particular role access to specific functions and data. For example, when attending to a patient, a nurse assigned the 'nurse role' may be allowed to retrieve that patient's

records from the SCM database *via* the SCM client, but may not be allowed to order a lab test or medication for that patient.

52. The SCM application supports the tagging of Very Important Persons ("**VIPs**") within its system. For these tagged patients, only selected users are allowed access to the medical records. Even when an authorised user seeks to access a VIP's visit record, a prompt will be displayed for the user to enter the reason for the access before he/she can proceed to view the record. All instances of access to VIP records are logged and an alert is generated daily to both the user and the hospital's Chief of Medical Board ("**CMB**") *via* email. The user is required to validate his/her access in response to the alert email. If more than a set number of records are accessed at the same time, an alert would be sent to the IHiS security team, and the cluster IT and Operations teams.

53. The SCM client does not have any functionality which allows for the bulk retrieval of records from the SCM database. There are reporting functions which allow users to print, download, or export data into Microsoft Excel. Reporting tools, or custom applications would be used for generating such reports.

8 PARTS OF THE SCM SYSTEM AND NETWORK RELEVANT TO THE CYBER ATTACK

8.1 Overview

54. This section covers the key parts of the SCM system relevant to the Cyber Attack. The attacker also compromised servers that were not part of the SCM system, such as the servers referred to in this report as the NCC Server and S.P. Server, which nonetheless played a role in the attack – see section 14.4.1 (pg 57), and section 14.7 (pg 70), below respectively.

55. Broadly speaking, users in the User Zone of the SingHealth network access the SCM database through an SCM client application hosted on Citrix servers. The Citrix servers serve as middleware between users' workstations and the SCM servers. The users' workstations communicate with Citrix servers,

which in turn communicate with the SCM database servers, where patient data is stored.

8.2 SCM application and database servers

56. The SCM database is hosted on a server in the HDC, and is accessible through the SCM client application. The SCM application and database servers are physically located at HDC.

8.3 Citrix servers

57. The SCM Application is not installed on individual users' workstations. Instead, users access the SCM Application through Citrix servers in the H-Cloud, on which the said application is 'published' (*i.e.* made available for access by multiple users through a process known as *virtualisation*). The HDC Citrix servers were protected by a firewall.

58. There is another set of Citrix servers critical to this Inquiry, which was located in Singapore General Hospital ("**SGH**"). The SGH Citrix servers were located at the SGH Local Data Centre ("**LDC**"). These servers were deployed in a sub-net that was *not* protected by a firewall. The SGH LDC also had another location at the SGH Academia building, which was protected by a firewall. This location contained other IHiS servers that are not relevant to the Inquiry.

8.4 Migration of the SCM to H-Cloud and open network connection between SGH and HDC

59. Prior to June 2017, the SCM infrastructure – including the Citrix servers and SCM application and database servers – were hosted at SGH premises, at the SGH LDCs. The SCM application and database were migrated from Citrix servers at SGH to HDC in July 2017. With the migration, the SCM application, database, and servers were all within the H-Cloud environment, and the SCM system at SGH was supposed to be decommissioned.

60. While the migration of SingHealth servers to H-Cloud has been largely completed for the key systems, there are some remaining non-business critical servers in this zone that are in the progress of migrating to H-Cloud. This includes the Citrix server farm, which still continued to operate at SGH premises, after June 2017 and as at July 2018. These Citrix servers were used to host applications for the SGH intranet, including SAP (which is enterprise software used to manage business operations and customer relations), pharmacy systems, Operating Theatre Management systems etc.

61. Notably, there was still network connectivity between the Citrix server farm at SGH and the SCM database server in the H-Cloud data centre. This connectivity between the SGH Citrix servers and the SCM database servers in the HDC proved to be a significant vulnerability that was exploited during the attack, as explained below at section 15.1 (pg 72).

9 IHIS TEAMS RESPONSIBLE FOR IT AND IT SECURITY ADMINISTRATION AND OPERATIONS

62. The IHiS organisational chart, highlighting the teams involved in the Cyber Attack is presented on the following page. The roles and responsibilities and key personnel from these teams will be detailed in section 9.1 and 9.2 below.

Figure 4: IHiS Organisational Chart

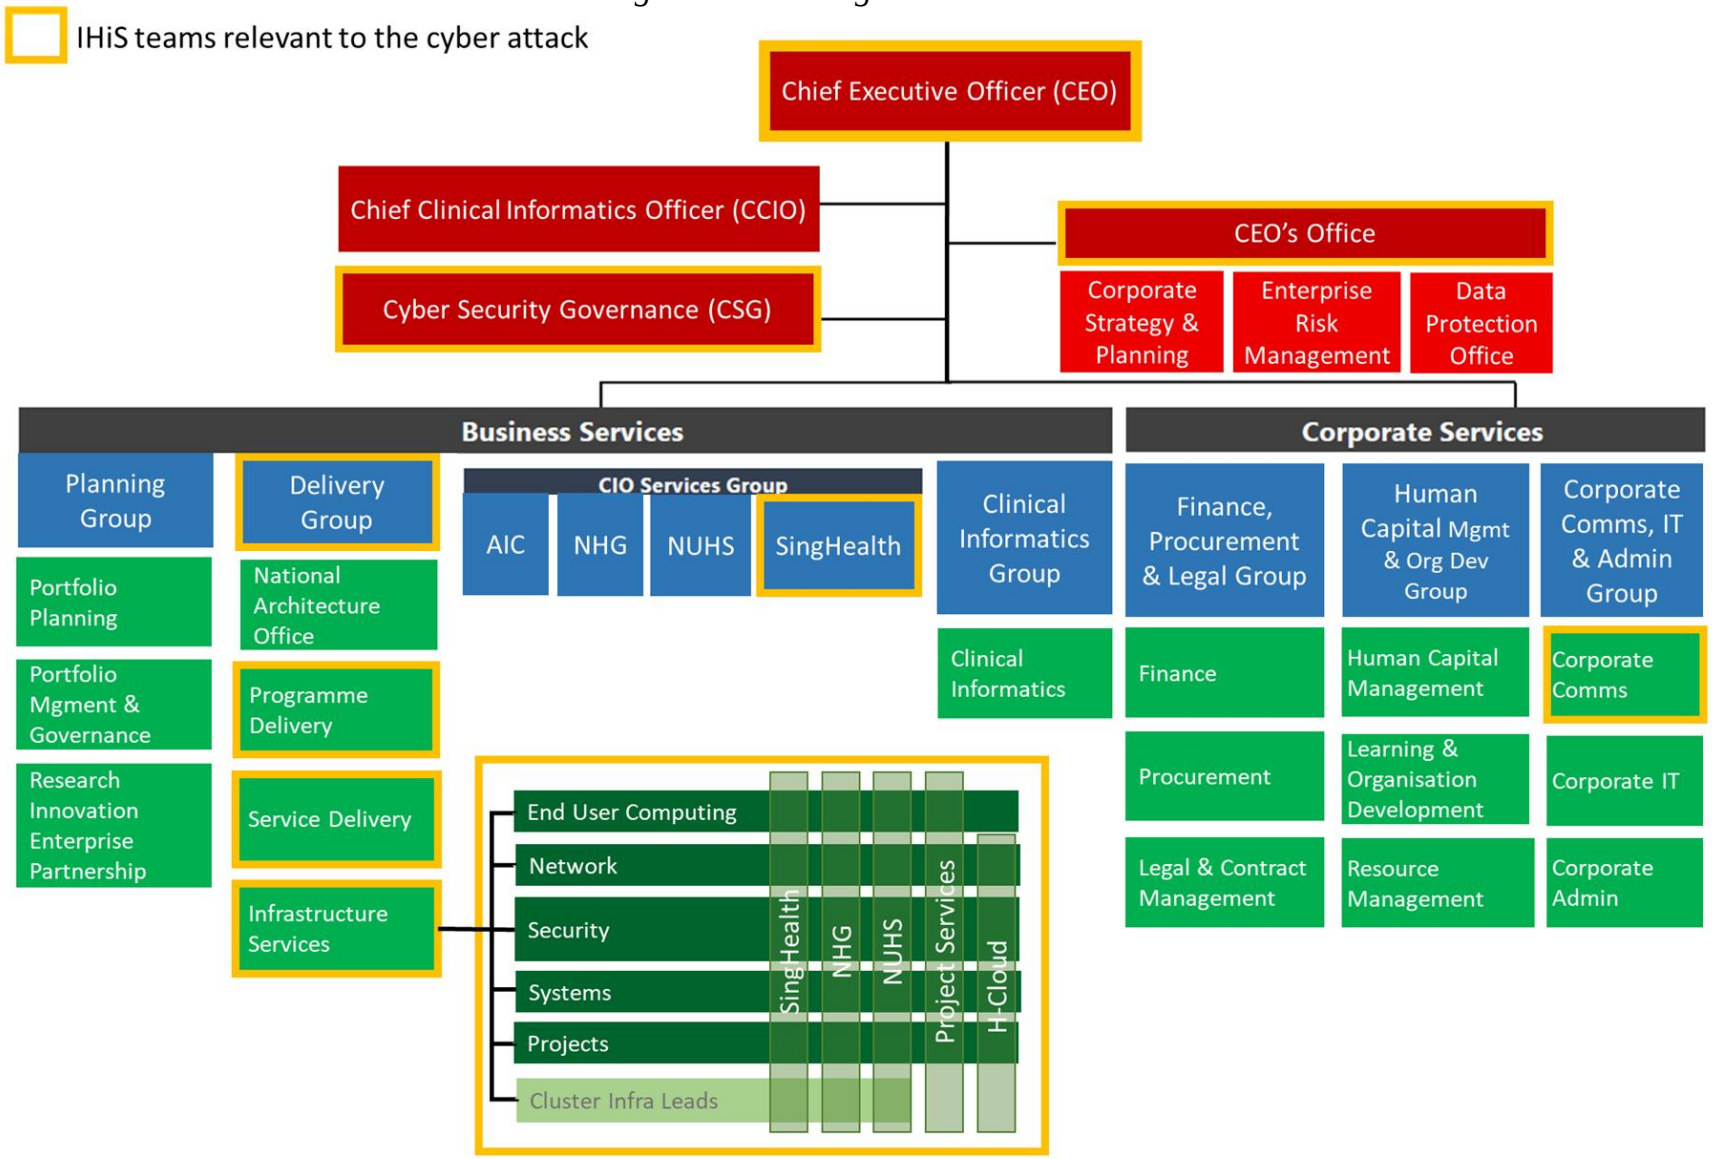

9.1 The Infrastructure Services Division

9.1.1 Overview

63. IHiS Delivery Group, and in particular its sub-group Infrastructure Services, encompass salient functions pertaining to the Inquiry, including the managing of the SCM system and SingHealth network, and responding and reporting of security incidents.

64. Delivery Group is headed by Director Ong Leong Seng ("**Leong Seng**"). Three teams within the Delivery Group are responsible for ensuring the functioning and integrity of the SCM system at different stages:

(a) Product Management and Delivery – Clinical Care ("**PMDCC**");

(b) Infrastructure Services; and

(c) Service Delivery

65. Descriptions of these three teams, as well as their roles and responsibilities, are set out below.

9.1.2 Product Management and Delivery – Clinical Care

66. Generally, the PMDCC team is responsible for the Development and Implementation Stages in the provision of IT systems. PMDCC includes the SingHealth SCM team headed by Programme Director Foong Lai Choo ("**Lai Choo**"), which manages the development and implementation of updates, upgrades and changes to the SCM system.

9.1.3 Infrastructure services

67. Infrastructure services is headed by Director Serena Yong ("**Serena**"). There are four 'towers' of competencies in the Infrastructure Services team, which each tower responsible for a separate IT domain:

(a) End User Computing ("**EUC**") Management, which focuses on the rollout and deployment of end point devices such as PCs, laptops, and printers, including software installation, configuration and administration of these devices. EUC's day-today operations include the physical transportation of end point devices, technical refreshes, trouble-shooting, and user support. The majority of EUC's operations are outsourced to third parties, and EUC works closely with these third parties to ensure smooth EUC operations.

(b) Network Management, which is in charge of network connectivity within the healthcare environment. The team manages the wired and wireless Local Area Network ("**LAN**") that end users connect to, the Wide-Area Network ("**WAN**") that connects the LAN to H-Cloud data centres, the server LAN within the data centres, and all connections to the internet and private third party networks. The team's day-to-day operations include continuous network traffic monitoring, network equipment technical refresh and upgrades, network equipment firmware patches, and network routing and configuration changes to support system rollouts and new facilities. This team is largely in-sourced within IHiS.

(c) Security Management, which is led by Assistant Director, Infrastructure Services, Han Hann Kwang ("**Hann Kwang**"). The Security Management tower comprises three Cluster security teams, that are in charge of security operations and projects, and security awareness and training for their respective clusters. The Security Lead for the SingHealth security team, also called the Security Management Department ("**SMD**"), is Tan Choon Kiat Ernest ("**Ernest**"), who has reporting lines to Hann Kwang, as well as SingHealth Cluster Infrastructure Lead Leong Woon Lan ("**Woon Lan**"). The security operations activities include the operation and administration of security tools such as firewalls, Virtual Private Networks ("**VPNs**"), Intrusion Detection and Prevention Systems ("**IDS/IPS**"); security incident investigation

and response; security assessments; and working with the Infrastructure and/or Application teams in IHiS to close escalated security tickets from IHiS' outsourced Managed Security Services ("**MSS**") provider. The team has doubled in size from 2016 to 2018, and is augmented by outsourced partner services. More information on SingHealth's Computer Emergency Response Team ("**CERT**"), which is under SMD, may be found at at paragraph 112 (pg 40) below.

(d) Systems Management, whose day-to-day operations include (but are not limited to) server operations and monitoring, database storage, server technical refresh and upgrades, and server patching. The Systems Management Department includes a team that provides infrastructure support for the SCM system and SCM Application Citrix servers (the "**Citrix Team**"), and a Database Management team whose responsibilities include managing the SCM database.

68. All members of the Infrastructure Services team are organised in a matrix reporting structure. In addition to the four 'towers', there is a horizontal Cluster Infrastructure Services grouping across the towers. Under the current structure, when an infrastructure services issue is picked up at one Cluster site and needs to be addressed, the tower lead of the IT domain into which the issue falls is supposed to ensure that the issue is addressed across all the Clusters. The current structure is intended to facilitate the propagation of information and holistic implementation of actions across Clusters.

9.1.4 Service Delivery

69. The Service Delivery team is generally responsible for the Maintenance and Support Stage of production systems including SCM. They work closely with the outsourced helpdesk (*i.e.* level 1 support) on all IT incidents, and provide level 2 support *i.e.* restore services to normalcy as quickly as possible. They also follow-through with relevant program teams to resolve identified issues (level 3

support). Of note is the Production Enhancement Team, also known as the SCM Application Team, which provides support for end-user issues with applications in the SCM system.

9.1.5 The SingHealth GCIO and Cluster ISO

70. Each of the Clusters, including SingHealth, has a Group Chief Information Officer ("**GCIO**") and an Information Security Officer ("**Cluster ISO**"), both of whom are IHiS employees. This arrangement has been in place since the formation of IHiS in 2008. The Cluster GCIOs are accountable to the Clusters for Chief Information Officer ("**CIO**") services, such as IT capability development, and systems resilience and security; and are concurrently accountable to IHiS' CEO for the quality of CIO services provided to the Clusters and other IHiS leadership responsibilities.

71. The Cluster GCIOs are accountable to the Clusters for Chief Information Officer ("**CIO**") services, such as IT capability development, and systems resiliency and security; and are concurrently accountable to IHiS CEO for the quality of CIO services provided to the Clusters and other IHiS leadership responsibilities.

72. The SingHealth GCIO is Benedict Tan Wee Bor ("**Benedict**"). GCIO Benedict has a reporting line to IHiS CEO Bruce; as well as to SingHealth management *via* SingHealth Deputy GCEO (Organisational Transformation and Informatics) ("**Dy GCEO**") Professor Kenneth Kwek ("**Prof. Kenneth**").

73. The SingHealth GCIO's roles and responsibilities include:

(a) Strategic IT planning to align IT to support SingHealth's business objectives, including IT capability development, systems resiliency and security (*i.e.* 'Keeping The Lights On' or "**KTLO**"), and IT cost-effectiveness.

(b) Working with MOHH's Group Internal Audit team ("**GIA**") in connection with yearly internal audits on SingHealth's IT systems.

(c) Ensuring that SingHealth's IT enterprise programs remain aligned with security requirements, ensuring compliance with prevailing security policies and standards, and overseeing SingHealth's IT risk assessment.

74. The SingHealth GCIO is supported by the SingHealth GCIO office, which comprises about 50 staff, who are mostly IT directors from SingHealth's PHIs and domain or business analysts.

75. SingHealth GCIO Benedict is assisted by Cluster ISO Wee Jia Huo ("**Wee**") in fulfilling his responsibility for cybersecurity in SingHealth. Wee is the only staff in the SingHealth GCIO office who has a portfolio specific to security, with no officers reporting to him. For cybersecurity matters, the GCIO office (including Wee) works collaboratively with IHiS CSG and IHiS Delivery Group. The SingHealth GCIO office is reliant on IHiS delivery group for both technical implementation of cybersecurity measures and compliance with cybersecurity policy and procedure.

76. The SingHealth Cluster ISO's roles and responsibilities include:

(a) Working on IT risk assessment (see section 12.3.2 (pg 44) below);

(b) Liaising with internal auditors GIA and on follow-up on any audit findings or observations;

(c) Being part of the security incident response and reporting process (see paragraph 108(a) (pg 38) below); and

(d) Assisting GCIO in raising end-user awareness of IT security in SingHealth.

77. The centralised IHiS team in the GCIO office supports the Clusters and GCIOs by delivering the necessary services. IT projects are articulated in the Clusters' annual work plans which are agreed between the GCIOs and the IHiS Delivery Group management, with resiliency projects having the highest priority. Members of the Cluster CIO office and the IHiS Delivery group meet regularly to synchronise demand and supply, and review projects and operations.

9.2 IHiS Cyber Security Governance ("CSG")

9.2.1 Overview of CSG

78. IHiS Cyber Security Governance ("**CSG**") comprises 12 staff who report directly to Director of CSG Kim Chuan, who provides both IHiS CEO and the CSC with a broad overview of security from the governance perspective. The formation of CSG was detailed at section 6.4.1 (pg 14) above.

79. CSG is in charge of (i) developing cybersecurity policies and standards; (ii) liaising with Clusters and IHiS Delivery Group about their implementation of cybersecurity policies for the Clusters; and (iii) tracking and providing compliance assurance on the implementation of cybersecurity policies. CSG acts as the Secretariat for CSC, and proposes policies and make recommendations for CSC's approval.

80. As mentioned, Kim Chuan has a dual appointment as Director of CSG, IHiS and MOH Chief Information Security Officer), which he said *"allows (him) to align IHiS' cybersecurity policies with broader Government standards and facilitates engagement with the Clusters on cybersecurity policies".*

81. CSG is separate from the Security Management Department (located within the Delivery Group) which provides security advisory services and subject matter expertise as well as the Cluster ISOs which spearhead cybersecurity in their respective Clusters. Both the Security Management Department and Cluster ISOs do not report to Kim Chuan.

82. CSG also performs the function of Sector Lead, which is explained in detail below.

9.2.2 CSG's healthcare Sector Lead role

83. CSA requires Sector Leads to oversee and regulate CII owners within their respective sectors. For example, the National Cyber Incident Response Framework ("**NCIRF**") places obligations on Sector Leads to report security incidents to CSA. CSG's role is also to ensure that there is proper incident response for security incidents within the healthcare sector.

84. To avoid any conflict of interest that may arise from its undertaking of the Sector Lead's operational activities, CSG does not have operational responsibilities for any CII systems in the healthcare sector. CSG is also independent of the Delivery Group in IHiS, which performs all functions relating to the operation of IT systems (including the CII systems).

85. CSG also communicates threat intelligence and any indicators of compromise from CSA *via* IT security-related circulars and directives to the Cluster CIOs and the Cluster ISOs in each of the healthcare Clusters, for them to carry out the necessary checks and follow-up.

9.2.3 Conducting compliance reviews and penetration tests

86. CSG performs the inhouse red teaming function for the public healthcare system. Red teaming refers to ethical hacking *i.e.* penetration testing to test the IT systems of PHIs for vulnerabilities. Since 2015, Kim Chuan's team (then at MOHH ISSD, now CSG in IHiS) has been conducting ethical hacking on PHIs' internet-facing systems, and reporting the results to the PHIs' management. Apart from this, CSG does not conduct any compliance assurance, *i.e.* going on the ground to check whether IT security policies and standards are being complied with by the PHIs.

87. In April 2018, CSG started to form a compliance and assurance team to carry out compliance reviews of PHI systems, as well as to move beyond ethical

hacking of internet-facing systems and conduct penetration tests of internal systems. This was in response to discussions at the November 2017 IHiS ARC meeting that CSG should take on such a role as an independent check on PHIs' compliance levels in respect of IT security policies and standards. IHiS is in the process of assembling this team.

88. The relative roles of CSG and GIA with respect to audit and compliance function is explained further at section 12.7 (pg 47) below.

*9.2.4 Conducting Table Top Exercises ("***TTXes***")*

89. CSA mandates that all CII sectors carry out cybersecurity exercises annually within their respective sectors. In 2016, 2017 and 2018, TTXes were conducted to understand the healthcare sector's, including the Cluster's and IHiS', effectiveness and preparedness in responding to cyber attacks. The TTXes were discussion-based sessions where team members met in a classroom setting to discuss their roles and responses during various emergency scenarios. A facilitator guided participants through a discussion of the scenarios and evaluated their responses.

90. Upon completion of an exercise, an After Action Report covering the key observations and areas of improvement shall be prepared, and CSG shall track the progress of the follow-up implementation plans on the areas for improvement.

10 NATIONAL INCIDENT REPORTING FRAMEWORK FOR CRITICAL INFORMATION INFRASTRUCTURE

91. Having established the parties involved and their relationships, we now turn to the incident reporting responsibilities as at the time of the Cyber Attack.

10.1 Identification of SCM as a CII system

92. The SingHealth Electronic Medical Records ("**EMR**") system was identified as CII in a review initiated by the Singapore InfoComm Security

Authority ("**SITSA**"), the predecessor to CSA, in 2011. The SCM, which is an integral part of SingHealth's EMR, is a CII under the charge of SingHealth.

93. As the Sector Lead for the healthcare sector, IHiS is responsible for reporting security incidents to CSA.

10.2 National Cyber Incident Response ("NCIRF")

10.2.1 Overview

94. The NCIRF is the framework for the reporting and management of cyber incidents affecting CII in Singapore. The NCIRF was approved by the Homefront Crisis Executive Group ("**HCEG**") in December 2015. The HCEG is part of the Homefront Crisis Management System, and is the executive body tasked with managing peacetime crises. Under the NCIRF, Sector Leads have to report all security incidents within their respective CII sectors to CSA.

95. At the time of the Cyber Attack, the NCIRF was the only relevant national-level security incident reporting framework. Accordingly, in considering the policies in place at the time of the Cyber Attack and the incident response, the Committee will make references to the NCIRF where appropriate. The Committee notes that the Cybersecurity Act 2018 (Act No. 9 of 2018) (the "**Cybersecurity Act**") came into force on 31 August 2018, and this act will apply to the SCM system, which has been designated a CII under the act. In making its recommendations, the Committee will refer to the Cybersecurity Act where appropriate.

96. The NCIRF has a three-tiered framework, as follows:

(a) CII Owner. CII owners are the entities that own CII assets. They are the first-tier cyber incident responders, and are responsible for immediate response to any cyber incidents that affect the CII assets.

(b) Sector Lead. Sector Leads oversee and regulate CII owners within their respective sectors. They are the Sectoral Cyber Incident Managers, providing second-tier response, as they are in the best position to assess the related business risks, and impact of such threats, to the sector.

(c) CSA. The national agency in charge of cybersecurity, CSA oversees 11 CII sectors [9] and is the National Cyber Incident Manager, which co-ordinates incident response efforts across the sectors. CSA provides third-tier response, supporting Sector Leads and CII owners when required.

10.2.2 Categories of security incidents

97. CII perform critical functions in order to provide essential services which, if disrupted, would have a debilitating impact on Singapore's national security, economy, or public health and safety. Incidents associated with the critical functions of CII or provision of essential services must be reported to CSA in a timely manner to facilitate investigations. The three categories of incidents that Sector Leads must report to CSA are:

(i) Category 1: Incident directly affecting CII.

(ii) Category 2: Incident occurring on systems or network that could put the CII at risk.

(iii) Category 3: Incident occurring on systems or network within CII sector that is not covered under Category 1 and Category 2.

[9] The 11 sectors are: Energy, Water, Banking and Finance, Healthcare, Land Transport, Maritime Transport, Aviation, Infocomm, Media, Security and Emergency Services, and Government.

98. For each category of reportable incident, the NCIRF also states the reporting flow and timing requirements for the Sector Lead or its Computer Emergency Response Team ("**CERT**") to report the incident to CSA.

11 IHIS' INTERNAL FRAMEWORK FOR INCIDENT REPORTING AND RESPONSE

99. The main policy document governing IT security in the healthcare sector, including in the Clusters, is the *Healthcare IT Security Policy and Standards Version 3.0* ("**HITSPS**")[10]. The HITSPS was developed under the charge of Kim Chuan (when he was Director of the Identity & Security Services Department within MOHH ISD) and Francis (the former IHiS Group Director (Technology Management)). Broadly, it prescribes IT security policies, technical security standards and processes to be implemented by the PHIs. Relevant to the Inquiry are policies within the HITSPS pertaining to user-ID management, password management, and technical vulnerability management (vulnerability and penetration tests).

100. The HITSPS states that the reporting timelines and escalation processes for all IT security incidents shall be as per two documents, namely (i) the Healthcare IT Security Incident Response Framework ("**SIRF**") and (ii) the Cluster IT Security Incident Response SOP ("**IR-SOP**").

101. It must be highlighted that the SIRF and IR-SOP are meant primarily for the sector-to-CII level, and it is for the Cluster GCIOs and their IT leads to develop lower level processes to comply with their requirements. There is also no written protocol for how IHiS staff who discover an IT security incident affecting a Cluster's assets are to assess and report the matter.

[10] IHiS plans to update HITSPS by issuing HITSPS Version 4.0, and provided in its evidence a draft of Version 4.0, dated October 2017.

11.1 The Healthcare IT Security Incident Response Framework ("SIRF")

102. The Healthcare IT Security Incident Response Framework ("**SIRF**") – translates the NCIRF requirements into how PHIs are to report IT security incidents to their management and to the Healthcare Sector Lead, for onward reporting to CSA. This was prepared under Kim Chuan's charge, and issued by MOHH and IHiS in February 2017.

103. From the present proceedings, there was no evidence that the SIRF had been circulated or otherwise communicated widely to staff, and was not known to most of the witnesses who were IHiS staff.

11.2 The Cluster IT Security Incident Response SOP ("IR-SOP")

104. The Cluster IT Security Incident Response SOP ("**IR-SOP**") details the various protocols for Clusters and their respective PHIs, for reporting and responding to specific scenarios of IT security incidents. This was created by Han Kwang, based on the SIRF.

105. The IR-SOP was shared in March 2018 with the Security Management team members, incident responders (*i.e.* the CERT), Serena and CSG.

11.3 Security incident reporting flow for SingHealth

106. Both the SIRF and IR-SOP categorise reportable security incidents in an identical manner to the NCIRF. These documents also dictate IHiS' internal reporting timelines to Healthcare Sector Lead (CSG) for each category of incident.

107. Information on the incident reporting flow for SingHealth is set out in both the SIRF and IR-SOP. While some information is in common between the two documents, there are different focuses and updated information in the IR-SOP, which was issued a year after the SIRF.

(a) The SIRF describes the framework for PHIs, also termed as "healthcare entities", in reporting security incidents to their management and the Healthcare Sector Lead. It includes the typical incident reporting flow, starting from the Cluster ISO (Security Officer), which then branches to multiple incident reporting chains – including to the Cluster GCIO and CSG as Healthcare Sector Lead; and also results in the incident being reported to CSA, and management of SingHealth, IHiS, MOHH and MOH.

(b) The IR-SOP covers incident reporting as part of the roles and responsibilities of the SIRT in mounting a broader response to security incidents. The IR-SOP includes the typical incident reporting flow mentioned in (a). Unique to the IR-SOP is the SIRT Reporting Structure and description of the SIRT members' roles, which provides for a linear incident reporting from Cluster ISO to Cluster GCIO, and then to Cluster senior management.

Figure 5: Reporting flow for IT security incidents for SingHealth

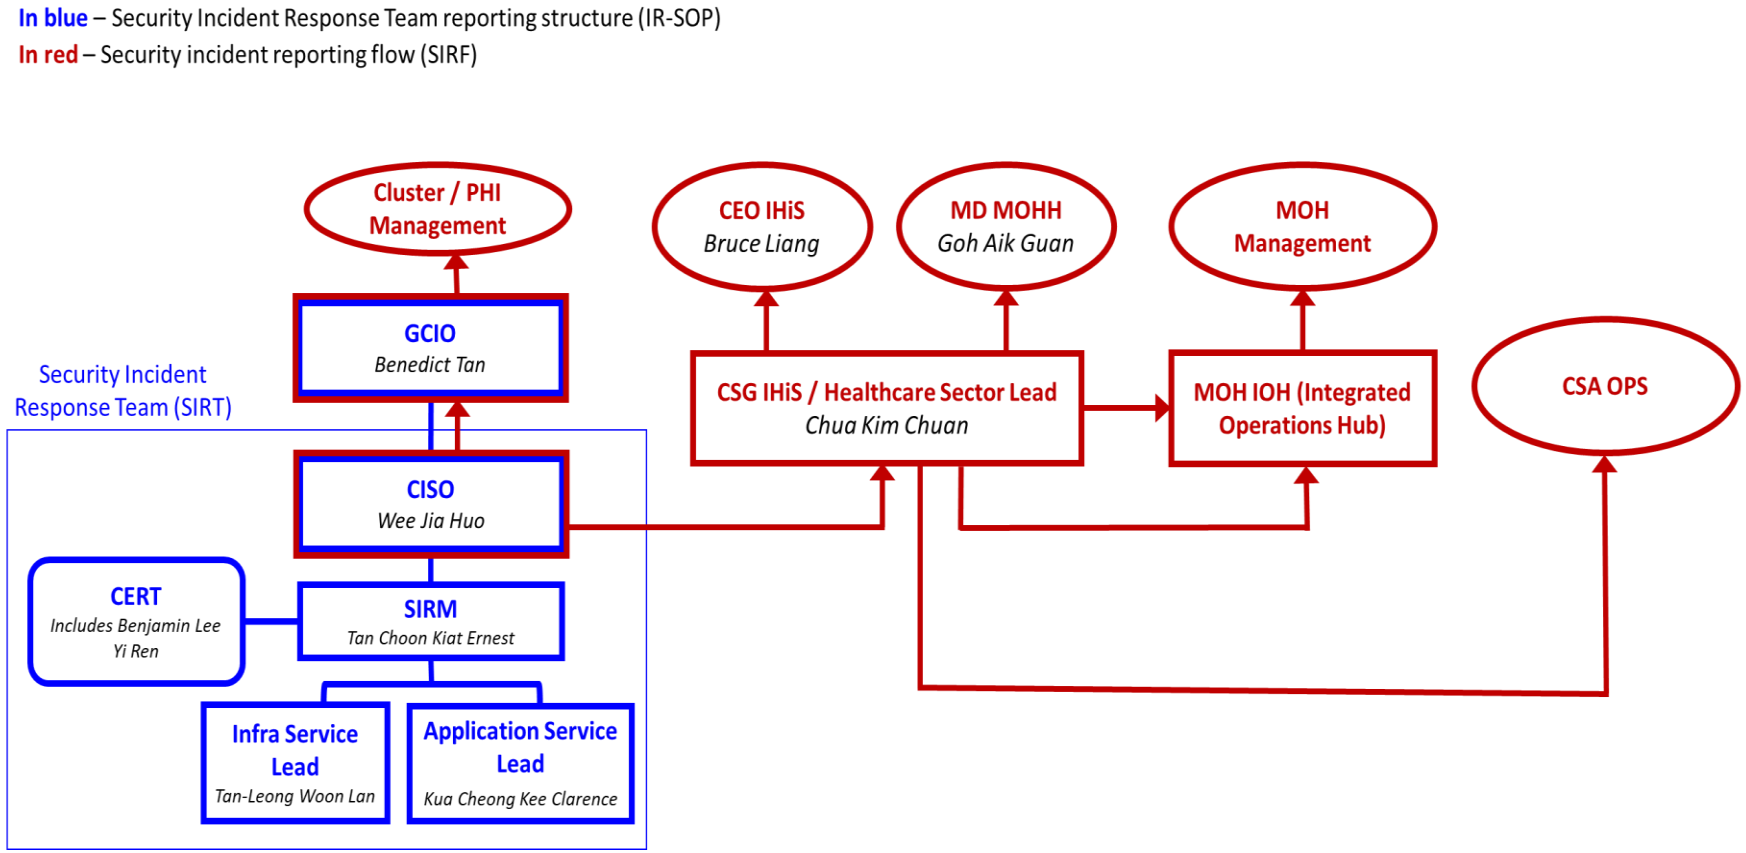

108. With reference to Figure 5 above, the bottom-up reporting flow for IT security incidents in the SingHealth Cluster would be as follows:

(a) Cluster ISO Wee is the first in the reporting chain, who is then to report IT security incidents upwards to multiple stakeholders, including GCIO Benedict and to CSG as Healthcare Sector Lead. Wee's reporting to CSG would be as per the NCIRF incident categorisation and IHiS' internal reporting timelines. Before Hann Kwang wrote the IR-SOP, both communications (including incident reporting) and technical incident handling roles were supposed to be done by the Cluster ISO. But Hann Kwang decided to split the roles in the IR-SOP such that Cluster ISO is in charge of communications (including incident reporting), and Security Incident Response Manager ("**SIRM**"), in this case Ernest, would lead technical incident handling. Nonetheless Cluster ISO Wee gave evidence of his close working relationship with Ernest and SMD in practice, for reporting security incidents:

(i) Wee would typically come to know about security incidents when informed by Ernest or SMD; and

(ii) Upon the receipt of this information, Wee would have a "*two-way conversation*" with Ernest who is the "*subject-matter expert*", to determine if the incident had been confirmed and the category of incident, before escalation to GCIO and Healthcare Sector Lead.

(b) GCIO Benedict is to report the incident to SingHealth senior management. GCIO Benedict does not usually have a *direct role* in the subsequent investigation, response or containment measures, but given that SingHealth is system owner, Benedict would be involved in incident tracking, oversight and management.

(c) Once CSG is alerted, Kim Chuan, who is Director of CSG and the Sector Lead point-of-contact, should review the information and determine if there is a reportable IT security incident, and if so, what the categorisation of the incident is according to CSA's framework. Kim Chuan is then to report the incident to CSA as per the NCIRF timelines, and update IHiS management and MOH. CEO, IHiS has acknowledged that "*while Kim Chuan is the point-of-contact for the Sector Lead, IHiS is the Sector Lead, and as the CEO IHiS, I have the ultimate responsibility to ensure that reporting to CSA is done appropriately*".

11.4 Technical incident response – the Security Incident Response Team ("SIRT"), Security Incident Response Manager ("SIRM") and Computer Emergency Response Team ("CERT")

109. The IR-SOP also details the methodology for incident response, which is to be undertaken by Security Incident Response Team ("**SIRT**"). The SIRT is responsible for investigating and verifying threats, and includes technical experts from various teams, who are also to trigger the necessary response to contain and remediate security incidents, and report services.

110. The SIRT consists of members from the following teams:

(a) Cluster Information Security Officer

(b) Security Incident Response Manager ("**SIRM**")

(c) Computer Emergency Response Team ("**CERT**")

(d) Infrastructure Service Lead

(e) Application Service Lead

111. SingHealth's SIRM is Ernest. The SIRT reports to the SIRM, and the SIRM plays a key role in leading and coordinating technical incident response, namely to "*lead the effort of the (SIRT) and coordinate activities between all of its respective groups*" and to "*receive the initial IR alerts and responsible for activating the IR team and managing all parts of the IR process*".

112. Of note is the SingHealth Computer Emergency Response Team ("**CERT**"), the first responders who are responsible for performing incident analysis to determine the scope and nature of the incident, collect forensic evidence, tracking or tracing the intruder, and providing on-site assistance to help with incident recovery. The three-man CERT was established in March 2018. Benjamin is the one member of the CERT who has attended an incident response course ("*Hacker Tools, Techniques, and Incident Handling*" by SANS Institute), while the other two members have not received any formal incident response training.

113. Also included in the IR-SOP is a set of Security Incident Response Plans, or 'playbooks', that provide a step-by-step guide on the SIRT's incident response for specific scenarios. Hann Kwang explained that the playbooks were targeted in terms of malware, ransomware and website defacement, as this was based on the threat intelligence for the healthcare sector "*for the last 1, 2 years*". There was no playbook on attacks by Advanced Persistent Threats, and the existing playbooks lacked details on the tactics, tools, and procedures of advanced threat actors.

12 IT AND IT SECURITY GOVERNANCE FOR SINGHEALTH

114. Relevant to the Inquiry are the structures and processes for IT and IT security governance pertaining to SingHealth. This includes oversight and decision-making for the policies, technical implementation, and IT security risk management.

12.1 Healthcare sector-wide platforms: The Healthcare IT Steering Committee and the Cyber Security Council

115. Of note are two platforms with sector-wide oversight over the development and implementation of IT strategies for the public healthcare sector, namely the Healthcare IT Steering Committee ("**HITSC**"); and the Cyber Security Council ("**CSC**").

116. The HITSC is a strategic-level forum for decisions on broad policies, strategies and issues relating to overall healthcare IT (including but not limited to cybersecurity). The HITSC is chaired by the Permanent Secretary of MOH, and its members include the Managing Director of MOHH ("**MOHH MD**") Goh Aik Guan ("**Aik Guan**"), Cluster Group CEOs ("**GCEOs**"). If key cybersecurity issues require discussion and consensus amongst all Clusters at the GCEO levels, they are surfaced to the HITSC for decision.

117. The CSC serves as a forum for discussion on the operational feasibility and implementation of cybersecurity policies and initiatives at the Cluster level. The CSC is chaired by MOHH MD Aik Guan and its members include Cluster GCIOs or Group Chief Operating Officers ("**GCOO**"). The CSC discussions are pitched at the operational level, for instance how a measure is to be phased in or how initiatives are to be prioritised. IHiS Cyber Security Governance ("**CSG**"; see paragraph 79 (pg 29) above) acts as the Secretariat for CSC, and proposes policies and makes recommendations for CSC's approval.

118. Bruce, who is the MOH CIO and IHiS CEO, is a member of both the HITSC and CSC.

12.2 Cluster-level platforms for SingHealth

119. For the SingHealth Cluster, there are four Board-level committees that have oversight of IT security matters for SingHealth.

(a) First, SingHealth Board. The SingHealth Board receives summaries of the discussions in Board committees, and if necessary, key matters are escalated to the Board for attention or decision. The Board meets quarterly.

(b) Second, the IT Committee ("**ITC**"), comprising Board members and co-opted members from external institutions who have IT expertise. Senior management representatives from SingHealth, such as GCEO Prof. Ivy, Dy GCEO Prof. Kenneth and Group Chief Information Officer ("**GCIO**") Benedict Tan ("**Benedict**"), attend ITC meetings, which are held two to three times a year. The terms of reference of the ITC include reviewing IT security; providing oversight and direction on IT infrastructure development; and making recommendations to the Board on Cluster IT development policies, plans and issues.

(c) Where audits and key risks relate to cybersecurity matters, these are also deliberated by the Audit Committee ("**AC**") and the Risk Oversight Committee ("**ROC**"). GCEO Prof. Ivy also attends the AC and ROC meetings, which are held two to three times a year. On an annual basis, the MOHH Group Internal Audit team ("**GIA**") identifies and prioritises the key risk areas (including for cybersecurity) together with input from SingHealth management, and comes up with the annual audit plan for the AC's review and approval.

(d) At SingHealth management-level, the Cluster IT Council ("**CITC**") is the overall governing body for IT across the SingHealth Cluster. The CITC reports to the ITC. The CITC is chaired by GCEO Prof.

Ivy, its members include all the CEOs and heads of the various public healthcare institutions ("**PHIs**") in SingHealth. and its secretariat is the office of the GCIO. The CITC's role is to ensure that IT strategy and investments are aligned with the business strategy and IT architecture of the Cluster, resulting in the effective and efficient use of IT in enabling SingHealth to achieve its goals. Each year, GCIO Benedict, with the assistance of SingHealth PHIs, consolidates the SingHealth cluster IT workplan which will be presented to the CITC (and thereafter the ITC) for approval. An IT workplan would typically include IHiS' direction for implementation of IT initiatives, including IT security initiatives for the financial year. The CITC also meets on a monthly basis to review and endorse SingHealth's Cluster-wide IT projects and initiatives which are presented by IHiS staff and sometimes together with relevant SingHealth staff who provide the user perspective.

12.3 IT security-related risk management

12.3.1 MOHH Audit and Risk Committee ("ARC") and Group Internal Audit ("GIA")

120. MOHH has an Audit and Risk Committee ("**ARC**"), which is chaired by an MOHH Board member. The MOHH ARC members include the respective Chairmen of the audit committees or audit and risk committees of the three Clusters and IHiS.

121. MOHH's GIA, which provides internal audit services to the MOHH Group, including the Clusters and IHiS, and helps improve their governance, controls and risk management. The GIA has a specialised IT audit team that conducts IT security audits, led by IT audit head Thng Chiok Meng. The GIA has a direct reporting line to the MOHH ARC, as well as direct reporting lines to the audit or audit and risk committees of the Clusters, IHiS and the Agency for Integrated Care.

122. An overview of audit findings, including cybersecurity findings for IT-related audits, is tabled for discussion at IHiS ARC meetings, which oversee and review systems of risk management within IHiS, including audit and business processes to manage risks. IHiS ARC also agrees on the appropriate follow-up actions to be taken to address the audit findings.

12.3.2 Internal IT security risk assessments

123. In accordance with CSA's requirements, all CII owners are to conduct risk assessment of their CII at least once every 12 months, and are to submit the risk assessment results, together with the risk mitigation plan and timeline, to the Sector Lead for tracking.

124. Cluster ISO Wee handles the IT risk assessment for SingHealth including the annual risk assessment of the SCM system. To prepare the risk assessment, Wee coordinates with the relevant teams in the IHiS Delivery Group *(e.g.* the Systems team, and the Security team) to obtain their views, and submits the risk assessment results to CSG (Healthcare Sector Lead) while also sharing the results with with GCIO Benedict for his information. If any new technical controls are required in response to the risks identified, Wee will coordinate with the relevant teams in the IHiS Delivery Group to ensure they provide and implement the necessary measures. CSG is to track the risk assessments of CII.

125. Relevant to the Inquiry is the FY2016 risk assessment report for the SCM system ("**FY16 CII Risk Assessment**") prepared by Wee with inputs from the Infrastructure and Application Teams, dated 3 January 2017. This will be discussed in section 18.3 (pg 104) below.

126. It is worth noting that the HITSPS, the internal IT policy document for the healthcare sector, also has a requirement for security risk assessments on mission-critical IT systems (which includes the SCM), but there is no fixed frequency for conducting risk assessments.

12.4 IT security audits

127. IHiS does not have its own internal audit department. Audits, including IT security audits, are carried out by MOHH GIA. Periodically, the GIA would conduct audits on the clusters' network and systems. The team that looks at IT within GIA conducts audits of both CII and non-CII systems.

128. The HITSPS states that independent audits of PHI's IT systems are to be conducted by the GIA periodically to evaluate and test the adequacy of, and the compliance to prevailing IT security policies and standards.

12.4.1 CII audits on the SCM system

129. Since the SCM system is a CII system, SingHealth as CII owner is to conduct an independent cybersecurity audit of the SCM system at least once every 12 months[11], with the scope of the audit conducted in accordance with CSA's requirements. These results are then to be submitted to Sector Lead CSG, together with mitigation/improvement plan and timeline. GIA would carry out the audit itself, while CSG as Sector Lead would follow up on the results to track the progress of action plans for reporting to MOHH management. Further details on CSG's role in follow up are at section 12.5 (pg 46) below.

12.4.2 Audits for non-CII systems and the FY16 H-Cloud Pen-Test

130. For non-CII systems, the GIA will prepare an audit workplan, with inputs from SingHealth management. These audits are typically conducted by the GIA, although the GIA may contract some audits to external auditors. Findings of these internal audits are reported to SingHealth's Audit Committee, and where risks are highlighted in the audit, will be surfaced to SingHealth's Risk Oversight Committee ("**ROC**"). The GIA keeps SingHealth updated on audit findings and the status of remediation plans in response to the audit findings, at Audit Progress

[11] This requirement has been superceded by the requirements of the Cybersecurity Act, which came into force on 31 August 2018.

Update ("**APU**") meetings which are held quarterly and at SingHealth's Audit Committee meetings.

131. Of relevance to the Inquiry is the FY2016 network penetration testing from SGH to H-Cloud, conducted by GIA in January 2017, as part of its internal audit activities for FY2016 (the "**FY16 H-Cloud Pen-Test**"). GIA had engaged an external consultant to conduct a set of network penetration tests from three PHI's systems (including SGH) to H-Cloud. By March 2017, certain high-risk weaknesses had been uncovered from these penetration tests, and IHiS senior management and MOHH ARC were notified that month. The findings from and response to the FY16 H-Cloud Pen-Test were reflected in an Internal Audit Report published in May 2017 (the "**FY16 GIA Audit Report**"), and will be discussed further in section 15.7 (pg 89) below.

12.5 Compliance reviews and tracking of progress on action plans from audits

132. CSG carries out annual compliance reviews of mission-critical IT systems (which includes the SCM system) for compliance with prevailing IT security policies and standards. Before the formation of CSG, the Cluster GCIOs were initiating such compliance reviews, but with the formation of CSG in November 2016, CSG has been coordinating compliance reviews for all Clusters.

133. CSG is also responsible for tracking the progress status of action plans arising from CII audits, for reporting to MOHH senior management. Specifically, CSG is to (a) compile all submitted audit results in an Audit Tracking Sheet; (b) collate updates from SingHealth on the progress of the mitigation/improvement plans for the SCM system every 6 months; (c) gather the corresponding evidence of the completion of mitigation/improvement plan for closures; (d) update the Audit Tracking Sheet accordingly; and (e) update the CSC on the results of the audit conducted, and the progress of the CII owner's mitigation/improvement plan, once every 6 months.

12.6 Follow-up for IT Security audits

134. Where the audit concerned the SingHealth network or systems, the Infrastructure Services Lead would ordinarily lead the follow up on the audit observations and findings. For follow-up of audit findings in 2017, including the follow-up for the FY16 H-Cloud Pen-Test, it would have been Serena Yong's role as Infrastructure Services Lead for SingHealth, to work with the Tower Leads. In 2017, the Tower Leads would have been:

(a) Nick Thoo for Network Services;

(b) Loh Khim Huat for End-User Computing;

(c) Ernest for Security Services; and

(d) Woon Lan for Data Centre Services.

135. Verification of audit findings, *i.e.* that follow up action has indeed been taken, is conducted:

(a) By CSG on a 6 monthly basis pertaining to the tracking of progress of action plans from CII audits, for updating to SingHealth management; and

(b) By GIA on a yearly basis, as part of the overall audit process for that financial year.

12.7 Relative roles of MOHH GIA and CSG

136. There have been various discussions on the role of MOHH GIA vis-à-vis CSG. IHiS ARC agreed in March 2017 on the following roles and responsibilities for CSG and MOHH GIA, with concurrence from IHiS CEO Bruce:

(a) CSG would perform all necessary checks of security implementation through its compliance programs and MOHH GIA would review the adequacy of the compliance programs carried out by CSG.

(b) MOHH GIA could also conduct independent tests including network penetration tests periodically to validate the effectiveness of controls.

137. Since 2017, there has been discussion at the IHiS ARC over the three lines of defence model for effective cyber risk management and control, which is being designed. In brief, this would comprise operations as the first line, compliance checks at the second line, and internal audit as the third line of defence. At the time of the Cyber Attack, the respective roles of GIA and and CSG were not yet finalised. The Committee will discuss this further when it makes its recommendations in section 36.1 (pg 235) below.

Part III – The attacker and the events and contributing factors leading to the Cyber Attack

TABLE OF CONTENTS – PART III

13 INTRODUCTION TO THIS PART

138. In this Part, the Committee presents its findings in respect of TOR #1, establishing the events and contributing factors leading to the Cyber Attack and the subsequent exfiltration of patient data.

139. Although TOR #1 refers to the Cyber Attack as having occurred on or around 27 June 2018, the evidence shows that the Cyber Attack in fact *began earlier*, with the earliest signs of compromise dating back to 23 August 2017. It was only the *querying* of the SCM database which began on 27 June 2018, continuing on until 4 July 2018. Thereafter, instances of malicious activity took place on 18 and 19 July 2018. No further instances of malicious activity were observed after internet surfing separation was implemented on 20 July 2018. Thus, taking a broader view, the Cyber Attack spanned a period from around 23 August 2017 to 20 July 2018. Accordingly, the Committee's findings in this Part will encompass all relevant events that took place in this period.

140. The Committee's findings in this Part comprise three main issues. First, reconstructing the events of the Cyber Attack; second, identifying the pre-existing vulnerabilities that were exploited or may have been exploited by the attacker in the course of the Cyber Attack; and third, profiling the attacker.

141. In considering the events of the Cyber Attack, it is useful to bear in mind the Cyber Kill Chain framework developed by Lockheed Martin, which identifies what adversaries must complete in order to achieve their objectives, going through seven stages starting from early reconnaissance to the final goal of data exfiltration. Having this framework in mind will facilitate understanding of the actions and the tactics, techniques and procedures ("**TTPs**") of the attacker in this case.

Figure 6: The Cyber Kill Chain developed by Lockheed Martin[12]

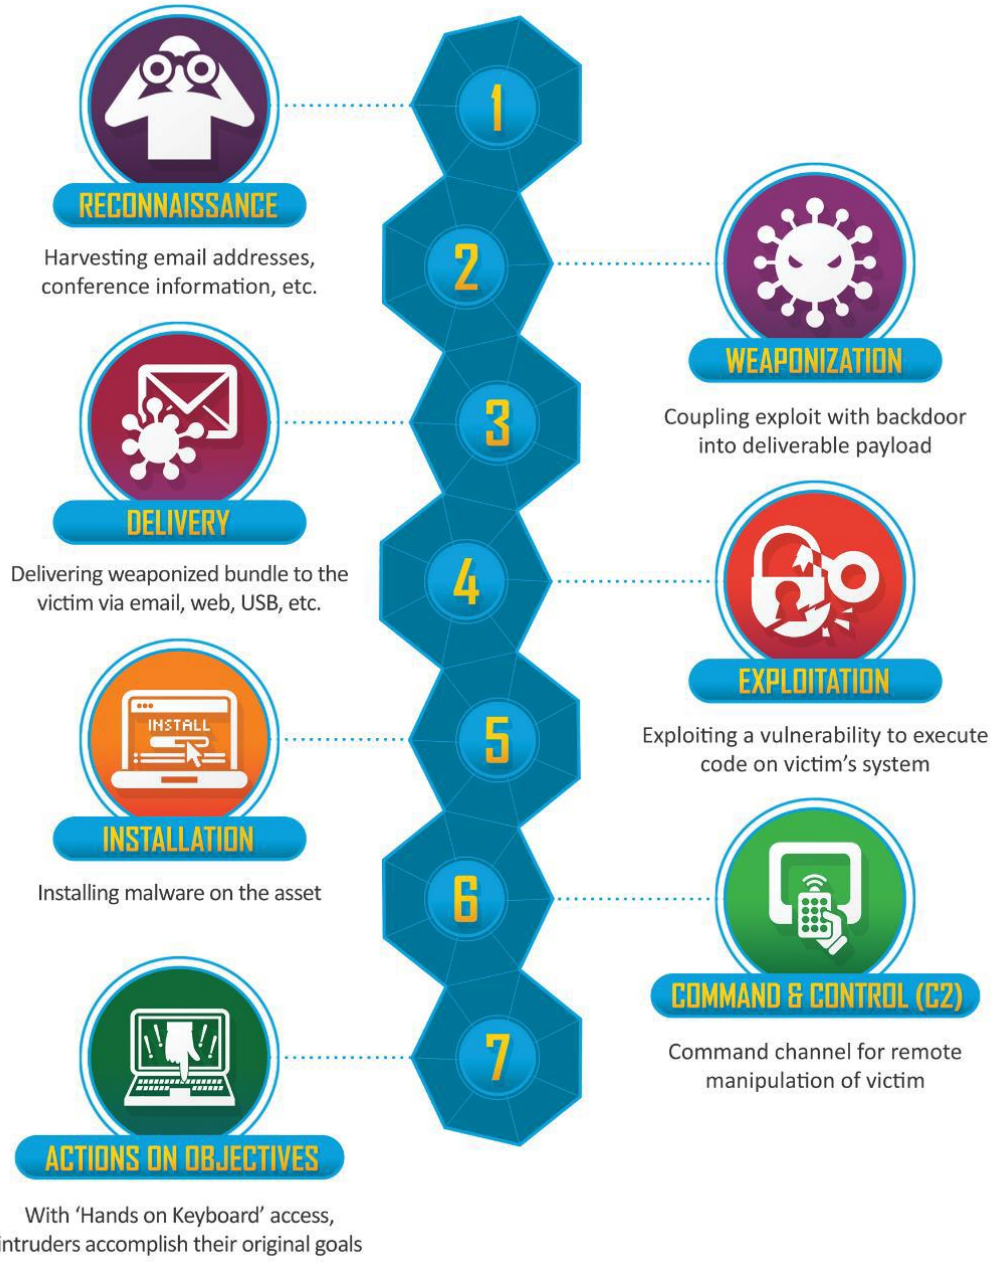

[12] Lockheed Martin Corporation, "The Cyber Kill Chain", 2018. <https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html>.

14 THE CYBER ATTACK

14.1 CSA's reconstruction of events

142. CSA's National Cyber Incident Response Team ("**NCIRT**") was able to substantially reconstruct the events of the Cyber Attack through thorough forensic analysis of machines suspected to have been compromised, network traffic flow data, and systems' logs. The initial batch of data was collected based on information provided by IHiS, and included forensic images provided by IHiS of some machines. As more information was revealed in the course of investigations, more forensic images and memory dumps of workstations and servers were collected. Proxy and network logs from various network segments, such as login logs and firewall logs, were also collected.

143. The NCIRT has provided a graphical summary of its findings:

Figure 7: Key events of the Cyber Attack

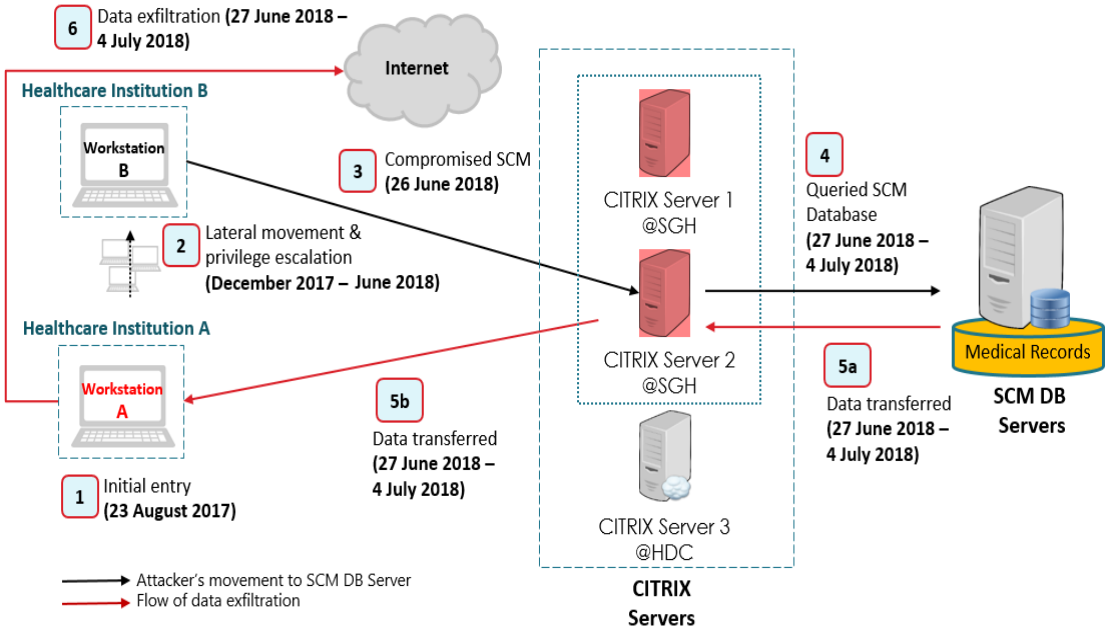

144. Having considered the evidence before it, the Committee accepts CSA's reconstruction of the sequence of the attack, and presents its findings below.

14.2 First evidence of breach and establishing control over Workstation A – August to December 2017

145. Forensic investigations uncovered signs of callbacks to an overseas command and control server[13] ("**C2 server**") from 23 August 2017. Callbacks refer to communications between malware and C2 servers, to either fetch updates and instructions, or send back stolen information. The computer that these callbacks originated from had been decommissioned in October 2017, and was not available for forensic analysis.

146. A different workstation, **Workstation A** began calling-back to the same C2 server on 24 August 2017, one day after the earliest-detected callback.

147. As will be shown subsequently, Workstation A went on to play a critical role in the Cyber Attack as a key pivoting point through which the attacker entered the network, and was also used for the exfiltration of the stolen patient and medical data between 27 June and 4 July 2018. In the course of investigations by the Criminal Investigation Department ("**CID**") of the Singapore Police Force, the user of Workstation A denied being involved in any way in the unlawful access of the SCM system in 2018. Investigations by the CID also did not reveal any evidence of the user's involvement in the Cyber Attack.

148. While not conclusive, there is some evidence to suggest that the initial intrusion was through a successful phishing attack, which led to malware being installed and executed on the workstation.[14]

---

[13] C2 servers are centralised devices operated by attackers to maintain communications with compromised computers within a target network.

[14] Phishing refers to a common technique used by hackers to trick people (typically through emails) into divulging personal information, transferring money, or installing malware.

149. CSA analysts discovered a number of malicious artefacts in Workstation A, including (i) a log file which was a remnant of a malware set; (ii) a publicly available hacking tool, (iii) a customised Remote Access Trojan[15] referred to in this report as "**RAT 1**". Pertinent details of these artefacts are as follows:

(a) The log file was a remnant file from a known malware which has password dumping capability;

(b) The publicly available hacking tool enables an attacker to maintain a persistent presence once an email account has been breached, even if the password to the account is subsequently changed. It also allows an attacker to interact remotely with mail exchange servers, perform simple brute force attacks on the user's email account password, and serve as a hidden backdoor for the attacker to regain entry into the system in the event that the initial implants are removed; and

(c) RAT 1 provided the attacker with the capability to access and control the workstation, enabling the attacker to perform functions such as executing shell scripts remotely, and uploading and downloading files.

150. The log file was created on Workstation A on 29 August 2017. The file contained password credentials in plaintext, which appeared to belong to the user of Workstation A. The malware was likely to have been used by the attacker to obtain passwords for privilege escalation and lateral movement.

151. The publicly available hacking tool was installed on Workstation A on 1 December 2017 by exploiting a vulnerability in the version of Microsoft Outlook ("**Outlook**") that was installed on the workstation. Although a patch for Outlook addressing this vulnerability was available at the material time, the patch was not

[15] A Remote Access Trojan is a type of malware that provides the attacker with access to and control of the victim system through a remote network connection.

installed on Workstation A then. The tool was thus successfully installed and was used to download malicious files onto Workstation A. Some of these files were masqueraded as .jpg image files, but in fact contained malicious PowerShell scripts, one of which is thought to be a modified PowerShell script taken from an open source post-exploitation tool.

152. Also on 1 December 2017, shortly after the installation of the hacking tool, RAT 1 was created on Workstation A.

153. With the introduction of the hacking tool and RAT 1 in December 2017, the attacker gained the capability to execute shell scripts remotely, as well as to upload and download files to Workstation A. Referring to the Cyber Kill Chain framework referred to in paragraph 141 above, it can be seen that the attacker was able to go through the 'Delivery', 'Exploitation', 'Installation' and 'Command and Control' phases by 1 December 2017.

14.3 Privilege escalation and lateral movement – December 2017 to June 2018

154. After the attacker established an initial foothold in Workstation A, it moved laterally in the network between December 2017 and June 2018[16], compromising a number of endpoints and servers, including the Citrix servers located in SGH, which were connected to the SCM database. CSA's assessment

[16] The Committee notes that in CSA's reconstruction of events, the period of "privilege escalation and lateral movement" is stated to be from December 2017 to May 2018, and the events of June 2018, where the attacker made unauthorised logins to the SGH Citrix servers and attempted to log in to the SCM system, are viewed as a different "phase". This conception of the events has the merit of clarity, with clearly defined "phases". At the same time, having regard to the Cyber Kill Chain and the specific facts of the Cyber Attack, the period between the 'Command and Control' stage (*i.e.* gaining control of Workstation A) and the 'Actions on Objectives' stage (*i.e.* retrieving and stealing records from the SCM database) may be viewed holistically as a period of "privilege escalation and lateral movement" – where the attacker moved from system to system within the network, and gained additional privileges by compromising more accounts and systems. Viewed in this light, the events of June 2018 may also constitute "privilege escalation and lateral movement".

is that the attacker moved in a targeted manner, planning his route in the network to reach his ultimate objective, the SCM database.

155. Evidence of the attacker's lateral movements was found in the proliferation of malware across a number of endpoints and servers. Malware samples found and analysed by CSA were either tools that were stealthy by design, or unique variants that were not seen in-the-wild and not detected by standard anti-malware solutions. Such malware included RAT 1, another Remote Access Trojan referred to in this report as "**RAT 2**", and the malware associated with the earlier-mentioned log file.

156. There was also evidence of PowerShell commands used by the attacker to distribute malware to infect other machines, and of malicious files being copied between machines over mapped network drives. These were clear indicators that the attacker had moved laterally around the network.

157. CSA has also assessed that the attacker is likely to have compromised the Windows authentication system and obtained administrator and user credentials from the domain controllers.[17] This meant that the attacker would have gained full control over all Windows based servers and hosted applications, all employee workstations, and underlying data, within the domain.

158. A number of notable events between December 2017 and June 2018 are set out in the following section.

14.4 Notable events between December 2017 and June 2018

14.4.1 Establishing control over the NCC server

159. The **NCC server** was located at a server room at the National Cancer Centre ("**NCC**"), and was part of the SingHealth IT network. In the context of

[17] The domain controller is a server that responds to, and validates, security authentication requests such as logging in and checking permission within a Windows domain.

the attack, uses of the the NCC server included use as a distribution point for malware, where malware was stored temporarily before being copied to other workstations or servers in the network.

160. Forensic analysis of the NCC server revealed the presence of malicious artefacts from as early as 29 September 2017. Malicious PowerShell scripts were also found to have been created on the server in January 2018, and it is likely that these malicious scripts were executed as part of the process through which the attacker strengthened its control over the server.

161. The NCC server was an IHiS asset. However, investigations have revealed that it was not in fact being managed by IHiS. Instead, it was managed locally by an NCC employee, Tan Aik Chin, since January 2016. This was a result of happenstance, and Aik Chin did not possess the necessary knowledge to administer the server. As a result, patches that would ordinarily be rolled out automatically for other servers under IHiS' care were not similarly rolled out to the NCC server. In fact, the server did not have an updated version of the antivirus program installed.

14.4.2 Callbacks to a foreign IP address in January 2018 from Workstation A and the PHI 1 Workstation

162. In January 2018, (i) a workstation from a SingHealth public health institution (in this report, this specific institution is referred to as "**PHI 1**", and the workstation is referred to as the **"PHI 1 Workstation"**), and (ii) Workstation A from SGH, were separately making callbacks to a foreign IP address. As will be shown in section 19 (pg 109) below, while IHiS staff were aware of callbacks from both workstations on 19 January 2018, action was taken only to block connections to the address from PHI 1, and not SGH. CSA's investigations have revealed that this foreign IP address was that of one of the key C2 servers used by the attacker throughout the entire period of the Cyber Attack. When CSA's incident response team was onsite at IHiS after 10 July 2018, there was still ongoing communications with this C2 server from compromised computers.

14.4.3 Obtaining credentials of the L.A. local administrator account

163. A local administrator account, referred to in this report as the "**L.A. account**", was an account found on all the Citrix servers at the SGH data centre. The account has full administrative privileges to login to the Citrix server, including logging in interactively[18], and logging in remotely *via* RDP. The attacker obtained and used the credentials of the L.A. account to log in to at least two SGH Citrix servers (referred to in this report as "**Citrix Server 1**" and "**Citrix Server 2**" respectively) on multiple occasions in May and June 2018.

164. Investigations have revealed at least two possibilities of how the attacker obtained the password for the L.A. account:

(a) First, the L.A. account had a weak password, 'P@ssw0rd', that would produce a common password hash that could easily be decrypted with free online tools.[19] Attackers who are experienced in network intrusion techniques would be familiar with the use of such weak password hashes. From the numerous domain user profiles observed in Citrix Server 1, CSA deduced that the attacker could have logged in using a domain user account, obtained the password hash of the L.A. account, and then decrypted it with ease.

(b) Second, the credentials to the L.A. account were found to be reflected in clear-text on a batch file on Citrix Server 1. It is possible that the attacker had first achieved access to the file system of the Citrix server, and then accessed this file and obtained the credentials.

---

[18] An interactive log in is a process whereby the user gains access to the network by entering a username and password in response to a dialog box on the local console.

[19] Using a publicly available online tool, CSA was able to decrypt the password hash within seconds to reveal the actual password in plaintext.

14.4.4 Obtaining credentials of the S.A. service account

165. The attacker compromised a system level service account, referred to in this report as the "**S.A. account**". The S.A. account has full administrative privileges to login to the Citrix server, including logging in interactively, and logging in remotely *via* RDP. In the context of the attack, the attacker used this account to log in to Citrix Server 2 on multiple occasions in June 2018.

166. IHiS did not have any operational use of the service for which the S.A. account was created. CSA has observed that the attacker could have acquired the credentials to the S.A. account through the malware it used.

14.4.5 Obtaining credentials for the D.A. domain administrator account

167. The attacker also compromised a domain administrator account, referred to in this report as the "**D.A. account**". A domain administrator account is a member of the administrators group on all domain controllers, all domain workstations, and all servers that are members of the domain. An administrator account gives the user full control of the files, directories, services, and other resources that are under the control of the servers in the domain. In the context of the Cyber Attack, compromising the D.A. account allowed the attacker to access and control the SGH Citrix servers.

168. The D.A. account was subsequently used in attempts to log in to the SCM database, and in connecting from Citrix Server 2 in SGH to Citrix Server 3 in the H-Cloud.

14.4.6 Establishing control over Workstation B on 17 April 2018

169. On 17 April 2018, the attacker gained access to Workstation B, a workstation in the SGH, and planted a copy of RAT 2, thus gaining control of the workstation. Workstation B was a workstation which had access to the SCM application.

170. In the context of the attack, Workstation B was used to log in remotely to the SGH Citrix Servers 1 and 2. It is also suspected that Workstation B, or a spoof of it, was used to host virtual machines[20] (referred to in this report as "**VM 1**" and "**VM 2**") used by the attacker to log in to the SGH Citrix servers.

14.4.7 Attempts to log in to the SCM database from Citrix Server 1 from 24 May to 12 June 2018

171. The events discussed in this section are summarised in the following figure:

Figure 8: Attempts to log in to the SCM database from Citrix Server 1

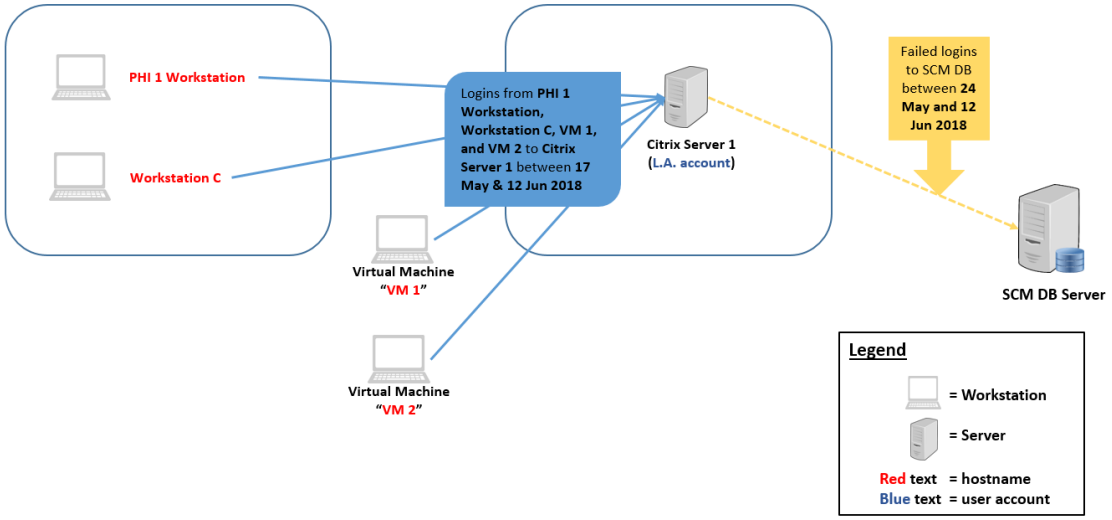

Unauthorised access to Citrix Server 1 from 17 May to 12 June 2018

172. From 17 May 2018 to 11 June 2018, the attacker used the L.A. account to remotely log in to SGH Citrix Server 1 on numerous occasions. The L.A. account is a local domain administrator account not ordinarily used for day to day operations.

[20] A virtual machine ("**VM**") is an emulation of a computer system that, like a physical computer, runs an operating system and applications. A VM allows one to run two operating systems alongside one another on a single machine.

173. The unauthorised logins to Citrix Server 1 were also made *via* Remote Desktop Protocol ("**RDP**") from workstations which would not ordinarily use the L.A. account, including (i) the PHI 1 Workstation; (ii) a SGH workstation referred to in this report as "**Workstation C**"; (iii) VM 1; and (iv) VM 2.

174. On 11 June 2018, IHiS staff became aware of the unusual logins to Citrix Server 1 using the L.A. account, and they changed the password for the L.A account that same evening. This was based on the understanding that (i) the L.A. account is not ordinarily used for day to day operations; and (ii) the unauthorised logins to Citrix Server 1 were made from workstations with hostnames which would not ordinarily use the L.A. account.

175. On 12 June 2018, the attacker attempted to log in to Citrix Server 1 using the L.A. account, but was unable to do so. It then used another account to access the server.

Failed attempts to log in to the SCM database from 24 May to 12 June 2018

176. Starting from 24 May 2018, the attacker made a number of failed attempts to log in to the SCM database from Citrix Server 1. These attempts failed because the attacker either used invalid user-IDs. The latter group included the user-ID of the user account of Workstation A. The failed logins prior to 11 June 2018 were not noticed by IHiS staff at the time.

177. On 11 June 2018, the attacker made a number of failed attempts to log in to the SCM database from Citrix Server 1. Most of these attempts failed because the attacker used invalid user-IDs. The attacker also attempted to use the D.A. account to log in to the SCM database, but this was unsuccessful because the account was not granted permission to the access the SCM database. It was on 11 June 2018 that Katherine, an IHiS database administrator, noticed some of the failed logins from that day.

178. The Citrix system event log for Citrix Server 1 was also deleted in the evening of 11 June 2018. The system event log is a set of Windows generated

logs categorised into 'Application', 'Security' and 'System'. These logs record events such as system boot-up, processes that have been started or stopped, and logins. In particular, the security event log would have captured the details of all the accounts that had logged in to Citrix Server 1. The deletion was not performed by any IHiS staff. It was presumably done by the attacker to cover its tracks.

179. On 12 June 2018, there were further failed logins from Citrix Server 1 to the SCM database. The error logs show that for most of these, the logins failed because they were from untrusted domains. For other attempts, the attacker used accounts which had not been granted access to the SCM database.

180. Based on the incidents of 11 and 12 June 2018, IHiS' Citrix administrators disabled logins to Citrix Server 1 on 12 June 2018, and shut down the server on 13 June 2018.

14.4.8 Attempts to log in to the SCM database from Citrix Server 2 and Citrix Server 4 on 13 June 2018

181. The events discussed in this section are summarised in the following figure:

Figure 9: Attempts to log in to the SCM database from Citrix Servers 2 and 4

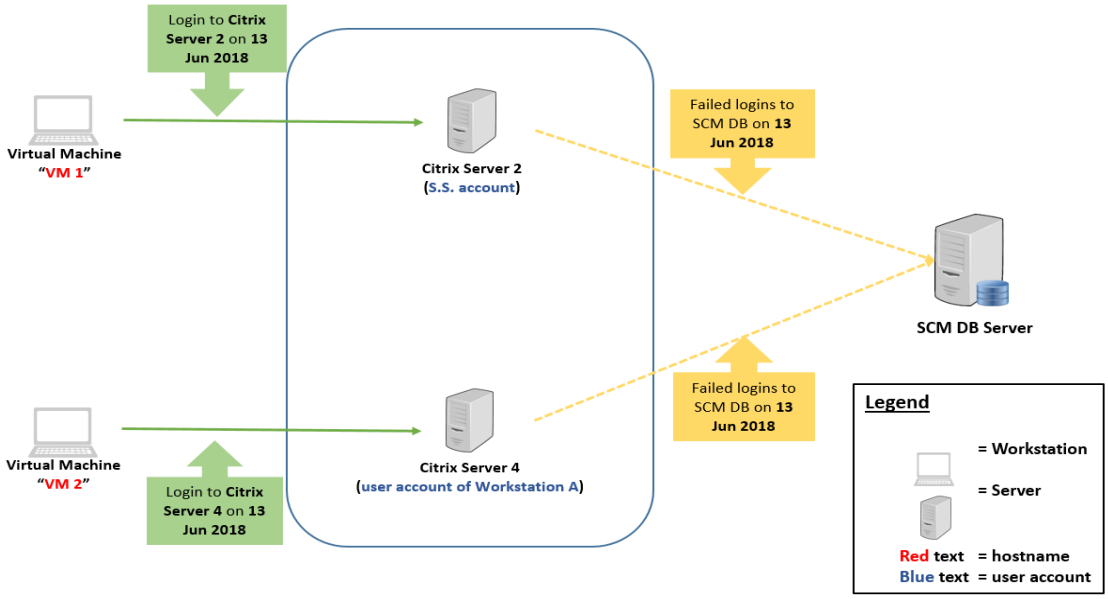

Citrix Server 2

182. On 13 June 2018, the attacker used a compromised local service account, the S.A. account, to remotely log in to Citrix Server 2, which was an SGH Citrix server. VM 1 was used to log in to Citrix Server 2, and these were not legitimate logins.

183. In the afternoon of 13 June 2018, a number of failed attempts were made to login to the SCM database from Citrix Server 2. These attempts failed because invalid user-IDs were used. In one attempt, the server name for a H-Cloud Citrix server (referred to in this report as "**Citrix Server 3**"), was used as a user-ID. Other attempts were made using the invalid user-IDs.

184. Later in the afternoon of 13 June 2018, another round of failed attempts was made to login to the SCM database from Citrix Server 2. Again, the server name for Citrix Server 3 was used as a user-ID in one attempt. The user-ID in another attempt was the name of a service account which would not ordinarily be used for the purposes of logging in to the SCM database. In yet another attempt, the attacker used a user-ID that it had used in a prior attempt to connect to the SCM database from Citrix Server 1 on 12 June 2018.

Citrix Server 4

185. In the afternoon of 13 June 2018, after the attempted logins from Citrix Server 2, the attacker used the account belonging to the user of Workstation A to remotely log in to another SGH Citrix server (referred to in this report as **"Citrix Server 4"**) from VM 2. A few minutes later, the attacker attempted to access the SCM database from Citrix Server 4, but this failed because the account used was not granted access to the SCM database.

14.4.9 Attempt to log in to the SCM database from Citrix Server 2 on 26 June 2018

186. On 26 June 2018, a failed attempt to connect to the SCM database from Citrix Server 2 was made using the account belonging to the user of Workstation A, from VM 2. Once again, this failed because the account was not granted access to the database.

14.4.10 Obtaining credentials of the A.A. account from Citrix Server 3 on 26 June 2018

187. The events discussed in this and the following section 14.5 are summarised in the following figure:

Figure 10: Obtaining credentials to the A.A. account and querying the SCM database

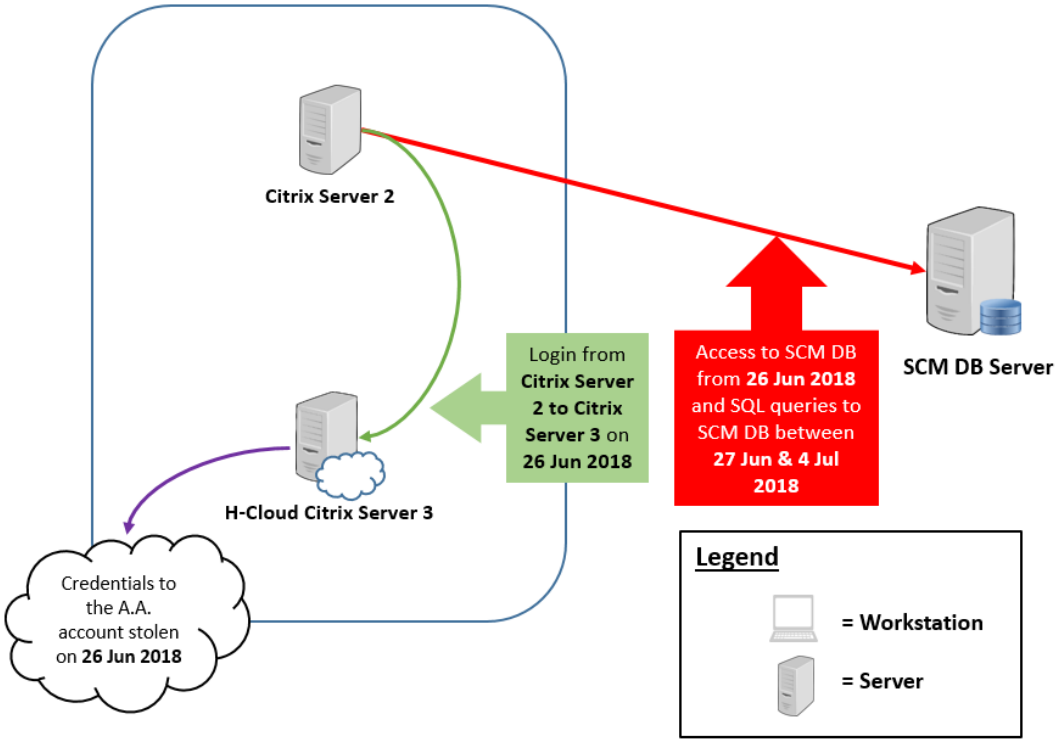

188. On 26 June 2018, the attacker remotely logged-in to Citrix Server 2 from Workstation B using the S.A. account. From Citrix Server 2, the attacker used the D.A. account to access a H-Cloud Citrix server, Citrix Server 3. While there is no conclusive evidence to show this, CSA assesses that it is probable that whilst logged into Citrix Server 3, the attacker stole credentials to an account referred to in this report as the "**A.A. account**". Obtaining the credentials to the A.A. account allowed the attacker to cross the last-mile to the SCM server, as it could be used to make SQL queries to the database.

189. CSA's assessment is that there was a coding vulnerability in the SCM application, and it is highly probable that this vulnerability allowed the attacker to easily retrieve the credentials of the A.A. account. Further details of this vulnerability will be discussed in section 15.6 (pg 86) below.

190. The lateral movement to Citrix Server 3 was significant because credentials of the A.A. account could not be obtained from the SGH Citrix Servers 1 and 2. This arose from the fact that the SGH servers were no longer being used actively to connect to the SCM database following the migration of the SCM application to H-Cloud Citrix servers in July 2017.

191. Lum has explained that connectivity between Citrix Server 2, which was an SGH server, and Citrix Server 3, a H-Cloud server, was present since June 2017 when the SCM system was migrated to the H-Cloud. The plan was to have all Citrix servers in both SGH and the H-Cloud form one logical farm, and the planned upgrade was scheduled for completion in September 2018. Only the ports that were required for the Citrix servers to communicate were left open. It was through this connection that the attacker was able to connect from Citrix Server 2 to Citrix Server 3.

192. With the credentials to the A.A. account, the attacker began the 'Actions on Objectives' phase as described in the Cyber Kill Chain, retrieving and exfiltrating patient data from the SCM database.

14.5 Queries to the SCM database from 26 June to 4 July 2018

193. From 26 June 2018, the attacker began querying the database from Citrix Server 2 using the A.A. account. Based on the evidence available, it appears that there were three broad types of Structured Query Language[21] ("**SQL**") queries which the attacker ran: (i) reconnaissance on the schema of the SCM database, (ii) direct queries relating to particular individuals, and (iii) bulk queries on patients in general. In total, the attacker performed over 200 SQL queries on the SCM database between 26 June 2018 and 4 July 2018.

194. The programs used to make the queries included programs that were legitimately used by IHiS, and also programs not used by IHiS and which were installed by the attacker. The hostnames from which the queries were logged as being made from were those of VM 1, VM 2, and Workstation B.

Reconnaissance on the schema of the SCM database and test queries

195. From 26 June 2018, the attacker began with reconnaissance queries which returned information relating to the schema of the SCM database, including information on database tables and views, stored procedures, and predefined SQL codes and functions. The purpose of this has been assessed by CSA to be to understand the SCM database and its design, before making queries on the data.

196. The attacker also executed test queries to understand the types of information in the database, and to confirm its findings from its reconnaissance work.

Direct queries relating to particular individuals

197. Thereafter, the attacker made a number of direct queries on specific NRIC numbers, including that of the Prime Minister Mr Lee Hsien Loong. The Prime

[21] Structured Query Language (SQL) is the standard language for relational database management systems, and is used to communicate with a database.

Minister's personal and outpatient medication data was specifically targeted and repeatedly accessed.

Bulk queries on patients in general

198. The attacker then made queries relating to patients in general, where no particular NRIC numbers were specified. IHiS staff detected the unusual queries on 4 July 2018. IHiS staff then terminated any subsequent bulk queries made on 4 July 2018, and took steps to prevent any similar malicious queries from being run against the SCM database. There was thus no further unauthorised access of the SCM database after 4 July 2018.

14.6 Exfiltration of data to overseas C2 servers

199. Between 27 June 2018 and 4 July 2018, the attacker was able to retrieve the following information from the SQL queries:

(a) The Prime Minister's personal and outpatient medication data;

(b) The demographic records of 1,495,364 unique patients, including their names, NRIC numbers, addresses, gender, race, and dates of birth; and

(c) The outpatient dispensed medication records of about 159,000 of the 1,495,364 patients mentioned in sub-paragraph (b) above.

200. From 27 June to 4 July 2018, the data was exfiltrated by the attacker *via* Workstation A to the attacker's C2 servers, as shown in the following figure:

Figure 11: Data exfiltration route

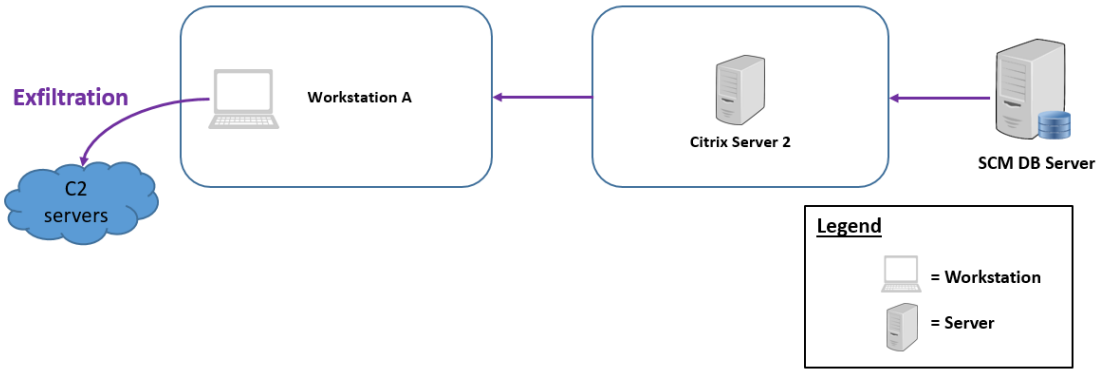

201. IHiS simulated the queries executed by the attacker and was able to approximate the data volume of the results. This was compared against data on the outgoing network traffic from Workstation A to an overseas C2 server from 27 June 2018 to 4 July 2018. The two data-sets clearly correspond to each other, and strongly indicates that part of the outgoing data contained the patient records accessed by the attacker.

202. IHiS has also confirmed that the size of the database query returns corresponds to the approximate size of 1.5 million patients' personal particulars and 159,000 outpatient dispensed medication records.

203. There is no evidence to show that patient records had been amended, deleted, or otherwise tampered with. Similarly, there is no evidence that other patient records, such as diagnoses, test results, or doctors' notes, were accessed. There was no disruption to healthcare services and patient care was not compromised.

204. The copying and exfiltration of data from the SCM database was stopped on 4 July 2018, after staff from IHiS discovered the unusual queries and took steps to prevent any similar queries from being run against the SCM database.

14.7 Attempts to re-enter the SingHealth Network on 18 and 19 July 2018

205. Although no data queries to the SCM database or exfiltration of patient records were detected after 4 July 2018, there was malicious activity in the SingHealth network on 18 and 19 July 2018, which suggested that: (a) that the attacker was trying to establish a fresh pathway into the network; and (b) that the attacker had established multiple footholds in the network and had re-entered the network through one of these hitherto unknown footholds.

206. On 18 July 2018, phishing emails were sent to a number of recipients in various SingHealth institutions. One of the recipients of the email was the user of a previously infected workstation – the PHI 1 Workstation. The email contained content similar to the earlier mentioned publicly available hacking tool, and would run automatically when the mail was previewed or read. It was also configured to lead to callbacks to a C2 server. IHiS discovered and informed CSA of the phishing emails on 1 August 2018, and the emails were assessed by CSA to be a possible attempt by the attacker to re-enter the network. The form and content of the emails also support the hypothesis that the initial breach could have been executed through a phishing email.

207. On 19 July 2018, IHiS informed CSA that a server, referred to in this report as the "**S.P. server**", was detected trying to connect to a C2 server, but the attempts were blocked by the firewall. On the S.P. server, malicious files were discovered.

208. There is no evidence of any callbacks to any known C2 servers from the S.P. server before 19 July 2018. The malicious files were created on the S.P. server on 19 July 2018, and the attacker would have required remote access to the SingHealth network in order to create these files. These facts indicated two things:

(a) First, the attacker had established multiple footholds in the SingHealth network, and had re-entered the system undetected

through one of these hitherto unknown footholds to create the malware, even while IHiS was actively implementing measures to contain the Cyber Attack and to monitor the system for malicious activity; and

(b) Second, the attacker was still active and trying to re-establish control of the network.

209. After detection of malware on and communications from the S.P. server, CSA recommended that internet surfing separation should be implemented, because this would be effective in preventing communications between elements in the SingHealth IT network and the attacker's C2 servers, thus preventing the attacker from exercising command and control over any remaining footholds it may have in the network. Internet surfing separation was implemented at 12:00am on 20 July 2018. No further signs of malicious activity were detected thereafter.

15 CONTRIBUTING FACTORS LEADING TO THE CYBER ATTACK

210. In the course of the enquiry, the Committee has heard of a host of pre-existing vulnerabilities, weaknesses, and misconfigurations that contributed to the Cyber Attack, in the sense that they were exploited or may have been exploited by the attacker in the course of the Cyber Attack. The Committee also heard evidence on circumstances which gave rise to or otherwise contributed to some of these vulnerabilities, weaknesses, and misconfigurations, and earlier opportunities in which some of them could have been remedied by IHiS prior to the attack. The Committee will present its findings on these matters in this section.

15.1 Network connections between the SGH Citrix servers and the SCM database were allowed

211. At the time of the Cyber Attack, network connections between SGH Citrix server farm to the SCM database server at HDC were allowed (this network connectivity has been referred to in the proceedings as the "open network connection"). The network connection was a critical pathway to the SCM database, over which the attacker was able to make SQL queries to and retrieve data from the SCM database. The Committee accepts the Solicitor-General's submission that but for this open network connection, the SCM database was adequately protected within the H-Cloud perimeter defences, and the attacker would not have been able to access the SCM database as easily.

212. These facts raise the issue of why the network connection was maintained. The Committee has heard evidence that during migration of the SCM system to the H-Cloud in June 2017, network connectivity between the SGH Citrix servers to the SCM database was required. After the migration in June 2017, the SCM infrastructure at SGH was decommissioned, but the network connection remained. This was because the SGH Citrix servers were used to host (i) administrative tools used for administering and managing SQL databases, including the SCM database in H-Cloud, and (ii) custom applications used by staff to query and retrieve data from the SCM database. These administrative tools and custom applications made use of the open network connection to perform their functions.

213. The administrative tools were hosted on the SGH Citrix servers as a matter of operational efficiency and not necessity. These tools were not used solely to administer the SCM database, but were also used to administer other SQL databases servers that were hosted in SGH and not H-Cloud. By hosting the tools on the SGH Citrix servers and maintaining network connectivity with the SCM database, the same set of tools could be used by administrators across all relevant SQL databases. Lum has clarified that this was not strictly necessary, as separate sets of tools could have been hosted on the H-Cloud Citrix servers (to service the SCM database), and on the SGH Citrix servers (to service the other SGH

databases). In elaborating on the inefficiencies that may result with having different sets of tools, Lum mentioned that "*the database administrator may have to manage different tools and may get confused at which one to launch*", and may end up being blocked by firewalls when attempting to use the wrong tool.

214. As for the custom applications, Lum has clarified that there were a few such custom applications. These applications were developed in-house and were not part of the Allscripts product. From a technical standpoint, the code base of some of these applications were dated, and some time would be required for their compatibility to be assessed before they could be migrated to the H-Cloud servers. While the applications could not have been migrated together with the SCM system in June 2017, there were plans to migrate these applications by September 2018. This was a deadline driven by the end-of-support for the software on the SGH Citrix servers. However, Lum has confirmed that with proper planning and resources, the applications could have been migrated earlier.

215. The Committee also notes that initially, IHiS had informed CSA that the SGH Citrix servers and the network connectivity were maintained to provide back-up connectivity to the SCM database. It was however clarified subsequently that this would not be technically possible, as the SGH Citrix servers, without the necessary upgrading, were not compatible with the latest version of the SCM application installed on the H-Cloud Citrix servers. The need for back-up connectivity was thus <u>not</u> a reason for maintaining the network connectivity between the SGH Citrix servers and the SCM database.

216. A basic security review of the network architecture and connectivity between the SGH Citrix servers and the SCM database could have shown that the open network connection created a security vulnerability. However, no such review was carried out. Woon Lan has confirmed that following the migration of the SCM system to H-Cloud, the network architecture of the SGH data centre was not redesigned. Ong has explained that network architecture "*is reviewed when there is a major change in infrastructure or needs*", and that "*(t)he SCM migration in June 2017 would not have involved a change in infrastructure in the SingHealth Sector*". It is surprising to the Committee why the migration of

the SCM database, a CII, to the H-Cloud, and accompanying migration of the SCM front-end application from the SGH Citrix servers to H-Cloud Citrix servers, was not seen as a "major" change meriting review of network architecture and connectivity.

15.2 Lack of monitoring at the SCM database for unusual queries and access

217. From 26 June to 4 July 2018, the attacker ran queries on the SCM database, including bulk queries. The attacker was able to do so unchallenged because of a lack of monitoring at the SCM database for unusual queries and access in at least two respects.

(a) First, there were no existing controls to detect bulk queries being made to the SCM database. While bulk queries are not uncommon as they are used for generating reports, the queries run by the attacker were anomalous in a number of ways. However, without controls in place to detect bulk queries and to identify anomalous queries, the the attacker was able to retrieve large amounts of data undetected.

(b) Second, one of the applications used by the attacker to query the SCM database was not a program that was legitimately used in the IHiS environment, and was not installed by IHiS on the SGH Citrix servers. This reveals a gap that was exploited by the attacker, namely, that there were no controls in place at the time of the attack to detect or block any queries to the SCM database made using illegitimate applications.

218. In the course of proceedings, the Committee has heard evidence on database activity monitoring ("**DAM**") solutions available on the market which could address some or all of the three gaps highlighted above. DAM was not implemented by IHiS at the time of the attack.

219. Mr David Koh, Chief Executive of CSA ("**CE, CSA**"), stated in his evidence that at the time of the attack, DAM was not common in the healthcare sector, but was common in both the security sector, and the banking and finance sector. Based on this, counsel for IHiS has submitted that the lack of DAM should not be viewed as an "*inherent weakness*" in SingHealth's network architecture, in light of the prevailing security posture in the healthcare sector at the time of the attack. The upshot of IHiS' submissions on this point is that it was not unreasonable for IHiS not to have implemented DAM at the time.

220. As discussed in the course of proceedings, the 'reasonableness' of IHiS' conduct in this respect is not in issue. What the Committee is concerned with is in (i) identifying the contributing factors (*i.e.* the lack of monitoring at the SCM database for unusual queries and access), (ii) identifying whether there was anything that could have been done better to address the vulnerability (*i.e.* implementing DAM), and (iii) the reasons, if any, why such steps were not taken.

221. It is in respect of this third issue that CE, CSA's evidence becomes relevant. The Committee notes that CE, CSA goes on to state in his evidence that the security and banking and finance sectors are *"(sectors) where database monitoring is commonly in place because of the mindset of the network designers*". The Committee is inclined to accept the Solicitor-General's view that the lack of security measures at the database-level to monitor for unconventional querying and access demonstrates that the need for such measures was *not* part of the consciousness of the network designers and operators for the SCM system at the time of the Cyber Attack.

15.3 SGH Citrix servers were not adequately secured against unauthorised access

222. The compromise of the SGH Citrix servers was critical in giving the attacker access to the SCM database. The Committee has heard evidence of a number significant security weaknesses concerning access to the SGH Citrix servers, some of which will be considered below.

15.3.1 Privileged Access Management was not the exclusive means for accessing the SGH Citrix servers, and logins to the servers by other means without 2-factor authentication were possible

223. Privileged Access Management ("**PAM**") is a means for organisations to restrict access to critical systems by privileged users, such as system administrators. As at the time of the Cyber Attack, PAM was implemented for both H-Cloud Citrix servers and SGH Citrix servers. This implementation required administrators to use 2-factor authentication ("**2FA**") in order to access servers.

224. In an internal risk assessment conducted by Wee and the IHiS Infrastructure and Applications teams around the end of 2016, "*Unauthorised Access and Account Theft (e.g. Stealing of Admin/User Accounts and Passwords)*" was listed as a threat or risk with a 'medium' likelihood of occurring, and a 'high' impact to business operations. The implementation of PAM was identified as an additional means of controlling this threat or risk.

225. However, the actual effectiveness of PAM was however severely undermined by the fact that it was not enforced as the exclusive means by which administrators could log in to the SGH Citrix servers. Even after PAM was implemented, IHiS' Citrix administrators were able to log in through an alternative route not requiring 2FA, and in fact preferred to do so.

226. Having less secure alternative routes would defeat the purpose of implementing PAM, as an attacker would simply exploit such alternative routes without having to concern itself with 2FA. Had PAM been the exclusive means of logging in to the SGH Citrix servers, the need for 2FA would have made it significantly more difficult for the attacker to move laterally and to gain privileged access to the Citrix servers.

227. The Committee has not heard any compelling reason why the alternative route was kept open. It is also of serious concern to note that IHiS Citrix administrators not only were aware of this alternative route, but knowingly made

use of it as a matter of operational convenience, when it would have been clear that this undermined the very purpose of implementing PAM and 2FA.

15.3.2 Lack of firewalls to prevent unauthorised remote access using RDP to the SGH Citrix servers

228. CSA's reconstruction of events show that the attacker had moved laterally using RDP to remotely access multiple SGH Citrix servers. This was done from compromised workstations and suspected virtual machines, and by using compromised user credentials. After compromising the SGH Citrix servers, the attacker was able to connect to Citrix Server 3 in the H-Cloud. The attacker also queried the SCM database from Citrix Server 2, a SGH server.

229. If RDP access from end-user workstations to the SGH Citrix servers had been disabled or restricted, it would have made it harder for the attacker to move laterally[22] and to compromise the SGH Citrix servers. However, at the time of the attack, there were no firewalls in place to prevent unauthorised remote access to the SGH Citrix servers using RDP.

230. This was not an unknown risk to IHiS. First, the HITSPS states that unecessary services including remote administrative access to servers and network devices should be disabled. Second, the need to enhance network segregation for administration access was in fact flagged-up in the FY16 GIA Audit Report of May 2017 (which stated the findings from and response to the FY 2016 H-Cloud Pen-Test) as a 'High^{H} Priority' issue, which in IHiS' risk classification framework meant that it was of a 'High' severity of impact, and had a 'High' likelihood of occurrence. The audit finding pointed to the possibility

[22] For completeness, CSA has clarified based on forensic findings that the attacker had also used other means to move laterally to the SGH Citrix servers. This means that even if RDP access from user workstations to the SGH Citrix servers were disabled or restricted, it would only have made it harder for the attacker to move laterally.

of accessing critical servers, including the Citrix servers that were connected to CII, from user workstations using RDP without any restriction imposed.

231. In response to the audit finding, IHiS decided that a combination of hardware and software firewall rules would be used to restrict RDP connections from the end-user segments to the SingHealth servers.

232. The SGH Citrix servers were deployed on a subnet which was not protected by a hardware firewall. Woon Lan has explained that while a hardware firewall was operational since January 2017 where the relevant Citrix servers were sited, the Citrix servers were not placed behind the firewall in view of plans to migrate them to H-Cloud. This was scheduled to be done by the end of FY2018 (*i.e.* April 2019). The security risk from not placing the servers behind the hardware firewall was recognised by IHiS, and the interim plan was to turn on the software firewalls in the servers.

233. However, software firewall rules to restrict RDP access were not enabled on the SGH Citrix servers either. Lum has explained that this was because the SGH Citrix servers were used to host a wide range of applications, some of which had complex requirements in terms of the ports they needed to access. If the built-in software firewall was enabled, it would be very difficult for staff to configure and manage the ports that had to be allowed in order for the various applications to function. Woon Lan has clarified that she was not aware that the software firewalls were not turned on for the SGH Citrix servers. While there is no written record evincing this, Woon Lan's evidence is that a decision had been taken around April 2017, further to a discussion between her and Nick Thoo (the IHiS Tower Lead for Network Services at the time), for the software firewalls to be enabled for the SGH Citrix servers. It is not clear whether this decision was communicated to the relevant staff, or if any steps were taken to confirm that the instructions were duly carried out.

234. In any case, the fact remains that as at the time of the attack, RDP access from user workstations to the SGH Citrix servers were not restricted by any hardware or software firewall. A person with the necessary account credentials

could remotely log on to the SGH Citrix servers from any workstation in any medical institution under SingHealthwithout 2FA or any other form of restriction on access. This was in spite of the known security risks, and the stated intent to remedy the specific risks identified in the FY16 GIA Audit Report.

15.3.3 Weak controls over and inadequate monitoring of local administrator accounts

235. As explained above, the password to the L.A. account was 'P@ssw0rd', which is easily cracked, and it is possible that the attacker gained control over the account by cracking the password. The L.A. account was also considered a 'dormant' account, which meant that it was an account that has been used before, but has not been logged into for the last 183 days.[23]

236. The weak password and the fact that the attacker was able to use the dormant account to access Citrix Server 1 were in spite of three relevant IHiS policies:

(a) First, the HITSPS states that user passwords are to be changed periodically. However, the password to the L.A. account was first set manually in 2012, and remained the same until it was changed on 11 June 2018.

(b) Second, in 2017, IHiS instituted a policy under which administrators were required to have more complex passwords. This policy applied to the L.A. account, but its password remained unchanged.

(c) Third, in-line with paragraph the HITSPS, dormant or unused accounts should be identified and disabled, in order to prevent

[23] More than 183 days had passed since the last legitimate use of the L.A. account in Citrix Server 1 on 13 October 2017, and the first instance of unauthorised use by the attacker on 17 May 2018.

usage in unauthorised activities. However, this was not done in the case of the L.A. account.

237. The issue of weak passwords for domain or privileged user accounts was flagged-up in the FY16 GIA Audit Report as a 'High^{H} Priority' issue, which in IHiS' risk classification framework meant that it was of a 'High' severity of impact, and had a 'High' likelihood of occurrence. In fact, one of the weak passwords identified in the course of the H-Cloud Pen-Test was the same "P@ssw0rd", which was used for another account. The password policy in paragraph 236(b) above was also instituted in response to this audit finding.

238. The management response from IHiS to the FY16 GIA Audit Report finding included a comment that passwords for *active directory* administrator accounts had been changed in-line with the new password policy by 21 March 2017. Both Lum and Woon Lan have recognised that there was no explicit mention of the need to change the *local administrator* account passwords to meet the new requirement, explaining that it did not occur to them at the time the management response was being discussed.[24]

239. On 21 March 2017, Woon Lan sent an email to the then-System Management Department, which included the Citrix administrators, directing recipients to change passwords for their "privileged accounts". Once again, there was no explicit mention of the need to change all local administrator account passwords. Likewise, in subsequent follow-ups with the GIA, the issue of local

[24] The Committee notes that Lum has stated in his conditioned statement that he had "*instructed Ping Hai and Ji Han to change the local admin password*", through an email. This email was sent in March 2017. The relevant section reads "*As mentioned this morning to all of you, we need to immediately "clean up" those password things that were flagged up. As a precaution, please reset your individual Citrix admin password and also the local admin password that we have exposed due to our own negligence.*" Viewed in context of the FY16 H-Cloud Pen-Test, the local admin password in question was an account belonging to a Citrix administrator. This direction does not appear to be a direction for all local administrator account passwords to be changed, and does not clearly indicate the Lum had in fact specifically considered applying the new password policy to the local admin accounts.

accounts did not arise, and there was no verification on whether the policy was implemented in respect of these accounts.

240. Evidently, the need to apply the same password policies to local administrator accounts was overlooked by the Citrix Team and Woon Lan. In addition to this oversight, IHiS' usual approach of implementing and enforcing password policies did not apply to local accounts for the SGH Citrix servers.

241. Password policies are usually effectuated in IHiS through the use of the Group Policy Object ("**GPO**"), which automate the implementation and enforcement of policies. GPOs should apply to all servers by default, except for groups of servers which have the 'block policy inheritance' setting applied. Applying 'block policy inheritance' prevents group policies from being inherited from these servers. The SGH Citrix servers were part one such group of servers which had group policy inheritance applied. As such, the GPOs implementing the complex password policy and policy for the deactivating of dormant accounts was not applied to the L.A. account.

242. Lum has explained that the password was not meant to expire because it was the local administrator account that would be used as a last resort for accessing the server if administrators were unable to use their active directory domain administrator accounts for whatever reason. It is not apparent to the Committee how any of the password policies mentioned above would necessarily prevent the use of the account as a back-up means of access, since all that is required is a proper process to be put in place to manage the change in passwords or disabling of the account due to it being dormant.

243. It also bears mention that the L.A. account was last legitimately used on 13 October 2017, after the institution of the new password policy. While no evidence has been led on this particular point, it appears that the administrator who had used the account and presumably keyed in the password paid no heed to the fact that the password was against IHiS' policies.

15.3.4 Lack of sight over and mismanagement of the S.A. service account

244. As explained above, the S.A. account was used by the attacker to access Citrix Server 2, including when querying the SCM database. The existence of and privileges attached to the account facilitated this use. From the evidence, the Committee finds three points that are relevant in this regard.

(a) First, to begin with, there was no real need for the S.A. account to exist, as there was no actual use in IHiS of the relevant service for which it was created. Yet it existed on all Citrix servers in which the service had been installed, and the account had full administrative privileges to login to the server, including logging in interactively.

(b) Second, the Citrix Team did not know of this account. Lum's evidence is that he had first come to know of the account on 13 June 2018, after the Citrix Team discovered that the account was used in unauthorised logins to Citrix Server 2.

(c) Third, the S.A. account was an unused account that should have been identified and disabled in accordance with IHiS' policies. An 'unused account' refers to accounts that were created but never logged into. As mentioned above, unused accounts should be identified and disabled in-line with the HITSPS, in order to prevent usage in unauthorised activities. This however was not done. The GPOs for password policies also did not apply to the account as 'block policy inheritance' was applied.

15.3.5 Observations on the overall management of SGH Citrix servers

245. A number of weaknesses in respect of securing the SGH Citrix servers against unauthorised access have been identified above. As the Solicitor-General has submitted, such failures likely stem from a failure to recognise the SGH Citrix servers as being part of a mission-critical system. While IHiS recognised the SCM system to be a mission-critical system, it did not regard the Citrix

servers as being part of this mission-critical system. Citrix servers were instead viewed simply as a *means* by which a mission-critical system is accessed, but are not themselves part of that system. This mindset on the lower criticality of the Citrix servers may have indirectly resulted in the vulnerabilities listed above. In addition, this mindset was expressed in the following two facts as well:

(a) The SGH Citrix servers were not monitored for real-time analysis and alerts of vulnerabilities and issues arising from these servers.

(b) Vulnerability scanning, which was carried out for mission-critical systems, was not carried out for the SGH Citrix servers. Vulnerability scanning is an inspection of the potential points of exploit on a computer to identify gaps in security. In the context of IHiS, the rules prescribed in vulnerability scanning included their internal security policies on issues such as minimum password lengths. Thus, if vulnerability scanning of the SGH Citrix servers had been carried out, the fact that the L.A. had a weak password that did not comply with IHiS' password policies would have been identified. In a similar vein, the S.A. account would have been detected as an unused account.

246. There are also clear indications of poor cyber hygiene and a lack of security consciousness on the part of the Citrix administrators. This is clearly seen in examples such as failing to change the password for the L.A. account, and the deliberate use of alternative methods to avoid PAM when logging in to the Citrix servers. Further examples evincing poor cyber hygiene and a lack of security consciousness will be covered in section 15.7 (pg 89) below, where the Committee presents its findings in respect of other weaknesses that were identified in the FY16 H-Cloud Pen-Test.

15.4 Internet connectivity in the SingHealth IT network increased the attack surface

247. The SingHealth network's connection to the Internet, while serving their operational needs, created an avenue of entry and exit for the attacker. This allowed the attacker to make use of an internet-connected workstation (Workstation A) to gain entry to the network, before making his way to the SCM database to steal the medical data.

248. The Committee has also heard examples of security concerns arising from internet connectivity on certain network elements that were involved in the attack:

(a) SGH Citrix servers: At the time of the attack, a user who accessed pre-configured internet websites through the SGH Citrix servers would be able to access websites other than the pre-configured sites simply by keying in the internet URL in the address bar of the web browser. If such other websites were malicious, it would be possible that malware would be downloaded onto the SGH Citrix server.

(b) The S.P. server: As mentioned in section 14.7 (pg 70) above, the S.P. server was detected trying to connect to a C2 server on 19 July 2018. Investigations revealed that the S.P. server was put to two uses: first as an intranet document repository for SGH users; and second as an internet web server hosting SGH websites. Leong Seng was unable to explain why the S.P. server was used both as a web server and an intranet server. The placement of the server in the local server zone was also a cause for concern – Leong Seng has clarified that intranet servers should be placed in an internal server zone with no connection to the internet. The implication of this appears to be that if the attacker fully compromised the S.P. server, it would have gained a foothold within the local server zone.

249. The background relating to the formulation of the healthcare sector's internet access strategy and the steps taken towards its implementation will be discussed in greater detail in section 48.1 (pg 390) below. For present purposes, it is sufficient to note the following:

(a) The security risks arising from internet-connectivity in the SingHealth network were raised by CSA to MOH from as early as August 2015;

(b) By June 2017, the healthcare sector had determined, among other things, that (i) internet access would be removed for staff that did not require the internet for work, and (ii) for staff that required the internet for work, access would be through a secure internet access platform which, at that time, was to take the form of a 'remote browser'.

(c) When the Cyber Attack occurred, the remote browser solution was not yet rolled out. IHiS was on the cusp of awarding the tender for the remote browser solution in July 2018 when the Cyber Attack occurred, and the award of the tender was consequently put on hold.

250. Thus, at the time of the Cyber Attack, while an internet access strategy to reduce and mitigate the risks posed by internet connectivity in the SingHealth IT network had been *formulated*, it had not been *implemented*.

15.5 Versions of Outlook used by IHiS were not patched against a publicly available hacking tool

251. A publicly available hacking tool played an important role in the compromise of Workstation A (see section 14.2 (pg 54) above). The attacker was able to install the hacking tool on Workstation A on 1 December 2017 by exploiting a vulnerability in the version of the Outlook application installed on the workstation.

252. A patch that was effective in preventing the vulnerability from being exploited (and thus to prevent the installation of the tool) was available since late-2017. Leong Seng has explained that software security patches are applied on SingHealth and IHiS issued endpoints based on a specified posting cycle, except for critical patches addressing serious vulnerabilities, which would be applied as soon as possible. The patch was scheduled to be rolled out as part of IHiS' regular patching cycle, but the patch had not been applied to Outlook on Workstation A as at 1 December 2017.

253. Counsel for IHiS has submitted that IHiS' conduct in respect of the patching cycle for Outlook was "*reasonable*", and that it was "*entirely fortuitous*" for the attacker to have executed the hacking tool within the period between the release of the patch and its application. Once again, the 'reasonableness' of IHiS' conduct in this respect is not in issue. What the Committee is concerned with, and has found, is that the hacking tool was installed on Workstation A by exploiting a vulnerability on Outlook, that a patch was available since late-2017 but was not applied at the time the hacking tool was installed on 1 December 2017, and that the patch was scheduled to be rolled out as part of the regular patching cycle. The Committee's recommendations on improving software upgrade policies are found in section 47 (pg 381) below.

15.6 Coding vulnerability in the SCM application

254. CSA's analysis of the SCM application showed that there were signs of insecure coding practices, giving rise to a vulnerability that was likely exploited by the attacker to obtain the credentials to the A.A. account.

255. Sometime in September 2014, a then-employee of IHiS, Zhao Hainan ("**Zhao**"), discovered a method of exploiting the vulnerability. Zhao informed his immediate superior, Angela Chen ("**Angela**"), about some of his findings on or about 15 September 2014. There is some inconsistency in the evidence as to the specifics of what Zhao told Angela. It is however clear that Zhao did *not* inform Angela about the technical details of his findings, or the precise fact that credentials could be obtained.

256. Angela also gave evidence that she asked Zhao to log a case with Allscripts, but she did not follow-up with him on whether he had in fact done so. Zhao's evidence is that Angela had asked him to provide feedback to the "architecture team", but he did not know who she meant by this. More pertinently, Zhao did in fact know that he could log a case with Allscripts, but presumed that Allscripts "*(would) not realise the importance of all this*", and thus did not log a case with Allscripts.

257. Zhao's actions must be viewed in context of his other action of independently sending an email to Epic Systems Corporation ("**Epic**"), a competitor of Allscripts, on 17 September 2014. The subject of the email was "*Allscripts Sunrise Clinical Products can be hacked easily*", and the email read:

> Dear Epic,
> There's a loophole in Allscripts Sunrise Clinical Manager products, where user can gain admin control of the whole database easily. The user can be just a medical student, nurse, pharmacist. This lies in their architecture of the product. Note the market share of Sunrise Clinical Manager in US hospitals, this could lead to a serious medical data leak, or even a national security threat.
> As a competitor, I am not sure whether you can leverage on this to gain more market share. Contact me if you guys are interested.
> Regards,
> HZ

258. On 18 September 2014, David Chambers, who is in charge of Allscripts' businesses in Asia, wrote to Dr Chong Yoke Sin ("**Dr Chong**"), the CEO of IHiS at the time, informing her of Zhao's email, and impressing that Allscripts was "*treating this as a very serious matter*".

259. Dr Chong tasked Foong Lai Choo (the then-Director of the Core Apps 1 Department, which Zhao was part of) ("**Lai Choo**") and Kua Cheong Kee Clarence (the Applications Service Lead for SingHealth systems, including the SCM system) ("**Clarence**") to verify whether Zhao was in fact the one who sent

the email to Epic. Having ascertained that he was the sender, Dr Chong immediately terminated Zhao's employment by 5:00pm on 18 September 2014.

260. There is some inconsistency in the evidence on whether Zhao communicated additional details about his findings to Lai Choo and Clarence, when they met him to ascertain if he sent the email. But what is undisputed is that no action was taken by IHiS to formally investigate, assess, or rectify the alleged vulnerability.

261. Dr Chong's evidence is that she "*considered this [matter concerning Zhao] to be primarily a disciplinary issue, and not an IT security issue*". On the alleged vulnerability, Dr Chong's evidence is that she, Clarence, and Lai Choo thought that the alleged vulnerability would be "*irrelevant*" following recent upgrades to the SCM system architecture, or that the alleged vulnerability was in fact a "*well-documented*" problem with Microsoft's SQL server and not the SCM itself, and which "*could be addressed by additional layers of security*". Since no steps were taken to investigate further, these views were unverified assumptions.

262. Later in the evening on 18 September 2014, Dr Chong wrote back to David Chambers, informing him that Zhao had been dismissed. Dr Chong also stated that "*My technical people have investigated the subject mentioned and concluded that the 'exposure' is a normal programming of codes to extract data from the database, which is done as a normal course of work.*" Dr Chong has explained that the "technical people" she referred to were in fact Lai Choo, Clarence and their staff. Dr Chong has also confirmed that the explanation given in her email was an expression of opinion, and there was in fact no formal inquiry conducted. No further steps were taken by IHiS in relation to this incident after this email was sent.

263. While the SCM vulnerability was not the sole contributing factor in the Cyber Attack, it likely played a pivotal role in allowing the attacker to obtain the SCM database credentials and cross the last mile to gain access into the SCM database. IHiS has accepted that if further queries and investigations had in fact been carried out, the coding vulnerability could have been discovered. In this

respect, the Committee agrees with the Solicitor-General's submission that the events concerning Zhao in September 2014 was a missed opportunity.

264. Investigations by the CID did not reveal any evidence of Zhao being involved in the Cyber Attack.

15.7 Other vulnerabilities in the network that were identified in the FY16 H-Cloud Pen-Test which could have been exploited by the attacker for privilege escalation and lateral movement

15.7.1 Administrator credentials were found on network shares

265. The FY16 H-Cloud Pen-Test revealed that administrator credentials were found in network shares. A Citrix administrator password was also found in a Windows batch file. The implication of this was that attackers having access to such files, or with physical or network access to shared folders, could read this sensitive information and further use it to perform enhanced focused attacks.

266. In the course of investigations, Citrix Server 1 was found to contain a batch file with administrator credentials in it. The batch file was created on 9 April 2017 and contained the administrator credentials of the L.A. account in cleartext. This remained available on the server until the server was taken offline for forensic imaging on 13 June 2018. CSA has given evidence that it is a reasonable hypothesis that the attacker gained initial access to the file system of Citrix Server 1, and obtained the credentials for the L.A. account, which were saved in the batch file in this server.

267. Similarly, during a scanning process done after the Cyber Attack, a script file containing credentials for an administrator account was found, which had the password 'P@ssw0rd'. This was in fact the very same account flagged by the penetration testers during the FY16 H-Cloud Pen-Test.

268. Back in March 2017, after being informed of the findings from the FY16 H-Cloud Pen-Test, Lum sent an email to the Citrix Team, directing them to

"*clean up*" any existing files containing admin credentials. He also instructed the team to enforce stringent controls such files and the folders in which they were stored. Finally, he impressed on the team that they should take these matters seriously, and that everyone in the team had to take ownership of the issues raised. Evidently, his exhortations went unheeded, given that the batch file discussed in paragraph 266 above was created shortly after on 9 April 2017.

269. Similarly, in March 2017, Woon Lan instructed all administrators to "*comb through*" their files to "*ensure there is no hardcoded password*". Woon Lan has explained that by "*combing through*", she had in mind the administrators checking through every server. Her thinking was that if the administrators had developed such scripts, they would know where the scripts were saved on the servers.

270. IHiS' management response, as stated in the GIA Internal Audit Report from May 2017, was that IHiS had "*Completed housekeeping of scripts in the server*". Woon Lan has explained that in making this response, she meant that the specific server flagged-up in the pen-test had undergone housekeeping. However, this response was given in spite of the fact that neither Woon Lan nor Lum had taken any steps to verify if their directions above had in fact been performed by the Citrix Team across *all* Citrix servers.

15.7.2 The Citrix virtualisation environment was not configured adequately to prevent attackers from breaking out into the underlying operating system

271. The penetration testers uncovered that the Citrix virtualisation environment used was not configured adequately to prevent attackers from breaking out of the virtualisation and into the underlying operating system. Exploiting the vulnerability allowed the penetration testers to access files and execute arbitrary commands. CSA's hypothesis is that this vulnerability could have been the means by which the attacker gained initial access to the file system of any of the compromised SGH Citrix servers.

272. In relation to this vulnerability, IHiS had indicated in its management response that it would lock down the Citrix server farm. However, the lock down was only carried out for the new Citrix farm in H-Cloud, and not for the SGH LDC. This meant that the vulnerability continued to be exploitable for the SGH Citrix servers at the time of the Cyber Attack.

15.7.3 Observations on the remediation of vulnerabilities identified in the FY16 H-Cloud Pen-Test

273. The FY16 H-Cloud Pen-Test was conducted in early 2017, and a number of vulnerabilities were identified. The vulnerabilities identified by the penetration testers should have been remediated at the time of the Cyber Attack, given that IHiS had been informed of the observations from the penetration test as early as March 2017, well before the various weaknesses were exploited in the Cyber Attack. Unfortunately, the remediation process undertaken by IHiS was mismanaged and inadequate, as is evident from the findings on issues such as(i) weak domain/privileged users' passwords; (ii) administrator credentials found on network shares; (iii) poor network segregation for administration access; and (iv) the Citrix environment compromise issue.

274. To make matters worse, some issues were reported by the management of the IHiS Infrastructure Services Division at the time (*e.g.* the Citrix Team, led by Lum and the Data Centre Services Tower Lead, Woon Lan, and Security Services Tower Lead, Ernest) to the GIA as having been resolved by the time the Internal Audit Report was published on May 2017, without first taking steps to verify if they were in fact resolved, or considering carefully if the steps taken were adequate. Clear examples are the cases involving weak domain/privileged users' passwords, and administrator credentials found on network shares, where the remediation that was done for these items were limited to the particular accounts or servers that were identified by GIA, and no thought was given to implement the same measures on all other local accounts and across all other Citrix servers.

275. In spite of the inadequacy of the measures taken, these audit items were marked in the Internal Audit Report as having been completed. The Internal

Audit Report was sent to Bruce and members of IHiS' and SingHealth's senior management. The understanding given at the SingHealth Audit Committee and IHiS' Audit Risk Committee meetings was also that these audit items had been resolved. No questions were raised at any level about the adequacy of the measures taken. Likewise, no major questions were raised at any level about the adequacy of any other measure which the management of the IHiS Infrastructure Services Division had proposed for the purposes of addressing the other audit findings.

276. As a result, from May 2017 to the time of the attack, organisationally, IHiS and SingHealth held the mistaken belief that some of the audit items had been adequately resolved, and that the remaining items would likewise be adequately resolved. As the findings above show, this was not the case.

277. It also bears mention that similar vulnerabilities were surfaced in further penetration tests conducted by the GIA in FY2017 at three local sites. The IT systems of these three sites are managed by IHiS as well. The repeated findings of similar weaknesses are particularly concerning given that these penetration tests were conducted in FY2017, *after* the findings of the FY16 H-Cloud Pen-Test were published. Evidently, the lessons learnt were not applied.

278. In sum, the internal audit discovered a number of vulnerabilities in the SingHealth network, and several of these vulnerabilities were present during the Cyber Attack, as IHiS had failed to properly implement adequate remediation measures. CSA found that these vulnerabilities could have been exploited by the attacker, and also noted that these were not *necessarily* the vulnerabilities exploited, given that the attacker could have achieved its ends through other means as well. Nevertheless, the fact remains that the failure to properly remediate these vulnerabilities, gave the attacker these *additional opportunities* through which it could compromise the SingHealth network. The failure to remediate likely made the attacker's path through the SingHealth network to its ultimate objective, the SCM database, *easier*.

16 THE ATTACKER – TOOLS AND COMMAND AND CONTROL INFRASTRUCTURE

279. In the preceding section, the Committee presented its findings on the contributing factors which allowed the attacker to achieve its objectives more easily. In the next two sections, the Committee will present its findings on the attacker – its tools, command and control infrastructure, and profile as a skilled and sophisticated threat actor.

16.1 Customised and stealthy malware

280. The attacker made extensive use of advanced, customised, and stealthy tools throughout the attack, which effectively overcame and evaded the antivirus software and conventional security defences that were in place. The malware samples CSA analysed were either (a) unique variants that were not seen in-the-wild, and had not been detected by the standard anti-malware solutions deployed by SingHealth, or (b) a mix of open source tools that were modified to provide stealth for the attacker.

281. A variety of custom web shells, tools, and unique malware were used in the attack. Early-stage tools were used to gain a foothold within the network. Intermediate-stage tools, including some custom tools, were used to perform various tasks such as reconnaissance, privilege escalation and lateral movement. Remote Access Trojans, such as the abovementioned RAT 1 and RAT 2, were used to provide the attacker with full control over specific infected systems and to serve as backdoors to re-enter the network. The wide range of tools and the fact that many of them were customised indicates that the attacker was well resourced, and possessed or was supported by developmental capabilities.

282. Notably, during the incident response, malware samples were given a cybersecurity company to develop malware signatures. The firm's software was initially unable to detect the samples as being malicious. After CSA shared their initial malware analysis findings with the company, it was able to develop malware signatures in their antivirus solution for mass network-wide scanning.

16.2 Extensive C2 infrastructure

283. CSA's forensic analysis of the exhibits revealed a number of network Indicators of Compromise ("**IOCs**") which appeared to be overseas C2 servers. CSA has explained that generally, the C2 servers were used for:

(a) Infection: where the server is used as a means of dropping malware into the system it is trying to infect;

(b) Data exfiltration: there were indications of technical data (and not medical records) being sent to the servers; and

(c) Beacon: infected machines may have connected to C2 servers to establish a 'heartbeat', which refers to a slow, rhythmic communication meant just to sustain communications.

284. The CSA furnished the details of a number of overseas network IOCs to the CID for follow-up to determine if the subscribers' information could be ascertained. Direct requests were made to foreign law enforcement agencies for the relevant information.

17 PROFILING THE ATTACKER

285. CSA has assessed that the attacker was a "*skilled and sophisticated*" threat actor, that had "*characteristics that are typical of an Advanced Persistent Threat* ("**APT**") *attack*". CSA has also provided the following description of an APT:

> APT refers to a class of sophisticated, usually state-linked, cyber attackers who conduct extended, carefully planned cyber campaigns, to steal information or disrupt operations. APT attackers are known to be extremely persistent in finding ways to get into a network/system once a target had been identified.

286. The Committee agrees with CSA's assessment of the attacker as skilled and sophisticated attacker bearing the characteristics of an APT group, having regard in particular to the following attributes seen from the evidence presented before the Committee:

(a) **The attacker had a clear goal in mind**, namely the personal and outpatient medication data of the Prime Minister in the main, and also that of other patients. CSA has assessed that the attacker's actions were targeted and specific, conducting reconnaissance in the network targeted at reaching the SCM database, and compromising only selected computers necessary to access, copy, and transfer data from the SCM database. The attacker also avoided secondary targets that might have drawn attention to its presence. The attacker also effected a quick turnaround time between access to the SCM database and exfiltration of data from the SCM database, showing both technical competence and mission-orientation.

(b) **The attacker employed advanced tactics, techniques, and procedures**, as seen from the suite of advanced, customised and stealthy malware used, generally stealthy movements, and ability to find and exploit various vulnerabilities in SingHealth's IT network and the SCM application. CSA has highlighted that network intrusion techniques with low attack signature are a hallmark of an advanced threat actor. Apart from evading detection for almost 10 months from 23 August 2017, the attacker was conscientious in erasing logs on compromised workstations and servers. Notably, the attacker even re-entered the network after being detected, to erase system and program logs.

(c) **The attacker was persistent**, having established multiple footholds and backdoors, carried out its attack over a period of over 10 months, and made multiple attempts at accessing the SCM database using various methods. It is particularly noteworthy that

even after its attack was thwarted on 4 July 2018, the attacker re-entered the system on 19 July 2018 through an earlier established foothold and sought to re-establish control over the network (see section 14.7 (pg 70) above).

(d) **The attacker was a well-resourced group**, having an extensive C2 network, the capability to develop numerous customised tools, and a wide range of technical expertise.

287. Our cyber defences will never be impregnable. The skill and sophistication of the attacker has been recognised by the Solicitor-General, CSA, and all the interested parties. The expert witnesses also noted that an APT, given enough time, will breach the perimeter of any network. However, it is vital to note that while it may be difficult to prevent an APT from breaching the *perimeter* of a network, the success of the attacker in obtaining and exfiltrating the data in this attack was not inevitable. In this regard, the Solicitor-General has rightly pointed out two key considerations:

(a) First, the attacker was stealthy but not silent, and signs of an attack were observed. As will be discussed in the next Part, these signs were not acted upon either because of: (i) the relevant staffs' inability to recognise that an attack was ongoing; or (ii) inaction on the part of the staff responsible for responding to attacks. Had they taken appropriate action, the attacker could have been stopped before it achieved its objectives.

(b) Second, as explored in this Part, there were vulnerabilities, weaknesses, and misconfigurations in the SingHealth network and SCM system that contributed to the attacker's success in obtaining and exfiltrating the data, many of which could have been remedied before the attack. Doing so would have made it more difficult for the attacker to achieve its objectives.

Part IV – Incident response by IHiS up to 10 July 2018

TABLE OF CONTENTS – PART IV

18 PRELIMINARY MATTERS

18.1 Introduction to this Part

288. In this Part, the Committee presents its findings in respect of TOR #2 for events up until 10 July 2018, when CSA was notified. Although TOR #2 refers to establishing how IHiS *and* SingHealth responded to the Cyber Attack, the facts show that the incident response up until 10 July 2018 was within the domain of IHiS, and there was no involvement of SingHealth in this period.

289. The Committee's findings on the events are largely set out in a chronological fashion in order to better reflect the sequence of events and the state of mind of the persons involved, and to better contextualise their acts and omissions. Following the account of the events of each day or period of days, and where appropriate, the Committee will also provide its assessments of the incident response by the persons involved. In the course of making its findings and assessments, the Committee will highlight facts and issues that would subsequently inform the recommendations that the Committee makes in respect of TORs #3, #4, and #5.

290. In making its findings, the Committee will also highlight, based on CSA's evidence, the various points prior to 10 July 2018 where CSA ought to have been informed in accordance with the NCIRF. Had these 'missed opportunities' been taken up by IHiS, CSA could have been involved before the unauthorised access to the SCM database began on 26 June 2018, and the attack could have been prevented or its impact significantly mitigated. These missed opportunities are similarly instructive for the recommendations that the Committee makes in respect of TORs #3, #4, and #5.

18.2 Key witnesses from IHiS and SingHealth

291. At this juncture, it is useful to set out the key witnesses involved in the response to the Cyber Attack, grouped according to their roles in the incident response.

Witness marking	Name	Designated Role	Roles in response to the Cyber Attack
IHiS IT administrators			
W2	Lum Yuan Woh	Assistant Director, Systems Management Department, IHiS	Citrix Team lead, led efforts in investigating into suspicious activities on the Citrix servers
W3	Katherine Tan Seang Lim	Senior Manager, Systems Management Department, IHiS	SCM database administrator, noticed attempts to log in to SCM database on 11 June 2018
W4	Chai Sze Chun	Assistant Lead Analyst, Production Enhancement Team, IHiS	Member of the SCM Application Team, Detected unusual query to SCM database on 4 July 2018
IHiS Security Management Department			
W7	Benjamin Lee Yi Ren	System Engineer, Security Management Department, IHiS	Investigated suspicious activity detected in a PHI 1 Workstation in January 2018, later also investigated the incidents taking place in June/July 2018
W8	Tan Choon Kiat Ernest	Senior Manager, Security Management Department, IHiS	SIRM for SingHealth, was meant to lead and coordinate IT security incident response
W9	Wee Jia Huo	Cluster Information Security Officer for SingHealth, IHiS	Cluster ISO for SingHealth, was accountable for the actions of the incident response team

Witness marking	Name	Designated Role	Roles in response to the Cyber Attack
IHiS management in-charge of matters concerning the SCM system			
W12	Yong Cheng Pei Serena	Director, Infrastructure Services Division, IHiS	Together with Clarence, escalated incident to IHiS senior management on 9 July 2018
W17	Kua Cheong Kee Clarence	Application Service Lead for SingHealth's clinical systems[25] (employed by IHiS)	Together with Serena, escalated incident to IHiS senior management on 9 July 2018
W26	Ong Leong Seng	Director, Delivery Group, IHiS	In-charge of the War Room set-up to deal with the Cyber Attack
IHiS Senior Management			
W27	Benedict Tan Wee Bor	Group Chief Information Officer, SingHealth (employed by IHiS)	Escalated incident to Bruce, Kim Chuan, and SingHealth senior management
W28	Chua Kim Chuan	Director, Cyber Security Governance, IHiS, and concurrently Chief Information Security Officer, MOH	Reported incident to CSA
W29	Bruce Liang Chwee Bock	Chief Executive Officer, IHiS, and concurrently Chief Information Officer, MOH	Reported the incident to MOH and MOHH, and oversaw the technical response to the attack

[25] Clarence is concurrently the Deputy Director of the Chief Information Officer's Office in SingHealth. He is an IHiS employee.

Witness marking	Name	Designated Role	Roles in response to the Cyber Attack
SingHealth Senior Management			
W31	Kwek Yung Chiang Kenneth	Deputy Group Chief Executive Officer (Organisational Transformation and Informatics), SingHealth	Took direct charge of patient outreach and communications efforts.

18.3 Knowledge of and preparedness against APTs as at June 2018

292. In order to properly assess the incident response, it is necessary to first ascertain the extent of knowledge that IHiS and SingHealth had of APTs at the time of the attack, and who had such knowledge.

293. IHiS has informed the Committee that they were alive to the risk of APTs from as early as August 2016, and had begun sourcing for an Advanced Threat Protection ("**ATP**") solution at around that time to address this threat. Bruce has explained that the deployment of ATP was originally scheduled for FY2017, but they faced delays in finding a suitable vendor. Eventually, the vendor was identified in June 2018, but the ATP solution was not yet implemented throughout the period of the Cyber Attack.

294. Towards the end of 2016, the Cluster ISO for SingHealth, Wee, prepared a risk assessment report for the SCM system. This risk assessment report, titled "*SHS & EHA IT Security Risk Assessment for Critical Information Infrastructure System*" (the "**FY16 CII Risk Assessment**") was dated 3 January 2017. The threat of APTs was flagged in two respects:

(a) First, at Item 7, the threat of "*Malware Attacks (Virus, Worms, Trojans, Rookits, Advanced Persistent Threats, etc.)*" was

identified[26], and it was noted that there was "*(l)imited protection against advanced persistent threats*" in place.

(b) Second, at Item 13, the threat of "*Cyber Extortion/Ransom e.g. theft of patient's medical record*" was identified[27], and it was noted that there was "*(l)imited protection against advanced persistent threats*" in place.

295. In response to the above two threats, the proposed control measure was for "*(i)nfrastructure Services – Cluster Infra[structure] Services to implement the client APT, advanced persistent threats [sic] protection, in stages and to be completed by end FY19*".

296. Wee has informed that the initial draft of FY16 CII Risk Assessment was sent to Serena Yong, Henry Arianto, Foong Lai Choo, and Clarence Kua. These were all senior members of IHiS' Infrastructure Services Division or the CIO Office. The FY16 CII Risk Assessment was also presented at a number of meetings in January 2017, including the SingHealth CITC (Cluster IT Council) meeting which, as described in paragraph 119(d) (pg 42) above, was chaired by SingHealth GCEO, Prof. Ivy. Subsequently, on 5 April 2017, the CSG also shared the results of the risk assessment with the CSC (Cyber Security Council), which was chaired by MOHH MD, Aik Guan.

297. In end-2017, the next risk assessment was conducted, and a "*2017 Risk Assessment*" was published on 31 December 2017. The same threats posed by APTs and the proposed implementation of ATP by FY2019 were repeated. On 31 January 2018, the CSG updated the CSC on this risk assessment.

---

[26] This threat was described as having a 'Medium' likelihood of occurring if there are no controls in place, and with a 'High' impact to business operations.

[27] This threat was described as having a 'Low' likelihood of occurring if there are no controls in place, and with a 'Medium' impact to business operations.

298. On 5 June 2018, Kim Chuan, in his capacity as CISO for MOH, presented to the IHiS ARC (Audit and Risk Committee) on the "*Cybersecurity Threat Landscape for Public Healthcare*". Pertinently, the following were identified:

(a) Remote Access Trojans used for "*steal(ing) confidential data from within organisations' network through backdoor on compromised PCs*" was identified as a threat;

(b) "*State-backed, highly-skilled cyber hackers who target national Infrastructure and systems for espionage*" was identified as a one of the profiles of cyber attackers; and

(c) IHiS and PHIs would have to remain vigilant against the "*potential threat*" of "*Advanced persistent threats (APT)s, stealth attacks to attack endpoint systems, exfiltrating data and/or facilitating backdoor access*", while also stating that there were no incidents as yet. A slide showing the "Anatomy of an APT Attack" was also included.

299. It is unfortunate to note that what was described in the 5 June 2018 IHiS ARC meeting as a "*potential threat*" was already a real and present danger unfolding at the time.

300. The overall picture that emerges from the above facts is that IHiS' senior management had knowledge of and was alive to the threat of APTs from as early as August 2016, and had some familiarity with the "Anatomy of an APT Attack". Senior management of SingHealth and MOHH may also have had some awareness of the threat of APTs based on discussions of the FY2016 and FY2017 CII risk assessments. However, as the Committee's findings on IHiS' incident response demonstrates, this knowledge did not effectively percolate down to the IT administrators, security personnel and line management in IHiS. The ATP system was also not yet implemented throughout the Cyber Attack.

18.4 Timeline of events

301. A timeline showing the main events in IHiS' incident response, and the relevant sections in which they are discussed in this Report, is as follows:

Figure 12: Timeline of events pertaining to IHiS' incident response

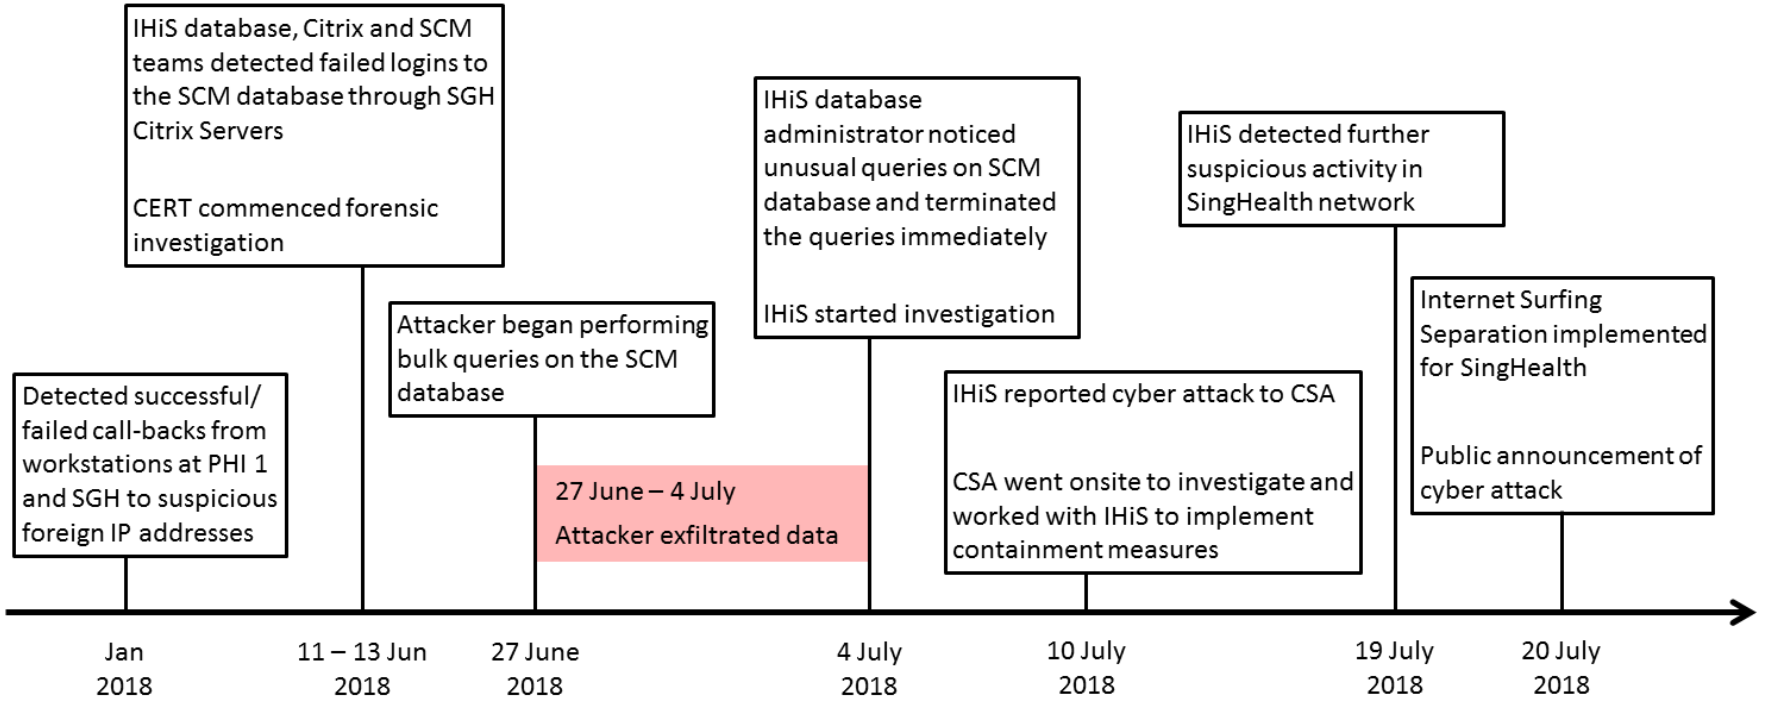

19 EVENTS OF JANUARY 2018

19.1 Detecting malware on the PHI 1 Workstation and callbacks to suspicious IP addresses – 18 January 2018

302. On 18 January 2018, Benjamin Lee ("**Benjamin**"), a System Engineer from the IHiS Security Management Department ("**SMD**"), was performing a routine check and noticed an alert about suspicious activity detected on a workstation located in a SingHealth public healthcare institution (referred to earlier as "**PHI 1**" and the "**PHI 1 Workstation**"). The alert provided him with the filename of the suspected malware found on the workstation, and the date of infection was stated to be 18 January 2018. Benjamin decided to investigate the matter, and informed Tan Choon Kiat Ernest ("**Ernest**"), Senior Manager of the SMD, of the same.

303. In the course of investigations, Benjamin determined that the PHI 1 Workstation was:

(a) attempting to communicate with what he understood to be a **foreign IP address** and an associated URL; and

(b) sending commands to two other IP addresses.

304. The foreign IP address was in fact one of the key C2 servers used by the attacker throughout the entire period of the Cyber Attack.

305. As for the other two IP addresses, Benjamin found that public IP addresses beginning with those numbers were associated with a different foreign country, and thus believed that the commands were being sent to IP addresses in another country. This view would subsequently be proved to have been erroneous.

306. While the file name of the suspected malware was that of a legitimate program, the program should not be located in the file path where it was found.

However, it appears that Benjamin or other members of the SMD did not notice this fact.

19.2 Blocking and monitoring of suspicious IP addresses and re-imaging the PHI 1 Workstation – 18 January 2018

307. Acting on his findings above, Benjamin blocked network traffic between the PHI 1 network and all three IP addresses on 18 January 2018. However, he did not take any action to block network traffic to those IP addresses from the rest of the SingHealth network.

308. Benjamin then informed the site engineer to disconnect the PHI 1 Workstation from the SingHealth network. Checks with the user of the workstation did not reveal any useful information.

309. Benjamin then quarantined the suspicious file, and reconnected the machine to the SingHealth network. He then confirmed that the workstation was no longer attempting to connect to the foreign IP address. However, commands were still being broadcast to the other two IP addresses.

310. Benjamin then instructed the site engineer to re-image the PHI 1 Workstation to eradicate any malware, and this was done overnight. He also recorded the suspected malicious IP addresses and URL and asked IHiS' outsourced vendor for MSS (Managed Security Services) to continue monitoring network traffic to the suspicious IP addresses and URL, in all the public healthcare clusters.

19.3 Discovering multiple attempts from Workstation A to communicate with the same suspicious foreign IP address – 19 January 2018

311. On 19 January 2018, Benjamin obtained a set of network logs (comprising proxy logs and firewall logs). The search range for both the proxy logs and the firewall logs was from 1 to 19 January 2018.

312. On reviewing the network logs, Benjamin noticed that there were many instances of access to the foreign IP address. All the *successful* instances of access were from a single IP address, and involved either a particular SGH user-ID, or the hostname of a SGH workstation. The IP address and hostname were in fact that of Workstation A, and the user-ID was that of the user of Workstation A. As discussed in Part III above, Workstation A played a significant role in the Cyber Attack.

313. In the afternoon of 19 January 2018, Benjamin sent an email to the SMD, including Ernest, titled "*Hits to IOCs*" ('IOCs' refer to indicators of compromise), attaching the network logs. In this email, Benjamin informed his colleagues that he had arranged for scans for hits involving the malicious IPs and URLs that he was aware of at that time.

314. There was no immediate reply to Benjamin's 19 January 2018 email. Ernest stated that he "glanced" at the logs when he received the email. He could see that there were multiple attempts to communicate with the foreign IP address from one or more workstations in SGH and PHI 1. However, no steps were taken by Ernest or Benjamin to (i) identify the owner of the user-ID shown in the logs (*i.e.* the user of Workstation A); (ii) identify the physical location of Workstation A; (iii) investigate into the callbacks from Workstation A, including whether it was infected with malware; or (iv) to block connections to the suspicious IP address from SGH or the rest of the SingHealth network. There is also no indication that there was any follow-up from any other members of the SMD.

315. Apart from the URL related to the foreign IP address, Benjamin's email also mentioned another URL. Benjamin has explained that this URL was also flagged up as having been accessed by the PHI 1 Workstation, but he had inadvertently forgotten to include this URL in his subsequent reports on the matter. Although Benjamin had flagged this other URL in his email, no further action was taken by Benjamin or Ernest to block the domain name for the SingHealth network. This URL was in fact one of the attacker's C2 servers.

19.4 Further steps taken in respect of queries to the other two IP addresses – 19 January 2018

316. As mentioned above, Benjamin found that commands were being made to two other IP addresses, which he believed erroneously to be from another foreign country. On 19 January 2018, Benjamin conducted further checks and found that a number of other workstations across PHI 1 were sending queries to similar IP addresses. Investigating further, he discovered that the Windows printer settings on the affected workstations were configured to send the queries. He instructed site engineers to remove these printer settings from the workstations, and there was no further traffic to these IP addresses thereafter.

317. At the material time, Benjamin had the understanding these IP addresses were foreign IP addresses. However, investigations into the Cyber Attack have since revealed that the queries to these IP addresses resulted from legacy printer settings that had not been removed. The queries to the IP addresses were thus not malicious.

318. It appears that Ernest had known at the time that the PHI 1 network previously used such IP addresses, and was thus of the view that these IP addresses were not suspicious. However, neither Ernest nor any other members of the SMD informed Benjamin of this at the material time.

19.5 Analysing process dump of the suspected malware – 20 January 2018

319. On the evening of 20 January 2018, Benjamin performed an analysis of a process dump[28] of the suspicious file that had been identified on the PHI 1 Workstation. The analysis was performed through an online service which

[28] Benjamin has explained that this refers to a memory dump of a particular process that is running on a computer.

analyses suspicious files and facilitates detection of viruses, Trojans, worms and malware. This was done on his own initiative and without Ernest's knowledge.

320. The online analysis returned a benign result. CSA has explained that the malware signature was not available publicly at the time, and the online check would thus not have been able to flag the process as malicious.

19.6 Concluding investigations without further escalation – 22 January 2018

321. Benjamin continued to monitor the network traffic until 22 January 2018. In this period, he found no further malicious outbound traffic from the PHI 1 Workstation. He then recorded his investigation process and findings into a deck of slides and sent the slides to Ernest on the afternoon of 22 January 2018.

322. That night, Ernest replied stating that it was an "*informative report*", and "*as we do not [sic] overall impact of this malware, I'm just wondering whether to share this out* ☺". In saying this, Ernest had in mind sharing this with the Security team. The next morning, Benjamin replied agreeing that they did not know "*the true impact of the malware*", and suggesting that the information be shared with IHiS security teams, "*(j)ust to check if other institutions/clusters have the same malware or printer misconfiguration*".

323. Ernest stated that this was, to his mind, not a reportable security incident as the malware on the PHI 1 Workstation had been contained. He also cited the IR-SOP, which states that malware infections that have been detected, contained, and cleaned, without network propagation, need not be reported. He did not inform Wee "*because suspected malware infection of a workstation is a very common occurrence*". He also stated that they did not file an Incident Reporting Form ("**IRF**") because they would not typically file an IRF for cases involving suspected malware infections.

19.7 Assessment of IHiS' incident response in January 2018

324. Based on the evidence presented, the Committee finds a number of deficiencies in the incident response by Ernest and Benjamin:

(a) First, no steps were taken whatsoever by Ernest or Benjamin to investigate Workstation A, despite the evidence of communications between Workstation A and what was understood to be a suspicious IP address. The Committee finds this unsatisfactory, since such investigations would have been the next logical step following the containment of malicious activity on the PHI 1 Workstation, and was clearly not beyond their technical experience or expertise.

(b) Second, Ernest did not take steps to verify whether the malware found on the PHI 1 Workstation was propagated across the network. Vivek has highlighted that this was a failure to implement the IR-SOP which, as mentioned in paragraph 323 above, states that malware infections are not reportable if the malware is cleared and if there has been *no network propagation*. In this case, however, the callbacks from Workstation A were indicative of network propagation, and ought to have prompted further investigations.

325. These deficiencies arose from a failure by Ernest and Benjamin to appreciate the significance of the findings. Benjamin has explained that while he had previously dealt with malware infections, this was his first time dealing with malware which was communicating with C2 servers. Ernest has also explained that this was the first time he had dealt with a case where there were multiple attempts to communicate with a foreign IP address, from one or more workstations, and where one of those workstations was found with a suspicious file. As Ernest himself stated in his email to Benjamin on 22 January 2018, they did not know the "*overall impact of this malware*".

326. However, instead of carrying out further investigations to better understand the circumstances, the matter was concluded without further investigations or reporting. Ernest has stated that he did not take any further action because the foreign IP address was not a known C2 server, in that he had not received any information positively identifying the IP address as a C2 server. Unless he had received such information, no steps would be taken to block other suspicious IP addresses that had not been flagged as C2 servers. This completely passive attitude towards the identification and addressing of potential security risks is, in the Committee's view, fundamentally inconsistent with the roles and responsibilities of the SIRM.

327. The failure to block the suspicious IP address across the whole network and to investigate Workstation A constituted a significant missed opportunity to prevent the attack. CSA is of the view that the fact that callbacks were being made to a suspicious URL and IP address from a workstation within the CII sector, and which was suspected to be infected with malware, should have been reported to CSA as a security incident. Had this been done at the time, it is possible that the Cyber Attack could have been detected earlier, and the appropriate actions could have been taken.

328. Separately, the Solicitor-General has submitted that Benjamin's attempt at self-help by performing an analysis of process dump through the online service was "*resourceful but inadequate*". The Committee has also heard that as at January 2018, Benjamin neither had the training nor the tools to analyse the process dump himself, and he was only trained in digital and memory forensics in March 2018. In view of this, the Committee agrees that Benjamin displayed a good sense of initiative and resourcefulness. However, there were security implications with the use of the online service. Unfortunately, at the material time, Benjamin did not have the proper training to appreciate the consequences of his actions.

20 EVENTS OF 11 JUNE 2018

20.1 Detecting failed logins to the SCM database and changing of passwords for the D.A. account

329. At 3:10pm on 11 June 2018, Katherine received a system-generated email showing a number of failed logins to the SCM database within a very short period of time earlier that day. Shortly after that, she was notified of a few more failed logins earlier that same day.

330. At 5:08pm and 5:09pm on 11 June 2018, Katherine forwarded the details of the failed logins *via* email to Robin Seah ("**Robin**"), Kelvin Chong Wee Kiat ("**Kelvin**") and Reynaldo Delgado Francisco ("**Rey**") from the IHiS Service Delivery (Clinical Care) Department, asking if they had any idea what was going on. Shortly after, Katherine also forwarded this email chain to Vicky Boh, Thota Veerendra Naidu ("**Veerendra**") and Joanne Lim Shan Shan ("**Joanne**"), who are Citrix administrators, to ascertain whether a particular IP address was that of a Citrix server, and to follow-up with any further investigations. She also copied Lum Yuan Woh (**"Lum"**) in this email, in his capacity as the Assistant Director of the Citrix Team.

331. These were in fact part of the failed logins discussed in paragraph 177 (pg 62) above. All the failed logins were shown to have originated from one IP address, which was subsequently determined to be the IP address of Citrix Server 1.

332. Katherine noticed that a number of different account names had been used to attempt to log in to the database. The log in attempts generally failed because they were invalid user-IDs. Applying basic common-sense, the obvious inference to be drawn was that someone was guessing user-IDs, and therefore, the attempted access to the SCM database was likely to have been unauthorised. Yet, this did not occur to Katherine, and she initially thought that some IHiS staff might have been "*testing the system*".

333. The D.A. account was also used in the log in attempts. This was an account belonging to an IHiS domain administrator. The account was not authorised to access the SCM database. On the evening of 11 June 2018, Katherine called the domain administator asking if he had attempted to access the SCM database, and he confirmed that he did not attempt to do so. Later that same evening, the domain administator changed the passwords to the D.A. account.

334. At this stage, Katherine realised that it was not a test but was something unusual, because "*[the domain administrator]'s user-ID would not be used in testing the system*", and the domain administrator had confirmed that he had not tried to access the SCM database. She surmised that someone was trying to access the SCM database. However, as she had already escalated the matter to the Citrix administrators, she left it to them to follow-up with further investigations.

20.2 Detecting unusual logins to Citrix Server 1 using the L.A. account

335. Upon reviewing the contents of Katherine's emails, Lum and his team of Citrix administrators determined that: (i) the attempted log ins were made at the database-level and not through the SCM front-end application, and (ii) the IP address in question was assigned to Citrix Server 1.

336. On the understanding that a user would have to log in to Citrix Server 1 before attempting to log in to the SCM database, Lum tried to trace who had logged into Citrix Server 1 on 11 June 2018, in order to identify the person or persons who had attempted to log in to the SCM database.

337. While there were many logins on that day, there were two logins which Lum saw as "*unusual*" – these were the logins using the L.A. account. These were unusual to Lum as the L.A. was not an account that staff would use in day-to-day operations.

338. The logs also showed that the logins on 11 June 2018 using the L.A. account to Citrix Server 1 came from a workstation bearing the hostname of VM 2. This appeared unusual to Lum as it was not a valid hostname.

339. Lum then filtered the logs from Citrix Server 1 to find all logins to the server using the L.A. account. He found that the last legitimate login into the server using the L.A. account was on 13 October 2017.

340. Lum also found that after 13 October 2017, there had been numerous logins to Citrix Server 1 between 17 May 2018 to 11 June 2018 using workstations bearing hostnames which should not normally have been logging into the Citrix server using the L.A. account. Lum noticed the use of workstations VM 1 and VM 2, but did not know where these workstations were located. He also felt that the names of these two workstations were unusual. Lum's hypothesis was that these were virtual machines running on legitimate workstations that had already cleared IHiS' network access control measures.

20.3 Discovering that Citrix system event logs for Citrix Server 1 were deleted

341. In the evening of 11 June 2018, Vicky also discovered that the Citrix system event log for Citrix Server 1 had been deleted. As discussed at paragraph 178 (pg 62) above, these logs would have captured the details of all the accounts that logged in to Citrix Server 1. The Citrix Team however had access to another set of logs.

342. IHiS staff noted that the record of the log being cleared was reflected as having been carried out by the "System" account. However, they were unable to explain how the "System" account had been used in this way, nor identify the person who had deleted the event log.

343. Ordinarily, if the system event log has been deleted, there would be no other record of who had logged into the server. However, IHiS staff had access

to a separate set of logs, and was thus able to identify the various logins to Citrix Server 1 using the L.A. account.

20.4 Changing passwords to the L.A. accounts on all SGH Citrix servers

344. By the evening of 11 June 2018, Lum was of the view that the password for the L.A. account had been stolen, and had been used to access Citrix Server 1 without authorisation on numerous occasions. However, at this stage, Lum did not report the incident to anyone from the SMD (Security Management Department).

345. The password for the L.A. account on Citrix Server 1 was changed in the evening of 11 June 2018. Thereafter, there were no subsequent unauthorised logins to Citrix Server 1 using the L.A. account.

346. Local administrator accounts named "L.A." also existed in all other Citrix servers in the SGH Citrix server farm and in H-Cloud, and Lum was of the view that it was highly possible that other "L.A." accounts in the SGH Citrix servers used the same password, 'P@ssw0rd'. As such, later in the night of 11 June 2018, the passwords for the L.A. accounts on all the other SGH Citrix servers were changed.

20.5 Discovering that malware was detected earlier on Citrix Server 1

347. On the night of 11 June 2018, Lum was concerned that Citrix Server 1 may have been infected with malware. He checked the antivirus software logs on Citrix Server 1 and found that malware had been deleted from the server on 8 June 2018. At that point, he did not check further as to the details of the malware, but he provided this information to the SMD the next day.

348. Lum has explained that given what he had learned on 11 June 2018, his first thought was that there was "*some kind of audit/penetration testing or*

scanning of the SCM database to test for vulnerabilities". However, after checking with the Infrastructure Services and the SMD the next day, he realised that this was not the case.

20.6 Assessment of IHiS' incident response on 11 June 2018

349. The responses of IHiS' staff to the events on 11 June 2018 could have been improved in a number of ways:

(a) First, the Citrix Team became aware of multiple failed logins to the SCM database being made from a number of workstations through Citrix Server 1, including workstations which appeared to be unauthorised (*i.e.* workstations VM 1 and VM 2). The presence of unauthorised devices on a network is a serious issue which merits immediate follow-up, but no such action was taken until 13 June 2018.

(b) Second, the Citrix Team became aware that Citrix Server 1 had been accessed on multiple occasions dating as far back as 17 May 2018 using the L.A. account, which should only have been used by the Citrix Team. No one in the Citrix Team had accessed Citrix Server 1 using this account, and the accesses were clearly unauthorised. CSA is of the view that this, and other subsequent unauthorised logins to Citrix servers detected by IHiS staff, should have been reported to CSA as a security incident.

(c) Third, Katherine and the Citrix Team became aware that the D.A. domain administrator account was used in an attempt to access the SCM database. The owner of the account verified that he had not attempted to access the database. Besides changing the account password, there was no further follow-up to ascertain how the account was misused in this way.

(d) Fourth, the Citrix Team became aware that event logs to Citrix Server 1 had been deleted. This would mean that the attacker had

sufficient privileges to delete system logs. CSA is of the view that this, and other instances of deleted server logs detected by IHiS staff subsequently, should have been reported to CSA as a security incident.

(e) Fifth, the Citrix Team became aware that Citrix Server 1 was previously infected with a malware (this was in fact a hacking tool used by the attacker). This indicated that there was malicious activity occurring in a server directly connected to a CII system, and CSA is of the view that this should have been reported to CSA as a security incident.

(f) Sixth, Lum was aware that he ought to have informed the SMD (*e.g.* Ernest, head of the SMD team for SingHealth and the SingHealth SIRM) should he encounter a security incident, however, he did not do so on 11 June 2018, despite believing that the L.A. account had been compromised, and all the other findings identified above.

350. Vivek has also observed that the Citrix Team's action of resetting the compromised passwords during the investigations in a hurry could have hampered the investigations – doing so would have flagged to the attacker that his presence had been discovered. In such a scenario, attackers usually respond by moving over to use other passwords that are not yet flagged as compromised, as was the case in the Cyber Attack. In the process, an investigation team can lose track of the attacker at least temporarily until further compromised passwords are discovered. In Vivek's expert opinion, a better practice would be to put the compromised passwords on active monitoring and use them to learn more about the attacker's behaviour as well as presence across other systems within the network.

351. In respect of the events of 11 June 2018, the responses of Katherine, Lum and the Citrix Team were inadequate on the whole. They could not fully appreciate the security implications of their findings, and were unable to co-relate

these findings with the TTPs of an advanced cyber attacker. That said, this was not necessarily their fault. There is no evidence that they had been given the requisite training to do so, despite the fact that IHiS' Senior Management had appreciated the risk of APTs from as early as August 2016 (as explained in section 18.3 (pg 104) above).

352. Likewise, Katherine and Lum were not familiar with relevant policy documents such as the IR-SOP and the SIRF. They were therefore not in a position to understand that the suspicious instances of attempted access to the SCM database, which is a CII, in fact constituted a potential security incident, requiring an urgent response and reporting all the way to CSA. Furthermore, IHiS did not have an incident reporting framework for line staff, and there was no clarity on how incidents were to be reported to management.

353. This lack of appreciation by the line-staff of the security implications of their findings and the need to report security incidents all the way to CSA would prove to be a consistent feature of IHiS' incident response up until 4 July 2018.

354. This was a weakness which was rightly observed by Vivek – in his expert opinion, representatives from all IT teams (*e.g.* database teams, network teams etc.) should be involved in IT security training, including tabletop exercises. This recognises the fact that aside from formal incident responders, the persons who would first experience the signs, and who would need to be equipped with the ability to detect signs of a cyber attack, are the operational IT staff in an environment such as IHiS.

355. That said, there were a few positive aspects from the Citrix Team's response. These were observed by Vivek: (a) the availability of logs (even after they were deleted by the attacker); (b) detecting the presence of malware; and (c) being able to identify the host names that were connecting to Citrix Server 1 and track historical logins.

21 EVENTS OF 12 JUNE 2018

21.1 Discovering failed logins to SCM database from Citrix Server 1 dating back to 24 May 2018

356. In the morning of 12 June 2018, Katherine provided Kelvin with the logs of failed logins to the SCM database dating back to May 2018, further to a request made by Kelvin earlier. These logs were subsequently forwarded by Kelvin to Lum.

357. Katherine reviewed these logs, and noticed that in addition to the unusual failed attempts on 11 June 2018, there were a number of unusual failed attempts to log in to the SCM database from Citrix Server 1 beginning from 24 May 2018. These were in fact the failed attempts discussed in paragraph 176 (pg 62) above. While Katherine would have received notifications of these failed attempts around the time they happened, she had not noticed them earlier.

358. Although Katherine noticed now that there was a pattern of unusual failed attempts to login to the SCM database dating back to 24 May 2018, she did not take any further steps to report the matter or discuss this matter with anyone. She did not see a need to, given that Kelvin and Lum had a copy of the logs, and she assumed that they would look into the matter.

21.2 Detecting further failed logins to the SCM database from Citrix Server 1 on 12 June 2018

359. In the afternoon of 12 June 2018, Katherine received system-generated database alerts showing a number of failed attempts to login to the SCM database from Citrix Server 1 earlier that day. As described in paragraph 179 (pg 63) above, these included the attempts made using accounts which had not been granted access to the SCM database. These emails were forwarded to Kelvin, Robin and Lum shortly after Katherine received them. At that time, Katherine had in mind that these errors involved end-user accounts, and it was thus appropriate for

Kelvin, as part of the Applications team that has responsibility for end-user accounts, to look into them.

360. In his reply, Kelvin directed a query at Lum, mentioning that a password had been changed, and asking to check who was logged in to Citrix Server 1. The Citrix Team checked the active sessions on Citrix Server 1 at the time of the failed logins to the SCM database, and it did not appear that the L.A. account was used to log in to Citrix Server 1 at the time. IHiS staff were unable to identify which account had been used to log in to Citrix Server 1.

361. Katherine has explained that in view of the unusual failed attempts at logging in to the SCM database on 11 June and 12 June 2018, she "*was concerned that something was wrong*". While she "*did not know exactly what was taking place*", she "*knew that it was unusual*". Thus, on 12 or 13 June 2018, she called Lum, and Lum told her to inform him every time she received any notice of failed attempts at logging in to the SCM database.

362. On Lum's part, he thought the multiple attempted log ins to the SCM database "*could be somebody attempting to gain unauthorised access to the SCM database*".

21.3 Discovering numerous instances of suspicious folders in Citrix Server 1

363. Having found out the night before on 11 June 2018 about the earlier infection of Citrix Server 1 with malware, the Citrix Team was on the look-out for any other suspicious files. On 12 June 2018, they found that many users' profile folders contained a folder with a particular name. These folders did not contain any executable programs. Lum checked online and found out that the name of the folders was that of an open source SQL tool. The tool is not a software that is used in the SingHealth IT environment.

21.4 Disabling logins to Citrix Server 1 and informing the CERT and Wee

364. In view of the above circumstances, the Citrix Team disabled logins to Citrix Server 1. Thereafter, in the afternoon of 12 June 2018, the Citrix Team sent an email to Sean Navin from the SMD, informing the latter (i) of attempts to connect to the SCM production database from Citrix Server 1 on 12 June 2018, and (ii) that they found the suspicious folder in all user sessions, and seeking Sean's help to "*gather any information suspicious about this abnormal behaviour*". The Citrix Team also provided a screenshot of the log entry showing the presence of malware on Citrix Server 1, but did not make clear which computer or server this malware had been detected, or what its significance was.

365. Subsequently, Sean forwarded the email to Benjamin at 6:02pm on 12 June 2018, seeking the latter's assistance on the matter in his capacity as a member of the CERT. Benjamin replied at 8:53pm on 12 June 2018, copying Ernest, Wee, and two other members of both the SMD and CERT, Zac Lim Zi Yang ("**Zac**") and Muhammad Azzlan Bin Zainuddin ("**Azzlan**").

366. In their subsequent correspondence on 12 June 2018, Benjamin and Veerendra agreed to meet at SGH the next morning. Benjamin also clarified the following: (i) that in order to install the suspicious folder in every user's profile, administrative rights are required, and (ii) that it will be possible to suspend Citrix Server 1.

367. Notably, Ernest and Wee were copied in Benjamin's 8:53pm emails. Ernest did not read this email as he was overseas at the time. Wee states that he "*glanced through*" the emails the next morning and "*do(es) not recall looking in detail at the logs and screenshots in the first email*" from the Citrix Team. Likewise, for subsequent emails in the thread received by him on 13 and 14 June 2018, he explains that he "*may have briefly gone through the details of these emails, but (he) cannot remember them now.*" In any case, Wee did not take any follow-up action in spite of the information he had received.

21.5 Assessment of IHiS' incident response on 12 June 2018

368. The Committee notes that the Citrix Team had acted appropriately to raise the matter to the SMD. However, the response could have been improved by quicker action (*e.g.* immediately upon learning at around 1:30pm about the failed attempts at logging in to the SCM database), and by providing clearer explanations of their findings and their views to the SMD.

22 EVENTS OF 13 JUNE 2018

22.1 Meeting to update Benjamin on the events of 11 and 12 June 2018 and sharing of information with the CERT and Wee

369. At around 10:00am on 13 June 2018, Benjamin met with Veerendra and Vicky from the Citrix Team. Veerendra and Vicky showed Benjamin some logs, and explained the following:

(a) That attempts had been made to access the SCM database from Citrix Server 1, most recently on 11 and 12 June 2018;

(b) That multiple usernames had been used in attempts to login to the SCM database;

(c) That there had been unauthorised access to Citrix Server 1 using the L.A. account on multiple occasions dating as far back as 17 May 2018;

(d) The hostnames of the workstations used in the abovementioned instances of unauthorised access to Citrix Server 1, which included (i) the PHI 1 Workstation; (ii) Workstation C (a SGH workstation); (iii) VM 1; and (iv) VM 2; and

(e) That the L.A. account should only have been used by the Citrix Team.

370. Benjamin's evidence is that based on the above, it appeared to him that the SCM database, which he knew to be a CII, was being targeted.

371. Shortly after the meeting, Vicky forwarded Benjamin an email from Katherine containing screenshots of the alerts received by Katherine showing the failed login attempts to the SCM server. The CERT, Ernest, and Wee were copied in Vicky's email to Benjamin. While Ernest was still overseas and did not read the email at the time, Wee was at work and would have received the email. Once again, however, Wee "*cannot now recall*" if he had read Vicky's email or the attached email from Katherine.

22.2 Follow-up action in respect of workstations used in unauthorised logins to Citrix Server 1

372. Throughout and amidst the events on 13 June 2018, Benjamin and the Citrix Team took steps to investigate Citrix Server 1 and the workstations involved in the unauthorised access to Citrix Server 1. They started first with trying to find the physical locations of the workstations by pinging them.

373. That morning, Benjamin also purchased three external hard disks to store forensic images of the workstations and Citrix Server 1.

Pinging workstations VM 1 and VM 2

374. There was no response when Benjamin and the Citrix Team pinged workstations VM 1 and VM 2. In light of this, the Citrix Team suspected that these were virtual machines.

Seizing the PHI 1 Workstation on 13 June 2018

375. When Benjamin was informed that the PHI 1 Workstation was one of the workstations from which unauthorised logins to Citrix Server 1 were made, he noted that it appeared to be the same workstation which he had re-imaged in January 2018. He suspected that the PHI 1 Workstation may have been

compromised by malware again. He also suspected that the matter raised by the Citrix Team on the morning of 13 June 2018 was a security incident.

376. There was a response from the PHI 1 Workstation when pinged, showing that it was in PHI 1. On 13 June 2018 itself, Benjamin contacted his IHiS colleagues who were based in PHI 1 and asked them to locate the PHI 1 Workstation, and to unplug the power cable. Thereafter, Benjamin acquired the forensic image of the PHI 1 Workstation on the same day.

Arranging for the seizure of Workstation C on 18 June 2018

377. Sometime in the afternoon of 13 June 2018, Benjamin had ascertained the user to whom Workstation C was assigned. Benjamin's colleague contacted the user on the phone, and they learnt that he was overseas. The user consented to the workstation being seized for forensic investigations.

378. In an email conversation starting from the afternoon of 13 June 2018, Benjamin explained to the user's head of department that Workstation C was found to be attempting to connect to the SCM database using several different username and password combinations over several days from 22 May to 4 June 2018, and requested for permission to seize the workstation. The head of department gave permission for the workstation to be seized, and further informed Benjamin that the user was overseas and had not used the workstation after 29 May 2018. Thereafter, arrangements were made for Workstation C to be seized on 18 June 2018, with a view to taking a forensic image.

379. This email correspondence was copied to an email group which included Ernest, who was still on holiday at the time.

Shutting down Citrix Server 1

380. On Benjamin's advice, the Citrix Team exported the server image of Citrix Server 1 for the CERT to conduct further investigations, and shut down Citrix Server 1 thereafter.

381. Wee was copied in an email from Benjamin to Veerendra at 10:58am on 13 June 2018, in which Benjamin mentioned the obtaining of the server image of Citrix Server 1, which was referred to by its hostname. Wee has explained that this was the first time he recalls a server image being acquired for forensics. He has explained that he "*could not tell what kind of server Benjamin was investigating, based only on the hostname mentioned*", and that he was waiting for Benjamin's investigation findings.

22.3 Setting-up the TigerConnect chat group

382. After meeting with the Citrix Team on the morning of 13 June 2018, Benjamin created a chat group called "Citrix-SCM Incident" on the TigerConnect messaging platform to discuss the events of 11 and 12 June 2018, and to update members of the chat group on his findings and actions taken. Persons added to the chat group included Ernest, Zac, Azzlan, Alvin Chua Yu Long, and Benjamin from the SMD; Lum from the Citrix Team; and Cluster ISO Wee.

383. On the chat group, IHiS staff discussed the events that occurred on 11 and 12 June 2018, including the failed logins to the SCM database, and the use of workstations with unusual hostnames to access Citrix Server 1 from May 2018 onwards.

384. Ernest has explained that he was overseas until 18 June 2018, and did not read the messages on the TigerConnect chat group before then.

385. Wee has explained that his recollection of the updates on the chat group regarding the multiple attempts to login to the SCM database on 11 and 12 June 2018 using multiple usernames was "*very vague*", and he "*cannot recall if (he) saw this update*". However, he also represents that had he seen the update, he "*would have questioned more about how this had occurred*." At the same time, Wee "*vaguely remembers*" Benjamin mentioning workstation VM 2, and the collection of images of affected endpoints. He was of the impression that

Benjamin had already identified and isolated the affected endpoints and was in the process of conducting forensics.

386. The TigerConnect chat logs show that on the morning of 13 June 2018 when the chat group was first created, Wee had in fact read the messages, and sent some queries, but did not give any instructions. Wee's evidence is also that he took no steps to confirm with Benjamin if the situation had been isolated and contained.

22.4 Detecting failed logins to the SCM database from Citrix Server 2

First series of failed logins

387. At 3:50pm on 13 June 2018, a system-generated email alert was sent to Katherine, notifying her of a number of failed attempts at logging in to the SCM database within a short period earlier that same day. All the attempts were made from an IP address which was different from the IP address from which the attempts on 11 and 12 June 2018 were made. At this point, Katherine was not aware that this IP address was that of Citrix Server 2.

388. All the attempts had failed because invalid user-IDs were used. From the email alert, Katherine noticed the following:

(a) That one of the invalid user-IDs was also used in the earlier attempts on 11 June 2018;

(b) The server name for Citrix Server 3 was being used in an attempt to log in to the SCM database, which was very unusual; and

(c) Access was being attempted through a different Citrix server, Citrix Server 2, after Citrix Server 1 was shut down.

389. Katherine surmised that the failed attempts at logging in to the SCM database were evidence of someone attempting to gain unauthorised access to the SCM database. At 3:52pm, Katherine forwarded it to Kelvin, Robin and Lum, with the subject of her email being "*Login failed – new server.*', and with the IP address of Citrix Server 2 in the body text.

Second series of failed logins

390. At 4:10pm on 13 July 2018, another system-generated email alert was sent to Katherine, informing her of a few more failed logins to the SCM database over a short period of time earlier that day. Once again, all the attempts were made from Citrix Server 2.

391. In one attempt, the server name for Citrix Server 3 was again used as a user-ID. The user-ID in another attempt was the name of a service account which would not ordinarily be used for the purposes of logging in to the SCM database. In yet another attempt, the attacker used a user-ID that it had used in a prior attempt to connect to the SCM database from Citrix Server 1 on 12 June 2018.

392. Katherine was of the view that "*these failed log-in attempts were even more unusual as compared to the others*", and when she saw these errors, coupled with her knowledge of all the other failed login attempts, she realised that someone was repeatedly trying to gain unauthorised access to the SCM database. She forwarded these alerts to Lum as well.

Lum's reply at 4:19pm

393. Upon receiving Katherine's 3:52pm email, Lum identified the hostname of Citrix Server 3 to be that of a H-Cloud Citrix server. He also identified the IP address from which the attempted logins to the SCM database were made as being associated with Citrix Server 2. Lum checked the login logs for Citrix Server 2, and ascertained that (i) VM 1 was logged into Citrix Server 2 at the time of the failed logins to the SCM database, and (ii) the account used to log in to Citrix Server 2 was the S.A. account.

394. At 4:19pm, Lum replied to Katherine's 3:52pm email. This reply was addressed to Katherine, Kelvin and Robin. Lum also copied Veerendra, Vicky Boh, Yu Ping Hai, and Joanne (who were Citrix administrators). The contents of the email included:

> Please do not disclose this information to any other people.
> Veerendra – Need to catch this [VM 1] machine.
> We found a malicious login to Citrix from a machine [VM 1].
> […]

395. Lum has explained that at the point he sent this email, he had already determined that the login to Citrix Server 2 using the S.A. account was malicious.

396. Lum has further explained that he asked that recipients "*not disclose this information to any other people*" because they had just discovered new information, namely, that a different Citrix server was being accessed (*i.e.* Citrix Server 2) using a different account (*i.e.* the S.A. account), and he wanted to look further into the matter and understand how the account could be used in that way; prior to this, he had not even heard of the S.A. account.

397. Katherine has explained that by the time she received this email, she was of the opinion that IHiS was dealing with what could be classified a security incident. However, she did not report this to the security team or to her head of department, in view of Lum's statement in his email "*not to disclose this information to any other people*". At this point, Katherine did not know that the Citrix Team had contacted the SMD.

398. It is pertinent to note that Benjamin was not informed of the use of the S.A. account to login to Citrix Server 2 until 26 June 2018.

22.5 Removing the S.A. account from the admin group

399. At 4:30pm on 13 June 2018, Lum removed the S.A. account from the admin group, thereby disabling remote access to the Citrix servers using this account. At the time, the Citrix Team suspected that access from the workstations to the Citrix servers must have been through an RDP client. They also noted that the Windows Remote Desktop Connection application was installed by default on all workstations. By removing the S.A. account from the admin group, RDP access into Citrix servers using the S.A. account would be disabled from all workstations.

22.6 Detecting failed logins to the SCM database from Citrix Server 4

400. At 4:50pm on 13 June 2018, a system-generated email alert was sent to Katherine informing her of one failed attempt at logging in to the SCM database that same afternoon from another IP address which was different from the earlier observed attempts. Although Katherine did not know this at the time, this IP address was that of Citrix Server 4. The user-ID used in the failed login attempt was that of the user of Workstation A.

401. Katherine forwarded the email to Kelvin, Robin, Lum and Joanne immediately, highlighting in the email title that there was a new server involved. Further to Lum's directions, Katherine also forwarded the email to Veerendra for investigations.

22.7 Investigations into the account used to log in to Citrix Server 4 and resetting the account password

402. Having received Katherine's second email at 4.55pm, Lum determined that the IP address was associated with Citrix Server 4, a SGH server. By reviewing the login logs to Citrix Server 4, the team found that the account belonging to the user of Workstation A was used to log in to Citrix Server 4 on a few occasions, including on 13 June 2018 from VM 2. This log in to Citrix

Server 4 on 13 June 2018 took place a few minutes before the failed attempt to log into the SCM database from Citrix Server 4. The Citrix Team then identified the user of the account.

403. The Citrix Team provided the above information to Benjamin on 13 June 2018 itself. At 6:51pm on 13 June 2018, Benjamin asked SingHealth's outsourced vendor for IT services to reset the user's password, giving the reason that the credentials were being abused. At 8:24pm that same day, the vendor replied, stating that they had contacted the user and gotten his approval for the password reset, and indicating that the password reset had been carried out.

22.8 Determining that VM 2 was not a workstation issued by SingHealth

404. Earlier at 5.41pm on 13 June 2018, Benjamin emailed SingHealth's outsourced IT vendor, asking for their help in finding the location of VM 2, and stating that the workstation was "*attempting to attack (SingHealth's) servers*". By this time, Benjamin thought that they could be facing a "*genuine cyber attack*".

405. Shortly after, the vendor replied confirming that based on their checks, VM 2 was not a machine that was joined to the SingHealth domain. The vendor also suggested that VM 2 could be a personal computer.

406. After receiving this reply, Benjamin emailed a SingHealth system administrator, explaining that VM 2 was trying to attack one of the Citrix servers, and asked the administrator to check if there was an IP address assigned to VM 2. The administrator replied shortly after, stating that there was no record of an IP address having been assigned to VM 2. Benjamin then updated the members of the TigerConnect chat group on the above.

22.9 Assessment of IHiS' incident response on 13 June 2018

407. The Committee finds that Benjamin's response was, on many counts, timely and appropriate. He recognised that the incidents were suspicious, took the initiative by investigating leads and acquiring forensic images, and even purchased extra external hard drives to do his work. He also kept his superiors, Ernest and Wee, informed at all times. However, his response could have been improved in some respects:

(a) First as observed by Vivek, some of Benjamin's actions demonstrated inexperience and poor judgment. In particular, Benjamin's focus on shutting down systems that were exhibiting suspicious behaviour (*e.g.* Citrix Server 1 and the PHI 1 Workstation) led to loss of potentially valuable forensic evidence. A better practice would have been to put the systems on a quarantine network without turning off the power, for further study.

(b) Second, the fact that Benjamin was communicating over both email and TigerConnect was not ideal, as it led to fragmentation of information and confusion for recipients. Additionally, it made it hard to keep records of the information flow, as TigerConnect chats are deleted after 30 days. This was not Benjamin's fault, as no formal system of recording investigation findings was in place for use during incident response. Nonetheless, it would have been better for official modes of communication to be mandated and enforced to prevent confusion.

408. The Committee also notes that at the material time, there was no relevant playbook in the IR-SOP that could guide Benjamin in identifying the nature of and responding to the suspicious activities that IHiS staff had detected. The playbooks that were available lacked details on the tactics, tools and procedures of advanced threat actors. As stated in section 18.3 (pg 104) above, the senior levels of IHiS' management were alive to the risk of APTs from as early as August 2016. But this awareness was not reflected in the SOPs in place at the

time of the Cyber Attack. In this regard, Benjamin did well by applying himself to the problem at hand, and to come up with the appropriate responses to the best of his abilities.

409. In contrast, Wee's response was clearly inadequate. Under the IR-SOP, Wee was accountable for incident response. Despite being apprised of a series of investigations into what were, to Wee's own admission, circumstances involving a potential risk to a CII system, and by a channel used specifically for reporting of security issues and risk (*i.e.* TigerConnect), he did not make further inquiries, and instead passively waited for updates.

410. This also raises the issue of whether the incident should have been escalated by Wee to IHiS' senior management, and on to CSA. In the course of evidence, CSA has identified two facts which would, in themselves, have been sufficient reasons for escalation:

(a) First, the unauthorised logins to Citrix Servers 2 and 4, which are both systems that were directly connected to a CII system; and

(b) Second, the series of failed logins to the SCM database, which happened over short period of time and was indicative of persistent attempts at accessing a CII system.

411. The fact is that by 13 June 2018, Benjamin and Wee had been apprised of the events of 11 to 13 June 2018. Unlike Katherine and Lum who were unfamiliar with the IR-SOP and not trained in security matters, Benjamin and Wee had defined roles in security incident response under the IR-SOP, and were familiar with these roles. Wee, in particular, was both accountable for the incident response team, and responsible for escalation to the GCIO. The failure on the part of Wee and the CERT to even *consider* whether the incidents should be reported, is a cause for serious concern. It is also apposite to note at this point SMD Lead Han Hann Kwang's ("**Hann Kwang**") evidence that by 13 June 2018, the incidents were "*very suspicious*" and that he (assuming he was in Wee's

position) would have reported the matter to SingHealth GCIO Benedict on 13 June 2018.

412. Turning to Ernest, who was overseas at the time, the Committee finds that he should have nominated a covering officer. As Vivek observed:

> SIRM is a critical role and must be staffed at all times. Also, there must be a proper handover-takeover process in place to ensure the responsibilities are transferred back to the primary SIRM once he or she is back from leave. Not having a functioning SIRM for an extended period could significantly hamper the investigation as CERT team may struggle to take decisions and seek the necessary support from other external parties including the SIRT team members.

413. The issues identified by Vivek were clearly borne out in this case, where Benjamin was effectively left alone to carry out his own investigations and coordinate the incident response, to the best of his abilities and resources. The fact that Ernest remained contactable while overseas did little to help the situation: although Ernest was aware of the TigerConnect messages and able to read them, he did not do so. He simply opened the application to dismiss the notifications. If Ernest had paid some attention to the messages, he may have realised something was afoot, and could have delegated his duties to a specific officer while he was away.

414. The Committee further notes that this being a security incident, under the IR-SOP, the SIRT (Security Incident Response Team) should have been activated. The SIRT comprises not only security staff from the CERT (Computer Emergency Response Team), but also the infrastructure services lead and the application services lead. Had the SIRT been activated, there could, *at a minimum*, have been better coordination, resourcing, and leadership in the incident response. Senior management who were aware of the reporting framework and the need to escalate the matter may have done so, and CSA could have been informed at a much earlier stage.

415. However, the fact is that throughout the entire period of IHiS' response to the Cyber Attack, neither Wee nor Ernest, who each had responsibilities under the IR-SOP for leading the SIRT and coordinating the incident response, took any steps to activate the SIRT. Instead, coordination of the incident response was left to the CERT, with its staff of three relatively junior and inexperienced officers.

23 EVENTS OF 14 TO 25 JUNE 2018

23.1 Monitoring access to the Citrix servers and the SCM database

416. In the morning of 14 June 2018, Benjamin emailed Lum, Vicky, and Veerendra from the Citrix Team, his fellow CERT members, and Ernest and Wee, laying out an action plan. Vicky and Veerendra were tasked with monitoring access to the Citrix servers. Azzlan, a member of the CERT, was tasked with "*monitoring direct access attempts to the SCM database…[and] to identify rogue internal PCs*". Between 14 and 25 June 2018, IHiS staff did not detect any unusual logins or attempted logins to the Citrix servers or the SCM database. Wee, once again, "*cannot quite remember if (he) read this email*", and took no further action.

23.2 Forensic investigations into the PHI 1 Workstation and Workstation C

417. Further to Benjamin's 14 June 2018 action plan, the CERT commenced forensic investigations on the PHI 1 Workstation on 14 June 2018. On 18 June 2018, Workstation C was seized, and forensic investigations on the PHI 1 Workstation had to be stopped on that day in order for forensic investigations on Workstation C to begin. The team was unable to find any evidence of malware or suspicious activities or files on either of the workstations.

418. The CERT team was hampered by their inability to run forensic investigations of the workstations concurrently. Although the CERT had been set-up in March 2018, they had not yet been provided with workstations that were

suitable for forensic investigations. The forensic tools were in fact installed on Benjamin's personal laptop, and forensic investigations could only be done on this one computer.

23.3 Obtaining of Citrix server system event logs on 19 and 20 June 2018

419. On 19 and 20 June 2018, Benjamin worked with the Citrix Team to obtain the system event logs of a number of Citrix servers, including that of Citrix Servers 2 and 4, that were involved in the failed attempts at logging in to the SCM database on 13 June 2018. The logs for Citrix Server 1 were provided to Benjamin earlier.

23.4 Ernest's actions after his return to Singapore on 18 June 2018

420. Although Ernest was added to the TigerConnect chat group on 13 June 2018, he was on overseas leave from 9 to 17 June 2018, and did not participate in the discussions, or provide any directions to the SIRT in this time. While he received the messages sent by members of the chat group as they were being sent, Ernest simply opened the TigerConnect application in order to dismiss the notifications, and did not read the messages until 18 June 2018, when he was back in Singapore.

421. Having read the TigerConnect messages upon his return, Ernest was generally aware that the team was trying to locate workstations with unusual hostnames and had taken some forensic images, but the team was unsuccessful in their efforts to locate VM 1 and VM 2. Ernest saw the messages from Benjamin on 13 June 2018 stating that an "*incident [was] ongoing*", and that someone had "*obtained local admin credentials*" and was "*try[ing] to login to the SCM [production] database*". However, Ernest has explained that he was "*not concerned*", as "*there was nothing to be concerned about while awaiting the results of forensic analysis*" on the forensic images that had been collected.

422. Ernest was also copied in email conversations pertaining to the CERT's incident response starting from the night of 12 June 2018. However, he did not read them contemporaneously, and did so after 18 June 2018 when he returned to Singapore. Having "*glanced through*" each of the emails he was copied in, he had the following understanding:

(a) He understood that Benjamin had been communicating with Veerendra, and that an image had been taken of Citrix Server 1, and that the server was shut down.

(b) He did not notice the email from Vicky to Benjamin at 11:32am on 13 June 2018, in which Ernest was copied in, and in which Vicky informed Benjamin that she had "*attached the email from Katherine the DBA for SCM regarding the login failed attempt to the DB server*". As such, Ernest did not see the emails from Katherine containing details of failed logins to the SCM system.

(c) Ernest however did notice the email from Benjamin at 9:28am on 14 June 2018, in which he laid out an action plan (see paragraph 416 (pg 138 above). Ernest thus concluded that "*this was a case that required forensic investigations*".

(d) Ernest was also copied in the emails sent by Benjamin from 13 June 2018 onwards regarding the seizing of Workstation C for investigations (see paragraph 377 (pg 128) above). He noted that Benjamin had indicated in the first email at 4:34pm on 13 June 2018 that Workstation C had been "*involved in an IT incident*". He also noted that Benjamin had stated that the workstation "*was found to be attempting to connect to the SCM database using several different username/password combinations*."

423. Having read the TigerConnect messages and the emails he was copied in, Ernest realised that there were two workstations and one Citrix server being forensically examined at the same time in relation to suspicious attempts at

connecting to the SCM database. However, in his words, "*this did not ring any alarm bells in my mind*". Despite Benjamin having used the words "IT incident", Ernest did not view these facts as constituting a security incident.

424. It is apposite at this point to note Ernest's understanding of what constitutes a 'security incident':

> (M)y definition of a security incident is one where (a) there must be a 100% confirmation of malicious intent, and (b) the malicious activity must be successful *i.e.* an attempt is insufficient.

425. The emphasis here is on *confirmation*, both in terms of malicious intent and the success of the malicious act. Thus, to Ernest, collecting workstations for investigation was "*a fairly common occurrence*" that was "*based on suspicion*", and did not amount to a security incident because it was not yet 'confirmed'. Likewise, attempts to connect to the SCM database were mere *attempts* – unsuccessful, and therefore not security incidents.

426. On 19 June 2018, Ernest returned to work and checked with Benjamin verbally whether the forensic analysis of the endpoints was complete, to which Benjamin replied in the negative. On 20 June 2018, Ernest met with Lum and they discussed the events that had occurred. Lum informed Ernest that the L.A. account had a weak password and could have been compromised, and that the password had already been changed. Since the password was changed, Ernest did not consider there to be a reportable security incident.

427. In the period preceding 26 June 2018, it appears that there was no significant discussion between Ernest and Wee about the events of 11 to 13 June 2018. From 20 June to 3 July 2018, Wee was on medical leave, but he was able to look at the TigerConnect chat group "*on and off*".

23.5 Assessment of IHiS' incident response from 14 to 25 June 2018

428. The Committee notes the initiative shown by Benjamin and the CERT team to carry out forensic investigations. However, the pace of their investigations could have been improved with better resources. For instance, it took five days to locate and collect Workstation C, as the CERT did not possess the tools necessary to remotely collect forensic evidence. The CERT was also hampered by the fact that they only had one laptop, Benjamin's personal computer, which was capable of processing the forensic images.

429. The CERT team, being new and relatively inexperienced, would also have benefitted from firm and effective leadership from the SIRM and Cluster ISO, both in terms of the conduct of investigations and on the issue of whether to escalate the matter to the GCIO. However, no such leadership was forthcoming from Ernest and Wee.

430. Turning to Ernest, two aspects of his evidence stand-out in particular:

(a) First, Ernest stated that he gets "about 200" emails in a day. This presumably includes emails on security issues and, based on Ernest's evidence, also includes emails about issues such as end-users not getting their account IDs and passwords, and complaints or feedback about his staff or the outsourced IT administration team. Ernest has explained that he would give *equal priority* to his emails and not prioritise them based on urgency, *except* when it came to complaints.

(b) Second, Ernest's definition of a 'security incident', which requires 'confirmation' of both malicious intent and a successful malicious act.

431. In the Committee's view, it is unacceptable for someone in position of a SIRM to hold such views. Even as a matter of common sense, the difficulties are obvious: user complaints are prioritised over security matters, and a security

incident is only recognised and treated as such after the damage has been done. These are clearly misguided, and are in fact the direct inverse of their proper order.

432. At this point, one might think of looking to relevant policies and frameworks in place, such as the IR-SOP and SIRF, to identify ambiguities or deficiencies therein in order to better explain Ernest's misconceptions. While there are certainly aspects of these documents that can and should be improved, and the Committee will make its recommendations on these in Part VII below, one must not lose sight of the fact that the treatment of cybersecurity issues and incidents by staff and middle management is very much shaped by organisational culture.[29] A sense of this can be gleaned from the evidence of Hann Kwang, Ernest's reporting officer (emphasis added):

> (I)n my view, when a security incident is reported, this is not a trivial matter, and it activates a whole team, including the Cluster ISO, GCIO and senior management. Everyone will have to attend to the security incident. **If a security incident is declared when it turns out there is no security incident, this may look bad on the person who made the declaration.**

433. The Committee observes the alignment between this comment from Hann Kwang, and Ernest's emphasis on "confirming" security incidents and prioritising complaints over all other matters. The evidence suggests that the reluctance to escalate the matter may have come from a belief that it would not reflect well in the eyes of the organisation if the matter turned out to be a false alarm.

---

[29] The Committee also notes parenthetically that there is a logical difficulty with looking to the text of the IR-SOP and SIRF to account for Ernest's misconceptions, since his own evidence is that he was "*not very familiar with the contents*" of the SIRF, and that he was familiar with the IR-SOP "*but not in great detail*".

434. The above does not detract from the fact that Ernest's failure to follow-up on clear evidence of malicious activity targeting a CII and the systems connected to it was plainly unacceptable. He did not apply himself properly to the facts before him, and did not take any further action on the matter pending the completion of Benjamin's forensic investigations – which he did not provide any guidance on. These are clear failings in the discharge of his duties as the SIRM. Yet at the same time, his comments and failure to act are suggestive of deeper cultural issues within the organisation as to where priorities should lie.

24 EVENTS OF 26 JUNE 2018

24.1 Detecting a failed attempt at logging in to the SCM database from Citrix Server 2

435. At 3:10pm on 26 June 2016, a system-generated email alert was sent to Katherine, notifying her of one failed attempt at logging in to the SCM database that same afternoon. The user-ID used in this attempt was again that of the user account of Workstation A, and the attempt was made from Citrix Server 2.

436. Katherine recognised this as another attempt to gain unauthorised access to the SCM database that was being made from one of the same IP address (that of Citrix Server 2) as the attempts on 13 June 2018. She forwarded the email alert to Kelvin, Robin and Lum immediately.

24.2 Investigating further into the use of VM 2 and the S.A. account to log in to Citrix Server 2

437. Upon receiving the alert from Katherine, the Citrix Team was able to identify a suspicious login to Citrix Server 2 made earlier in the afternoon of 26 June 2018 using the S.A. account. Lum was surprised to learn that S.A. had been added back into the administrator role and could be used to log in to Citrix Server 2, despite the fact that the Citrix Team had removed the account from the administrator group on 13 June 2018. Given that adding the account back to the administrator group can only be done using an account which has administrator

privileges, and having confirmed that no one in IHiS had done so, Lum suspected that a privileged account had been compromised.

438. Shortly after on 26 June 2018, Lum logged in to Citrix Server 2 and found an active RDP session where the S.A. account was used to log in to the server. Lum checked the login logs and established that the S.A. account was logged in to Citrix Server 2 from VM 2. Lum then searched for active RDP sessions on the network and found only one connection. This connection was made from an IP address (referred to in this section as the "**first IP address**"), and Lum concluded that this IP address must have been related to VM 2.

439. At around 3:40pm on 26 June 2018, Lum sent emails to Benjamin containing screenshots showing the RDP session to Citrix Server 2 from the first IP address, and the suspicious login to the server using the S.A. account. Benjamin has confirmed that this was the first time he had heard of the S.A. account being used to log in to a Citrix server. He was not aware at this point of the earlier use of the account on 13 June 2018.

440. Shortly after at 3:47pm on 26 June 2018, Joanne replied to the same email thread with a screenshot showing the session information of an RDP session, with the client name being that of VM 2, and the user name of the S.A. account. However, the client address was another IP address (referred to in this section as the "**second IP address**"), which was different from the first IP address.

441. At 4:03pm on 26 June 2018, Benjamin forwarded the above emails to Ernest and the other members of the CERT. Wee was not copied in this email. The emails comprised entirely of screenshots, and did not contain any explanations.

442. Lum was of the view that the use of the S.A. account to access Citrix Server 2 *via* RDP was clearly unauthorised. Thus, on 26 June 2018, he removed the account from the administrator group once again.

24.3 Identifying and seizing Workstation B

443. While Benjamin was unable to understand why two different IP addresses appeared to be associated with the same RDP session, he conducted some searches and found that the first IP address was associated with Workstation B.

444. Benjamin informed Lum and the Citrix Team of the association of the first IP address and Workstation B at 3:46pm on 26 June 2018. Lum has explained that based on his own checks and the information provided by Benjamin, he thought at that point in time that the RDP connection to Citrix Server 2 came from a virtual machine (*i.e.* VM 2) running from Workstation B. This opinion was shared by Benjamin, who later informed the members of the TigerConnect chat group of his "*guess*" that the Workstation B was "*used as a victim PC to host a virtual machine*", and that the second IP address was that of the virtual machine.

445. Benjamin identified the user of Workstation B, and with the user's permission, Workstation B was seized for forensic investigations on 26 June 2018. Both a memory dump and a forensic image of the hard disk were acquired.

24.4 Imposing firewall blocks for the IP address range for the second IP address

446. As mentioned in paragraph 441 above, Benjamin forwarded Lum and Joanne's emails to Ernest and the other members of the CERT at 4:03pm on 26 June 2018. Given the lack of explanation in Benjamin's email, Ernest could not understand the emails, and asked Benjamin for clarifications. He instructed Benjamin to try tracing the source of the second IP address, which was associated with VM 2. They determined that the IP address range was not part of the SingHealth network, and they were also unable to determine conclusively if this was a valid IP address. In the circumstances, Ernest arranged for firewall blocks for the IP address range for the second IP address to be imposed as a precaution.

447. Ernest has explained that even though the IP address range was not part of the SingHealth network, he did not think that this was a security incident because he had already taken action to impose firewall blocks for the IP address range, thus preventing any access to any of the SingHealth servers from this IP address range.

24.5 Discovering background processes being run on Citrix Server 2

448. Curious about what the S.A. account was doing when logged in to Citrix Server 2, Lum reviewed the system task-list and noticed some background processes being executed. However, he was unable to ascertain what scripts were being run. Lum forwarded a screenshot of the task-list *via* email to Benjamin and the Citrix Team at 3:44pm on 26 June 2018.

24.6 Discovering the use of the D.A. account to access Citrix Server 3 from Citrix Server 2 and that the system event logs for these servers were deleted

449. On 26 June 2018, the Citrix Team reviewed the security event logs for Citrix Server 2 and discovered that the D.A. account was used to access the H-Cloud Citrix Server 3 from Citrix Server 2. As explained in Part III above, it is probable that the attacker had stolen SCM database credentials from Citrix Server 3 at this time.

450. As mentioned above in section 20.1 (pg 116), the password for the D.A. account had been changed on 11 June 2018. When contacted, the domain administrator confirmed that he had not logged into Citrix Servers 2 and 3 on 26 June 2018. On 26 June 2018, the domain administrator changed the passwords to the D.A. account again.

451. The Citrix Team also discovered on 26 June 2018 that the Windows event logs for Citrix Servers 2 and 3 were deleted earlier that afternoon. This was further evidence of malicious activity.

24.7 Discussions between Ernest and the CERT on the events of 26 June 2018

452. The TigerConnect chat logs show that around 4:20pm to 4:50pm on 26 June 2018, members of the chat group were discussing the use of the S.A. account. Ernest was enquiring why the account could be used to log in to the server *via* RDP, stating that such a means of access was "*weird*", and asked if "*even tat acct got prob?*". Benjamin replied that it was "*possible the attacker guessed the password*". Shortly after, Ernest replied stating "*guys pl secure yr citrix accts…please…they oredi know yr passwords*".

453. While the face of the chat logs strongly indicates that there was awareness amongst the CERT and Ernest that they were dealing with an attacker, Ernest has sought to explain that (i) Benjamin was merely raising the possibility that an attacker guessed the password, but Ernest himself did not believe this, and (ii) Ernest's own statement that "*they oredi know yr passwords*" was "*made up*" by himself, "*just to pressurise the Citrix Team to quickly secure the accounts*". Ernest's account is that he did not think they were dealing with a security incident at that point in time.

454. Likewise, Wee, who was then on medical leave, occasionally glanced at the updates sent by Benjamin in the TigerConnect group, but took no further action.

24.8 Assessment of IHiS' incident response of 26 June 2018

455. To begin with, the events prior to 26 June 2018 were already highly indicative of a security incident. In the face of the events of 26 June 2018, it should have been abundantly clear that IHiS was facing a serious security incident that had to be reported. CSA has identified three facts which, in particular, underscore the seriousness of the events of the day:

(a) First, the S.A. account, which had been removed from the administrator group on 13 June 2018, was added back to the

administrator group and used to log in to Citrix Server 2 over RDP. This was a strong indication of malicious activity as an administrator account had been compromised to perform this function.

(b) Second, the D.A. account was used to access Citrix Server 3 from Citrix Server 2. This was the first time that a H-Cloud Citrix server had been accessed without authorisation (as the account owner confirmed that he had not accessed the server). Furthermore, the D.A. account had been used in this manner even though the account holder had changed the password on 11 June 2018. Given the circumstances, this was clearly an unauthorised access into Citrix Server 3, which was directly connected to a CII system, and, in CSA's view, should have been reported to CSA as a security incident.

(c) Third, the Windows event logs for Citrix Servers 2 and 3 were deleted using the D.A. account. The natural inference would be that the entities behind the malicious activities was deliberately trying to cover its tracks. Yet the significance of this seemed to have been lost on Ernest and the IHiS teams: other than changing the password again for the D.A. account, no further steps were taken to investigate into the compromise of the D.A. account or the deleted logs.

456. IHiS CEO Bruce's own evidence is that the fact that the SCM database was facing a deliberate attack would have been "*firmly established by 26 June 2018*", and that on that basis Ernest and Wee should have reported the incident accordingly.

457. In respect of the events of 26 June 2018, Wee may be given the benefit of the doubt in light of the fact that he was on medical leave.

458. However, the Committee finds that Ernest was wholly irresponsible in his attitude towards the cumulative facts before him. Ernest's failure to report the matter on 26 June 2018 proved to be the last and most significant 'missed opportunity' to prevent the attack. Having obtained the SCM database credentials from Citrix Server 3, the attacker began stealing and exfiltrating patient data from the SCM database over the course of the next few days.

25 EVENTS OF 27 JUNE TO 3 JULY 2018

25.1 Further investigations into Workstation B

459. In the days following 26 June 2018, Benjamin forensically analysed the memory dump of Workstation B, and Zac examined the hard disk image. Zac was unable to locate any evidence of malware or other suspicious activities or files on the hard disk. Antivirus programs also did not indicate any malware infection or suspicious files on Workstation B. He did, however, notice something unusual about the way certain background processes were being run on the workstation.

460. At 4:47pm on 28 June 2018, Benjamin sent an email to Ernest and the rest of the SMD; Wee was not copied in the email. This email contained some of Benjamin's forensic findings on Workstation B as described above in this section, and a set of Microsoft PowerPoint slides containing screenshots supporting these findings. There was, however, little explanation on how to interpret the screenshots.

461. Ernest has informed that he received the email, but he did not understand the information provided by Benjamin in his email. Ernest also did not understand the screenshots in the slides, which appeared to him to be "forensic-related", and he was not trained in digital forensics. However, he did not ask Benjamin for clarifications.

25.2 Assessment of IHiS' incident response from 27 June to 3 July 2018

462. It is important to note that in this period, the attacker was in fact making SQL queries on the SCM database and exfiltrating the stolen patient records. However, there is no evidence that IHiS staff had detected any suspicious activities in this period. As discussed in section 15.2 (pg 74) above, this was because of a lack of monitoring at the SCM database for unusual queries and access.

463. Benjamin's investigations into Workstation B and his identifying of unusual processes were steps in the right direction. Unfortunately, he was unable to fully appreciate the security implications of his findings, or to associate them with earlier findings. It is likely that this was due to the limited training and experience that he had. It also did not help that his reporting officer, Ernest, similarly lacked the necessary technical knowledge and experience, and did not take any steps to find out more about the findings which Benjamin presented to him.

464. The above clearly illustrates the importance of timely reporting and escalation. Had the matter been escalated to a level which could provide effective leadership, and which possessed the appropriate resources and technical expertise, it may have been possible to determine, from all the suspicious activities to-date, that the attacker was targeting the SCM database and sought to exploit the open network connection between the SGH Citrix Servers and the SCM database. With timely action, the attack could have been detected and contained, minimising the damage caused. Unfortunately, the matter was not escalated and valuable time was wasted. On the facts, further suspicious activity was only discovered on 4 July 2018, eight days after the attacker first began querying the database.

26 EVENTS OF 4 JULY 2018

26.1 Discovering queries to the SCM database

465. In the afternoon of 4 July 2018, Chai Sze Chun ("**Sze Chun**"), an Assistant Lead Analyst in the Production Enhancement Team supporting the SCM application, received some text alerts triggered by scripts that he had put in place to monitor the SCM database server.

466. Intending to make sure that there was no persistent issue with the SCM database server, Sze Chun decided to look into what triggered the alert. In the course of investigations, Sze Chun checked the queries that were currently running at the time, and a particular query caught his attention. He checked back after a while, and the particular query was no longer running. He also did not receive any significant alerts from his scripts thereafter for the rest of 4 July 2018.

467. Prior to 4 July 2018, Sze Chun had not seen queries similar to this particular query, which was in fact one of the bulk queries run by the attacker. Although the query was no longer running, Sze Chun decided to investigate further.

468. Noticing that Citrix Server 2 and Workstation B were associated with the query, Sze Chun thought that he could find out the user-ID of the user that was logged in to Citrix Server 2 at that point. However, Sze Chun was unable to find the necessary information. Sze Chun then asked Robin to provide the logs of users who had logged in to Citrix Server 2 on 4 July 2018. Sze Chun received and reviewed the logs, but found no record of Workstation B having logged in to Citrix Server 2 on 4 July 2018.

469. Sze Chun also noticed that the A.A. account was associated with the query, which appeared unusual to him. The query also being run from a program which was unusual (referred to in this section as the **"first program"**). At this point in time, Sze Chun thought that the unusual query may not have been run from the SCM application, and had in mind four possibilities:

(a) First, the query could be a legitimately run query which originated from some other automated services;

(b) Second, the query could be a legitimately run query which originated from a SQL-linked server;

(c) Third, the query could have been run by some external program that he did not have knowledge of and that could run on the Citrix servers; and

(d) Fourth, some individual had run the query from the Citrix server.

470. The third and fourth possibilities would have meant that there was some possible misuse of the A.A. account. However, as Sze Chun was still trying to figure out what was happening and could not be sure that the A.A. account was being misused, it did not occur to him to report the matter to the SMD at that point.

26.2 Informing Katherine and the Citrix Team

471. Having reached a dead end, Sze Chun emailed Katherine, Kelvin, Lum, Loo, and Robin to seek their advice in the afternoon of 4 July 2018. In his email, Sze Chun stated that there was a query made using the A.A. account from Workstation B, and provided the text of the query along with some screenshots supporting his findings.

472. Several minutes later, Sze Chun replied to the same email thread, stating that he "*may be contacting assistance to make a visit to that PC*", referring to Workstation B. Sze Chun has explained that there was "*a sense of urgency in (his) mind"*, because if the query was not run with permission, it would mean that the A.A. account had been misused. By then, he had learnt that Workstation B was deployed in SGH, and he thought to ask someone there to find out if there was an explanation for the query.

26.3 Detecting active queries to the SCM database

473. Upon receiving Sze Chun's email, Katherine noticed that the Citrix server in question was Citrix Server 2. She also noted that the account, workstation, and program that were involved were unusual.

474. Katherine then logged in to the SCM database to look at the current active sessions running on the database. Her intention was to check if the query mentioned by Sze Chun was still running. She found that very similar queries were being run. The active sessions reflected the hostname of VM 1, and the queries were being run using a different program (referred to in this section as the "**second program**"). She was of the view that this was indicative of abnormal activity, and called Benjamin to inform him of the active sessions.

475. Following the call, Benjamin checked the active sessions and found the same queries described by Katherine. He thought that perhaps it could be some new modules that were deployed or being tested, and which used the second program to run queries on the SCM database. Benjamin called some colleagues, who confirmed that they were not running any queries, and that they were unaware of the second program.

476. At 3:21pm, Katherine emailed Sze Chun with a screenshot showing the SQL sessions on the SCM database involving the second program and running from VM 1. Katherine asked Sze Chun in her email why the hostname was that of VM 1. Shortly after at 3:27pm, Katherine also pointed out that that the query had been running since 3:06pm, and was still running. At 3:30pm, Sze Chun replied Katherine, confirming that the query she identified was still running, and that there was a second query that was running at the time.

477. By this point, Sze Chun had become more concerned, as the probability of the A.A. account being misused appeared to be higher, in view of all the unusual circumstances.

26.4 Terminating unusual queries to the SCM database

478. By this point, Sze Chun and Katherine, being unsure who was running the ongoing queries, decided they should immediately terminate the queries, and wait and see if they received any calls from any affected users or colleagues.

479. Later on 4 July 2018, further to discussions between IHiS staff including Henry, Kelvin, Sze Chun, and Katherine, it was also decided that they would immediately terminate any similar queries as they may arise. A few more such queries were terminated over the course of 4 July 2018. Katherine and the Applications Team never received any calls from any users or colleagues complaining about terminated queries or sessions.

26.5 Attempts to locate Workstation B and linking up with Benjamin

480. At around the same time on 4 July 2018, Sze Chun took steps to ascertain the exact physical location of Workstation B. He was informed that Workstation B had been confiscated by the SMD, and was directed to Benjamin.

481. Sze Chun then met with Benjamin, who informed Sze Chun that the workstation was not connected to the network, and was with the Security Management Department for investigations. Sze Chun then informed Benjamin that Workstation B had been detected as having executed SQL queries to the SCM database on 4 July 2018, and showed Benjamin the details of the SQL queries that had been run from Workstation B and VM 1.

482. Seeing as Workstation B was with the CERT and could not have run the SQL queries, Benjamin thought that "*this could not be happening*", and wanted to immediately escalate the matter to Ernest.

483. The TigerConnect chat logs show that at around 4:32pm on 4 July 2018, Benjamin addressed Ernest in the TigerConnect chat group, stating that he had met with Sze Chun, and that "*we really need to escalate into incident ... Seems*

like someone managed to get into SCM db already ... Attack is going on right now".

484. At Sze Chun's workstation, Benjamin collected some screenshots and pasted them in a set of Microsoft PowerPoint slides and named the file "*SCM Breach.pptx*". These screenshots showed the details of the suspicious queries, including the hostname, program name, query run, and user-ID used. Copies of SCM Breach.pptx were shared with Benjamin, Ernest, Wee, the rest of the CERT, Lum, Kelvin, and Katherine on 4 July 2018. However, Benjamin did not provide any explanation of the slides.

26.6 Comparing and drawing links between the uses of Workstation B in June 2018 and 4 July 2018

485. At 4:45pm and 4:52pm on 4 July 2018, Lum called Benjamin and spoke with him briefly about the suspicious queries to the SCM database that Sze Chun had identified. They also discussed Workstation B, which had been identified as the machine that was used to run the suspicious queries on that day. Lum and Benjamin observed that Workstation B:

(a) Was the same workstation that had been discovered to be used to log in to Citrix Server 4 earlier without authorisation;

(b) This was the same hostname that they suspected to have been used to run a virtual machine, VM 2, that was used in connection with an unauthorised RDP session into Citrix Server 2 on 26 June 2018; and

(c) Had been seized by the SMD earlier on 26 June 2018.

486. At that point in time, Lum guessed that Workstation B was being spoofed to run as a virtual machine. It was also unclear to him how Citrix Server 2 was being accessed, as the login logs did not show either the L.A. or S.A. accounts being used.

487. On his part, Benjamin had sent the SCM Breach.pptx to Ernest (see section 26.5 (pg 155) above). There is, however, no evidence showing that Benjamin had specifically informed Ernest at this point of their observations regarding the role of Workstation B.

26.7 Further investigations by Ernest into the SQL query and the use of the A.A. account

488. It happens that at around 5:00pm on 4 July 2018, Wee was at Ernest's desk discussing another matter when Ernest received the email enclosing a copy of SCM Breach.pptx from Benjamin. They looked through and discussed the slides, trying to understand the information contained therein. Wee did not notice the program that had been used to run the query or the user-ID that had been used to access the SCM database. He also did not understand the syntax of the SQL query. On the other hand, Ernest had a rough understanding that the query was seeking to select a large number of records from a particular database table.

489. Ernest and Wee then went to see Katherine in person, knowing that she was a SCM database administrator. Ernest asked Katherine to explain the SQL query, including what the query could do and the records it could receive. Katherine directed Ernest to Vida Junitha ("**Vida**"), an IHiS IT administrator whom Katherine knew would have knowledge of the table that was being queried.

490. Vida informed Ernest that the database table contained data which was "*obsolete*". Ernest did not clarify further what she meant by this. Vida also informed Ernest that the database table contained information about dispensed medication. She performed a "*sample query*" to retrieve one record from the database table in question. Ernest reviewed the retrieved record, and realised that it "*contained staff information*".[30]

[30] Evidence was not led on what this meant. Based on the context, it appears likely that Vida's "sample query" was made using the identfiers of a member of staff, and that Vida and Ernest had found that the test query did in fact retrieve the records of this staff member.

491. Looking at the screenshots in SCM Breach.pptx in greater detail, Ernest realised that the queries were run using the A.A. account. Ernest also noticed from the screenshots that the query was made from the second program, which was unusual.

492. At this point, Ernest only knew of the query reproduced in SCM Breach.pptx, and did not know that there were similar queries being run repeatedly and being terminated by Katherine as they were being run. But based on the use of the A.A. account and the second program, "*alarm bells started ringing*" for Ernest, and thought that the query was suspicious. He and Wee thus arranged for a meeting to be held on the morning of 5 July 2018.

26.8 Ernest's reasons for not reporting the incident

493. Ernest has explained that while he had received Benjamin's messages on the TigerConnect group (see paragraph 483 (pg 155) above), he "*did not see any reason to report the incident upwards*", and "*did not agree with Benjamin that the matter needed to be escalated*".

494. As explained above, Ernest held that view that there must be 'confirmation' of both malicious intent and a successful malicious act before a matter is considered a 'security incident'. He has further elaborated that a security incident would be 'reportable' only after obtaining all of the following additional information:

(a) All the information about the impact of the attack;

(b) The identity of the attacker;

(c) Where the attack is coming from;

(d) Whether the attacker is an 'internal' or 'external' attacker, *i.e.* whether the attacker is a SingHealth user, or whether the attacker is from outside of SingHealth;

(e) Whether data in the SCM database had in fact been accessed; and

(f) Whether there was more than one instance of access to the SCM database.

495. Ernest has explained that despite the fact that access to the SCM database would have meant that patient data had been accessed (*i.e.* item (e) above), the events of 4 July 2018 "*just aroused (his) suspicions*" and "*still did not rise to the level of a reportable security incident*", as he had to obtained all other necessary information in (a) to (f) above.

496. As such, upon receiving Benjamin's messages, Ernest did not agree that the matter had to be escalated, and simply told Benjamin to "*continue to investigate and isolate*". This remained his view even after he found out more about the use of the A.A. account and the second program.

26.9 Wee's reasons for not reporting the incident

497. Wee was with Ernest when they spoke with Katherine and Vida (see section 26.7 (pg 157) above). Although he saw that the test query had returned a record, and heard that the database table had "*something to do with medication*", Wee erroneously thought that the record did not contain any sensitive or up-to-date data, and assumed that the records that the SQL query was seeking to retrieve similarly would not return any sensitive data. He also heard from Katherine that the second program was used, and this was not a tool that database administrators used.

498. Despite the above, Wee did not seek any clarifications from Benjamin on the matters raised in SCM Breach.pptx, or to take further steps to investigate or clarify what he saw as a "*potential breach*". Wee also did not did not make any suggestions to Ernest on the investigations.

499. In Wee's view, if there had been a breach in the SCM, it would have been a Category 1 security incident, and he would have to report the incident to the

GCIO and the Sector Lead (*i.e.* the CSG). Although Benjamin's slides were titled "*SCM Breach.pptx*", Wee was of the view that it was a "*potential breach*", as it was "*not confirmed*"; his understanding at the time was that based on the IR-SOP, only a 'confirmed' breach of CII would need to be escalated. Wee was also of the view that since that Ernest was still investigating, it would not have been appropriate for Wee to report the matter to the GCIO at the time.

26.10 Query from Katherine about reporting the matter

500. At about 4:23pm on 4 July 2018, Sze Chun created a WhatsApp chat group for quicker communication. The chat group was titled "*unknown access*", and members of this chat group were Sze Chun, Katherine, Kelvin, Robin, Lum, Loo and Sze Chun's reporting officer, Kuah Peng Ann Steven. Members of the SMD were not included in this chat group. The chat group was used for information sharing and coordination between members, including for the terminating of on-going unusual queries.

501. At about 4:30pm on 4 July 2018, Katherine informed her immediate superior, Teresa Wu Rong-Jang, about the unusual queries running on the SCM database. It occurred to Teresa that they may be dealing with a security incident, and showed Katherine a single PowerPoint slide titled "*IT Security Incident Management – Reporting Flow and Timeline*". This slide in fact reproduced the reporting timelines and reporting flow found in the SIRF. Katherine understood this as directing staff to inform their HOD and the Cluster ISO in the event that they encountered a security incident.

502. Further to Teresa's directions to check if a report should be made, Katherine sent a message to the "unknown access" Whatsapp chat group at 4:41pm on 4 July 2018, asking if "*Please decide if need to rpt?*", and attaching a copy of the slide which Teresa showed her earlier. At 4:44pm, amidst discussions on the chat group on logins to Citrix Server 2, Sze Chun replied asking "*the reporting is .?*" Thereafter, the conversation turned to the termination of queries and the use of the A.A. account, and there was no response to Katherine's query on reporting.

503. Katherine did not press for an answer to her query. By this time, she was of the view that IHiS was dealing with a security incident, but she did not personally report the matter to anyone else. Neither did she associate the active queries with the failed login attempts in June 2018, in spite of certain indications, such as the fact that the database being queried was the same database on which failed login attempts were made repeatedly on 11, 12,13, and 26 June 2018, and the fact that VM 1 was used in attempts to log in to the SCM database both on 13 June 2018 and 4 July 2018. After Ernest spoke with her that afternoon, she was of the view that there was no further need for her to make any additional report.

504. On Sze Chun's part, when he first saw Katherine's question and the screenshot, he did not know who the incident should be reported to, as he was not aware of any incident reporting framework. Furthermore, Benjamin from the SMD was already aware of the situation. Sze Chun explained that he "*did not make any conscious decision, one way or the other, as to whether the incident should be reported*", and that he "*was focused on trying to stop the queries at that point in time.*"

26.11 Preventing further queries to the SCM database from the SGH Citrix servers

505. On the night of 4 July 2018, IHiS staff decided that something had to be done to prevent further queries or access to the SCM database from the SGH Citrix servers. To this end, Lum took steps on the night of 4 July 2018 to prevent such access.

26.12 Implementing scripts on the SCM database to block malicious queries

506. On the night of 4 July 2018, Katherine and Sze Chun worked together on a script to be implemented on the SCM database to block malicious queries and to alert them when any such queries were made. Katherine has explained that ordinarily, a Change Request has to be made in order for a script to be made and

implemented, but Katherine decided to bypass the usual procedure as manually terminating the queries was too time-consuming.

507. The script was completed and implemented at around midnight on 5 July 2018. Following its implementation, there were no alerts from the script.

26.13 Changing the password of the A.A. account

508. On the evening of 4 July 2018, Kelvin recommended to Henry that the password for the A.A. account should be changed. Henry agreed, but on the condition that the password change be tested in the development environment of the SCM system before being implemented in the production environment.

509. Kelvin proceeded to change the password for the A.A. account on the night of 4 July 2018 in the development environment of the SCM system. The change was subsequently implemented in the production environment on 8 July 2018.

26.14 Assessment of IHiS' incident response on 4 July 2018

510. The Committee commends Sze Chun for his initiative. Notwithstanding the fact that that security was *not* within Sze Chun's usual job scope and that he was unaware of any written procedure for responding to and reporting security incidents, he responded quickly and thoroughly upon noticing something unusual. He went so far as to terminate the unusual queries, and to locate the workstations from which the queries were being run. Sze Chun's actions exemplify an important principle of cybersecurity as identified by Dr Lim Woo Lip – that cybersecurity is the problem of every member of an organisation, and should not just be left to the dedicated security staff to handle.

511. Benjamin too was acting with a high degree of initiative and autonomy. That said, the process by which he communicated the results of his investigations to the rest of the SMD could be improved. For instance, Ernest and Wee were unable to understand the slides sent to them, as the slides were sent without

explanation. Further, Benjamin was still communicating in a somewhat fragmented manner, over both email and TigerConnect chat. In Vivek's expert opinion, such problems could have been mitigated by consolidating communications in a single, formal channel, to prevent fragmentation of information and facilitate ease of understanding.

512. On the other hand, Ernest's response was severely inadequate. As the SIRM, Ernest was expected to lead and coordinate the incident response, and also to decide on whether to escalate the matter. Up to 4 July 2018, Ernest had not properly applied himself to the events. He was aware of Benjamin's investigations and updates *via* TigerConnect and email, but withheld any further action pending 'confirmation' of a security incident. On 4 July 2018, Ernest finally realised something out of the ordinary was happening. But even so, his response from this point onwards left much to be desired. As Vivek has observed:

(a) Under the IR-SOP, it is the responsibility of the SIRM to lead and coordinate activities during an incident response but there was no formal coordination happening between the different teams. This wasted valuable time without making any real progress.

(b) Under the IR-SOP, the SIRM needs to report the incident up the command line so a formal incident can be declared, and all available resources can be deployed or re-deployed to respond to the incident. However, no formal incident was declared and therefore key experts and stakeholders kept operating in silos (or remained un-informed), which significantly hampered the incident response.

513. In respect of (b) above, Ernest's view on the information that is required before a security incident is 'reportable' is equally, if not more, unacceptable next to his misguided view of what constitutes a 'security incident'. By his definition, it would be necessary to obtain *all* information about the attack, including its source and impact, and the identity of the attacker, before a security

incident is considered 'reportable'.[31] Common sense alone would inform us that this cannot be the case. As before, and without detracting from Ernest's clear failures in understanding and discharging his duties, the Committee questions the manner and extent to which his views have been shaped by the organisational culture in IHiS.

514. Turning now to Wee, his response was also clearly lacking, and displayed an alarming lack of concern. First, by this point it was already clear that a CII system had been potentially breached. Wee should have recognised this as a Category 1 reportable security incident and taken steps to escalate the matter immediately. Yet he did not do so, and effectively abdicated to Ernest the responsibility of deciding whether to escalate the incident. Second, under the IR-SOP, Wee was also accountable for the actions of the SIRT. Yet he did nothing, and simply left Ernest and the rest of the SIRT to their own devices in the investigation of the matter and remediation efforts.

515. To sum up, considerable initiative was shown by officers on the front line, including Sze Chun, Katherine, and Benjamin. It is a shame that such initiative was then smothered by a blanket of middle management mistakes, by the likes of Ernest and Wee. Despite the fact that "*alarm bells*" had started ringing, Ernest and Wee failed to take any further action to escalate the matter and seek further assistance, instead leaving the SMD personnel and the Citrix Team to continue their investigations and remediation in much the same manner as before over the next few days.

[31] Further details of Ernest's views on what information is necessary before a security incident is reportable can be found in paragraph 494 (pg 171) above.

27 EVENTS OF 5 TO 8 JULY 2018

27.1 Meeting at 9:00am on 5 July 2018 between the Security and Citrix Teams

516. At around 9:00am on 5 July 2018, a meeting was led by Ernest, with Wee and members of the CERT also in attendance, and with the Lum and the Citrix Team calling in. The attendees discussed the events of 4 July 2018, and the focus was on the security of the Citrix servers. Ernest was also trying to relate the events of 4 July 2018 with the earlier events of 26 June 2018.

517. With the understanding that RDP had been used to access Citrix Server 2 and that there was no hardware firewall between end-user workstations and Citrix Server 2, Ernest asked Lum if the built-in Windows firewall could be used to block RDP. Such a firewall rule was in fact instituted later on 5 July 2018.

518. The Citrix Team also changed their administrator passwords on the advice of the Security Management Department, out of concern that the passwords may have been compromised.

519. The fact that based on the logs neither the L.A. nor the S.A. accounts were used to log in to Citrix Server 2 was also discussed.

520. Thereafter, Lum left the meeting and the discussion turned towards to forensic investigations that were being carried out. The CERT informed Ernest that investigations were still ongoing, but they had not found anything suspicious. It was also recognised that a problem they faced was that they only had one computer, Benjamin's personal computer, on which digital forensic examinations were carried out.

521. On Wee's part, he recalled that the discussion was about the SGH Citrix servers and the use of the S.A. account to log in to the servers. He did not think the use of the S.A. account was a security incident, and "*did not probe further as to why there was a need for strengthening of the SGH server security*", simply

on the basis that the S.A. account "*was a valid one*". He also did not link the contents of this discussion with the SQL queries that he was informed of the previous evening. Accordingly, he took no steps to report the matter, and left it to Ernest and the team to investigate further.

522. Benjamin asked at the meeting whether the matter should be escalated to IHiS senior management, in light of everything that had happened in June 2018 and on 4 July 2018. However, Ernest took no such steps to do so, despite the fact that he was, by his own account, "*bordering on the conclusion that this was a security incident*".

27.2 Detecting an active login to Citrix Server 2 and disabling the S.A. account on the morning of 5 July 2018

523. Midway through the meeting on the morning of 5 July 2018, Joanne noticed that there was a 'live' active session by the S.A. account connecting to Citrix Server 2 *via* RDP from VM 1. Lum observed that:

(a) VM 1 was the same hostname that was discovered to have accessed Citrix Server 1 using the L.A. account as early as 8 June 2018; and

(b) the S.A. account ought not to have the privileges to login to the server following its removal from the administrator group on 26 June 2018.

524. Lum told Joanne to terminate the RDP session immediately. He also informed Ernest of what they observed, and made it clear that this was an unauthorised access to Citrix Server 2. After Joanne terminated the session, the same RDP session reappeared a few minutes later and was again terminated.

27.3 Implementing a firewall rule to block all connections to the SCM database from any SGH Citrix server on 5 July 2018

525. On 5 July 2018, the Citrix Team implemented a firewall rule which blocked all connections to the SCM database from any SGH Citrix servers to ensure that the SGH Citrix servers could not be used to access the SCM database.

27.4 Enforcing the use of Privileged Access Management to access the SGH Citrix servers from 5 July 2018

526. The Citrix administrators were also told to access the SGH Citrix servers using only Privileged Access Management ("**PAM**"). The use of PAM required 2-factor authentication.

27.5 Forensic examination of Workstation B

527. On 5 July 2018, Benjamin conducted further forensic investigations into the memory dump and hard disk image of Workstation B using forensic tools.

528. For the forensic investigation of the memory dump, Benjamin detected a suspicious process and file. He took a memory dump of the process and performed an analysis using an online service, which indicated that this was an unsafe file.

529. Benjamin also searched the memory dump for more unusual background processes, given what he had learnt previously. Again, he found that there were other suspicious background processes, and analysed them using online tools. The results of one tool indicated "*malicious_confidence_80%*", and another tool indicated that this was an unsafe file.

530. Benjamin prepared a report of his findings from the memory dump of Workstation B. He updated Ernest orally of his findings, and also showed Ernest the report. For the forensic investigations into the hard disk image of Workstation B, Benjamin also made a number of findings from this.

531. On 12 July 2018, Benjamin prepared a consolidated report containing his findings from the forensic examination of the hard disk image of Workstation B, and a summary of his earlier findings from the examination of the memory dump.

27.6 Sze Chun discovering on 5 July 2018 that SQL queries were made to the SCM database since 27 June 2018, and informing Ernest of the same

532. On the morning of 5 July 2018, Sze Chun decided to investigate further to determine the earliest date on which unusual queries had been run on the SCM database. He found that the earliest date on which such queries were made was 27 June 2018, and that there had been many such queries between 27 June and 4 July 2018.

533. On that morning, Sze Chun sent an email to Lum, Katherine, Kelvin, Loo, and Robin, containing details of the queries made from 27 June to 4 July 2018. Sze Chun asked the recipients for advice on his findings and information regarding the whereabouts of VM 1. He also stated that he was "*trying to understand when [Workstation B] was confiscated*".

534. No one from the SMD was copied in Sze Chun's 5 July 2018 email. However, on 5 or 6 July 2018, Ernest approached Sze Chun to discuss the events of 4 July 2018, and Sze Chun informed Ernest then of his findings, including the fact that the earliest query dated as far back as 27 June 2018. Ernest told Sze Chun to hold on to findings, and something along-the-lines of "*not to jump to conclusions yet*". Sze Chun's understanding was that the SMD was still investigating. Sze Chun also did not have any other contact with Ernest after this meeting.

27.7 Series of measures taken on 6 and 7 July 2018 to secure the domain administrator accounts and domain controllers

535. At about 6:00pm on 6 July 2018, Ernest called Raymond Sun Xiang ("**Raymond**"), the Assistant Director of the Data Centre Team, and told him about the issues faced by the Citrix Team and the measures that had been proposed to Lum. Ernest also arranged to meet Raymond on 9 July 2018. Following the call, the Active Directory Team undertook a series of measures to secure the domain administrator accounts and domain controllers. It does not appear that the members of the Active Directory Team were informed of the events of June and July 2018.

27.7.1 Creating a new set of domain administrator accounts and removing the old accounts from the administrator groups of their respective domains

536. At around 7:00pm on 6 July 2018, Roy created a new set of domain administrator accounts. These new accounts were added to the administrator group and given administrator rights. At that point in time, every domain administrator had two accounts, one old and one new. At around 10:00pm on the same day, the old domain administrator accounts were removed from their respective administrative groups, but were not deleted.

27.7.2 Performing full antivirus scans on all domain controllers

537. At around 8:00pm on 6 July 2018, Chan Chee Choong ("**Chee Choong**"), the manager of the Active Directory Team, performed full antivirus scans on all the domain controllers under his charge. This was occasioned by the fact that one of the domain controllers was found to have been infected by a virus. The result of the scans was that the domain controllers were "clean".

27.7.3 Creating and enforcing a GPO to block the access of domain administrator accounts to servers

538. On 7 July 2018, on Raymond's instructions, the Active Directory Team created a new GPO (Group Policy Object) to block the access of domain administrator accounts to servers in their respective domains. The intention was that domain administrator accounts were not to be used to log in to servers in their domain at all during that period. This, however, was distinct from removing the domain administrator accounts entirely, which was not done.

539. The team implemented the GPO and specifically selected the option to 'enforce GPO', which is not something that is done usually. This was done with the intention that the GPO should be implemented on all servers regardless of whether the server had been set to block policy inheritance. However, there was no way for the team to tell if the GPO was successfully implemented across all the servers, as there is no status report generated. The Active Directory Team simply sampled a few of the servers they managed in order to confirm that the GPO had been implemented. They did not sample the Citrix servers.

27.7.4 Creating and implementing a GPO to prevent remote connections to domain controllers

540. At around 1:00am on 7 July 2018, Chee Choong created another GPO to prevent remote connections to domain controllers from domain clients using domain privilege accounts. Prior to this GPO, a domain administrator would be able to connect remotely to the domain controller from any machine.

27.8 Ernest's continued refusal to escalate the matter on 6 July 2018

541. At around 10:50am on 6 July 2018, Benjamin again raised the issue of escalating the matter, stating on a TigerConnect chat group: "*Ernest, the scope of compromise is quite wide now..[a domain administrator's] account was compromised before Citrix servers were compromised. I would suggest getting a 3rd party at this point to come in*", and that, based on their observations relating

to Roy's accounts, that IHiS' "*entire infra has been compromised…Followed by Citrix, and successful login and queries to our scm…*"

542. In reply, Ernest stated "*as mentioned, we need to isolate, contain and defend first...our tightening by infra is not strong enough.. even if we report now bring down the experts, they'll say our tightening is not well done...once we escalate to mgt, there will be no day no night... everyone I meant everyone in IHiS will be working non-stop on this case...*" Ernest has given an explanation for his reply:

> When I referred to management in this message, I was referring to GCIO Benedict. At the time I sent this message on 6 July 2018, it had occurred to me that I should report the incident to management. Nevertheless, I did not report the matter. I did not report because my focus was on isolating, containing and defending. I was so busy with this that I did not escalate to management about the security incident. In fact, I thought to myself, "If I report the matter, what do I get?" If I report the matter, I will simply get more people chasing me for more updates. If they are chasing me for more updates, I need to be able to get more information to provide to them. The moment I report the security incident, the clock will start ticking as per the time lines indicated at p 11 of the IR-SOP… I avoided reporting the matter as soon as it occurred to me to report it, because the clock will start ticking. Having to provide these updates on these timelines puts a lot of pressure on my team - CSA, CSG, MOH, IHiS and SingHealth senior management, GCIO and CISO will all want more information, and all of this pressure will be on my team…"[32]

[32] In context of Ernest's oral evidence, the term "CISO" was intended to refer to Cluster ISO Wee Jia Huo.

27.9 Arranging to meet Woon Lan on 9 July 2018

543. It appears that by Friday 6 July 2018, Woon Lan had come to know that something was amiss, and asked Ernest on 6 July 2018 for a meeting the next day, 7 July 2018, to inform Serena of what was going on. Ernest had declined to meet on 7 July 2018 because he was "*too stressed to work that weekend*". In recounting this exchange with Woon Lan to Benjamin over TigerConnect, Ernest told Benjamin that he "*told (Woon Lan) not 2 kajiao me in d wkend cos I stressed up. Dun wanna meet on Sat*." In his oral evidence, Ernest explained that his stress arose from his mother being admitted to the Accident and Emergency Department of a hospital on the night of 6 July 2018. A meeting was instead scheduled for 1:00pm on 9 July 2018.

27.10 Assessment of IHiS' incident response from 5 to 8 July 2018

544. Based on Ernest's replies on TigerConnect and the evidence he has given in the course of the Inquiry, it is clear that by 6 July 2018 at the latest, Ernest had himself come to the view that they were facing a security incident that should be reported. However, Ernest wilfully delayed reporting on 6 July 2018 because he felt that additional pressure would be put on him and his team once the situation became known to management. In his own words, once the matter was escalated to management, "*there will be no day no night"*. Thus, he procrastinated and as a result remediation efforts were unnecessarily set-back by several days.

545. Wee's silence on the issue of reporting during this period is also deeply unsatisfactory. Based on his evidence regarding the meeting at 9:00am on 5 July 2018, Wee remained oblivious to the significance of events, viewing the unauthorised logins to Citrix Server 2 using the S.A. as not something of concern, and failing to see any link between such unauthorised logins to Citrix Server 2 and the SQL queries discussed on 4 July 2018. Once again, he displayed his characteristic passivity and aloofness.

546. There is also no evidence of Ernest or Wee providing any significant degree of leadership in sharing information and coordinating investigations and

remediation efforts across the various IHiS teams. The teams continued to operate largely in silos, and no meeting to consolidate all the findings and to decide on a concrete way forward was held until 9 July 2018, and, even then, apparently under the direction of Woon Lan.

28 EVENTS OF 9 JULY 2018

28.1 Shutting down Citrix Server 2

547. On 9 July 2018, at around 12:45pm, the Citrix Team shut down Citrix Server 2. Lum was unable to recall who gave the instruction to do so.

28.2 Meeting amongst various members of the Infrastructure Services Division at 1:00pm

548. At around 1:00pm on 9 July 2018, Ernest coordinated a meeting attended by members of the SMD, the Citrix Team (including Lum), the Active Directory Team, Raymond, Wee, and Woon Lan, in her capacity as Deputy Director of the Infrastructure Services Division. The purpose of the meeting was to discuss the events of June and July 2018, including the unauthorised access to the Citrix servers and attempts to login to the SCM database, to correlate findings, and to discuss measures to tighten the security of the SingHealth network.

549. At or before the afternoon of 8 July 2018, the SMD was putting together a list of action items. This list of action items was discussed at the meeting at 1:00pm on 9 July 2018 and was updated thereafter based on the discussion at the meeting. The updated list was distributed by Ernest to the relevant IHiS staff at 10:17pm that same day.

550. With reference to the list of action items, the focus of the discussion was on the list of technical and administrative measures to tighten the security of the network. However, in addition to these measures, the list of action items also recorded a number of ‘notable events’ and ‘considerations’, which is indicative

of what IHiS staff had known by that point in time. The updated version as of 10:17pm on 9 July 2018 that Ernest sent to the attendees included the following:

> **NOTABLE EVENTS**
>
> - A potential sucessful [*sic*] DB [database] dump […]
> - [a domain administrator's] account is already compromised even before Citrix breach
> - [the same domain administrator's] account is constantly being used to clear the logs
> - Priv account was removed from local admin group to deny RDP but it was added back again???
>
> **CONSIDERATIONS**
>
> - What if Domain Controllers are compromised?
> - Potentially all LDC citrix servers are compromised.
> - Consider checking with IHIS pentesting team, to check for common AD password dumping tools, and search for traces anywhere on SingHealth PCs/Servers […]

551. Wee has explained that the 'notable events' and 'considerations' listed above were shown during the meeting, but were not the focus of the meeting.

552. At the meeting, Woon Lan asked Wee for his assessment of the incident, and asked whether there was a need to escalate the matter to senior management.[33] Based on what was discussed at the meeting, he had "*some concerns that (he) was not able to confirm*", and that he "*still thought it may not be a security incident*".

553. Examples of such "concerns" include concerns over the fact that there were suspicious SQL queries on the SCM database, and that the A.A. application account was being used. On the deletion of the Windows event logs, Wee was

[33] This is contrasted against Ernest's claims in his conditioned statement that "*At this meeting, we did not discuss whether the matter should be escalated to IHiS senior management*".

also of the view that it was "*not a reportable security incident*", "*as it was not conclusively shown that the deletion had been malicious*". It also appears that the attendees at the meeting were discussing whether this was a scheduled clearing of logs or a manual clearing for malicious purposes, despite knowing that the domain administrator's account was "*constantly being used to clear the logs*", and that it was "*already compromised even before Citrix breach.*"

554. While Wee had these concerns and the relevant IHiS staff were present, he did not voice his concerns or seek clarification. Ultimately, no decision was made at the meeting to escalate the matter to the SingHealth GCIO, Benedict.

28.3 Raising the matter to Clarence Kua and Serena Yong

555. Sometime in the morning of 9 July 2018, Serena happened to meet Clarence, and informed him that there were issues with the SCM system, and asked that he find out more, in his capacity as the Applications Service Lead for SingHealth.[34] At around 3:00pm, Clarence called Henry, his point of contact for issues relating to the SCM system. Henry directed Clarence to the Security Team, knowing that they were looking into the matter.[35] Clarence then called Wee and arranged to meet him at IHiS' offices at ConnectionOne. It appears that separately, after the 1:00pm meeting, Ernest and Wee had discussed the matter and decided that they would consult Clarence.

556. Wee met with Clarence in person at around 5.00 pm, before inviting Ernest to join them. Ernest briefed Clarence on the events of June and July 2018, including the queries run on the SCM database on 4 July 2018, the failed attempts

[34] It appears that by around 6 July 2018, Woon Lan had some knowledge of issues concerning the SCM system – hence her request to Ernest on 6 July 2018 to meet Ernest the next day. Based on Serena's evidence, Woon Lan had told Serena on 8 July 2018 about an issue involving the SCM system. Evidence has not been led on how Woon Lan had come to know of the issue, or the precise extent of her knowledge. Based on the surrounding circumstances, it appears likely that prior to 9 July 2018, Woon Lan and Serena only had a very vague idea of what the issue was about.

[35] Henry has explained that around half an hour before Clarence contacted him, Ernest and Wee had called him to arrange for a meeting on 10 July 2018 to understand how the SCM application works.

to login to the SCM database from the Citrix servers, and some of the accounts that had been used, including the A.A. account. This was the first time Clarence had heard of these matters. The meeting lasted for around 1.5 hours.

557. At around 7:00 to 7.30pm on 9 July 2018, Clarence called Serena and informed her that there were SQL queries being run on the SCM database, that the source of these queries could not be identified, and that there were attempts to access the SCM database. Serena understood this to be a serious issue, and asked for an urgent meeting to be convened at ConnectionOne. Lyonel Cha (Director System Management, IHiS Infrastructure Services Division) and Han Hann Kwang (Head of the SMD), who happened to be with Serena at IHiS' offices at Serangoon North at the time of the call, followed her to ConnectionOne. Arrangements were made for relevant staff from the Infrastructure Services Division to attend the meeting as well.

28.4 Meeting at ConnectionOne and the decision to escalate the matter to Benedict Tan

558. Serena led the urgent meeting at ConnectionOne on the night of 9 July 2018. Attendees included Woon Lan, Han, Ernest, Wee, Lum, Clarence, Henry, Katherine, Teresa and Loh Khim Huat. The attendees began with piecing together information about incidents on 4 July 2018, and then worked backwards and started adding in information about incidents in June 2018.

559. The attendees of the meeting also attempted to interpret each of the SQL queries made to the SCM database. It appears that prior to the meeting, Kelvin informed Henry in a phone call that the query discovered on 4 July 2018 had returned zero results from the SCM database. Henry then informed Clarence of the same over a separate phone call, and subsequently informed all the attendees of the meeting as well. The evidence suggests that at the time of the meeting, no

steps were taken to verify if the queries did in fact return zero results.[36] Han also confirmed that there were no on-going audits or penetration testing that may have led to the queries.

560. By the end of the meeting, it was agreed that the matter should be escalated to Benedict. Serena has explained that to her mind, while there was no operational impact, it "*would be better to over-communicate*" and escalate the matter to Benedict. On Clarence's part, he had in mind a few possible explanations, namely, that the queries were run by internal staff; that there was an on-going audit; or that it was a security incident. "*Worried*" that they were unable to identify the source of the queries, Clarence agreed that they should inform Benedict to seek his advice.

561. At around 9:30pm on 9 July 2018, Serena, Clarence, Ernest and Wee called Benedict, informing him of the unusual activity detected on the SCM database, that no records had been extracted, and that IHiS staff had stopped the SQL queries on 4 July 2018. Serena and Clarence also informed Benedict that there were no on-going audits or penetration testing.

562. Benedict was of the view that they should have informed him earlier. He informed Serena and Clarence that he would inform IHiS CEO Bruce and Director CSG Kim Chuan, and also asked for a meeting to be held the following day to discuss matters further.

28.5 Informing Bruce, Kim Chuan and Prof. Kenneth

563. While Benedict felt that the information he received was "*still quite vague*", he was of the view that he ought to escalate the matter to Bruce. Benedict called Bruce on 9 July 2018 immediately after speaking to Clarence and Serena,

[36] Some support for this can be found in Clarence's conditioned statement, in which he states that "*I do not recall seeing any of the logs at the meeting, as we were focussed on gathering the information from the various teams*".

relaying the information that he had heard to Bruce, including the understanding at that point that the queries returned zero results.

564. Bruce told Benedict to inform Kim Chuan, and to set up a conference call the next day, 10 July 2018, at 1:00pm with Kim Chuan and the team from the Infrastructure Services Division. Bruce explained that the conference call was fixed at 1:00pm because he had other meetings scheduled on the morning of 10 July 2018.

565. After the call with Bruce, Benedict called Kim Chuan and relayed the same information, and informed him of the conference call to be held at 1:00pm on 10 July 2018. It was also understood that a decision would be taken at the 10 July 2018 conference call on whether there was a need to inform CSA and MOH. Bruce also called Kim Chuan, asking the latter to look into the matter urgently, and consider if it was a security event.

566. At the time, Kim Chuan considered whether the incident was a "deliberate adverse event", which could amount to an IT security incident under the SIRF, and which would ultimately have to be reported to the CSA. He also considered whether the incident would be considered a Category 1, 2, or 3 incident. At the time, he only had the information provided to him by Benedict. Pertinently, he did not know then that user accounts and the local administrator accounts for the Citrix servers had been compromised. He did however obtain confirmation that there were no on-going audits or red teaming exercises. Nonetheless, in view of the conference call with Bruce arranged for the next day, he did not report the matter to the CSA on the night of 9 July 2018.

567. On the part of Bruce, he did not consider what the categorisation of the incident should be, because he thought that the incident may not be a security event. He had in mind that there have been previous incidents of unauthorised access, and these incidents did not turn out to be security incidents.

568. At around 10:00pm on 9 July 2018, Benedict also called Prof. Kenneth (Deputy Group CEO (Organisational Transformation and Informatics) of

SingHealth), informing him that "suspicious activities" were detected on the SCM database, but that he was not very sure of the details. Prof. Kenneth asked if the matter was serious, and if they should report the matter to MOH. Benedict replied that it was too early to decide if there was a need to report to MOH. Benedict also stated that he had already informed Bruce, that they were working hard to find out more, and that he would give Prof. Kenneth and Prof. Ivy an update the next day. In view of Benedict's reply, Prof. Kenneth did not inform Prof. Ivy of the incident that night.

28.6 Assessment of IHiS' incident response on 9 July 2018

569. The Committee is troubled by the fact, having regard to paragraph 560 above, that even senior members of IHiS' management, such as Serena and Clarence, did not fully appreciate that there was a security incident and breach of the SCM system, even though the facts they were provided with relating to the events of June and July 2018 would have provided strong indications that they were facing an attack by an APT. Their decision to escalate the matter to Benedict was seen as "over communicating" (on the part of Serena), and tentative (on the part of Clarence, who had a few possibilities in mind and did not have a firm view). This indicates that the lack of training and security awareness observed by this Committee earlier in respect of the IT administrators was also present in the more senior members of IHiS' management.

570. Nonetheless, the Committee does note that in spite of their doubts, Clarence and Serena did escalate the matter to Benedict swiftly. Likewise, Benedict immediately informed IHiS CEO Bruce and Kim Chuan, the Sector Lead point-of-contact for the healthcare sector, despite the fact that, in Benedict's own words, information about the incident at that juncture "*was still vague*". This underscores the point, acknowledged by Benedict, that there is value in escalating potential incidents quickly to senior management, even as they are being investigated, so that the right judgment call can be made on how to respond to the incident. This was echoed by Vivek, when he explained that it is critical that incidents are reported to management so that management can "*give marching orders and realign the troops, realign priorities and get everybody else working*

on [the] problem so that the response can be very firm and aggressive", because "*[u]nless management is involved at a functional level, at an operational level, these calls cannot be taken*".

571. Under the SIRF, Kim Chuan is responsible for reporting security incidents to CSA. In the case of unauthorised access to the SCM database, which is a CII system, CSA would have to be alerted verbally within 2 hours. Under the SIRF, Kim Chuan was also responsible for reporting security incidents to Bruce, MOHH, and MOH.

572. No steps were taken by Kim Chuan or Bruce to escalate the matter further on the night of 9 July 2018 itself. The Committee notes that the information provided to them at this stage was very brief, hence both their evidence that they were each unable to determine at the time whether this was indeed a security incident. At the same time, however, even the limited information provided would have indicated that there was unauthorised access to the SCM database.

573. On the facts, Bruce fixed a conference call for 1:00pm the next day (10 July 2018), and the understanding was that a decision on escalating the matter would only be made then. Eventually, the matter was reported to CSA at 4:40pm on 10 July 2018.

574. The Committee is of the view that, knowing the urgent reporting obligation for a suspected Category 1 incident, Bruce and Kim Chuan should have acted with more urgency, instead of only convening the meeting on the afternoon of the next day.

29 EVENTS OF 10 JULY 2018

29.1 Discovering that the queries did result in data being returned

575. On the morning of 10 July 2018, members of IHiS management, including Serena, Clarence, Woon Lan, Henry, Hann Kwang and Teresa met at ConnectionOne to continue their discussion on the events of 4 July 2018. At 11:00am, they were joined by Kim Chuan.

576. While discussions were on-going, Henry decided to run one of the queries to double-check whether any data would be returned from the database. He was shocked to discover that the query did in fact result in data being returned – this was contrary to Kelvin's earlier representation that no results were returned – and he informed those present at the meeting.

577. Kim Chuan directed the team to ascertain the number of records retrieved from the SQL queries. At the time, the team estimated that there were around 600,000 records retrieved, and they found that the SQL queries could have been run since late June 2018. Separately, Kim Chuan was also informed that one of the logins to a Citrix server could be traced to at least one compromised PC in SingHealth.[37]

578. By this point, Kim Chuan "*thought that it could be an APT attack, and that the incident could be categorised as Category 1*". However, he did not inform CSA of this immediately as the conference call with Bruce and Benedict was scheduled at 1:00pm that same day, and Kim Chuan felt that it was important for Bruce to be briefed and be able to assess the situation and facts.

[37] Kim Chuan does not state which Citrix server this is, but in the circumstances, this must be a reference to Citrix Server 2.

29.2 Conference call with Bruce at 1:00pm

579. At around 1:00pm, a conference call was held and participants included Bruce, Benedict, Kim Chuan, Serena, Clarence and others. Bruce was informed that, based on the IHiS team's re-running the queries, more than 600,000 records could have been retrieved. The IHiS team also provided Bruce with a set of slides containing a summary of the facts of incident, investigation findings, actions taken to-date, and actions that were still in progress.

580. The call took about an hour, after which Bruce directed the IHiS team to return to the IHiS office at Serangoon North immediately to continue discussions in-person. He also wanted to study the logs of the SQL queries to get a more complete picture.

581. At the time when he was taking the call, Benedict was at a work lunch with Prof. Ivy and Prof. Kenneth. He stepped out to take the call, and was still on the call when Prof. Ivy and Prof. Kenneth were leaving at the end of the lunch. Benedict briefly informed them that there was some suspicious activity on the SingHealth database that looked serious, and that he would provide them with a further update later. In response, Prof. Ivy stated that if there was unauthorised access, they should report it to MOH right away. Benedict however asked for more time to find out more details.

29.3 Meeting at Serangoon North at 3:00pm

582. At around 3:00pm on 10 July 2018, a meeting was convened at IHiS' offices at Serangoon North. Attendees included Bruce, Kim Chuan, Benedict, Leong Seng, and members from the IHiS team. Bruce was shown the SQL queries, and he noted that queries had been made on 26 June 2018 for retrieval of the schema of the SCM database, and that queries had been made for retrievals from tables for dispensed medication and patient demographics from 27 June to 4 July 2018.

583. Bruce was also informed of the following in response to some of his queries:

(a) There were attempts to access the SCM database from "back-up" Citrix servers since 11 or 12 June 2018, and this was not a typical route that an end-user would take to access the SCM database.[38]

(b) One of the Citrix servers used to attempt access the SCM database had been taken down for investigation before 27 June 2018 (*i.e.* Citrix Server 1).

(c) The SQL queries were made from a Citrix server (*i.e.* Citrix Server 2), and that the name of the program used.

(d) The SQL queries were coming from authorised accounts and an authorised Citrix server.

(e) The SQL queries could not have been the activity of an internal attacker or a bad program running in the system.

584. In view of the above "*signs of strange activity*" that "*could not be accounted for*", Bruce decided that the matter should be reported to CSA and asked Kim Chuan to do so. Bruce also asked Kim Chuan how the incident should be categorised, to which Kim Chuan replied that it should be seen as a Category 1 incident, as the incident involved unauthorised to a CII system, the SCM database.

585. At the time of the meeting, Bruce did not ask the team why they did not report these events earlier, or why they said the night before that zero records

[38] As mentioned in paragraph 215 (pg 81) above, it was established subsequently in the course of investigations that the SGH Citrix servers could not, in fact, be used for back-up connectivity to the SCM database.

were retrieved. Bruce has explained that "*these were not priority questions (at the time of the meeting)*".

586. Kim Chuan recalls discussing at the meeting why staff had not escalated the matter earlier, but he did not receive any answers. However, in his view, IHiS security staff should have been able to recognise that the incident was a Category 1 incident based on the information that was presented on 10 July 2018 and their experience from the TTX conducted in March 2018.

587. At the meeting, Serena asked the team to start tabulating an event log, recording all staff observations and actions in relation to the events of June and July 2018. Bruce assigned Leong Seng to be in-charge of IHiS investigations into the matter. Benedict also told Bruce that he would inform SingHealth's management.

29.4 Informing SingHealth's management, MOH, the Chairman of the SingHealth Board, and the Chairman of the Risk Oversight Committee

588. At 3:57pm on 10 July 2018, Benedict emailed Prof. Ivy, Prof. Kenneth, Tan Jack Thian ("**Jack Thian**") (SingHealth's Group COO), and Loo Chian Min (SingHealth's Medical Informatics Officer), informing them that IHiS "*detected unauthorised accesses to the SCM production database*" on 4 July 2018, that the team "immediately terminated/blocked all the programs and access channels" on 4 July 2018, and that IHiS was "now doing forensics to determine the source/cause and if any data was compromised". Benedict also provided a summary of events relating to the incident known to IHiS up to that point, and asked for the recipient's advice on whether to inform the MOH Integrated Operations Hub (the "**MOH Ops Centre**") through Jack Thian. Prof. Ivy replied *via* email on 5:22pm, stating that "(t)his is very serious indeed", and asked that the MOH Ops Centre be informed in accordance with protocol. Thereafter, Benedict worked with Jack Thian to prepare the incident report to the MOH Ops Centre.

589. Separately, at 7:20pm on 10 July 2018, Bruce sent an email to (i) the Permanent Secretary of Health, Mr Chan Heng Kee; (ii) the MOH Director of Medical Services, Associate Professor Benjamin Ong; (iii) the Deputy Secretary (Policy) of Health, Ms Ngiam Siew Ying; and (iv) the Managing Director of MOHH, Aik Guan. In this email, Bruce informed the recipients of "*a potential EMR systems breach*", and provided an interim update on IHiS' investigation findings. In addition, Bruce analysed the situation as such in his email:

> Our Citrix servers and SCM EMR database servers are likely to have been attacked and breached by a highly sophisticated & intelligent hacking ops. The attacker demonstrated significant understanding of Citrix, SCM and our physical computing infrastructure. We noticed database retrieval commands (SQLs) to SCM database were made but we are trying to locate evidence that the commands were successfully executed and records accessed. There's likely a system security breach but we can't confirm a data breach. But if the data accesses were successful, it would be very serious as up to 621K dispense medication records could have been accessed.

590. At 9:29pm on 10 July 2018, SingHealth submitted a formal incident report to MOH Ops Centre *via* email. The email was titled "*Incident Report to MOH – 2018/02/01 (Initial Report) on "Unauthorized Access to SCM Production Database"*". The report stated that the incident was assessed to be a Category 1 incident, and contained a summary of the facts known to IHiS at the time.

591. On 11 July 2018, the Chairman of the SingHealth Board and the Chairman of the Risk Oversight Committee were informed of the Cyber Attack.

29.5 Informing CSA and setting-up the War Room at ConnectionOne

592. At around 4:40pm on 10 July 2018, Winston Chua ("**Winston**"), Deputy Director CSG, called CSA's hotline on Kim Chuan's instructions to inform them that a Category 1 incident had occurred. Kim Chuan also sent a text message to Douglas Mun of the CSA. Douglas and Kim Chuan met at IHiS' Serangoon

North office at around 7:30pm that night, and Douglas recommended setting up a War Room to coordinate investigations and recovery efforts. The War Room was set-up at ConnectionOne on the night of 10 July 2018 itself.

30 CONCLUDING OBSERVATIONS FOR THIS PART

593. IHiS' incident response up until 10 July 2018 was commendable in some respects, but was inadequate on the whole in preventing the attacker from stealing and exfiltrating the patient data. Two aspects stand out in particular:

(a) First, IHiS staff did not have adequate levels of cybersecurity awareness, training, and resources to appreciate the security implications of their findings and to respond effectively to the attack.

(b) Second, certain IHiS staff holding key roles in IT security incident response and reporting failed to take appropriate, effective, or timely action, resulting in missed opportunities to prevent the stealing and exfiltrating of data in the attack. Ernest delayed reporting because he felt that additional pressure would be put on him and his team once the situation became known to management The evidence also suggests that the reluctance to report may have come from a belief that it would not reflect well in the eyes of the organisation if the matter turned out to be a false alarm.

594. In a similar vein, the Committee recalls the discussion in Part III regarding the mismangement and inadequacies in remediating the vulnerabilities, weaknesses, and misconfigurations in the SingHealth IT network that had been identified prior to the Cyber Attack.

595. Taken together, it can be seen that there were multiple gaps and deficiencies in IHiS' cybersecurity posture and readiness. IHiS would have benefitted from better training for staff, and more effective processes that would ensure that senior management had better oversight of security incidents and

audit findings. In this regard, the Committee notes that IHiS CEO, Bruce Liang, had limited visibility over some of the matters raised above.

(a) In respect of the events of 11 June to 9 July 2018, Bruce did not have any sight over IHiS' incident response. Bruce's evidence is that due to the scale of IHiS' operations, he relied on the processes and frameworks in place for visibility over security incidents. However, on the facts, the reporting process had broken down, with a bottleneck resulting from Ernest and Wee's failure to escalate the matter. Without a sufficiently robust system for oversight and information flow in place, Bruce did not have visibility over the incident until nearly one month after the first signs of suspicious activities were discovered.

(b) Insofar as the vulnerabilities in the SingHealth IT network identified in the FY16 H-Cloud Pen-Test are concerned, Bruce's evidence is that he relied on various 'dashboards' presented at the ARCs (IHiS Audit Risk Committees) where issues were classified broadly into 'green', 'amber', 'dark amber', and 'red' categories, with general information on whether remedial measures were implemented or in progress. While recognising that the vulnerabilities identified in the FY16 H-Cloud Pen-Test were serious, he did not raise any specific queries as to the completion and adequacy of remedial measures. Instead, he relied on the processes that were in place for the remediation of vulnerabilities, which proved to be inadequate.

596. In order to prevent and respond effectively to future attacks, the cybersecurity posture and readiness of IHiS must be strengthened. In this regard, effective leadership from the CEO and other members of IHiS' senior management is essential, and the starting point must be to improve their visibility over all matters relating to cybersecurity.

Part V – Incident response after 10 July 2018

TABLE OF CONTENTS – PART V

31 INTRODUCTION TO THIS PART

597. In this Part, the Committee presents its further findings in respect of TOR #2, establishing how IHiS and SingHealth responded to the Cyber Attack. This Part follows on from Part IV above, and covers events that took place after CSA was informed of the Cyber Attack on the evening of 10 July 2018. Whereas the incident response till this point was conducted by IHiS alone, the response thereafter was a combined and concerted one, involving a range of parties, including CSA, IHiS, SingHealth, MOH, and MCI.

598. Three main topics will be covered in this Part. First, the joint investigation and remediation efforts by IHiS and CSA's National Cyber Incident Response Team ("**NCIRT**"); second, the public announcement on the Cyber Attack and efforts for patient outreach and communications by SingHealth; and third, the additional measures taken by CSA in its capacity as the national cybersecurity agency.

32 JOINT INVESTIGATION AND REMEDIATION BY IHIS AND CSA

32.1 Setting-up of War Room and sharing of information with CSA

599. On the night of 10 July 2018, CSA deployed the NCIRT to jointly investigate the incident with IHiS, and to implement measures to contain the attack. After the NCIRT was deployed, IHiS provided a detailed oral briefing on the evening of 10 July 2018, informing the NCIRT of their findings to-date.

600. The NCIRT was deployed onsite from 10 to 20 July 2018 at the War Room at Connection One. During this period, daily meetings were held for the NCIRT and IHiS to share findings. In the course of investigations, IHiS also provided the NCIRT with the relevant forensic artefacts and information, including forensic images, memory dumps, and proxy and network logs. CE, CSA recognised that IHiS' response during this period was competent, responsive and cooperative.

601. The NCIRT conducted forensic analysis of these artefacts to verify that data had been successfully exfiltrated, as well as to determine the sequence of attack, and the nature of the attacker. The NCIRT also correlated the investigation findings with information from partners and vendors, as well as research from open source information. This provided them with a better appreciation and understanding of the attacker and its tactics, techniques, and procedures.

602. The War Room was set-up to facilitate coordination between agencies, containment and recovery, investigation, impact analysis, situational updates, and public communications. Leong Seng was placed in-charge of the War Room, which was organised with five "working cells":

(a) Containment: This cell focussed on dealing with containing the Cyber Attack, and was led by Leong Seng.

(b) Investigation: This cell focussed on investigations on how the Cyber Attack happened, and was also led by Leong Seng.

(c) Patient Impact: This cell focussed on reviewing whose records had been accessed, and was led by Benedict.

(d) Communications: Communications were dealt with by the Ministry of Communications and Information ("**MCI**") and MOH and was supported by IHiS Director, Corporate Communications, IT & Admin Group, Loh Chee Peng.

(e) Reviewing security measures for other systems and other Clusters: This cell was led by IHiS Director, Service Delivery, Mark Winn.

603. In Vivek's expert opinion, setting up the War Room was an "appropriate" action and these five cells were appropriately tasked to cover the key areas on which focus was required.

604. From 11 July 2018, Benedict was based at the War Room. His role was to understand the extent of the breach and the data accessed and exfiltrated, in order to keep SingHealth updated for it to prepare its patient communications and outreach plans. At 3:22pm on 12 July 2018, Benedict also lodged a police report for the incident. Kim Chuan went to the War Room daily to maintain situational awareness and keep up with developments, but did not play any role in containment of the attack. From 11 to 23 July 2018, Bruce oversaw the technical response, focussing on the containment measures and addressing IT security weaknesses.

32.2 Ascertaining the queries run and data exfiltrated

605. From 11 July 2018, CSA and IHiS also worked on:

(a) Recreating the SQL queries run on the SCM database between 27 June and 4 July 2018 to find out what data had been extracted;

(b) Monitoring for fresh SQL queries made on the SCM database using the A.A. account; and

(c) Checking whether there were any queries prior to 27 June 2018 that were similar to those run from 27 June to 4 July 2018.

606. On 11 July 2018, Sze Chun found that the Prime Minister's data had been accessed. Henry reported this to the War Room, where Benedict, Clarence and Irene Kwek (an IHiS employee in-charge of providing updates to MOH) were present.

607. IHiS and CSA found that there were altogether more than 200 queries that had been run. Sze Chun also found that there were queries run on 26 June 2018 which involved the attacker looking at the schema of the SCM database, and that no other queries were made before 26 June 2018 or after 4 July 2018. They also tabulated the exact number of records extracted by each query, determined which tables in the SCM database were queried, and ascertained whether the data of

VVIPs had been queried. It was also confirmed that the records in the SCM database were not amended, deleted, or otherwise tampered with, and no other patient records such as diagnosis, test results, or doctors' notes, were accessed.

608. IHiS also simulated the queries that were made by the attacker and compared this against the data traffic patterns going to the C2 servers. Based on the similarities between the two, IHiS confirmed on 13 July 2018 that data had been exfiltrated.

32.3 Containment measures implemented

609. During the joint investigation, IHiS and CSA put in place several containment measures that were aimed at containing the existing threat, eliminating the attacker's footholds, and preventing recurrence of the attack. The measures implemented were meant to contain the immediate threat of the attack, and were not intended to provide a permanent solution for SingHealth and IHiS.

32.3.1 Resetting the Kerberos Ticket Granting Ticket account

610. IHiS' investigations revealed that the attacker had gained administrative privileges and moved across the network to access the Citrix servers. This was an indication that the KRBTGT account[39] could have been compromised.

611. The KRBTGT account is a service account in the active directory, and by obtaining the password hash to this account, the attacker would have been able to compromise every account within the active directory, possibly to the extent

[39] KRBTGT stands for the "Kerberos Ticket Generating Ticket Account". Kerberos is a network authentication protocol that works on the basis of tickets to allow computers and devices communicating over a non-secure network to prove their identity to one another in a secure manner. The KRBTGT account is a special hidden account that encrypts all other authentication tokens in the Kerberos authentication protocol used by Windows. An attacker who has compromised the KRBTGT account can create a Kerberos Golden Ticket to gain complete access to the entire domain.

where it would be able to access any service in the system without the need for a user-ID or password.

612. The KRBTGT account stores two of the most recent passwords in its password history. Resetting the KRBTGT account password twice in succession will invalidate any ticket granting ticket that an attacker might have accessed. Thus, as a precautionary measure, CSA recommended, on 10 July 2018, that IHiS perform a reset of the KRBTGT account twice in succession to invalidate the Kerberos tickets, which could have been stolen or forged by the attacker. IHiS checked with Microsoft, which confirmed that the recommended practice was to perform the second reset 12 hours after the first.

613. IHiS performed the first reset at 10:00pm on 11 July 2018, and the second reset at 10:00am on 12 July 2018.

32.3.2 Changing of passwords across all healthcare clusters

614. On the assumption that the domain administrator account had been compromised, IHiS forced password changes at two levels to further ensure that the attacker would not be able to reuse any user's existing password to breach the network:

(a) The first level was a forced password change for all the domain users at the next point of Windows login. The instruction was sent out on 12 July 2018, and the IHiS active directories team triggered a password reset for all SingHealth users at 1:00pm on 13 July 2018. All users would be prompted to set a new password when they next logged in. Users that were already logged in would be prompted to change their passwords, once their machines were rebooted or entered screen saver mode. Users who failed to reset their passwords by 20 July 2018 would have had their accounts disabled. IHiS also applied the same precautionary measure to the other two healthcare clusters, NHG and NUHS.

(b) The second level was scheduling changes to the passwords of all the privileged and database application IDs, and host IDs, which was completed by 19 July 2018.

32.3.3 Cleaning-up of network-based IOCs, instituting of firewall rules, and reloading of Citrix servers

615. IOCs (indicators of compromise) discovered by CSA in the course of their forensics and malware analyses were incorporated into IHiS' corporate antivirus system from 17 July 2018. From 13 July 2018, the IHiS network team also created firewall rules to block off malicious callbacks to the C2 servers identified by the CSA analyst team. CSA also shared the identified IOCs with the other CII Sector Leads for dissemination to their CII owners, so that they could scan for similar infections.

616. Upon discovery that the SGH Citrix server had been used by the attacker to access the SCM database, the IHiS network team added firewall rules to block access from the SGH Citrix servers to the SCM database on 11 July 2018.

617. However, as it was not possible to ascertain through detailed forensic examination whether each Citrix server was compromised (nearly a thousand such servers were running in the HDC), IHiS set out to reload each of the Citrix servers in the HDC Citrix server farm with a clean image on 14 and 15 July 2018. This ensured that no compromised Citrix server was left running after the clean images were reloaded. All Citrix servers were fully refreshed by 16 July 2018.

32.3.4 Disabling of PowerShell on endpoints

618. After learning from CSA that the attacker had made use of PowerShell malware in the attack, IHiS disabled PowerShell on 13 July 2018 on all end-user machines.

32.4 Heightened monitoring of IT network and implementation of Internet Surfing Separation on 20 July 2018

619. From 11 July 2018, IHiS was placed on heightened alert for any sign of the attacker in the network. IHiS actively monitored the network for security events from the active directory, internet proxy, and firewall, to detect signs of compromise or failed login attempts. They also actively reviewed network flow logs to determine if there were further signs of mass data exfiltration.

620. As a result of the active monitoring, IHiS detected on 19 July 2018 the attempts being made from the S.P. server to connect to a known C2 server that same day, enabling IHiS and CSA to respond quickly to investigate.

621. As explained in paragraph 207 (pg 70) above, the attempted callbacks indicated that the attacker still had access to SingHealth's network even while IHiS was actively implementing measures to contain the incident, and that the attacker was still active and trying to regain a foothold in the network. In these circumstances, CSA strongly advised IHiS to implement ISS, on the basis that ISS would be effective against this particular attack because it fully blocked the callbacks and disrupted the attacker's command and control in the network.

622. IHiS acted decisively, and on 12:00am of 20 July 2018, cut off user internet surfing and internal server access to the internet for the SingHealth Cluster. On 22 Jul 2018, IHiS also cut off user internet surfing and internal server access to the internet for the NHG and NUHS Clusters.

623. No further suspicious activity was detected after ISS was implemented.

33 THE PUBLIC ANNOUNCEMENT AND PATIENT OUTREACH AND COMMUNICATIONS

33.1 The public announcement

624. After being notified of the Cyber Attack, SingHealth's senior management, in consultation with MOH, IHiS, CSA, and MCI began making plans for a public announcement, and for patient outreach and communications.

625. SingHealth's senior management recognised that SingHealth had an obligation to inform, in the shortest time reasonably possible, all patients who may have been affected by the Cyber Attack. At the same time, it was recognised that any announcement should not compromise ongoing forensic investigations, and that information should not leak out in an uncontrolled way that may cause public panic.

626. SingHealth's senior management was also of the view that before a public announcement could be made, they had to first ensure that patient data was intact and secure, and to obtain more information about the attack, including the information that was accessed and whether there was any exfiltration. Such information was not available as at 10 July 2018, but would be necessary before concrete plans for the announcement and patient outreach and communications could be made, as they had to be able to address patients' concerns and anxieties.

627. On 12 July 2018, Prof. Kenneth attended a meeting called by MOH. There was agreement at the meeting that more information was required before making a public announcement. In particular, the number of patients affected was still in flux. Following this meeting, Prof. Kenneth started mobilising resources for SingHealth's patient outreach and communications plan by briefing SingHealth's Communications Team on the outline of the Cyber Attack.

628. On 13 July 2018 at 5:00pm, Prof. Kenneth attended a meeting with MOH. At this meeting, IHiS confirmed that data had been exfiltrated and the type of data affected: (a) 4,600 line items of dispensed medication records had been

exfiltrated; (b) around 160,000 patients had their dispensed medical records accessed; and (c) around 1.5 million patients had their demographic data accessed.

629. At the meeting on 14 July 2018 at 4:00 pm, CSA informed SingHealth that no patient data had been overwritten or changed. With this key piece of information, SingHealth further shaped its plan for patient outreach and communications. At this same meeting, it was agreed that there was a need to make a public announcement, but the timing of the announcement was in issue.

630. Eventually, 20 July 2018 was scheduled as the date for the public announcement, to ensure that there was sufficient time for proper containment of the Cyber Attack. As it turned out, the attempted callbacks from the S.P. server to the C2 server on 19 July 2018 indicated that the attacker could still be in the system, and led to the implementation of internet surfing separation at 12:00am on 20 July 2018. The public announcement of the Cyber Attack was made later that same day at around 5:30pm.

33.2 Patient outreach and communications

631. Prof. Kenneth took direct charge of patient outreach and communication efforts for SingHealth, in close consultation with Prof. Ivy. Upon receiving more complete information about the attack, SingHealth began making detailed plans on how to contact patients, which patients to contact, and how to prioritise the contacts.

33.2.1 Identifying the patients who should be contacted

632. On 18 July 2018, IHiS informed SingHealth that the attack only affected patients who visited SingHealth institutions from 1 May 2015 to 4 July 2018. SingHealth decided to expand the date range and reach out to patients who had visited SingHealth insitutions from 1 January 2015 to 4 July 2018, on the basis that patients might recall the year they had visited the institutions but might not be sure of which month they had visited. Further, drugs may sometimes be logged

into SingHealth's system as having been dispensed on a later date than the date of the patient's visit, for various reasons such as amendments to medication or charging.

633. SingHealth thus took the view that it was important to also include those patients who visited between January and May 2015 so as to reassure them that their data had not been accessed. In total, SingHealth intended to contact 2.16 million patients.

33.2.2 Modes and content of communications

634. SingHealth decided to use the following modes of communication in their patient outreach and communications:

(a) Sending SMS messages to all patients, reassuring those whose data were not affected, and informing those whose data were affected and what data had been accessed;

(b) Sending letters to patients for whom SingHealth were unable to contact *via* SMS messages;

(c) Setting-up telephone hotlines in addition to SingHealth institutions' general call centres, and informing the patients whose medication data had been accessed of the hotline numbers;

(d) Creating a dedicated email account for public queries; and

(e) Allowing the public to perform checks themselves on whether their data was accessed, by using the channels provided on the Health Buddy mobile application and SingHealth's websites.

635. SingHealth had, further to a suggestion by Prof. Ivy on 16 July 2018, decided to use SMS messages as the primary mode of communication with patients in view of the need for quick dissemination of information on a large scale. SingHealth engaged a third-party vendor, which was able to send the SMS

messages more quickly, better regulate the time the messages were sent, and to track in real-time the number or messages that were delivered or undeliverable.

636. The multiple channels of communications were designed to allow SingHealth to reach out to patients (*via* SMS and letters), and also for patients to reach out to SingHealth for further information (*via* the Health Buddy mobile application, SingHealth websites, telephone hotlines, and emails).

637. SingHealth also sought to anticipate the concerns and needs of the affected patients. Across all modes of SingHealth's communications with their patients, SingHealth apologised unreservedly for any anxiety or inconvenience caused by the attack. Patients received personalised communications where they were addressed by name, and were informed of the extent of their data that was accessed. SingHealth also sought to reassure patients of the following through their communications:

(a) Their care delivery was not affected.

(b) Their medical records were intact and had not been tampered with.

(c) Information on their diagnosis, medical conditions and investigations/test results had not been accessed.

(d) Their telephone numbers and financial details (*e.g.* credit card number) had not been accessed.

638. For patients who approached SingHealth's staff directly, SingHealth also prepared information leaflets in the four official languages, and made these available at all SingHealth outpatient clinics and polyclinics. These served to answer some FAQs for patients who enquired about the Cyber Attack or the status of their data.

639. Care was also taken to safeguarding the confidentiality of information in handling the communications. The text and email messages did not contain the patient's demographic or medical data. Likewise, the call centre portal, developed by IHiS, was set up in a manner which safeguarded patient information by limiting the information which operator using the portal could view to whether the patient was (a) not affected, (b) had demographic data access, or (c) had both demographic and medication data accessed; no demographic or medical data was visible by the operator.

640. MOH reviewed and agreed with the contents of all the above forms of communications.

33.2.3 Operationalising the outreach and communications efforts, and the role of SingHealth and IHiS staff

641. Up until shortly before the public announcement on 20 July 2018, SingHealth's senior management were unable, in light of the news embargo, to engage their staff to seek their assistance in operationalising the outreach and communications plan. In spite of this limitation, SingHealth had to ensure that there was sufficient manpower to put their communications and outreach plans into effect once the public announcement was made.

642. More than 1,000 staff from across SingHealth and IHiS were mobilised at short notice to assist in patient outreach efforts. Staff were engaged through the usual leadership platforms, townhalls, and memos from senior management. These served to inform SingHealth staff of the attack, to provide staff with the information necessary to inform and reassure patients, and as a means to request for staff assistance in patient outreach and communications.

643. Staff volunteers juggled their usual duties, worked around the clock, stayed back after their usual shift or over the weekend, thus enabling SingHealth to reach their patients in the shortest time possible. Staff volunteers included doctors, nurses, allied health professionals, as well as administrators from HR,

Ops and Finance. Frontline staff in the wards and clinics also helped address patients' queries and concerns in the usual course of their work.

644. SingHealth also worked closely with IHiS staff who supported their patient outreach plans. For instance, IHiS staff developed applications for the HealthBuddy mobile application and the call centre portal. SingHealth also worked with IHiS to increase the number of servers on which the Health Buddy mobile application was run, to prepare for a surge in the volume of users. In this regard, SingHealth also worked with the Government Technology Agency of Singapore to ensure that SingPass could cope with a surge in the volume of cases.

645. To minimise the likelihood of any patient anxiety arising from their being unable to access the communications, SingHealth ensured that there were sufficient resources and infrastructure available to manage the anticipated volume of queries efficiently and effectively. This included increasing the number of calls that the call centre could support concurrently from 90 to 270, in anticipation of an increased number of calls and to ensure that patients would be able to reach SingHealth without their calls being dropped. All channels of communication and outreach were ready and in place as soon as the news embargo was lifted.[40]

[40] SingHealth had only one-and-a-half hours before the public announcement to train the first batch of call centre staff on 20 July 2018, as there was an information embargo. They then trained the rest of the call centre staff over the next few days at 7:30am each day. The hotlines were operational from 6:00pm on 20 July 2018, within 15 minutes of the public announcement. In total, 104 volunteer staff were trained and organised into two teams working 12-hour shifts to man the hotlines.

646. As at 11:00am on 25 July 2018, the extent of SingHealth's public outreach was as follows:

Figure 13: Results of SingHealth's outreach as at 11:00am on 25 July 2018

Communication channel	**Numbers # (cumulative from 20 – 25 July)**
SMS messages to patients with valid mobile numbers in SingHealth records	2.03 million
Letters to patient with no mobile numbers in SingHealth records	86,700
Self-check on HealthBuddy and SingHealth website	215,600
Telephone calls to hotlines and SingHealth's general call centers	13,400
Email enquiries received at dedicated account	3,100
Leaflets	36

The numbers in each category may overlap as multiple channels may have been used to reach some patients

647. From 26 July 2018, SingHealth began sending letters to patients for whom SMS messages to their mobile numbers on record had failed. In total, around 434,000 letters were sent as at 30 July 2018. Over 400 volunteer staff supported the operation to sort, check and print letters daily from 21 to 28 July 2018. All printing was done in-house by SingHealth.

33.2.4 Guarding against deliberate falsehoods and phishing risks

648. SingHealth also closely monitored for fake news, fake websites, and scams. Within a few hours of the first batch of SMS messages being sent on 20 July 2018, SingHealth received information that there were fake SMS messages being sent. SingHealth made a police report about the fake messages, and alerted the public through social media, mass media, and SingHealth's websites.

SingHealth initially used a bit.ly[41] link in its SMS messages (bit.ly/cyber-attack18), to improve readability for recipients. However, as a bit.ly link can be generated by anyone and carries a phishing risk, some recipients of the SMS-es had concerns as to the authenticity of the SMS messages. These concerns were realised when the fake SMS messages emerged. Fortunately, as SingHealth was closely monitoring the emergence of fake communications, it was able to quickly alert the public to the fake SMS messages. SingHealth also changed the bit.ly link to www.singhealth.com.sg/cyberattack (in full) in subsequent SMS messages.

33.2.5 Patient satisfaction

649. To track the sentiments of patients who called the hotlines, SingHealth introduced a callers' emotion survey at the call centres on 22 July 2018. Call centre staff were provided with a chart which showed a happy face, neutral face and unhappy face, and were asked to indicate on the chart after every call how they gauged the caller's sentiment. The chart was intended to be simple and easy to use, and was based on the staff's assessment. For calls, 82% were assessed to be satisfied, 16% were neutral, and 2% were unhappy.

650. A similar chart was introduced for staff to assess the sentiments of each person who sent in emails to the dedicated email account that was set-up, check@singhealth.com.sg. 85% of the persons who emailed were assessed to be neutral, 8% were unhappy, and 7% were satisfied.

33.3 Assessment of SingHealth's incident response

651. The efforts of SingHealth, with the assistance of its partners, in patient outreach and communications are commendable. A large number of patients were able to receive and obtain the necessary information in a timely and

[41] Bit.ly is a URL shortening service. URL shortening is a technique in which a URL may be made substantially shorter and still direct to the required page. This is achieved by using a redirect which links to the web page that has a long URL.

effective manner, through multiple modes of communication. The Committee notes that the scale of the outreach was unprecedented, and was planned and operationalised over a span of just 11 days, from when SingHealth's management was apprised of the situation on 10 July 2018. The dedication shown by the staff volunteers from SingHealth in assisting their patients is especially heartening.

652. The fact that the public announcement was made on 20 July 2018 is also well regarded. CE, CSA has noted that the general consensus among professionals, both in Singapore and around the world, is that the Singapore Government publicly announced the Cyber Attack in a "*remarkably short time*", and that this is contrasted against the "*long runways*" between discovery and public disclosure in many other cases of data breaches.

653. The use of multiple channels of communications, with SMS messages being the primary means of informing patients, proved to be effective. The Committee notes in this regard the submission by counsel for SingHealth, that SingHealth's approach may be contrasted with the experiences of the UK's National Health Service ("**NHS**") during the 2017 WannaCry Ransomware cyber attack. In that case, the NHS was found to have been over-reliant on email communications, and the need for alternative communication channels and multiple communication routes to support incident response was identified as a learning point.[42]

654. Nonetheless, there remains room for improvement in respect of the collecting and updating of patient contact details. The Committee has heard that 15% of the SMS messages failed to be delivered. At the time of the Inquiry, SingHealth was still unable to contact 2.9% of the affected patients. despite having utilised all these modes of communication. In this regard, the issues faced

[42] William Smart (UK Chief Information Officer for Health and Social Care), "Lessons learned review of the WannaCry (Ransomware Cyber Attack" (February 2018), at p33, [5.13], <https://www.england.nhs.uk/wp-content/uploads/2018/02/lessons-learned-review-wannacryransomware-cyber-attack-cio-review.pdf>.

by SingHealth in respect of the sending of SMS messages are illustrative of the issues posed by incomplete or inaccurate contact records:

(a) First, in some cases, SingHealth simply had no record of some patients' mobile numbers.

(b) Second, several mobile phone numbers in SingHealth's database had multiple names associated to each number. Some of these numbers belonged to human resources staff of a company or agency, and may have been used to register several company employees.

(c) Third, some mobile phone numbers in SingHealth's database did not belong to the names against which they were recorded, resulting in SMS messages being sent to incorrect recipients. This might have been because some pre-paid phone numbers were recycled by telecommunications providers; or patients might have inadvertently provided incorrect mobile numbers at the point of registration; or mobile numbers might have been inadvertently wrongly entered by staff at the point of registration.

655. The Committee has been informed of a number of improvements that SingHealth has implemented or are undertaking to improve the collecting and updating of patient contact details and to ensure their integrity and accuracy. A few of the more significant measures are reproduced here:

(a) Patients who were uncontactable following the Cyber Attack will be identified in SingHealth's system, such that the next time they visit a SingHealth institution for appointments, it would be flagged to staff that the patient's contact particulars need to be updated, and they will be directed to do so;

(b) Since 5 November 2018, SMS reminders are sent to all patients on the day of their outpatient appointments, reminding them to

approach counter staff to update contact details if they have made any changes to their personal details;

(c) A Patient Identity Integrity Workgroup, comprising representatives from across the SingHealth cluster including the Chief Operating Officers of SingHealth institutions, was set up in October 2018 to look into measures to improve the obtaining and updating of contact particulars.

656. The Committee agrees that these are positive steps in the right direction in improving SingHealth's capabilities to respond to similar incidents.

34 ADDITIONAL MEASURES TAKEN BY CSA

34.1 Putting CII sectors on alert

657. Over the course of the investigation, CSA issued alerts and disseminated newly discovered IOCs to the other ten CII sectors, to scan and monitor their networks and systems based on the IOCs for signs of the attacker. CSA also provided information about the attacker's MO to CII sectors to enable them to review their own security posture, and to implement appropriate security measures. These recommended security measures included the review of domain administrator accounts, monitoring for unauthorised remote access, and disabling the unnecessary use of PowerShell. Following the first alert issued on 16 July 2018, CSA sent a total of six addendums, between 17 and 25 July 2018, to the CII sectors.

34.2 Briefing entities hosting large amounts of personally identifiable information

658. On 19 July 2018, CSA organised a briefing for relevant stakeholders in all CII sectors, and recommended that these stakeholders review their PII protection measures.

34.3 Raising of National Cyber Threat Alert Level

659. The National Cyber Threat Alert Level ("**NCTAL**") provides the national level of alert in the cyber domain in Singapore, which is derived from the associated threats and the corresponding required responses. In anticipation of potential opportunistic attacks on sensitive systems by individuals or groups stemming from the media release about the Cyber Attack, CSA sought approval from the Chairman of CMG(Cyber)[43] on 19 July 2018 to raise the NCTAL on the day of the press conference, and CII sectors were instructed to adopt heightened defence measures as a precautionary measure.

34.4 Publishing advisories on protection and precautionary measures

660. SingCERT published two advisories on 20 July 2018. The first was a technical advisory on measures for the protection of customers' personal data. This was tailored to companies and incorporated specific recommendations for companies to adopt, in order to protect their systems and networks from the MO of the attacker and the vulnerabilities that had been exploited. The second advisory on precautionary measures to take, in view of the SingHealth incident, was tailored to members of the public to encourage them to take precautionary measures to protect themselves from the misuse of the personal data that had been exfiltrated from SingHealth. In this advisory, SingCERT recommended that members of the public enable two-factor authentication (especially for users of e-government services and i-banking transactions), change their passwords (if their passwords had been derived from PII), and check for possible fraudulent transactions.

661. After SingHealth started sending out SMS messages to notify affected citizens, SingCERT received feedback that there were phishing SMS messages

[43] Crisis Management Groups ("**CMGs"**) support the Homefront Crisis Executive Group's ("**HCEG"**) management of crises across different sectors of the nation. CMG(Cyber) is responsible for specifically managing cybersecurity incidents and implementing incident mitigation efforts during significant cyber incidents in Singapore.

going around that spoofed the SingHealth SMS ID, and contained links that directed citizens to fake websites designed to collect their personal data. The second advisory was updated on 23 July 2018 to inform the public that fake SingHealth text messages leading to phishing sites were being circulated. The updated advisory contained an infographic on precautionary measures for easy reference, and also reminded members of the public to visit the SingHealth website by keying the web address directly into their browser's address bar.

662. MCI also developed an infographic for the precautionary measures advisory so that members of the public could easily absorb the information. This was published with the updated advisory on 23 July 2018.

34.5 Requesting IMDA to issue blocking order on IOCs

663. CSA sent a request to IMDA on 21 July 2018 for the ISPs to block the domain and IP addresses of the IOCs that had been discovered. The ISPs confirmed that the domain and IP addresses were blocked, as directed by IMDA, on 28 July 2018. This effectively blocked any communications made by the attacker, through local ISP networks, with the C2 servers.

Part VI – Key Findings of the Committee on TORs #1 and #2

664. The Committee's findings in respect of TORs #1 and #2 have been set out in Parts III, IV, and V of this Report. From these findings, the Committee has identified five Key Findings.

Key Finding #1: IHiS staff did not have adequate levels of cybersecurity awareness, training, and resources to appreciate the security implications of their findings and to respond effectively to the attack

- A number of IHiS' IT administrators are commended by the Committee for their vigilance in noticing suspicious activity, such as unauthorised logins to the Citrix servers, suspicious attempts at logging in to the SCM database, presence of unauthorised software, and suspicious queries being run on the SCM database.
- However, these same IT administrators could not fully appreciate the security implications of their findings, and were unable to co-relate these findings with the tactics, techniques, and procedures ("**TTPs**") of an advanced cyber attacker.
- They were also not familiar with the relevant IT security policy documents and the need to escalate the matter to CSA. There was also no incident reporting framework in place for the IT administrators.
- Members of the Security Management Department, Computer Emergency Response Team, and senior members of IHiS' management were similarly unable to fully appreciate the security implications of the findings.

Key Finding #2: Certain IHiS staff holding key roles in IT security incident response and reporting failed to take appropriate, effective, or timely action, resulting in missed opportunities to prevent the stealing and exfiltrating of data in the attack

- The Security Incident Response Manager ("**SIRM**") and Cluster Information Security Officer ("**Cluster ISO**") for SingHealth, who were responsible for incident response and reporting, held mistaken understandings of what constituted a 'security incident', and when a security incident should be reported.
- The SIRM delayed reporting because he felt that additional pressure would be put on him and his team once the situation became known to management.
- The evidence also suggests that the reluctance to escalate the matter may have come from a belief that it would not reflect well in the eyes of the organisation if the matter turned out to be a false alarm.
- The Cluster ISO did not understand the significance of the information provided to him, and did not take any steps to better understand the information. Instead, he effectively abdicated to the SIRM the responsibility of deciding whether to escalate the incident.

Key Finding #3: There were a number of vulnerabilities, weaknesses, and misconfigurations in the SingHealth network and SCM system that contributed to the attacker's success in obtaining and exfiltrating the data, many of which could have been remedied before the attack

- A significant vulnerability was the network connectivity (referred to in these proceedings as an "open network connection") between the SGH Citrix servers and the SCM database, which the attacker exploited to make queries to the database. The network connectivity was maintained for the use of administrative tools and custom applications, but there was no necessity to do so.
- The SGH Citrix servers were not adequately secured against unauthorised access. Notably, the process requiring 2-factor authentication ("**2FA**") for administrator access was not enforced as the exclusive means of logging in as an administrator. This allowed the attacker to access the server through other routes that did not require 2FA.
- There was a coding vulnerability in the SCM application which was likely exploited by the attacker to obtain credentials for accessing the SCM database.
- There were a number of other vulnerabilities in the network which were identified in a penetration test in early 2017, and which may have been exploited by the attacker. These included weak administrator account passwords and the need to improve network segregation for administrative access to critical servers such as the domain controller and the Citrix servers. Unfortunately, the remediation process undertaken by IHiS was mismanaged and inadequate, and a number of vulnerabilities remained at the time of the Cyber Attack.

Key Finding #4: The attacker was a skilled and sophisticated actor bearing the characteristics of an Advanced Persistent Threat group

- The attacker had a clear goal in mind, namely the personal and outpatient medication data of the Prime Minister in the main, and also that of other patients.
- The attacker employed advanced TTPs, as seen from the suite of advanced, customised, and stealthy malware used, generally stealthy movements, and its ability to find and exploit various vulnerabilities in SingHealth's IT network and the SCM application.
- The attacker was persistent, having established multiple footholds and backdoors, carried out its attack over a period of over 10 months, and made multiple attempts at accessing the SCM database using various methods.
- The attacker was a well-resourced group, having an extensive command and control network, the capability to develop numerous customised tools, and a wide range of technical expertise.

Key Finding #5: While our cyber defences will never be impregnable, and it may be difficult to prevent an Advanced Persistent Threat from breaching the perimeter of the network, the success of the attacker in obtaining and exfiltrating the data was not inevitable

- A number of vulnerabilities, weaknesses, and misconfigurations could have been remedied before the attack. Doing so would have made it more difficult for the attacker to achieve its objectives.
- The attacker was stealthy but not silent, and signs of the attack were observed by IHiS' staff. Had IHiS' staff been able to recognise that an attack was ongoing and take appropriate action, the attacker could have been stopped before it achieved its objectives.

Part VII – Recommendations by the Committee on TORs #3, #4, and #5

TABLE OF CONTENTS – PART VII

.

35 PREAMBLE

665. In this Part, the Committee makes its recommendations on TORs #3, #4, and #5. The recommendations are made in light of the Committee's findings, the testimony of expert witnesses and CSA, and submissions from the public. The Committee has also taken into consideration the comprehensive, careful, and thoughtful recommendations by the Solicitor-General, and the collective recommendations from MOH, MOHH, SingHealth, and IHiS.

35.1 Terminology

666. The importance of the recommendations and the seriousness with which we take their implementation are denoted by the use of the following terms:

(a) The term "**MUST**" indicates requirements to be followed strictly and from which no deviation ought to be permitted. The use of "MUST" reflects our view that the degree of necessity for implementation of these recommendations is particularly high.

(b) The term "**SHOULD**" indicates that, among several possibilities, one is recommended as particularly suitable, without mentioning or excluding others, or that a certain course of action is preferred but not necessarily required, or that (in the negative form) a certain possibility or course of action is discouraged but not prohibited. There may exist valid reasons in particular circumstances to choose a different course, but the full implications must be understood and carefully weighed before choosing a different course.

35.2 Recommendations for TORs # 3, #4, and #5

667. The Committee makes 16 recommendations, comprising seven Priority Recommendations and nine Additional Recommendations.

668. The recommendations have been categorised by borrowing broadly from the U.S. National Institute of Science and Technology ("**NIST**")[44] Framework for Improving Critical Infrastructure Cybersecurity (the "**NIST framework**"), with necessary supplementation to address broader issues that are relevant to the Cyber Attack. The broad areas that the recommendations will address are:

(a) PREVENTION – Prevention requires maintaining fundamental security capabilities, and implementing appropriate safeguards to stop or limit the impact of a cyber attack;

(b) VIGILANCE – Being vigilant means putting in place procedures and solutions to identify vulnerabilities and misconfigurations, and to anticipate emerging threats;

(c) DETECTION – Detection means putting in place measures to discover potential cyber attacks and alert responders to their existence;

(d) RESPONSE – Response entails being prepared to react to cyber attacks, contain the impact, repair damage to operations, and return to normal operations;

(e) GOVERNANCE – Governance involves creating a culture of security to mitigate risk and better protect the company's critical

[44] NIST is a non-regulatory federal agency within the U.S. Department of Commerce. Its mission is to promote innovation and industrial competitiveness by advancing standards, and technology. NIST guidelines often become the foundation for best practice recommendations across the security industry and are incorporated into other standards.

infrastructure, by adopting and maintaining a proper posture in the area of cybersecurity; and

(f) PEOPLE DEVELOPMENT – People development involves training and equipping staff at all levels with cybersecurity knowledge and skills, to increase the cyber resilience of the organisation.

669. The broad areas that each of the Committee's recommendations relate to and addresses will be identified at the beginning of each relevant section with hashtags ('#').

670. Out of the 16 recommendations, the Committee proposes seven Priority Recommendations. They relate to certain strategic and operational measures to uplift the cybersecurity posture of SingHealth and IHiS, and steps must be taken to implement them immediately.

671. Given that the next attack may not follow the same attack pattern as the Cyber Attack and may also target different assets, the first six Priority Recommendations address areas for improving cybersecurity policies and capabilities as well as embedding cybersecurity awareness into daily operations. The senior management of SingHealth and IHiS must provide effective and agile leadership for the timely and effective implementation of these recommendations, allocating adequate resources, and keeping a close and careful watch. There must also be appropriate oversight over and verification of their implementation, including by external entities where appropriate. As CE, CSA has pointed out, "*from a technical, professional perspective…certain things need to be done. They have to be done and perhaps they should have been done yesterday*".

672. The seventh Priority Recommendation addresses the issue of collective security, and builds on the first six Priority Recommendations to bring our cyber defences to a new and higher level. This is imperative given the high degree of digitalisation and interconnectivity in Singapore, and the risks at the national level.

673. The nine Additional Recommendations relate to the specific issues raised in the course of this Inquiry, including technical, organisational, training, and process-related issues. The measures, which are similarly aimed at uplifting the cybersecurity posture of SingHealth and IHiS, must be implemented or seriously considered.

674. Collectively, the 16 recommendations serve to (i) build a culture of security; (ii) secure particular aspects of the system; (iii) improve incident response capabilities; (iv) improve post-incident recovery capabilities; and (v) promote collective security.

675. All 16 recommendations are made in respect of TORs #3 and #4, and apply equally to TOR #5[45]. In this regard, the experts confirmed to the Committee that their recommendations were *not* limited to IHiS or SingHealth and were applicable generally to all organisations responsible for large databases of personal data. Some of the recommendations also relate to enhanced measures for CII systems (*i.e.* recommendations #2, #4, #7, and #8).

676. How the recommendations should be adopted in practice by organisations responsible for large databases of personal data will depend on the existing policies, processes and personnel in each of these organisations.

677. Cybersecurity threats are constantly evolving, and will continue to increase in sophistication, intensity, and scale. Similarly, while implementing the recommendations is a necessary and vital first step, organisations must constantly renew, review, and refresh their security structures, technology, and readiness.

---

[45] TOR #5 reads: "*In light of the cybersecurity attack and the findings above, recommend measures to reduce the risk of such cybersecurity attacks on public sector IT systems which contain large databases of personal data, including in the other public healthcare clusters.*"

35.3 Key considerations for recommendations

678. In drawing up the recommendations, we agree with the Solicitor-General that these should be guided by the following key considerations:

(a) **First, in the current landscape, it must be acknowledged that attackers are increasingly sophisticated and will find a way to breach your network**. While this means that one should adopt an 'assume breach' mindset, it does not mean sitting back and waiting to be attacked. Instead, organisations and in particular those responsible for large databases of personal data, must adopt a "defence-in-depth" strategy. This involves: (i) arming themselves with sophisticated security systems and solutions which can facilitate early and accurate detection, *e.g.* by adopting emerging technologies such as database activity monitoring ("**DAM**"), endpoint detection and response ("**EDR**"), managed EDR ("**MDR**"), NetFlow analysis and advanced behaviour-based analytics; and (ii) complementing such security systems and solutions with the right people and processes, *e.g.* having dedicated and trained IT security personnel reporting to the right level within the organisation, engaging external expertise as required and having staff that have the right levels of cybersecurity awareness.

(b) **Second, at a practical level, the push towards a defence-in-depth strategy will no doubt be met with challenges given the current cybersecurity maturity levels in many organisations and the trade-offs that will need to be made *vis-à-vis* operational requirements and costs**. Hence, we acknowledge that the transition to a defence-in-depth strategy cannot happen overnight. However, even during the transition phase, there must be prioritised efforts to adopt certain strategic and operational measures to uplift security immediately – these measures are discussed below in the specific context of IHiS and SingHealth. In addition, it is an important priority that even during the transition

phase, cybersecurity is managed at the right level of leadership *i.e.* cybersecurity issues are deliberated at the right level within the organisation, by senior management who have oversight of the operational and business imperatives.

(c) **Third, having regard to the cyber threat actors of today – many of whom are state-linked, it must be recognised that the battle cannot be fought and won solely at the organisation-level**. It is therefore important that we operate a 'networked defence' with an emphasis on collective security. In this regard, the Singapore Government must play a role and the establishment of CSA is an example of the Singapore Government's commitment to this. For example, CSA's framework for critical information infrastructure ("**CII**") incident reporting provides the agency with awareness of cyber threats affecting Singapore, so that it can take the necessary actions to mitigate and respond to the threats, as well as provide early warning and alerts to other non-affected sectors. CSA must not only continue its good work with the CII but it must, as the national authority on cybersecurity, actively work to build a resilient cyberspace in Singapore.

35.4 Priority Recommendations

679. These recommendations have been informed by the unique circumstances of the Cyber Attack. Some recommendations may seem axiomatic, but unfortunately they had not been practised or implemented effectively by IHiS at the time of the attack. Each recommendation contains specific considerations and implementation details to guide IHiS, SingHealth, and other organisations on how to protect their crown jewels against similar attacks and respond to such attacks.

680. The seven Priority Recommendations are as follows:

Recommendation #1: An enhanced security structure and readiness must be adopted by IHiS and Public Health Institutions

- Cybersecurity must be viewed as a risk management issue, and not merely a technical issue. Decisions should be deliberated at the appropriate management level, to balance the trade-offs between security, operational requirements, and cost.
- IHiS must adopt a "defence-in-depth" approach.
- Gaps between policy and practice must be addressed.

Recommendation #2: The cyber stack must be reviewed to assess if it is adequate to defend and respond to advanced threats

- Identify gaps in the cyber stack by mapping layers of the IT stack against existing security technologies.
- Gaps in response technologies must be filled by acquiring endpoint and network forensics capabilities.
- The effectiveness of current endpoint security measures must be reviewed to fill the gaps exploited by the attacker.
- Network security must be enhanced to disrupt the 'Command and Control' and 'Actions on Objective' phases of the Cyber Kill Chain.
- Application security for email must be heightened.

Recommendation #3: Staff awareness on cybersecurity must be improved, to enhance capacity to prevent, detect, and respond to security incidents

- The level of cyber hygiene among users must continue to be improved.
- A Security Awareness Programme should be implemented to reduce organisational risk.
- IT staff must be equipped with sufficient knowledge to recognise the signs of a security incident in a real-world context.

Recommendation #4: Enhanced security checks must be performed, especially on CII systems

- Vulnerability assessments must be conducted regularly.
- Safety reviews, evaluation, and certification of vendor products must be carried out where feasible.
- Penetration testing must be conducted regularly.
- Red teaming should be carried out periodically.
- Threat hunting must be considered.

Recommendation #5: Privileged administrator accounts must be subject to tighter control and greater monitoring

- An inventory of administrative accounts should be created to facilitate rationalisation of such accounts.
- All administrators must use two-factor authentication when performing administrative tasks.
- Use of passphrases instead of passwords should be considered to reduce the risk of accounts being compromised.
- Password policies must be implemented and enforced across both domain and local accounts.
- Server local administrator accounts must be centrally managed across the IT network.
- Service accounts with high privileges must be managed and controlled.

Recommendation #6: Incident response processes must be improved for more effective response to cyber attacks

- To ensure that response plans are effective, they must be tested with regular frequency.
- Pre-defined modes of communication must be used during incident response.
- The correct balance must be struck between containment, remediation, and eradication, and the need to monitor an attacker and preserve critical evidence.
- Information and data necessary to investigate an incident must be readily available.
- An Advanced Security Operation Centre or Cyber Defence Centre should be established to improve the ability to detect and respond to intrusions.

Recommendation #7: Partnerships between industry and government to achieve a higher level of collective security

- Threat intelligence sharing should be enhanced.
- Partnerships with Internet Service Providers should be strengthened.
- Defence beyond borders – cross-border and cross-sector partnerships should be strengthened.
- Using a network to defend a network – applying behavioural analytics for collective defence.

35.5 Additional Recommendations

681. The nine Additional Recommendations are as follows:

Recommendation #8: IT security risk assessments and audit processes must be treated seriously and carried out regularly

- IT security risk assessments and audits are important for ascertaining gaps in an organisation's policies, processes, and procedures.
- IT security risk assessments must be conducted on CII and mission-critical systems annually and upon specified events.
- Audit action items must be remediated.

Recommendation #9: Enhanced safeguards must be put in place to protect electronic medical records

- A clear policy on measures to secure the confidentiality, integrity, and accountability of electronic medical records must be formulated.
- Databases containing patient data must be monitored in real-time for suspicious activity.
- End-user access to the electronic health records should be made more secure.
- Measures should be considered to secure data-at-rest.
- Controls must be put in place to better protect against the risk of data exfiltration.
- Access to sensitive data must be restricted at both the front-end and at the database-level.

Recommendation #10: Domain controllers must be better secured against attack

- The operating system for domain controllers must be more regularly updated to harden these servers against the risk of cyber attack.
- The attack surface for domain controllers should be reduced by limiting login access.
- Administrative access to domain controllers must require two-factor authentication.

Recommendation #11: A robust patch management process must be implemented to address security vulnerabilities

- A clear policy on patch management must be formulated and implemented.
- The patch management process must provide for oversight with the reporting of appropriate metrics.

Recommendation #12: A software upgrade policy with focus on security must be implemented to increase cyber resilience

- A detailed policy on software upgrading must be formulated and implemented.
- An appropriate governance structure must be put in place to ensure that the software upgrade policy is adhered to.

Recommendation #13: An internet access strategy that minimises exposure to external threats should be implemented

- The internet access strategy should be considered afresh, in the light of the Cyber Attack.
- In formulating its strategy, the healthcare sector should take into account the benefits and drawbacks of internet surfing separation and internet isolation technology, and put in place mitigating controls to address the residual risks.

Recommendation #14: Incident response plans must more clearly state when and how a security incident is to be reported

- An incident response plan for IHiS staff must be formulated for security incidents relating to Cluster systems and assets.
- The incident response plan must clearly state that an attempt to compromise a system is a reportable security incident.
- The incident response plan must include wide-ranging examples of security incidents, and the corresponding indicators of attack.

Recommendation #15: Competence of computer security incident response personnel must be significantly improved

- The Computer Emergency Response Team must be well trained to more effectively respond to security incidents.
- The Computer Emergency Response Team must be better equipped with the necessary hardware and software.
- A competent and qualified Security Incident Response Manager who understands and can execute the required roles and responsibilities must be appointed.

Recommendation #16: A post-breach independent forensic review of the network, all endpoints, and the SCM system should be considered

- IHiS should consider working with experts to ensure that no traces of the attacker are left behind.

36 RECOMMENDATION #1: AN ENHANCED SECURITY STRUCTURE AND READINESS MUST BE ADOPTED BY IHIS AND PUBLIC HEALTH INSTITUTIONS

#VIGILANCE #GOVERNANCE #PEOPLE DEVELOPMENT

682. All organisations, whether commercial, non-profit or governmental, need to build a secure organisation to ensure long-term success. This means that organisations must implement and maintain a strong security posture, including in relation to cybersecurity. This is particularly relevant to organisations like IHiS and the public health institutions ("**PHIs**"), which own and/or maintain public sector IT systems which contain large databases of personal data – failing to secure the organisation can lead to potentially devastating consequences beyond the four walls of the organisation.

683. Over the course of the Inquiry, the evidence showed that certain aspects of the public healthcare sector's cybersecurity posture were poor, in particular on the sector's mindset towards cybersecurity. This was the case even at the MOHH level. At the same time, even for aspects of the public healthcare sector's cybersecurity posture that are adequate, there is scope to further improve. The recommendations below aim to address this.

36.1 Gaps between policy and practice must be addressed

684. A comprehensive IT security policy, on its own, serves little purpose. For such a policy to be effective in fact (as opposed to in theory), the practice on the ground must comply with policy. Hence, any gaps between policy and practice must be addressed.

685. As regards the public healthcare sector, Dr Lim Woo Lip ("**Dr Lim**")'s opinion is that the existing IT security policy framework appears relatively well-established. However, there are obvious gaps between policy and practice – for example:

(a) Password management policies were not properly implemented;

(b) Incident reporting policies were not followed;

(c) Security hardening policies were not properly implemented (*e.g.* Remote Desktop Protocol access was not disabled, and there were patching delays); and

(d) User-ID management policies were not properly implemented (*e.g.* unused or dormant accounts not disabled).

686. As part of enhancing the public healthcare sector's security posture, these gaps must be addressed.

687. To achieve this, CE, CSA has recommended "*deliberate efforts to improve training and adherence to SOPs, as well as raising the level of awareness and cyber hygiene of the healthcare sector's personnel*". The Committee agrees and recommends the following:

(a) Training and Table Top Exercises ("**TTXes**"). There should be greater emphasis on training and TTXes for IT staff so as to build familiarity with policy, and to reveal weaknesses and gaps in practice. One of the greatest security assets is an organisation's own employees, but only if they have been properly trained to comply with security policies and to identify potential security problems.[46] The benefits of training and realistic TTXes will be discussed further in section 38 (pg 269) below, in the context of improving incident response processes.

(b) Audit and compliance. Regular audits and compliance checks are also important. They help to identify non-compliance, and if

[46] *Network and System Security* (John R. Vacca) (Elsevier Inc., 2010) ("***Network and System Security***") at p17.

findings are properly remediated, help to bridge any gaps between policy and practice. To this end, it is important that the ongoing discussions between IHiS and the GIA on the three lines of defence are properly reviewed and that an appropriate decision is taken soon. IHiS acknowledges that the three lines of defence model is a good target model. The key benefit of implementing an effective three lines of defence model is that it improves coverage of risks and controls by identifying and refining the population of risks and controls, and it appropriately allocates the ownership and performance of these risks and controls across the lines of defence. As a result, any unintended risks and gaps in controls can be avoided, and unnecessary duplication of work should be avoided by removing layers of redundant controls. An effective model of three lines of defence will, therefore, better address the gaps between policy and practice.

36.2 IHiS must adopt a "defence-in-depth" approach

688. The adoption of a "defence-in-depth" approach was recommended by CE, CSA, Gen. Alexander and Dr Lim. Defence-in-depth is not a new strategy. The basic idea behind the strategy is to hinder an attacker as much as possible with multiple layers of defence, even though each layer might be surmountable.[47] As Vivek Chudgar ("**Vivek**") has pointed out, "*the enterprise must have full visibility of their internal network, their controls, strengths and weaknesses, their exceptions, it is all about having visibility of your backyard*". In particular, more valuable assets are protected behind more layers of defence.[48]

[47] *Network and System Security* at p92.

[48] *Ibid.*

689. To properly implement defence-in-depth, active steps must be taken to identify and secure vulnerabilities that are "out there", particularly in legacy systems, to protect against future exploitation. Defence-in-depth also involves people, technology and operations[49]:

(a) People. Trained security personnel should be responsible for securing the network and systems;

(b) Technology. A variety of technological measures should be used for layers of protection; and

(c) Operations. Preventative activities (*e.g.* penetration testing, software patching, access controls, *etc*.) and reactive activities (monitoring, detection, blocking isolation, *etc*.) required to maintain security should be put in place. Several measures for this purpose will be set out below.

690. In the context of cybersecurity, *one cannot protect against vulnerabilities that one is unaware of*. IHiS should study and adopt the measures discussed in this report, and consciously layer them to adequately protect its systems. The following measures contain a particular emphasis on the review of systems, assets and networks.

36.2.1 Reviewing legacy systems

691. CE, CSA explained that legacy systems (such as the SCM) are *not* unique to the public healthcare sector and many system owners across the board (*e.g.* public transport, banking and finance and the Government) have re-looked their legacy systems through a new lens of potential vulnerabilities which did not exist at the point when the systems were put in place.

[49] *Network and System Security* at p93.

692. As regards the principle that more valuable assets should be protected behind more layers of defence, it is imperative that stronger, multi-layered security mechanisms should have been in place around SingHealth network's crown jewels – the electronic medical records of all SingHealth patients. This includes safeguards in the system to trigger alarms when abnormal activities are attempted or executed on the crown jewels.

693. An issue was raised in the Inquiry on whether it is realistic to expect a legacy system such as the SCM to have such in-built safeguards. The experts' view on this issue is clear: for legacy systems, there should be a regular process to constantly review such systems and penetration testing should be built-in as part of safety review. CE, CSA is also of the same view.

694. Hence, all legacy systems in the public healthcare sector must be reviewed as a matter of priority. This must involve a thorough review and assessment of legacy systems/applications, including penetration testing and consideration of whether such systems/applications should be isolated or decommissioned (if hardening them is not possible). In this regard, IHiS can consider commissioning an independent external expert to conduct an initial review of all the legacy systems in the public healthcare sector. This will ensure that the review will be objective and provides assurance that the systems have been thoroughly reviewed. Thereafter, subsequent regular reviews can be conducted internally.

36.2.2 Reviewing all assets including lower-priority assets

695. While the defence-in-depth strategy envisages that more valuable assets are protected behind more layers of defence, this is not to say that lower-priority assets are ignored. Vivek's expert opinion is that ignoring such lower-priority assets would be a mistake as such assets are targeted and regularly exploited by APTs. As regards the Cyber Attack, two instances of this were seen:

(a) NCC server: This is a server located at the National Cancer Centre ("**NCC**"). The Committee heard evidence that although the server was an IHiS asset, it was not being managed by IHiS in practice

and was being locally managed by an NCC employee, Tan Aik Chin, since January 2016 by happenstance. As a result, patches that would typically be rolled out automatically for other servers under IHiS' care, were not similarly rolled out to the NCC server. As it turned out, the NCC server was used used by the attacker as a point of distribution for malware, to infect other computers in the network.

(b) S.P. server: This server was a dual-use server, that functioned both as a web server hosting SGH websites accessible from the internet, and as an intranet server for SGH users to store documents. In fact, Director of the Delivery Group, Leong Seng, did not even know that the server had two functions, and could not explain why it was located in the local server zone. As it turned out, the S.P. server was compromised by the attacker and was used on 19 July 2018 in an attempt to regain access to the SingHealth network.

696. The above examples show the real security implications if assets (even lower-priority ones) are "forgotten". Hence, as part of the defence-in-depth strategy, IHiS should regularly review all its systems comprehensively to ensure that the necessary security and mitigation measures are in place across both higher-priority and lower-priority assets. This means, for example, that *all* assets must be identified and centrally managed to ensure that they meet IT security requirements, and are subject to periodic review. This is consistent with the recommendations of the experts who have explained the importance of inventorying all hardware and software assets and having "*full visibility of the IT assets that are added to or removed from the networks*".

697. CEO, IHiS Bruce has explained that there are two exercises currently being conducted in IHiS and the Clusters to achieve this: (i) an asset reconciliation exercise whereby the Infrastructure Services Division is checking the list of devices connected to the network against the inventory of devices managed by IHiS; and (ii) a Ministry-led exercise whereby the Clusters are inventorying their own assets, with a specific emphasis on biomedical equipment.

However, these are manual processes which not only are error-prone but will also require constant updating.

698. Hence, experts such as Dr Lim have recommended the use of an asset discovery tool to automate the asset discovery and management process, as opposed to a physical asset register updated manually. In his expert opinion, such a tool should be adopted to augment a network access control solution (which is limited in its effectiveness as a tool to discover and manage assets in the network).

699. The Committee was informed that IHiS is planning to set up a central Public Key Infrastructure ("**PKI**") to issue digital certificates such that only authorised devices and applications with valid certificates can connect to IHiS' network, and intends for the central PKI to support key exchange for encryption purposes.

700. The Committee also notes that IHiS is working towards the implementation of posture checking, which will ensure that endpoints have necessary operating system ("**OS**") patches and antivirus/malware signature updates before they are allowed to connect to the corporate network. This measure will help to enhance network access controls.

36.2.3 Reviewing the network

701. In addition to the abovementioned regular reviews, rules that allow or limit network traffic between different network segments must be periodically reviewed to identify vulnerabilities. In particular, any changes to the network configuration or architecture must trigger a separate security review to check that the change has not created new gaps in the existing layers of defence. As regards the Cyber Attack, following the migration of the SCM system to H-Cloud, there remained an open network connection from the Citrix server farm at SGH to the SCM database server at H-Cloud data centre. The open network connection was a critical pathway exploited by the attacker.

702. This was a security gap that should and could have been plugged. Instead, the evidence led showed that some senior staff were not even aware of the open network connection until after the Cyber Attack. This was a result of IHiS' current practice of reviewing the network architecture only when there is "*a major change in infrastructure or needs*" – according to Leong Seng, the SCM migration to H-Cloud was not one such change. A more proactive approach, *i.e.* one that would have required a security review of the network following the migration, would likely have identified the gap and IHiS would have had the opportunity to address it in time. Woon Lan in her evidence has said that such a proactive approach is now being considered for the SingHealth network – she explained that she will be putting forth a plan whereby the SingHealth network will be reviewed annually and also reviewed each time there is any major upgrade or migration. It is recommended this proactive approach and plan for network review be enshrined in policy for *all* Clusters (*i.e.* in the HITSPS).

36.3 Cybersecurity must be viewed as a risk management issue, and not merely a technical issue – decisions should be deliberated at the appropriate management level, to balance the trade-offs between security, operational requirements and cost

703. Effective cybersecurity requires an "*acceptance that [cybersecurity] is an organisation-wide problem, not just an IT problem*".[50] As with all high level business risks, cybersecurity should be managed at the senior level of leadership. In any organisation, cybersecurity requires balancing and trade-offs between security, operational requirements, cost; and also patient safety in the case of the public healthcare sector. This requires judgment and accordingly, decisions need to be deliberated at the right level within the organisation – not the technical staff

[50] Mark Barmby, "*Cybersecurity: Moving from Awareness to Understanding*" in *Managing Cybersecurity Risk (Jonathan Reuvid) (Legend Business Books, 2nd Ed, 2018)* ("***Managing Cybersecurity Risk***") at p43.

but by senior management who have responsibility and oversight of the operational and business imperatives.

704. To this end, IHiS and the Clusters must review their organisational and reporting structure, to ensure that cybersecurity considerations and decisions are escalated to the appropriate decision-makers. Some examples are highlighted below.

36.4 Ensuring appropriate lines of reporting

705. On the issue of appropriate decision-makers, an issue that came up in course of the proceedings was whether the double-hatting of officers such as Bruce (as IHiS, CEO and MOH CIO) and Kim Chuan (as Director, CSG and MOH CISO) raised conflict of interest concerns. As mentioned by MD, MOHH, "*there will always be the real possibility that there is a conflict of interest because the person promulgating the policy is the one who implements, and the one checking is the person who promulgated the policy*". There was an attempt to explain this conflict of interest by showing that the double-hatting enables alignment between: (i) MOH's priorities in IT and cybersecurity strategy, policy and programmes; (ii) IHiS' planning and implementation of the same for MOH; and (iii) it ensures a channel for IHiS to provide to MOH feedback from the running of programmes on the ground, so as to inform MOH's policy-making. In an organisation, there needs to be alignment of organisational objectives and processes, but there should not be any conflict of interest. While the Committee can understand the need for alignment of organisational objectives and processes, alignment alone does not address conflicts of interest. The oversight of IHiS by entities such as the CSC does not fully resolve conflicts of interest for IT and cybersecurity strategy and programmes.

706. The Committee notes that MOH is considering setting up an independent CISO office within MOH. This is a step in the right direction in this matter. If such an office is set up, it should be independent of IHiS.

707. In addition, another issue as regards appropriate decision-makers and the potential for conflict which needs addressing is one observed by CE, CSA. The security team in IHiS (*i.e.* the Security Management Division) is embedded as a sub-unit of Infrastructure Services within the Delivery Group. This may result in a misalignment of objectives. Given that the core mission of the Infrastructure Services and the Delivery Group is to provide IT services to the Clusters, security-related workstreams might be overlooked in favour of service delivery objectives. Moreover, the SMD may be too far detached from the key decision makers such as the Cluster's GCEO and GCIO, and the IHiS CEO. This dilutes the authority and effective control of decision makers over the SMD, to ensure that their day-to-day functions are executed properly.

708. The experts have also raised concerns with the current structure. In Dr Lim's expert opinion, there is potential for conflict when IT implementation and IT security come under the same team and same reporting structure. Gen. Alexander recommended that the cybersecurity team in an organisation should have a *direct* reporting line to the CEO – "*[b]y elevating it to the CEO, what the CEO is made aware of is the risks that go beyond operations of the actual IT platform into the security of the platform*".

709. Hence, the current structure should be changed such that the SMD has a direct reporting line to CEO, IHiS.

36.4.1 Ensuring appropriate management visibility

710. Another example is that of ensuring appropriate management visibility when it comes to security incidents. Management visibility is important – only by being well-informed will management be able to react in time and appropriately. It is unrealistic to expect a leader to know *everything* and to know it *all the time*. However, processes and tools should be available to allow management to have as much visibility as possible over security incidents.

711. For example, Vivek's expert opinion is that it would be helpful to put in place a management dashboard that covers not only security incidents which

were responded to and reported, but also incidents which were responded to but did not meet the threshold for reporting. In many cases, as regards the latter group of incidents, even if they had been overlooked or were not correctly responded to, there is an opportunity for review and consideration at the right level, if management has visibility over them. These dashboards should be given visibility all the way up to the organisation's CEO and should be reviewed periodically. This way, management can understand what is going on on the ground, where the bottlenecks are, and if any resources need to be further assigned.

36.4.2 Ensuring appropriate cybersecurity resources at the Cluster senior management level

712. Current practice and policy require the GCIO to have responsibility over a number of functions in the Cluster:

(a) First, the GCIO is in charge of strategic IT planning and development for the Cluster, including the overseeing of project delivery for the Cluster.

(b) Second, the GCIO is also generally responsible for ensuring that the Cluster's IT enterprise programs are aligned with security requirements, ensuring compliance with prevailing security policies and standards, and overseeing the Cluster's IT risk assessment.

(c) The HITSPS also states that the GCIO is responsible for: (i) providing leadership and direction for the IT security program (including the establishment and maintenance of the program objectives, strategy, and near and medium term activities); and (ii) updating the Cluster Board on important IT security matters (including IT security incidents, security policy changes, and non-compliance with security policies resulting from internal audits or from self-compliance reviews).

713. However, the evidence shows that the SingHealth GCIO Benedict does *not* have the resources to properly fulfil his functions in respect of cybersecurity. His GCIO office comprises about 50 staff, but these were mostly IT directors from SingHealth's PHIs and domain/business analysts. Specifically for cybersecurity, Benedict is supported by a team of one – the Cluster ISO Wee. Benedict's evidence is that he and Wee "*collaborate*" with IHiS' Delivery Group and CSG on cybersecurity matters, but generally, they are reliant on IHiS because the technical and engineering capabilities are all centralised within IHiS. Benedict's own evidence is that his technical expertise is limited – for example, as regards approving proposed management responses to audit findings, he can only consider their adequacy "*to the extent of [his] technical knowledge*".

714. As a result, SingHealth's senior management, who rely on Benedict, are left dependent on the central IHiS team to manage SingHealth's cybersecurity risks. This was confirmed by SingHealth's Dy GCEO Prof. Kenneth, who explained that even at management level, SingHealth is "totally dependent" on IHiS for their oversight on cybersecurity risks.

715. This position is difficult to sustain in the light of the new Cybersecurity Code of Practice[51] ("**CCoP**") which requires CII owners (*i.e.* SingHealth) to establish and approve policies, standards and guidelines for managing cybersecurity risks and protecting CII against cybersecurity threats, and to also review the policies, standards and guidelines against the current CII cyber operating environment and cybersecurity threat landscape at least once a year, starting from the date of the last review or the effective date of each policy, standard or guideline.

716. SingHealth acknowledged that based on the present relationships, generally, domain expertise and resources lie with IHiS on the one hand, whereas risks and responsibilities lie with SingHealth on the other. In order to improve on the current governance and risk management framework, SingHealth requires

[51] The CCoP was issued on 1 September 2018.

the ability (through empowerment and domain expertise) to carry out the independent oversight function of IT operations in the cluster in three main areas:

(a) Strategic oversight: policy and project alignment with a cluster's strategic and business interests, and horizon scanning.

(b) Risk management: audit and risk assessments about IT projects and security risks, and checks and balances in decision-making and assessments.

(c) Project management: operations and implementation of IT projects; and pricing, terms, competitiveness and value of project proposals.

717. It is important that there is appropriate cybersecurity expertise at the SingHealth senior management level. One way to do this would be to give the GCIO the right personnel and resources to perform his cybersecurity functions effectively. This minimally would mean increasing manpower in the GCIO office specifically in the area of cybersecurity, and also ensuring that the additional manpower includes personnel with technical and IT security expertise. This way, the GCIO is better equipped to educate and advise SingHealth senior management on cybersecurity risks and the trade-offs that can or cannot be made. There are however, two potential challenges with this approach.

718. First, at a practical level, there may be a challenge in terms of being able to attract enough quality staff at each Cluster CIO office and there is also the concern of duplication of resources (*i.e.* staff with technical and IT security expertise being spread across IHiS and each of the Clusters). Second, there may be a challenge in terms of managing conflicts of interest, given that under the current structure, the GCIO:

(a) Has responsibility over a number of functions (as mentioned in paragraph 712 (pg 245) above), and the GCIO will have to balance between the imperatives of each function, and compromises may have to be made in the allocation of limited operational and

budgetary resources. Gen. Alexander explained that a conflict of interest may arise if an organisation's Chief Information Security Officer is made to report to its Chief Information Officer. Similar conflicts may arise if these functions are held by the same person, as appears to be the case with the SingHealth GCIO.

(b) Is an IHiS employee. Potential conflicts may arise from this fact as the GCIO functions are shaped by IHiS and his KPIs are measured by IHiS.

719. To address the above challenges and to ensure that SingHealth senior management has appropriate oversight on cybersecurity risks, an alternative would be to appoint an independent and dedicated CISO for SingHealth with a direct reporting line to SingHealth senior management. A dedicated CISO for SingHealth will not only have depth and breadth of knowledge about the threat landscape, protective approaches, tools and techniques to protect infrastructure and information, but a unique perspective on how to analyse and mitigate cybersecurity risks. This is consistent with Gen. Alexander's recommendation that each company should appoint a CISO. The Committee notes that MOH is undertaking a horizontal review and assessment on IT governance, policies, standards and processes of MOH, MOHH, SingHealth and IHiS, and this issue may be considered.

37 RECOMMENDATION #2: THE CYBER STACK MUST BE REVIEWED TO ASSESS IF IT IS ADEQUATE TO DEFEND AND RESPOND TO ADVANCED THREATS

#PREVENTION #DETECTION #RESPONSE

720. It is imperative for organisations to give sufficient prominence to technology when formulating and implementing an overall cybersecurity strategy. Of course, it is important that the correct governance structure and policies are in place – technology cannot replace those elements. However, no matter how sophisticated, no paper document or process will thwart an attack until you have strong IT security technologies in place.

721. In Gen. Alexander's expert opinion, a comprehensive cybersecurity capability should be deployed and implemented, as cybersecurity teams cannot protect against threats that they cannot see and that are not detected by the cyber tools they are using. His vision of such a capability is one that not only includes the current set of cyber tools, but also leverages an expert system, behavioural analytics (which is rigorously tested and proven in the networks) and a collective security capability. In his opinion, such a capability would have been important in detecting the theft of credentials, lateral movement in the network, and data exfiltration in the Cyber Attack. Dr Lim echoes the sentiment that organisations like SingHealth need to subscribe to more effective cyber tools to analyse and detect more advanced and sophisticated cyber attacks.

37.1 Identify gaps in the cyber stack by mapping layers of the IT stack against existing security technologies

722. The "cyber stack" is a construct that conveys the notion that IT security must be an integrated set of solutions. No fixed or universally accepted definition of the "cyber stack" is available, but it can be understood as the layers of security technology that an organisation puts in place to form an integrated defence to cyber attacks, by providing prevention, detection and response capabilities to an organisation.[52] The "IT stack" is a hierarchical framework for computing, where network infrastructure and endpoints[53] provide a foundation, with various layers of software and applications on top. Mapping the cyber stack, and the capability provided by security technologies, against the IT stack, provides a framework for gaining greater visibility of the extent to which existing technologies address risks, and allows for gaps in coverage to be identified. This is illustrated in the following figure:

[52] The completeness of the cyber stack is necessary but not sufficient for effective defense. The organisation must have the appropriate expertise and intelligence to effectively operate the cyber stack.
[53] The term "endpoint" as used in this Recommendation refers to both end-user workstations and servers.

Figure 14: Mapping of IT stack against cyber stack

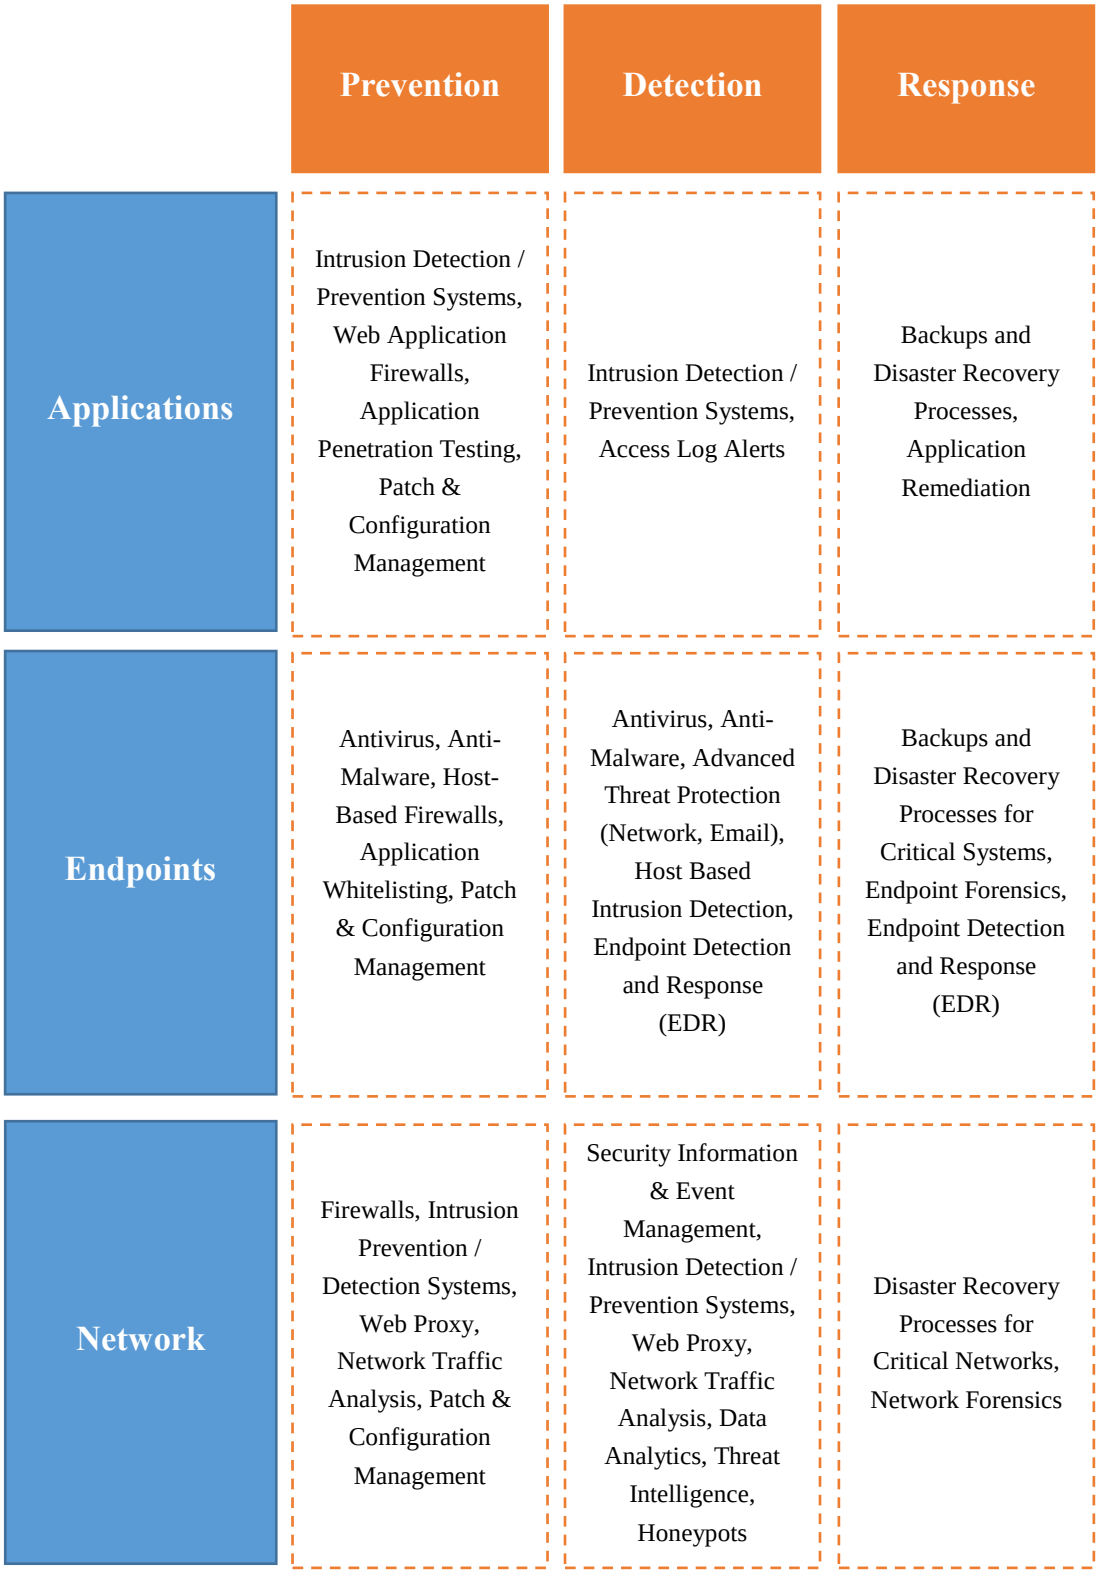

723. The Committee recognises that at the time of the Cyber Attack, IHiS had in place a range of enterprise-level security technologies including:

(a) preventive measures for endpoints, servers, network security, and applications; and

(b) detection measures such as continuous, real time monitoring.

724. Broadly speaking, IHiS had put in place a first line of defence to protect the "perimeter",[54] and several other common necessities such as antivirus and anti-malware systems, intrusion detection/prevention systems, and a SIEM (security information and event management) system.

725. However, as demonstrated in the Cyber Attack, there were gaps in the security framework which allowed the attacker to more easily enter the network, traverse and compromise wide-ranging systems, and make off with the crown jewels. The following measures address the aforementioned gaps.

37.2 Gaps in response technologies must be filled by acquiring endpoint and network forensics capabilities

726. While Leong Seng's evidence has addressed the technological measures in place to support prevention and detection measures, the silence in relation to technological systems in place to support the *response* to a cyber attack is telling. IHiS does not have such technological support in place. The "Response Measures" he has informed the Committee of relate only to processes.

37.2.1 Endpoint forensics

727. There is no enterprise-level forensics platform in place – IHiS uses only open source software for its forensics. These tools require IT security staff to

[54] The "perimeter" is the "outer wall" or the (logical) border line around an organisation's infrastructure and network, which separates it from an untrusted network such as the Internet.

physically go to individual machines to image them, or to take memory dumps – this is a process that simply takes too long, according to Vivek, and valuable time is lost in responding to a security incident.

728. In any event, even with the use of the open source forensics software, IHiS had no dedicated and suitably-equipped computers to run the desktop-based forensic software; Benjamin in fact used his personal laptop when running the software during forensic investigations. This being the only computer that could be used to carry out forensic investigations, the processing of digital forensic evidence required a "painfully long amount of time".

729. The gap in the response technologies available at IHiS undoubtedly hampered the response to the Cyber Attack, and their ability to limit the impact of the Cyber Attack.

730. We recommend the implementation of a centralised enterprise-level forensics platform for collection and analysis of digital evidence. Features of such a system would include:

(a) 360-degree visibility across all endpoints;

(b) Remote collection of forensic artefacts; and

(c) An ability to search and collect forensic evidence across multiple devices concurrently.

731. In the case of the Cyber Attack, the IHiS SMD did not have access to an endpoint detection and response ("**EDR**") system[55] that would have allowed rapid isolation and containment of the infected systems, and enabled the *rapid collection of forensic evidence from multiple systems at the same time*. Vivek testified that an EDR system would have allowed the team to fast-track the

[55] See section 37.3 below for further elaboration on EDR systems.

collection of forensic artefacts, speeding things up from taking several days to weeks, to taking just one day.

732. This centralised EDR can be monitored by an Advanced Security Operations Centre ("**ASOC**"), and integrated with the rest of the detection and incident response processes. Evidence can be collected remotely, consolidated with other inputs, and analysed for indicators of attack. MOH envisages that IHiS' planned ASOC provider will assist IHiS with forensic and threat hunting capabilities, development of security tools, and security threat analytics.

37.2.2 Network forensics

733. Almost all modern network equipment such as routers, switches, firewalls *etc*. support the ability to capture data regarding network traffic that flows in and out of such devices. While it appears that IHiS had tools to capture network traffic information, they did not have the means to analyse it effectively for forensic purposes.

734. IHiS has access to NetFlow data, which contains information about traffic that traverses the network. NetFlow can provide complete network visibility by providing the ability to collect and store network traffic metadata.[56] Network administrators typically analyse NetFlow data to determine the source and destination of traffic, the type of service involved, and the causes of congestion. In essence, it is information largely used for troubleshooting purposes. However, NetFlow is also valuable for network forensics as shown by its use after the Cyber Attack to determine whether the stolen patient data had been exfiltrated.

735. However, it appears that the ability to obtain forensically significant information from the massive volume of traffic data was hampered by the lack

[56] "Metadata" includes information such as username, source and destination IP, URL, start and end time and much more. See Plixer: NetFlow Version 9 <https://www.plixer.com/support/netflow-v9/>.

of data analytics tools. This should be addressed by acquiring the necessary technological solution to maximise the use of NetFlow information.

736. With the necessary analytical tools, NetFlow can provide anomaly detection and investigative capabilities that can be used in incident response, for example, to uncover behaviour that may have been occurring over a long period. When a security incident is being investigated, the flow database can be used to determine what IP addresses accessed a system, the times the system was accessed, as well as quantifying the impact on related systems that the host conversed with on the network, before and after the incident. Without automated analytics, trawling through huge volumes of flow-data would be nigh impossible to determine the actions of a long term threat actor residing within a network, who may have been dribbling out stolen data over a prolonged period. Vivek emphasises that NetFlow alone is insufficient – in the context of traffic leaving the network perimeter, he stated that analytical intelligence needs to be applied to help determine if the outbound traffic is suspicious, and to determine if the data is indicative of beaconing by malware. Without this analytical ability, NetFlow alone would result in an information overload.

737. However, it must be noted that NetFlow itself does not contain any content of the observed traffic. The Committee was informed that IHiS has begun efforts to enable NetFlow at routers and switches to collect traffic information for traffic profiling and intrusion detection, in particular, those relating to traffic moving laterally from server to server.

37.3 Effectiveness of current endpoint security measures must be reviewed to fill gaps exploited by the attacker

738. Endpoint security protects desktops, laptops, servers *etc*. from malicious internal and external threats. As security technology becomes more sophisticated, so do attackers' tools, tactics, and methods. Attackers are now adept at discovering the weak points in enterprise security strategy – and increasingly, endpoints are being targeted. However, asset classification is often still used as the means by which to prioritise risk, resulting in endpoints (assets of low priority

classification) being regarded as being of low risk. Consequently, endpoints have less coverage in terms of defensive, preventive and detection controls. Attackers know this, and exploit this vulnerability by targeting endpoints as part of their *modus operandi*.

739. Endpoints are the common points of ingress for attackers, and the platforms from which an attack is propagated, after initial breach is achieved. Further, multiple endpoints may be compromised during lateral movement, as the attacker navigates the network towards its end objective.

740. Given the nature of the advanced cyber threats that organisations now face, conventional signature-based and prevention-oriented solutions are insufficient. The conventional technique for detecting malware is to check to see if a program or process has been previously identified as being malicious. These checks depend on "signatures" that have been identified as being associated with the program or process – the name of the program or process, the size of the program, the date when it was created, a hash of the program *etc*. A signature-based approach to detection has two primary weaknesses. First, it is easy to alter the malware code without affecting what it can do. An unlimited number of functionally equivalent variants of the malware can thus be created with different signatures, thereby frustrating signature-based detection. Second, signature-based detection cannot identify a program as a virus or malware if the program has never been seen before.

741. Further, a new type of so-called fileless malware has emerged. Unlike attacks carried out using conventional file-based malware, intrusions using fileless malware do not involve attackers installing malicious programs on a victim's computer. Instead, tools that are built-in to Windows (for example, PowerShell) are abused by attackers and used for malicious purposes. The fact that conventional file-based malware is not used is significant, as this means that there is no signature for antivirus software to detect. Fileless malware can not only slip into a system without being detected by signature-based endpoint

security, but can also make itself persistent by manipulating Windows registry[57] entries. These entries will cause malware code to be reloaded into the computer's memory, even after the computer is rebooted, which would normally wipe out any purely memory-based malicious code.

742. It is therefore increasingly accepted that traditional anti-malware software is inadequate, and that a new strategy must be created to identify breaches at endpoints. Indeed, this was shown to be true in the Cyber Attack – while IHiS had enterprise-level antivirus and anti-malware protection for endpoints, the signature-based system was unable to prevent endpoints from being infected by fileless malware, nor could it detect the customised Remote Access Trojan deployed by the attacker.

743. To combat the sophisticated threats of today, modern endpoint security requires an endpoint security system with advanced security technologies and services, such as EDR, predictive analytics, and incident response. Advanced endpoint security solutions do not only address prevention, but also detection and response. The Committee notes that IHiS is in the midst of planning a roll out of EDR at all endpoints. Once rolled out, it will be able to detect IOCs (indicators of compromise), and record endpoints' system-level behaviours and events such as user or file processes, as well as registry, memory and network events.

744. Expert witnesses Dr Lim and Vivek have recommended the implementation of EDR systems. According to Gartner, EDR is a security technology "*created to satisfy the need for continuous detection and response to advanced threats – most notably to significantly improve security monitoring, threat detection and incident response capabilities.*"[58] Vivek recommends the use of EDR as it is a detection system that looks comprehensively at the overall network – the operating system, and the behaviour of the software operating on

[57] The Windows registry is a database of information, settings, options, and other values for software and hardware installed on Microsoft Windows operating systems.
[58] Business Wire, "Guidance Software Recognized as the Estimated Market Share Leader by Gartner in the Endpoint Detection and Response (EDR) Tools Market", December 2014.

that system *etc*. It provides more real-time information (as opposed to simply historical logs) of detectable and observable events in the network, and does not rely on the detection of known signatures only. He recommends that the EDR be centrally managed, bringing him in agreement with Dr Lim's recommendation for the implementation of what he refers to as "managed EDR" ("**MDR**"). MDR allows for the achievement of enterprise network visibility for more effective detection of advanced cyber threats. More than simply EDR, MDR collects, correlates and analyses all data obtained within an EDR, and can determine communications and movements *between* endpoints in different parts of the network. The system runs on two levels: there are software *agents* that run in the background on endpoints, and a centralised *endpoint security management system* that monitors and controls the agents. In essence, MDR allows a look at the bigger picture – a holistic look at data on a system level.

37.3.1 Detection

745. EDR tools work by monitoring endpoint and network events, and recording this data for analysis, detection, investigation, reporting and alerting. Such tools use sophisticated analytics that identify patterns and detect anomalies in the network, including rare processes, strange or unrecognised connections, or other risky activities that are flagged based on baseline comparisons. This monitoring process can be automated, and anomalies will trigger alerts for immediate action or further investigation. Instead of being a signature-based system, EDR systems use anomaly-based detection which compares definitions of what is considered normal activity, with observed events, in order to identify significant deviations. As explained by Vivek, this detection method can be very effective at identifying previously unknown threats.

746. Defending networks from cyber attacks necessitates a comprehensive EDR system which should meet the following criteria:

(a) Has comprehensive detection that (i) leverages on security analytics to identify threats, and (ii) automates threat detection across the Cyber Kill Chain.

(b) Offers total coverage that effectively secures all endpoints, including desktops, laptops, servers, and virtual environments.

(c) Has predictive security capabilities that use artificial intelligence and security analytics to predict potential threats, reduce false positives, and accelerate incident response.

37.3.2 Response

747. As discussed at paragraph 731 (pg 253) above, an EDR system would allow for rapid isolation and containment of infected systems, and enables the rapid collection of forensic evidence from multiple systems at the same time. In summary, an effective EDR system has the following capabilities, relevant to responding to a security incident:

(a) Allows for a complete response, as it validates, triages and remediates the effects of any threat with digital forensics.

(b) Seals off potentially compromised endpoints during investigations, and has the ability to do so remotely, without an IT security officer going to a compromised endpoint directly to physically disconnect it from the network.

(c) Allows for remote remediation of compromised systems by deleting malicious files and associated artefacts on all impacted endpoints.

(d) Conducts investigation and containment of suspicious events by sandboxing, quarantining, and retrieving endpoint process dumps.

748. Our recommendation is that IHiS/SingHealth (as CII owner) and other CII operators must implement advanced endpoint security solutions, given the clear evidence of how signature-based systems were thoroughly defeated in the Cyber Attack.

749. IHiS has fast-tracked the deployment and installation of an Advanced Threat Protection ("**ATP**") system in the aftermath of the Cyber Attack. The ATP system is an advanced endpoint protection system which is described as being able to replace a traditional antivirus system. Rather than try to keep up with the ever-growing list of known threats, it sets up a series of roadblocks that prevent the attacks at their initial entry points – where malicious access to the system is made through the abuse of legitimate executable files.

750. However, unlike EDR systems, the ATP system implemented by IHiS does not appear to have the response capabilities described above at paragraph 747(b)-(c) (pg 259). As such, IHiS must consider implementing a separate solution to fill these gaps.

751. At the end of the day, for effective detection and response to security incidents, the technical solution implemented must be able to send alerts and/or block the following attack methods that were observed in the Cyber Attack:

(a) the running of unauthorised applications;

(b) the use of system tools for malicious purposes (*e.g.* the solution must protect against fileless malware, the use of PowerShell, and methods of moving laterally in network); and

(c) the running of unauthorised Virtual Machines (*e.g.* as seen in the use of VM 1 and VM 2 to log in to Citrix servers).

37.4 Network security must be enhanced to disrupt the 'Command and Control' and 'Actions on Objective' phases of the Cyber Kill Chain

752. The SANS Institute defines network security as the process of taking physical and software preventative measures to protect the underlying network infrastructure from unauthorised access, misuse, malfunction, modification, destruction, or improper disclosure, thereby creating a secure platform for

computers, users, and programs to perform their permitted critical functions within a secure environment.[59]

753. According to Leong Seng, the SCM IT network had preventive measures securing network traffic at every tier and every access point, including within and across the various sectors. He added that prior to the Cyber Attack, IHiS had in place a range of preventive measures to address network security, including:

(a) Network firewalls, which segregate each network segment so as to ensure that only authorised network traffic is permitted to cross segments or zones, and which filter incoming and outgoing network traffic based on sets of rules;

(b) Intrusion Detection and Prevention Systems ("**IDS/IPS**"), which are used in SingHealth and H-Cloud to inspect network traffic in real-time, and to block and generate alerts for traffic associated with security risks and threats; and

(c) Proxy servers, which act as intermediaries between users and the internet.

754. However, the tools and technologies in place were shown to be inadequate during the Cyber Attack, in two respects:

(a) callbacks to C2 (command and control) servers went undetected for months; and

(b) lateral movement by the attacker through numerous systems similarly went undetected.

---

[59] SANS Institute, Network Security Resources.

755. These two aspects of the attacker's behaviour relate to the 'Command and Control' and 'Actions on Objective' phases of the Cyber Kill Chain.[60] Had the network cyber stack been adequate, the Cyber Kill Chain may have been disrupted at either one, or both, of these phases.

37.4.1 A solution must be put in place to better detect and block malicious outgoing traffic

756. C2 servers, to which callbacks were being made from compromised endpoints in the SingHealth network, were identified through malware and forensic analysis by CSA.

757. During the early stages of the Cyber Attack, outgoing communications with one C2 server were detected, but only by the fortuitous actions of Benjamin, who discovered the callbacks in January 2018 when investigating a malware infected workstation. However, human error on the part of Benjamin resulted in this C2 not being blocked. Worse still, according to Ernest, the Senior Manager of SMD, communications with the C2 server need not have been blocked, in any event, as it had not been "confirmed" as being a malicious C2. The failure to have an effective solution to automatically detect and block malicious outgoing traffic had dire consequences as the C2 server was actively used throughout the attack in June/July 2018.

758. It is precisely to avoid errors in judgment like this, that Vivek has recommended the implementation of advanced detection tools for malicious traffic on all outbound internet traffic. This is important because most attacker communications would have to traverse the internet and so can be spotted with the right level of monitoring. Alerts should be configured such that every detection of C2 traffic is treated with high priority.

---

[60] The Cyber Kill Chain reveals the phases of a cyber attack: from early reconnaissance to the goal of data exfiltration. See also paragraph 141 (page 51).

759. As regards the defences that IHiS had in place, firewalls can block malicious outgoing communications with C2 servers, but only if the C2s are *known* to be malicious – firewall rules can then be created to block outgoing communications with such servers. Firewalls, however, will be of little use in blocking outgoing communications with servers that are *not* known to be malicious.

760. As for intrusion detection systems such as IDS and IPS, these typically focus on the protection of local resources by identifying signs of malicious activity to help prevent a network intrusion and limit its effects. They are designed to prevent *incoming* attacks by checking all incoming traffic for security threats. IDS and IPS typically would not monitor outgoing traffic.

761. However, an IDS/IPS can be configured to monitor *outgoing* traffic to monitor and help mitigate compromised hosts on a network from reaching the internet, and this can prevent C2 functions. It is unclear if the IDP/IPS deployed in the SingHealth were so configured. What is clear is that monitoring of the SingHealth network did not flag the callbacks from compromised endpoints to multiple C2 servers.

762. It is recommended that IHiS review the effectiveness of current security technologies to detect and block malicious outgoing traffic. If no effective solutions are in place to detect callbacks to malicious C2s, such capability must be attained. Vivek has informed that solutions with such capability to provide real-time intelligence on callbacks to C2s are commercially available. These solutions can recognise the signature of malware calling back to a C2, since callback traffic has several fingerprints which can be tracked and caught.

763. Such solutions are known as ATP for networks, and typically use anomaly-based detection, which often relies on machine learning. When outgoing traffic deviates from parameters of traffic known to be benign, *i.e.* "good" traffic, the system takes this as evidence of malicious traffic and responds accordingly.

37.4.2 Modifications to network architecture and/or monitoring of east-west traffic within the network must be undertaken to limit the ability of attackers to move laterally within a network

764. Traditional security thinking prioritises preventing an initial intrusion into a network. However, the initial compromise is often only the beginning. Once an attacker gains a foothold, it would attempt to move around the network and access other systems. This was the case in the Cyber Attack.

765. Once the attacker had established an initial foothold, network logs indicate that the attacker moved laterally in the network between December 2017 and June 2018. Forensic analysis revealed clear indicators that the attacker had moved laterally around the network. For example, the PHI 1 Workstation was compromised and infected with malware on 18 January 2018. This infected workstation was also found to be communicating with foreign C2 servers. Moving laterally, the attacker also gained access to Workstation B and planted a customised Remote Access Trojan on 17 April 2018. After this workstation was compromised, the attacker was able to remotely log in to Citrix Servers 1 and 2 using the L.A. account and the S.A. account. The attacker had planned its route in the SingHealth network to reach its ultimate objective – the SCM database.

766. Given the risk of lateral movement in a future attack, IHiS must adopt measures to structure the SingHealth network in such a way to limit an attacker's opportunity to move laterally, or implement solutions to monitor, detect, and block lateral movement.

767. Network segmentation. Network segmentation in computer networking is the act of splitting a computer network into sub-networks, each being a network segment. In essence, groups of systems or applications are separated from each other. One of the advantages of splitting a network in this manner is improved security, as it makes it more difficult for an attacker to propagate an attack throughout the entire network. For example, there is a reduced attack

surface available for the attacker to pivot[61] in if one of the hosts on the network segment is compromised.

768. By contrast, a flat network infrastructure, while easy to manage, provides a greater prospect for malicious activity. In a flat network, all servers and workstation are on the same local area network ("**LAN**"), which may be unnecessary, as in most cases, the systems have no reason to talk to or "trust" each other. The principles of least privilege and need-to-know should be used. If a host, service or network does not need to communicate with another host, service or network, it should not be allowed to. If a host, service or network only needs to talk to another host, service or network on a specific port or protocol, and nothing else, it should be restricted to that. Allowing open communication between hosts, services or networks, when it is *unnecessary*, offers multiple pathways for an attacker to pivot from one system to another, and allows malware to propagate across the network.

769. Gen. Alexander, Vivek, Richard, and Dr Lim have all recommended network segmentation as a means of limiting an attacker's ability to move laterally in a network. In essence, according to these experts, network segmentation makes the attacker's job exponentially more difficult; it makes it much harder for an attacker to move laterally within a network as systems are not all inter-connected. An attacker would have to exploit segments one at a time, resulting in a far longer time to compromise the network as a whole. The additional time it takes an attacker to break into a network is valuable time given to the defenders to stop the attacker from succeeding.

770. In fact, in the wake of the FY16 GIA penetration test, GIA too had recommended network segmentation, in the form of a separate management virtual LAN ("**VLAN**") that should be established for administration access. This recommendation was made after the penetration testers observed that administrative access to critical infrastructure was possible from the employee

[61] Pivoting refers to the use of a compromised system to attack other systems on the same network.

endpoints through RDP, allowing for lateral movement throughout the network, without restriction.

771. Segmentation is a significant deterrent to attackers, and we recommend that a network architecture review be carried out to segment the SingHeath network as part of a defence-in-depth strategy.

772. Monitoring of east-west traffic. After an attacker has gained access to a network, seeing, detecting and tracking their actions is crucial to reducing the likelihood of their mission objective (*e.g.* exfiltration of data) being achieved. East-west visibility of network traffic refers to the ability to see malicious activity that is *contained within* the network.

773. Unless a specific solution is in place to monitor east-west network traffic, blind spots will exist where an attacker could be hiding. Without a solution for monitoring network communications between endpoints, a wide variety of malicious lateral movement will not be detected, and valuable forensic information will not be collected; information which could prove essential for analysis after an attack.

774. We recommend that an east-west security solution be implemented that can identify abnormal traffic types on the network. Leong Seng has testified that IHiS plans on enhancing network monitoring of east-west traffic to detect lateral movement within the entire network.

37.5 Application security for email must be heightened as it is the most common attack vector for cyber attacks

775. It starts with one email – malicious emails are, by far, the weapon of choice for cyber attackers. The Cyber Attack has reaffirmed the fact that emails are the most common intrusion vector,[62] and that stepped-up measures are essential to defend against this threat.

776. CSA's hypothesis was that the initial intrusion into the SingHealth network was *via* a phishing email. CSA was unable to determine conclusively what the source of the initial infection was, but based on a phishing email sent on 18 July 2018 when the attacker attempted to regain a foothold in the SingHealth network, CSA's hypothesis was that the attack vector was a phishing email containing malicious code.

777. While we acknowledge that no security solution can be 100% effective, the successful phishing attack in 2017, and the fact that in the Cyber Attack similar emails laden with malicious code passed through email security filters and reached the inboxes of a number of recipients in SingHealth institutions, necessitates an urgent review of email security measures that are in place.

778. According to Leong Seng, SingHealth email systems are managed centrally by IHiS with multi-layered preventive measures including:

(a) Antivirus, anti-spam, and attachment blocking technology, which filters emails that may pose security risks, analyses attachments, and scans macros in attachments;

(b) URL re-write technology to detect malicious URLs and render them benign; and

---

[62] An intrusion vector, or attack vector, is a path or means by which an attacker can gain access to a computer or network in order to deliver a payload or malicious outcome.

(c) ATP which analyses URLs and attachments in virtualised machines.

779. Vivek's expert opinion is that IHiS needs to implement advanced malware detection on incoming emails as emails remain one of the most preferred means by which advanced attackers target organisations. Phishing emails allow the attacker to get to a person with the right kind of message to lure that person into clicking on the email, attachment, link *etc.* In the process, the attacker gains a foothold in the network. Vivek also recommends that alerts be configured such that every advanced malware detection is treated with high priority. Protecting emails from advanced malware will go a long way in stopping cyber attacks at the point of entry.

780. We recommend that IHiS, together with CSA, review the efficacy of the email-protection measures that are currently in place, as testified to by Leong Seng. With such measures in place, the questions that need answering are: "Why did the phishing emails go undetected? Was there a failure in technology? Were the emails cleared as benign, when they were in fact malicious? Was the malicious code not detected because the systems currently in place are signature-based, and the code had not been seen in the wild before? Was it a process failure? Were the emails flagged as malicious, but alerts/blocking were not triggered as required?"

781. These are crucial questions that must be answered in order for IHiS, working with CSA, to ensure that adequate email protection measures are in place henceforth.

38 RECOMMENDATION #3: STAFF AWARENESS ON CYBERSECURITY MUST BE IMPROVED TO ENHANCE CAPACITY TO PREVENT, DETECT, AND RESPOND TO SECURITY INCIDENTS

#PREVENTION #DETECTION #RESPONSE #PEOPLE DEVELOPMENT

782. Employees can be the first line of defence in a cyber attack, but they can also be an organisation's Achilles heel. If employees do not understand security policies and procedures, how to mitigate risks, or are not prepared to respond to a security breach, they are potentially contributing, whether intentionally or not, to breaches in cybersecurity.

783. Even the best technological solutions can be circumvented by lax security practices by end-users. For example, in the case of the Cyber Attack, CSA's hypothesis is that the attacker gained its initial foothold *via* a phishing email.

784. It is thus important to inculcate in all staff a culture of good cyber hygiene, and the understanding that cybersecurity is everyone's responsibility, not just that of the IT department.

785. Having strong security technology is not enough. This is in recognition of the fact that cybersecurity is both a science and an art. Even if one is able to achieve the science (*i.e.* all the technical capabilities), it can be undermined by people who are un-trained in the art. Training employees in cybersecurity is therefore a priority. Adequate training for personnel can dramatically decrease the likelihood of a successful cyber attack.

38.1 The level of cyber hygiene among users must continue to be improved

786. Organisations cannot only focus on external cybersecurity threats – they must also focus on the role their employees may play in exposing vulnerabilities from within.

787. Despite efforts in cyber training and literacy, employees continue to engage in risky cyber behaviour. As aptly stated by CE, CSA:

> "[T]he Clusters and IHiS must continue to improve the level of cyber hygiene among all front-end users – doctors, nurses, pharmacists and administrators – in the public healthcare clusters. Front-end users are often the weakest link in cybersecurity. Increasingly sophisticated social engineering techniques, combined with human error, give threat actors the means to establish their initial footholds onto a network. The vast majority of cyber-attacks are not that technically sophisticated, and can be averted by raising the basic level of cyber hygiene throughout the organisation… Promulgating basic security practices, such as the use of strong passwords and being able to spot signs of phishing, can greatly improve the level of cybersecurity in an organisation."

788. Empowering people with good cyber defence habits can significantly increase readiness. It is not just IT staff who needs to practise good cyber hygiene habits, it is a responsibility that falls on everyone in an organisation.

789. IHiS' and SingHealth's efforts in training their staff in this area can be summarised as follows:

(a) Efforts in relation to SingHealth staff:

(i) IT security training conducted by IHiS for: all new staff; staff newly promoted to managerial-level; as well as junior doctors, trainees and personnel on attachment;

(ii) Security alerts from IHiS' IT security team through email broadcasts to all staff (*e.g.* alerts on the Ransomware attack on the National Health Service in the United Kingdom and seasonal threats such as malware infection *via* e-greeting cards during festive periods, *etc*.);

(iii) Memos from management on significant cybersecurity risks and incidents;

(iv) Talks by IHiS' IT security team and external experts at town halls and healthcare conferences organised by SingHealth; and

(v) Phishing exercises conducted by IHiS on all SingHealth staff to create awareness and promote vigilance. These phishing exercises have been conducted regularly every year since 2015, and according to SingHealth, the proportion of staff who responded to the test phishing emails decreased significantly from 14% in the first exercise in 2015 to 3.8% in the most recent exercise in 2018. Staff who responded to phishing emails twice or more, are also given additional attention. They are requested to attend IT security briefings to become more aware of the risks and in the recent exercise in February 2018, such staff also received a formal letter, with a copy to their direct report, signed off by both SingHealth GCIO Benedict and Dy GCEO Prof. Kenneth, to strongly remind them on the need for vigilance.

(b) IHiS' efforts in relation to their own staff:

(i) IT security updates are shared with IHiS staff through a Chief Information Security Officer ("**CISO**") blog created by Kim Chuan;

(ii) CSG sends out weekly email blasts sharing the latest news in IT industry security trends;

(iii) Email blasts to inform IHiS staff of security policies and responsibilities, as well as to alert staff of security vulnerabilities; and

(iv) Provision of IT security information on the IHiS intranet.

790. However, these efforts failed to equip IHiS staff, in particular the SMD, to respond effectively to the Cyber Attack.

791. Current efforts at increasing cybersecurity awareness by SingHealth and IHiS have focused on employee on-boarding, and periodic dissemination of cybersecurity best practices *via* various channels, as highlighted above. Although the existing measures reflect effort and good intentions on the part of management, it is telling that at least in the area of creating awareness about the risks of phishing, a disturbing number of SingHealth staff fell prey to the phishing emails twice or more.

792. Aside from the phishing exercises conducted on SingHealth staff, there was no way to assess if IHiS and SingHealth staff absorbed and understood the cyber hygiene habits required of them. The bare efforts by IHiS in relation to their own staff in particular, were not operationalised in a manner that ensured that information disseminated was *in fact* even read by any of the staff.

793. The Cyber Attack has demonstrated that it only takes one employee to trigger a potentially disastrous cyber incident. In order to ensure that each and every member of staff is educated sufficiently, to identify and report cyber incidents, current efforts in SingHealth and IHiS must be improved upon.

38.2 A Security Awareness Programme should be implemented to reduce organisational risk

794. Providing security awareness training is a reliable way to reduce the insider threat and alter user behaviours. However, current efforts at creating security awareness must be improved in line with best practices.

795. It is recommended that a security awareness programme for all staff be implemented, which must be completed on a regular basis, to ensure they understand and exhibit the necessary behaviours and skills to help ensure the security of the organisation. This is in-line with best practice standards set out in the Center for Internet Security ("**CIS**") Controls[63].

796. A comprehensive security awareness programme would educate staff not only on how to identify a cyber incident (by educating them on what to look for in a cyber incident and what the threats are), but on how to react in case of a cyber incident (*e.g.* by reporting it). The understanding in this day and age must be that incidents can and will occur frequently, and it is therefore critical that all staff, at all levels, know how to react. One means by which to then assess staff understanding of this information, is to implement an online questionnaire testing the staff's ability to recognise indicators of a cyber attack, and their awareness of reporting lines and procedures.

797. The training must express the idea that cybersecurity is everyone's responsibility – not just the IT department. Staff across all domain areas, not just those in the security team, must be trained in cybersecurity detection and response. Vivek's expert opinion is that the training could be a two-day programme, where staff are sensitised to how cyber attacks have evolved, how attacks play out, the *modus operandi* of cyber attackers today, and the kind of weaknesses that may be exploited by an attacker. In addition to that, TTXes and

[63] CIS Controls Version 7 at sub-control 17.3. <https://www.cisecurity.org/blog/cis-controls-version-7-whats-old-whats-new/>.

simulation exercises *must* be conducted, as they are crucial in providing staff with real experience in dealing with cyber attacks.

798. Make it ongoing – The training should be conducted for all incoming staff and existing employees, and should be regularly updated[64] based on evolving policies and threats. This regular staff engagement ensures that security is at the forefront of their minds, and increases the likelihood that they will understand and adhere to security policies. To assess their understanding of policy, they should be given case studies that allow them to practise identifying and dealing with security threats. In addition, staff should be updated on topical security issues.

799. Sharpen the focus – Instead of trying to tackle dozens of security topics, there should be focus on themes that matter most to the organisation, and which will result in the greatest reduction of risk – keeping in mind that different classes of staff, and different departments, face different risks. If necessary, external vendors can be engaged to help customise training materials for specific needs.

800. Best practice guides indicate that a security awareness programme should train employees on[65]:

(a) secure authentication;

(b) identifying social engineering attacks;

(c) sensitive data handling;

(d) causes of unintentional data exposure; and

(e) identifying and reporting incidents.

---

[64] *Ibid* at sub-control 17.4.

[65] *Ibid* at sub-controls 17.5 – 17.9.

801. Use real-life incidents – Concrete examples of breaches and other security headlines should be used in the programme, to add realism and legitimacy to security awareness efforts. As explained by Vivek, sending staff generic messages about cybersecurity and awareness is ineffective – no one reads these messages in the way they should. What works far more effectively is the use of narrative and "storytelling" – using real experiences and examples of cyber incidents to illustrate the key learning points for staff is more effective as the staff are more likely to read and understand the information provided to them. Moving forward, the Cyber Attack itself can be used as a useful "storytelling" implement to educate users on many aspects of cyber attacks, including APTs.

802. Test effectiveness of training – Simply waiting for a security breach to test employee readiness cannot be the right strategy. Mock attacks staging simulated social engineering campaigns should be executed, to assess whether the number of staff falling for them is decreasing. Apart from phishing emails that seek to entice users to click on malicious links, simulated attacks should cover other social engineering scenarios, such as requesting users to divulge user credentials to the 'helpdesk'. In addition, organisations should go beyond using emails in their simulations, for example by employing impersonation phone calls to employees *etc*.

803. Reward good performance – As previously mentioned, in the wake of simulated phishing attacks, SingHealth staff who responded to phishing emails twice or more, are also given additional attention. They are requested to attend IT security briefings to become more aware of the risks and in the recent exercise in February 2018, such staff also received a formal letter, with a copy to their direct report, signed off by both SingHealth GCIO Benedict and Dy GCEO Prof. Kenneth, to strongly remind them of the need for vigilance. Aside from the use of brickbats, staff who perform well in the training and simulation exercises should be recognised and rewarded. Incentives help encourage behavioural changes, and some companies have turned to using gamification to make security awareness education more compelling *e.g.* points and prizes may be awarded to employees who flag a phishing message.

804. Reinforce the message – Training courses are unlikely to have a lasting effect if they are one-off or only conducted infrequently. It is important for key points to be reinforced and this can be accomplished through refresher sessions, as well as through media such as blogs, posters and newsletters.

38.3 IT staff must be equipped with sufficient knowledge to recognise the signs of a security incident in a real-world context

805. Detection is a state of mind. Detection can only happen where there is awareness on the part of the staff. IT staff must be equipped to go beyond recognising obvious cyber attacks such as ransomware and website defacement. *All* IT staff must be equipped with sufficient awareness of cyber threats and signs of a security incident to be able to respond effectively should the need arise. In fact, as recommended by Vivek, this awareness should go even further – training in cybersecurity fundamentals must be provided to *all* IT staff to equip them to play a role as a member of the SIRT team in the event of a cyber incident. This would entail *all* IT staff participating in cyber crisis simulation exercises that simulate real-life scenarios related to advanced cyber attacks. The importance of proper training is echoed by Gen. Alexander – training does not need to be limited to personnel of a certain level of capability. Everyone should be trained and there must be a culture of constant learning.

806. When a cyber incident occurs, IT staff might be the first to notice. An organisation will be at significant risk if there is a lack of initial security incident cognition on the part of its IT staff. It is not safe to assume that IT administrators are prepared and equipped to identify and manage a security incident, and to respond in the initial stages, before security personnel enter the picture. The skills, training, and mindset of an incident responder are different from that of a system administrator. To better respond to security incidents, operational-level IT staff should be provided with a sufficient degree of cybersecurity training.

807. When employees fail to report, or delay in reporting security incidents, it can lead to dire consequences and increase the overall extent of the attack. A single, unreported threat could lead to a large breach. This lesson was learnt in a

painful way in the Cyber Attack. The attacker was sophisticated but was not silent – signs of the Cyber Attack were seen by a number of IHiS staff over a substantial duration of time. Unfortunately, these were not correctly recognised as signs of malicious activity. The experiences of Katherine, Lum and Sze Chun are most illustrative – each of them observed these signs at various times, but each of them assumed these signs to be indicative of operational issues, rather than evidence of a security incident. Katherine, for example, thought that the failed logins to the SCM database were an indication of IHiS staff "*testing the system*". She similarly thought that the active queries to the SCM database on 4 July 2018 were queries being run by other IHiS staff. Sze Chun, who first caught the unusual queries to the SCM database on 4 July 2018, also did not think them suspicious initially – he thought that these were legitimate queries that were being run for a number of other operational reasons. Lum, too was unable to perceive the signs he observed in June 2018 as evidence of a cyber attack. Having observed the use of multiple suspicious login-IDs to attempt access to the SCM database, he initially thought that there was some sort of audit or penetration testing that was being conducted on the SCM database. These IHiS staff were unable to ascertain from the evidence before them that they were observing a cyber attack in motion. The consequence was indeed dire – an unprecedented amount of data was exfiltrated from the SCM database by the attacker.

808. The capability of employees to detect, alert and respond to indicators of system compromise must therefore be enhanced – the only thing worse than having your network penetrated is having it penetrated and not knowing it. Early identification of a security incident is paramount at all levels and across the various divisions and groupings in IHiS – operational staff, IT security staff, senior management *etc*.

809. All staff must be trained to recognise suspicious activity that may point to a cyber breach. Suspicious activity can include a number of different observables such as abnormal access patterns, database activities, file changes, and other out-of-the-ordinary events that can indicate an attack. Being able to recognise these activities is important. Employees should be trained to recognise common examples of suspicious activity:

(a) Unusual database activity – Abnormal database activity can be caused by either internal or external attacks. Signs to watch for include changes in users, changes in permissions, bulk queries, and unusual data content growth.

(b) Account abuse – The abuse of privileged accounts is a common sign of an attack. Signs to watch for include modified audit trails, deleted logs, unauthorised access, and unnecessary accessing of sensitive information.

(c) Changes in account privileges – Unexplained changes in account privileges are a sign that an attacker is trying to gain access to the network using a user's credentials. Other signs include users accessing accounts at odd hours, accessing remotely, having multiple failed attempts to log in, and deviations from the usual pattern of usage between a user and a particular device.

(d) File changes – Changes in file configuration, including files being replaced, modified, added, and deleted, without explanation, are classic signs of a data breach, as it indicates that somebody has infiltrated the network.

(e) Suspicious network behaviour – Another sign of an attempted infiltration from external sources is unusual network behaviour. Employees must be able to identify traffic with odd origins or targets, unusual ports or protocols being accessed, unexplained changes in network performance, and unauthorised scans.

810. Further details on how all IT staff need to be involved in incident detection and response is also found in our recommendation on improving incident response processes to effectively respond to cyber attacks (Recommendation #6).

39 RECOMMENDATION #4: ENHANCED SECURITY CHECKS MUST BE PERFORMED, ESPECIALLY ON CII SYSTEMS

#PREVENTION #VIGILANCE

811. A pro-active strategy of discovering security vulnerabilities, misconfigurations, potential attack vectors, and even the presence of attackers lurking within the network, must be implemented, especially in relation to CII and mission-critical systems. Such a strategy should involve the use of five measures: (a) vulnerability assessments; (b) safety reviews, evaluation and certification of vendor products; (c) penetration testing; (d) red teaming; and (e) threat hunting.

39.1 Vulnerability assessment must be conducted regularly

812. According to the Cybersecurity Code of Practice [66] ("**CCoP**"), vulnerability assessment is the process of identifying, assessing and discovering security vulnerabilities in a system.

813. In turn, the CCoP defines:

(a) "architecture review" as "*a process of reviewing and analysing the design of the application and network architecture to identify critical assets, network design weaknesses, sensitive data stores and business* critical *interconnections for potential attack vectors and potential vulnerabilities in the network and application architectures*";

(b) "host security assessment" as "*a process of security assessment on a host to assess the host security configuration that cannot be seen*

[66] The CCoP was issued on 1 September 2018.

from the network, to identify *additional exposures and configuration weaknesses. It checks if the host's systems and applications are hardened effectively. Host, in this context, includes operating system, database server, firewall, router/switch, virtualisation implementation, load balancer, IDS, web proxy, web server, application server, mail server and wireless devices*"; and

(c) "network security assessment" as "*a process to identify and evaluate security weaknesses of the network and the network perimeter of a computer or computer system*".

814. Under the CCoP, the concept of a vulnerability assessment on a system is a broad one, requiring thorough review of the architecture, host security and network security of the system. Against this backdrop, we turn to discuss our recommendations.

39.1.1 Vulnerability assessments must be conducted regularly and following specified events on all CII, mission-critical, and/or internet-facing systems

815. We recommend that vulnerability assessments must be conducted on all CII, mission-critical, and/or internet-facing systems:

(a) prior to the commissioning of the system, or any new systems connected to the system;

(b) after any major changes have been implemented to the system, such as adding on application modules, system upgrades and technology refresh, as well as after any system migration; and

(c) in any event, at least annually.

816. This is in fact a *requirement* imposed on CII owners in respect of CII, under the CCoP. The CCoP also requires CII owners to, if requested by the Commissioner for Cybersecurity (the "**Commissioner"**), submit a copy of the report of any completed vulnerability assessments or penetration tests to the Commissioner within 30 working days of receiving the request.

817. In respect of mission-critical systems and internet-facing systems (assuming these are not also CII), such requirements for vulnerability assessments to be conducted are also important, and were in fact part of IHiS' policy under the HITSPS.

39.1.2 The scope of the vulnerability assessment should extend to all assets and systems connected to the CII, mission-critical and/or internet-facing system in question

818. In relation to the SCM system, which is both a CII and mission-critical system in the healthcare sector, vulnerability assessments were not conducted on the Citrix servers which are critical assets connected to the SCM database. Leong Seng testified that the Citrix servers were not considered part of the mission-critical SCM infrastructure and were not treated as "*the same level*" as the SCM infrastructure, although he acknowledged that all servers should be considered critical assets to be protected. On the other hand, Benedict considered that systems connected to internet-facing systems, although not directly internet-facing themselves (such as the Citrix servers), should be treated as internet-facing systems for the purposes of the vulnerability assessment and penetration testing requirements under the HITSPS.

819. We recommend that:

(a) First, the scope of vulnerability assessments to be conducted should extend to key assets and systems connected to the CII, mission-critical and/or internet-facing system in question. As seen in the Cyber Attack, the attacker exploited access to the SGH Citrix servers as a key part of its attack route to the SCM database. It is

thus important for key assets and systems connected to CII, mission-critical and/or internet-facing systems to also be subject to vulnerability assessment. Such a measure would also cohere with the CCoP's requirement for CII owners to ensure that the scope of each vulnerability assessment includes: (a) a host security assessment; (b) a network security assessment; and (c) an architecture security review.

(b) Second, there must be clarity within the organisation on what IT infrastructure would be considered connected to or part of CII, mission-critical and/or internet-facing systems, and therefore subject to vulnerability assessments. This could be achieved by way of drawing up an inventory of assets comprised in and connected to each system, such inventory to be regularly reviewed and communicated to the persons within the organisation responsible for conducting and overseeing the results of the vulnerability assessments.

39.1.3 Vulnerability assessments should also be conducted regularly on other critical assets which may not be part of or connected to CII, mission-critical or internet-facing systems

820. Leong Seng testified that IHiS' intention going forward was for all applications and servers (but not endpoints) to be subject to vulnerability scanning on a periodic and perpetual basis. In this regard, IHiS intends to use an Enterprise Vulnerability Management tool to perform regular vulnerability scans to detect and prioritise vulnerabilities found for remediation. We concur with this intended practice, as all servers are critical assets, as acknowledged by Leong Seng (see paragraph 818 (pg 281) above). We recommend that IHiS should carefully consider what would be considered critical assets, and perform vulnerability assessments on these assets at regular periodic intervals.

39.1.4 A process must be established to track that vulnerabilities identified in a vulnerability assessment are addressed

821. The CCoP requires CII owners to establish a process to track and address vulnerabilities identified in a vulnerability assessment and in a penetration test, and validate that all identified vulnerabilities have been adequately addressed. We further recommend that IHiS/SingHealth state clearly as part of the process:

(a) who (organisation/department/team) will have ownership for the respective tasks of drawing up action plans to address the vulnerabilities; reviewing and/or approving the action plans; implementing the action plans; tracking the progress of the action plans; validating that the vulnerabilities have been addressed; reporting on the progress/status of the action plans; and overseeing the process; and

(b) where feasible, what timeframes would be applicable for the respective tasks.

39.2 Safety reviews, evaluation and certification of vendor products must be carried out where feasible

822. One of the factors that CSA assessed had contributed to the Cyber Attack was that there were signs of insecure coding practices, and it was likely that the attacker had exploited this vulnerability to retrieve the credentials of the A.A. account. This incident underscores the importance of ensuring the security of vendor applications and systems which are used by an organisation, particularly where they relate to CII.

823. Indeed, the CCoP provides that the CII owner shall establish processes for validating vendors' compliance with cybersecurity requirements in terms of contract (for example, third party review) and product validation.

39.2.1 Code reviews and safety reviews

824. As part of CSA's technical recommendations, CSA also recommended that organisations should conduct code review of applications that are installed on critical systems and ensure that such reviews have been performed to their satisfaction. This is to verify that there are no instances of insecure programming or security flaws that may present vulnerabilities or backdoors that could be exploited by cyber attackers. As to how an organisation could go about procuring the conduct of such code reviews, Dan Yock Hau ("**Dan**") elaborated in oral testimony on a few options:

(a) As a customer purchasing critical software, the organisation could try to exercise its customer's rights to see what access it could get to the source code to conduct its own review.

(b) The organisation could also consider leveraging government reviews/certifications. Dan cited the example of how some companies, *e.g.* Microsoft, had set up transparency centres in certain countries and allowed governments to, as the proxy/agent at the national level, go through the source code, verify the source code and certify it at the national level so that others could use it.

(c) Alternatively, the organisation/customer could list down the standards and criteria by which it would have conducted a source code review, and ask the vendor to conduct an internal review based on those criteria and give a declaration that those criteria were met.

825. Dan explained that it was better to verify the application before buying it, and that such verification was an important function to be done during the tender, and before the organisation signed the contract with the vendor; rather than to test the application after purchase.

826. Where, however, it is simply not feasible to inspect the source code or have it reviewed (*e.g.* due to availability issues with off-the-shelf proprietary software, or where the code size is simply too large), Dan and Vivek recommended performing a penetration test on the application as part of the safety review to make sure there were no vulnerabilities in the application.

827. Flowing from the foregoing, we recommend that:

(a) Prior to the purchase of critical applications/systems, in particular those used for storing, processing or accessing sensitive information, an organisation should:

(i) ask to conduct a source code/safety review to ensure security of the application/system, if possible and feasible;

(ii) alternatively, ask the vendor to conduct a review based on the organisation's security requirements and standards;

(iii) further and/or alternatively, have a third party conduct an independent evaluation and certification for security;

(iv) in any event, build into the contract the security requirements and standards which the organisation expects the application/system to meet; and

(v) in any event, periodically conduct penetration testing on critical applications/systems (which will be elaborated on in section 39.3 (pg 288) below).

(b) In respect of legacy critical applications/systems which have already been purchased prior to any of the above steps being taken, it is all the more important that penetration testing be conducted on such applications/systems. Dan emphasised that there was urgency to conduct reviews and penetration testing of older, legacy systems,

for which the organisation did not have sight of the source codes and performance criteria.

(c) As suggested by Dan, there should be consistent safety reviews of applications and systems throughout their life cycle and use, with penetration testing built in as part of the safety review. This would enable necessary mitigation measures to be taken once vulnerabilities are found.

39.2.2 Evaluation and certification

828. In addition, we recommend that when it comes to the acquisition of:

(a) new software or products for CII systems; and

(b) critical applications and systems used for storing, processing or accessing sensitive information such as patient data,

829. CII owners must require the vendor/developer to obtain security certification for the products/systems in accordance with international, national or industry-recognised standards such as ISO/IEC 15408[67], FIPS 140-2[68], IEC 62443[69] *etc*. This could be done by way of, for example, IHiS including in their tender specifications that the health informatics applications/systems which protect patient data are certified in accordance with ISO/IEC 15408.

[67] ISO/IEC 15408-1:2009 establishes the general concepts and principles of IT security evaluation.
[68] FIPS (Federal Information Processing Standard) PUB 140-2 is the benchmark for validating the effectiveness of cryptographic hardware. If a product has a FIPS PUB 140-2 certificate, it has been tested and formally validated by the U.S. and Canadian Governments.
[69] ISO/IEC 62443 specifies the process requirements for the secure development of products used in industrial automation and control systems.

830. ISO/IEC 15408 is the international standard for evaluating and certifying products/applications, which implements security functions ranging from authentication to access control to encryption. Briefly, the standard requires a product/application to be subject to, *inter alia*, examination of the design of the security functions, functional testing, vulnerability assessment and penetration testing [70] Using the knowledge gained about the product/application, the evaluators would find whatever creative ways they can to compromise it. If a product/application has been certified ISO/IEC 15408, it would provide assurance that the product/application has undergone such rigorous security testing and evaluation.

831. In similar vein, Richard Staynings ("**Richard**") proposed in his report implementing stronger third-party vendor risk management requirements for applications and other systems that have access to electronic medical records, personally identifiable information, or other confidential/non-public information.

832. In this connection, Richard commented that in its marketing and documentation materials, Allscripts claimed that it was ISO 27001 and SOC 2 certified. We note that ISO 27001 does not relate to whether products have been securely developed and tested, but rather, whether an organisation has policies and controls in place to safeguard information. ISO 27001 would thus not be relevant as a security standard for the SCM application product.

833. Further, according to Richard, the materials did not stipulate the frequency with which both assessments were updated, the control objectives of the SOC 2 attestation or whether a SOC 2 Type I report (where the controls are described and evaluated at a point in time to determine if they are functioning as they are

[70] The penetration testing referred to here is specifically for a product/application; *c.f.* for example, the network penetration testing that the GIA had conducted in FY16, which involved network sniffing, running scripts to harvest credentials, lateral movement to compromise domain controllers, but not penetration testing of the SCM application and SCM database (which are not in the typical scope of a network penetration test). Penetration testing of a product/application and the product evaluation are complementary.

described by management in the design) or Type II report (where the assessment is made over a period of time, and is thus much more detailed and valuable in understanding the actual security controls of the application) was produced. Richard suggested that SingHealth should evaluate the attested or certified security controls provided by its application vendors to ensure that control objectives aligned with SingHealth's internal security objectives, and should supplement vendor assessment with its own administrative and technical penetration testing of the systems. This would provide improved alignment of risk analysis objectives and may help to identify previously unknown weaknesses or vulnerabilities in applications or services, such as the coding vulnerability in the SCM application.

834. The above recommendations for security evaluation and certification of vendor products would serve as new safeguards in an ever-changing security threat landscape faced with growing risks.

39.3 Penetration testing must be conducted regularly

835. "Penetration testing" is defined in the CCoP as:

> "an authorised process of evaluating the security of a computer system, network or application by finding vulnerabilities attackers could exploit and includes the process of:
> (a) gathering information about the target;
> (b) identifying possible entry points;
> (c) attempting to break in (either virtually or for real); and
> (d) reporting the findings".

836. The purpose of performing a penetration test is to verify that new and existing applications, systems and networks are not vulnerable to a security risk that could allow unauthorised access to resources.

837. CSA's Douglas Mun explained how the coding vulnerability in the SCM application was revealed during a penetration test conducted by CSA on the SCM

application in July 2018. The H-Cloud network penetration tests conducted by the GIA in FY16 revealed vulnerabilities and misconfigurations, several of which were present during the Cyber Attack, and which overlapped with CSA's investigation findings as to the vulnerabilities and contributing factors exploited by the attacker in the Cyber Attack. These are just two tangible demonstrations of the value in conducting penetration tests on critical applications, systems and networks.

838. Indeed, the CIS recommends that "*[i]n a complex environment where technology is* constantly *evolving, and new attacker tradecraft appears regularly, organizations should periodically test their defenses to identify gaps and to assess their readiness by conducting penetration testing*".[71]

839. We will elaborate on the following recommendations for the conduct of penetration tests.

39.3.1 Penetration tests must be conducted regularly and following specified events on all CII, mission-critical and/or internet-facing systems

840. We recommend that penetration tests must be conducted on all CII, mission-critical and/or internet-facing systems:

(a) prior to the commissioning of the system, or any new systems connected to the system;

(b) after any major changes have been implemented to the system, such as adding on application modules, system upgrades and technology refresh, as well as after any system migration; and

(c) in any event, at least annually.

[71] CIS Controls Version 7 at control 20: Penetration Tests and Red Team Exercises.

841. This is in fact a *requirement* imposed on CII owners in respect of CII, under the CCoP. The CCoP also requires CII owners to, if requested by the Commissioner, submit the penetration test results to the Commissioner.

842. Regrettably, IHiS did not conduct penetration tests on the SCM application or system (although it was a mission-critical and CII system) prior to the Cyber Attack. IHiS' policy under the HITSPS was for penetration tests to be conducted on internet-facing systems, which IHiS staff interpreted as not applying to the SCM application or database as they were not "internet-facing". Under the CCoP, it is clear that penetration tests must be conducted on CII, and moving forward, IHiS should not exclude the SCM application or database (or connected systems and networks, as to which, see the recommendation at section 0 below) from penetration testing.

843. In respect of mission-critical systems and internet-facing systems (assuming these are not also CII), it would also be important for the above requirements for penetration tests to apply, and similar requirements in respect of internet-facing systems are in fact part of IHiS' policy under the HITSPS.

844. IHiS should review its written policy on penetration testing to ensure that all the requirements set out in these recommendations are comprehensively captured (*c.f.* HITSPS which only refers to penetration testing of internet-facing systems). In formulating the policy, regard can be had to the Association of Banks in Singapore's "Penetration Testing Guidelines For the Financial Industry in Singapore".[72]

[72] The Association of Banks in Singapore, "Penetration Testing Guidelines for the Financial Industry in Singapore", July 2015.

39.3.2 The scope of the penetration tests should extend to key assets and systems connected to the CII, mission-critical and/or internet-facing system in question

845. Similar to the recommendation at section 817 (pg 281) above, we recommend that:

(a) The scope of the penetration tests should be extended to key assets and systems connected to the CII, mission-critical and/or internet-facing system in question. In other words, all essential components of a system (such as in the case of SCM, the application, database and middleware such as the Citrix servers) should be included in the scope of the penetration test. This would cohere with the CCoP, which provides that CII owners shall ensure that the scope of a penetration test includes penetration tests of the CII's hosts, networks and applications.

(b) There should be clarity and clear communication within the organisation on the IT infrastructure which are to be subject to penetration tests as part of the penetration tests conducted on CII, mission-critical and/or internet-facing systems.

39.3.3 Penetration tests should also be conducted regularly on applications, systems and networks which may not be part of or connected to CII, mission-critical or internet-facing systems

846. Dan recommended that organisations should conduct regular and vigorous penetration tests to ensure that vulnerabilities within their systems and networks are discovered and fixed, especially for mission-critical systems. This indicates that, more generally, penetration tests should be conducted periodically even for non-mission-critical applications, systems and networks, and we would recommend this. As mentioned in paragraph 827 (pg 285) above, penetration testing should also be built in as part of safety reviews conducted on systems, especially older, legacy systems.

39.3.4 Penetration tests should be conducted outside of the regular schedule if a need to do so is indicated

847. In line with taking a pro-active strategy towards testing defences and detecting vulnerabilities, IHiS should also consider conducting penetration testing outside of any regular schedule if and when a need is indicated, *e.g.* when prompted by threat intelligence.

39.3.5 Penetration tests should be conducted by persons with the appropriate levels of expertise

848. Dan testified that there are various "levels" of penetration testing that can be done, referring to penetration testing expertise ranging from in-house (*e.g.* in IHiS' case, by GIA or CSG) to independent accredited commercial penetration testers to CSA. He explained that:

(a) There is nothing wrong with relying on in-house penetration testing as the "first-cut". However, there may be residual risks as this means there is no external view of system vulnerabilities. These risks may be mitigated by engaging third-parties to conduct penetration testing. In respect of penetration tests which the CCoP requires CII owners to conduct on their CII, CSA's requirement (under the CCoP) is that if such penetration testing is done by third-party penetration testing service providers, the service providers and their penetration testers must have the requisite industry-recognised accreditation and certification.

(b) CSA is building up advanced penetration testing teams, but the resources will be limited and CII owners cannot all rely on CSA to do penetration testing for them.

849. In this regard, we therefore recommend that:

(a) For CII systems, IHiS must engage independent third-party penetration testing service providers to conduct the penetration testing. These external penetration testers must fulfil CSA's accreditation and certification requirements under the CCoP. For non-CII systems which are nevertheless critical, IHiS should also consider periodic penetration testing by accredited independent third-party service providers.

(b) In addition, there should be a strong in-house penetration testing capability, which would include having the in-house penetration testers regularly trained, accredited and certified.

(c) Such in-house penetration testing function should be independent of IHiS in nature, and could be parked in GIA or another department reporting directly to MOH. For clarity, there should not be any double-hatting in this process, and the person responsible for this function should not be an IHiS employee.

(d) There should also be clarity on what in-house penetration tests are being conducted and by whom and when, to avoid overlap, or the inadvertent omission of applications/systems/networks for testing. In this connection, the penetration testing department pursuant to (c) above could consider drawing up schedules to track regular penetration testing and coverage of all relevant applications/systems/networks.

(e) IHiS should consider whether there is any serious need (*e.g.* prompted by any particular threat intelligence, or alerts from its monitoring and detection systems) for any particular application/system/network to be subject to an advanced penetration test by CSA, and if so, to engage CSA.

39.3.6 A process must be established to track that vulnerabilities uncovered by a penetration test are addressed

850. IHiS and SingHealth (as the CII owner) must own the remediation process. The Committee recommends that there needs to be a process to address and track vulnerabilities uncovered in a penetration test, and to validate that all uncovered vulnerabilities have been adequately addressed. This process mirrors the requirement for CII under the CCoP, and as set out in paragraph 821 (at pg 283) above. We note in this regard that IHiS has, since April 2018, set up a centralised audit liaison team to track all audit issues and remediation actions (across Clusters), and IHiS could build on this in formulating its processes for tracking and addressing other vulnerabilities that are discovered *via* other security checks such as vulnerability assessments and penetration tests.

39.3.7 A more comprehensive penetration test of the SCM application should be conducted

851. Given that (a) the SCM application is used for SingHealth's mission-critical EMR system, (b) the protection of SingHealth network's 'crown jewels', *i.e.* the patient database, is critically dependent on how secure or not the SCM application is, and (c) the basic insecure coding vulnerability already shown to be inherent in the SCM application, the penetration testing department referred to in paragraphs 0(b) and (c) above should consider conducting a more comprehensive and advanced penetration test of the SCM application to see if any other vulnerabilities will be detected.

39.4 Red teaming should be carried out periodically

852. As explained by Dan, red teaming is a more advanced measure that goes beyond penetration testing. Red teaming is conducted by an independent external group that assumes an adversarial role and can simulate an APT attack on an organisation, and includes vulnerability assessment, penetration testing, bug hunting and more. By providing an end-to-end and full-scope attack cycle, red

teaming can be more effective in validating the people, processes and technology of an organisation.

853. According to the CIS, "*Red Team exercises take a comprehensive approach at the full spectrum of organization policies, processes, and defenses in order to improve organizational readiness, improve training for defensive practitioners, and inspect current performance levels. Independent Red Teams can provide valuable and objective insights about the existence of vulnerabilities and the efficacy of defenses and mitigating controls already in place and even of those planned for future implementation*".[73]

854. In this Cyber Attack, CSA found that the attacker was a skilled and sophisticated APT actor who employed advanced network intrusion techniques and customised malware to evade security measures. Given that APT attacks are likely to become more prevalent, Dan recommended that organisations may consider red teaming to fully appreciate the vulnerabilities present in their networks. This is because there are limitations to penetration testing. Penetration test teams have a limited amount of time with a system, and would look for the easiest or most time-effective way to gain access to the system. In contrast, APT attackers could wait patiently for months or years in a network. The vulnerabilities identified in a penetration test may thus not be comprehensive or indicative of all the vulnerabilities present in a network that could be exploited.

855. Kim Chuan testified that the Clusters' internet-facing systems are subject to internal "ethical hacking", which can be considered to be red teaming activities. He also explained that efforts were being taken to conduct similar activities on Clusters' internal systems.

856. Dan explained that where such a function is done internally, the term is either "blue team" or "white team". The key is that red teaming is to be done by someone who is not involved in the daily operations, as well as not involved in

[73] CIS Controls Version 7 at control 20: Penetration Tests and Red Team Exercises.

the design of the system, so that there is a fresh perspective on how an attacker would come in. Dan suggested that the red team could augment the results of the blue or white team, for a more robust solution.

857. In light of the foregoing, we recommend that:

(a) CSG should continue to build up its "blue" or "white" teaming capabilities, and regularly carry out blue/white teaming on key systems and networks.

(b) Beyond that, IHiS should engage independent third-party service providers to periodically conduct red teaming exercises on key systems and networks.

(c) A clear policy on the conduct of red teaming exercises should be set out in the HITSPS.

39.5 Threat hunting must be considered

858. Many of the experts have testified that it is a matter of when, not if, the security of a system will be breached. An "assume breach" mentality places the focus not merely on prevention, but critically, on detection as well.

859. It is thus apposite that Dan recommended building up threat hunting capability. Threat hunting entails proactively searching through networks to hunt for and detect advanced cyber threats that evade existing security safeguards, before they manifest into major security incidents. This is recommended for high value systems on a regular basis, based on the risk management analysis of the organisation, to ensure that the systems are clean and uncompromised.

860. Dan also explained that there were not many mature offerings of such threat hunting services at this time, but over time, commercial companies would probably build up their own competency and offer such services. For now, CSA could fill that gap. Once threat hunting services become more commercially

available, Dan's view was that it should be done regularly, especially for critical systems.

861. Relatedly, on the issue of timing for organisations to start deploying threat hunting, Vivek cautioned that an organisation's security set-up would have to reach a certain level of maturity, which could take several years, before threat hunting would be of meaningful benefit. In his opinion, the focus areas should first be on developing threat intelligence and the Security Operations Centre, with threat hunting and cyber range exercises to follow several years down the line.

862. At the same time, however, it is encouraging that Bruce, Leong Seng and IHiS have identified threat hunting as an area that IHiS will move into. Leong Seng testified that IHiS was looking to set up an Advanced Security Operation Centre ("**ASOC**") which would provide proactive services such as active threat hunting.

863. In light of the foregoing, we recommend that IHiS together with its ASOC evaluate the value of conducting threat hunting in the public healthcare institutions' systems presently, and as soon as practicable, move to ensure that threat hunting is regularly carried out on high value systems, including CII like the SCM system, by an independent third-party service provider with the expertise to do so.

40 RECOMMENDATION #5: PRIVILEGED ADMINISTRATOR ACCOUNTS MUST BE SUBJECT TO TIGHTER CONTROL AND GREATER MONITORING

#PREVENTION #VIGILANCE

864. Privileged accounts on a network are prime targets for malicious exploitation. According to the CIS[74]:

> "The misuse of administrative privileges is a primary method for attackers to spread inside a target enterprise…[a] common technique used by attackers is elevation of privileges by guessing or cracking a password for an administrative user to gain access to a target machine. If administrative privileges are loosely and widely distributed, or identical to passwords used on less critical systems, the attacker has a much easier time gaining full control of systems, because there are many more accounts that can act as avenues for the attacker to compromise administrative privileges."

865. The abuse of privileged access is therefore at the core of many cyber attacks because privileged accounts have more authority and access to resources, which simplifies the achievement of an intruder's goals. Windows domain administrator credentials potentially allow an attacker to gain access to all servers in a domain, while server local administrator accounts provide unrestricted access to individual servers.

866. Compromised privileged credentials have been revealed as a primary attack vector in the Cyber Attack. Privileged credentials were used by the attacker to move about in the network, after the initial intrusion, in its hunt for valuable assets.

---

[74] CIS Controls Version 7, at CIS Control 4.

867. Examples of the compromise and abuse of privileged accounts include the use of local administrator accounts the S.A. account and the L.A. account to log in to Citrix Servers 1 and 2. Furthermore, the D.A. account was compromised, since it was observed to have been used in an attempt to log in to the SCM database, when it was not being used by its authorised user.

868. IHiS was aware that their systems were vulnerable to the risk of privileged passwords being compromised. The FY16 GIA Audit Report had, in fact, highlighted the vulnerability created by weak control of privileged accounts in the SingHealth network. The report stated that the penetration testers had successfully exploited the vulnerability and obtained full domain administrator control of the servers in the SingHealth network domain. In the FY16 GIA Audit Report, GIA had highlighted the dire consequences when this vulnerability is exploited (see also paragraph 1072 (pg 368) below).

869. GIA had highlighted that the weak control of privileged accounts stemmed from bad password compliance policies – passwords being used were very simple; the non-complex passwords used could be easily guessed or cracked with readily available password cracking tools. Recommendations were made to IHiS for remediation, but unfortunately, these were inadequately complied with.

870. The following are a series of measures to mitigate the risk of privileged account abuse.

40.1 Inventory of administrative accounts should be created to facilitate rationalisation of such accounts

871. The CIS Controls require that an inventory of administrative accounts be maintained, including domain and local accounts, to ensure that only authorised individuals have elevated privileges.

872. Over time, "privilege creep" may have occurred where too many users, and too many accounts have undocumented privileges. The L.A. account is an example of a dormant local administrator account. Although not used for day-

to-day operations, it remained in the system with full administrator privileges, and was eventually exploited in the Cyber Attack. The S.A. account is yet another example – it was an inactive service account, that had full administrator privileges although there was no real reason for its existence. This too was exploited in the Cyber Attack. Although the SMD was responsible for the periodic review of user-IDs to identify and disable dormant accounts, this was not done.

873. Policies in relation to the management of accounts are laid out in the HITSPS. HITSPS policy requires that user-IDs in the IT system be reviewed periodically to identify unused or dormant accounts. Unused user-IDs should be disabled to prevent them from being used for unauthorised activities. This was not done, as evidenced by the eventual abuse of the L.A. and S.A. accounts, dormant and unused accounts, respectively, which had not been identified.

874. It is recommended that the number of IT staff who have administrator privileges, and the number and nature of privileged accounts on the network should be reviewed as there may be scope for rationalisation to adhere to the principle of least privilege,[75] maintain system integrity and reduce the attack surface for privileged accounts to be compromised.

40.2 All administrators must use two-factor authentication when performing administrative tasks

875. The risk of active directory ("**AD**") administrator accounts being compromised must be mitigated. Windows server administrators need to use domain administrator accounts to perform standard administrative tasks but, ideally, domain administrator accounts should only be used when privilege is required. Administrators should not be granted domain administrator privileges for their regular AD accounts, which they use for carrying out day-to-day tasks,

[75] The principle of least privilege is the idea that at any user, program, or process should have only the bare minimum privileges necessary to perform its function. Following the principle of least privilege is considered a best practice in information security.

such as accessing email, and they should only be used by a single administrator (*i.e.* not shared) for accountability.

876. AD accounts are susceptible to being compromised by an attacker who has already gained a foothold within the network. Further, AD administrator accounts are susceptible because their passwords are not frequently changed.

877. The attacker gained unauthorised access to numerous SingHealth servers by stealing the passwords for privileged accounts. These servers used single-factor authentication, in the form of a password. Relying solely on the strength of passwords is insufficient to protect critical servers against the risk of compromise.

878. Given these vulnerabilities, a system of Privileged Access Management[76] ("**PAM**") using 2FA must be put in place, and enforced for administrator access to servers. Experts Dr Lim, Gen. Alexander, Vivek and Richard all concur with this recommendation. With 2FA, users must input two distinct identification methods — such as a password and a one-time-use PIN — to verify their permission to access a restricted system. A second factor of authentication would significantly secure access to privileged accounts, and the risk of unauthorised access to mission-critical servers would be reduced. An attacker who obtains compromised credentials would *not* be able to access a server, as it would not be in a position to provide the second form of identification, to complete the authentication process.

879. The Committee notes that the use of multi-factor authentication for all administrative account access is recommended in the CIS Controls.[77] We also

[76] PAM is a solution that helps organisations restrict privileged access within an existing Active Directory environment.
[77] CIS Controls Version 7 at sub-control 4.5.

highlight NIST's recommendations in SP 800-63-3 Revision 3[78] (table 6-1) and SP 800-63B[79] (table 4-1) requiring multi-factor authentication as a minimum, for systems and online services that process personally identifiable, sensitive or classified information (*i.e.* Assurance Level 2 or 3).

880. Any implementation of PAM by IHiS must be accompanied with strict controls ensuring that the PAM-based access is the *exclusive* means by which administrators access servers. If not, administrators are likely to use less secure means to access restricted systems, to avoid the perceived tedium of using PAM. This would nullify the effectiveness of 2FA. For instance, PAM had in fact been implemented by IHiS for servers in H-Cloud, and thereafter for servers in the SGH Local Data Centre ("**LDC**"). However, even after PAM had been implemented, administrators were not limited to accessing servers in the SGH LDC and H-Cloud only by using PAM. Administrators preferred to use an alternative method to access the servers, which did not require 2FA, because they found usage of PAM tedious – IHiS administrators found that the PAM sessions timed out quite quickly resulting in their having to re-enter credentials and 2FA to reconnect to the servers, while carrying out their administrative tasks.

881. The Committee stresses that the implementation of a technical solution is not enough. The use of security-related technical solutions must be *enforced*, and less secure authentication methods must be closed-off. As noted by Vivek, if all other means of access are not closed off when 2FA is introduced, the whole purpose of PAM would be defeated, as it could easily be circumvented by administrators, for a variety of reasons.

882. The Committee recognises that there are certain circumstances in which exceptions may be granted to certain administrators. However, as stressed by Vivek, where these exceptions are granted, they must be carefully monitored.

---

[78] NIST.SP.800-63-3.

[79] NIST.SP.800-63B.

Attackers will target persons who are granted these exceptions and use one of their credentials to break into restricted systems.

40.3 Use of passphrases instead of passwords should be considered to reduce risk of accounts being compromised

883. Passwords have long been the preferred method of user authentication, but poor password practices cause security issues. Attackers have developed sophisticated and effective methods to "brute force"[80] passwords. This means passwords can be compromised if they are weak or easy to guess.

884. Passwords appear to be significant weaknesses in IHiS' cyber defences. The evidence shows that employees used passwords that met the most basic requirements of the password policy, but were not strong enough to resist compromise. IT administrators used simple passwords that were too easily decipherable. That "P@ssw0rd" was a commonly used password for privileged accounts, is deeply concerning. It is notable that weak passwords appear to have been a perennial problem for IHiS' cyber defences – they were identified as vulnerabilities after penetration testing by the GIA at three local sites in FY17. They were in fact also identified as persisting vulnerabilities, as they were not only identified in FY17, but had been previously highlighted for remediation in the FY16 GIA Audit Report.

885. Reliance on passwords, and the ease with which attackers can defeat those passwords, has resulted in a negative feedback loop where users have been subjected to increasingly complex composition rules (upper case, lower case, numerals and special characters), increasing length requirements, and password expiry requirements.

---

[80] A brute force attack consists of an attacker trying many passwords with the hope of eventually guessing correctly. The attacker systematically checks all possible passwords until the correct one is found

886. It is recommended that IHiS adopt a better approach by moving from the use of passwords to passphrases.[81] Passwords, even those with complex combinations of letters, numbers and symbols, no longer afford sufficient protection. Passwords that were once considered almost unbreakable can now be cracked in a matter of hours or days. Passphrases are longer but need not necessarily contain numbers or symbols, which makes them easy to remember, eliminating the need for them be written down or stored. By using passphrases, brute force attacks can be rendered impractical.

887. In June 2017, NIST released new standards for password security entitled "*Authentication & Lifecycle Management*".[82] In these guidelines NIST recommends using long passphrases instead of seemingly complex passwords. NIST observed that the "*memory burden*" on users could be lightened, and recommended encouraging users to create unique passphrases they could more easily remember. The switch to passphrases has also been recommended by a number of other reputable institutions.[83]

888. The Committee notes that the Singapore public sector's IT policy has very recently encouraged the use of passphrases instead of complex passwords. The policy now requires the use of longer passwords, with fewer complexity requirements; implicitly encouraging users to switch to the use of passphrases.

889. It is also pertinent to note that the NIST guidelines also recommend that[84]:

(a) When processing requests to establish and change memorized secrets, verifiers shall compare the prospective secrets against a list that contains values known to be commonly-used, expected, or

---

[81] A secure passphrase can be as simple as a short sentence with proper punctuation, *e.g.* "IAmUsingAPassphraseOnThisComputer".
[82] NIST.SP.800-63B.
[83] Australian Cyber Security Centre, "Passphrase Requirements", November 2017; SANS Institute, "OUCH! Newsletter: Passphrases", April 2017.
[84] NIST.SP.800-63B at 5.1.1.2 Memorized Secret Verifiers, p14

compromised. For example, the list may include, but is not limited to:

- Passwords obtained from previous breach corpuses.
- Dictionary words.
- Repetitive or sequential characters (*e.g.* 'aaaaaa', '1234abcd').
- Context-specific words, such as the name of the service, the username, and derivatives thereof.

(b) If the chosen secret is found in the list, the Credential Service Provider ("**CSP**") or verifier shall advise the subscriber that they need to select a different secret, shall provide the reason for rejection, and shall require the subscriber to choose a different value.

890. The move to passphrases can be achieved by:

(a) Educating employees on the benefits of moving to passphrases as part of their security training, emphasising both the personal and organisational benefits of improved network security; and/or

(b) Deploying third party tools on domain controllers that enforce the use of passphrases, blacklist guessable complex passwords, blacklist leaked passwords *etc*.[85]

[85] *Specops Password Policy*, a tool which can target any GPO level, group, user, or computer with dictionary and passphrase settings is an example of one such tool – see www.specopssoft.com for details.

40.4 Password policies must be implemented and enforced across both domain and local accounts

891. Typically, when restrictions are implemented on the administrators group in GPOs, Windows applies the settings to members of a computer's local administrators group, in addition to the domain's administrators group. However, in the case of the SGH Citrix servers, a setting called "block policy inheritance" had been applied at the servers, meaning that domain level policies could not be 'inherited' (*i.e.* they were blocked) and would not apply to the SGH Citrix servers. Accordingly, although password policies had been implemented at a domain level, they were not applied on these servers.

892. We recommend that a technological solution be found to ensure that updated password policies will be pushed down for server local administrator accounts, *without exception*. If no such solution can be found, steps must be taken to individually implement the updated policies at the local servers, or an alternate solution must be implemented to centrally manage server local administrator accounts.

40.5 Server local administrator accounts must be centrally managed across the IT network

893. A server local administrator account has access to every file and application on the server. If an attacker can get a foothold in a system, it often looks for this privileged local administrator account as part of its attack roadmap. It will then use these accounts as it starts moving laterally across the network.

894. In short, that attacker guesses or acquires the local administrator's account password, grabs the hashes of domain-level users with password dumping tools, and then moves around the network.

40.5.1 Establish clear policies in relation to the use and management of server local administrator accounts

895. Server local administrator accounts are a security problem because one set of login credentials is typically used by many IT administrators. This can make it difficult or even impossible to implement an identity-based access management policy because the specific person gaining access to a server cannot be tracked at any given time.

896. The password for the L.A. account was compromised, with the same account and the same password being used across all Citrix servers. Such local privileged accounts must not be configured with the same credentials across systems. The use of the *same local admin password on every server* helped the attacker to move laterally within the network. One server 'taken-over' meant that all of them were 'owned' by the attacker. Since the local administrator account can control everything that can be performed on a server, if the single password is compromised on any server, all systems are susceptible to compromise.

897. We note that HITSPS makes no express reference to account management or password policies specific to the management of local administrator accounts (*e.g.* there is no policy that the same password cannot be used to local administrator accounts across multiple servers).

898. Specific policies addressing server local administrator passwords must be formulated, with the necessary tools put in place to enforce and ensure compliance with these policies. Examples of such policies include[86]:

(a) Change Default Usernames and Passwords - change all default usernames and passwords for local admin accounts;

[86] These policies are drawn from the CIS Controls Version 7.

(b) Use Unique Passwords – local admin accounts must use passwords that are unique to that system;

(c) Disable Dormant Accounts - Automatically disable dormant accounts after a set period of inactivity;

(d) Log and Alert on Unsuccessful Administrative Account Login - Configure systems to issue a log entry and alert on unsuccessful logins to a local administrator account;

(e) Monitor Attempts to Access Deactivated Accounts - Monitor attempts to access deactivated accounts through audit logging; and

(f) Alert on Account Login Behavior Deviation - Alert when users deviate from normal login behavior, such as time-of-day, workstation location and duration.

40.5.2 Access to server local administrator accounts should be made available on a needs-only basis

899. Local accounts are harder to manage than domain accounts. Changing a domain account is easily done in one place, affecting all computers where the account is used. A local account is modified on the workstation where it exists. Using Group Policy Preferences, *some* local account settings (*e.g.* password length) can be managed centrally with Group Policy, but as highlighted above, this may not be a fool proof approach.

900. As the local administrator password can cause a major security issue in any network, a best practice to follow would be to have unique and random passwords per server *and* distributed in a secure manner that still allows IT staff to know and use the passwords. This will prevent an attacker abusing such credentials by reusing the same credentials across the network. However, regular monitoring would still be required, to make sure local passwords are not reset or defaulted to a weak password.

901. As a more comprehensive solution, it is recommended that a solution such as an enterprise Password Vault should be implemented. Broadly speaking, this would prevent unauthorised users from accessing privileged account credentials, and still ensure that authorised users have the necessary access for legitimate purposes. A Password Vault serves to protect all privileged account passwords in a secure central repository to prevent the theft or unauthorised sharing of these credentials. Administrators will "check out" server local administrator credentials each time access using such an account is required. Further, such a system would ensure that the credentials "checked out" would meet password length and complexity requirements, be constantly changed, and be unique to each server.

902. Implementing such a solution would significantly reduce the risk of weak passwords leading to the compromise of local administrator accounts, and would slow down lateral movement in a network if a breach happens, as it would require each server to be compromised separately.

903. IHiS has in fact implemented a Password Vault solution in the wake of the Cyber Attack. As testified to by Woon Lan and Leong Seng, IHiS has procured a software to manage all local administrator accounts. This ensures that IHiS is no longer reliant on the administrators to change the passwords themselves – the Password Vault mandates that it is constantly changed.

40.6 Service accounts with high privileges must be managed and controlled

904. A service account is a "non-human" account that is used to run services or applications. A service account is not an administrative account, nor is it a "human" user account, used by administrators or other employees. These accounts are the target of many malicious actors because they are often implemented in such a way that they have privileged access.

905. The S.A. account is a key example of this in the Cyber Attack. The S.A. account has full administrative privileges to login to the Citrix server, including

logging in interactively, and logging in remotely *via* RDP. In the context of the attack, the attacker used this account to log in to Citrix Server 2 on multiple occasions in June 2018.

906. As detailed in the HITSPS, this account, being an unused account, should have been identified and disabled in order to prevent usage in unauthorised activity. Moving forward, there must be a recognition that such accounts with such high privileges need to be managed and controlled.

40.6.1 Establish clear policies in relation to the use and management of service accounts

907. The compromise and use of the S.A. account in the Cyber Attack clearly illustrates the real risk that presents when service accounts with high privileges are not properly managed and controlled. We note however that HITSPS is silent on the specific policies and measures in relation to the management of service accounts.

908. Because service accounts are not tied directly to a human, they must be treated differently from other accounts. A specific policy should be formulated in respect of service accounts. Examples of such policies include:

(a) Longer Password Length – A policy requiring very long and complex passwords for service accounts is appropriate, as there is no 'memory burden' on the part of a human user to remember such passwords.

(b) Longer Password Expiration – It is hard to set password expiration policies that are short because resetting a service account password may break an application. However, a policy requiring the password to be changed, albeit at a longer interval, should still be imposed. This is necessary as, in the event a password is compromised by an attacker, he would otherwise have perpetual access to the service account.

(c) Privilege Management – It is a best practice to implement the principle of least privilege. Only provide the minimum necessary privileges to service accounts. For example:

- Restrict Interactive Logins – An interactive login is a process whereby the user gains access to the network by entering a username and password in response to a dialog box on the local console. Service accounts are not used by human users, and interactive login privileges are unnecessary.

- Restrict Remote access – Service accounts are typically used to log in to services or applications on the host server itself, and privileges allowing the service account to remotely login to a server, from elsewhere in the network, are unnecessary.

(d) Disable Dormant or Inactive Accounts – Automatically disable dormant or inactive service accounts after a set period of inactivity.

(e) Log and Alert on Unsuccessful Service Account Login – Configure systems to issue a log entry and alert on unsuccessful logins to a service account.

(f) Frequent Privilege Review – Automated checks should be carried out at fixed intervals to ensure that the privilege levels of service accounts have not been inadvertently or intentionally elevated beyond that which was granted.

(g) Monitor Attempts to Access Deactivated Accounts – Monitor attempts to access deactivated service accounts through audit logging.

(h) Alert on Account Login Behavior Deviation - Alert when service accounts deviate from normal login behavior, such as time-of-day, workstation location and duration.

(i) Not hardcoding or including administrator credentials in cleartext in scripts on servers - In addition to having carried out a scan to identify all scripts containing administrator credentials in text files stored in shared folders on servers, we note that IHiS plans to continue to conduct such scans periodically and take disciplinary action on administrators who are found to not comply with security policies on the creation of such scripts.

40.6.2 Create and maintain an inventory of service accounts, and disable accounts which are unnecessary

909. Locking down service accounts must be a basic component of the hardening strategy for servers. An inventory of all existing service accounts must be created, and existing privileges should be reviewed with the view to granting the least privileges necessary. When new servers are provisioned, specific regard should be given to reviewing the service accounts that are created, and whether such service accounts (and the underlying service) are necessary. Unnecessary accounts should be disabled as part of basic account administration hygiene.

41 RECOMMENDATION #6: INCIDENT RESPONSE PROCESSES MUST BE IMPROVED FOR MORE EFFECTIVE RESPONSE TO CYBER ATTACKS

#DETECTION #RESPONSE #GOVERNANCE

910. An effective incident response plan can reduce the extent and impact of an attack by identifying its source and shutting it down quickly. In the event of a cyber attack, warnings may come at short notice and the pace at which an attack escalates may be rapid. The correlation between the effectiveness of an incident response plan and recovery is evident, with organisations recovering from attacks proportionally to their incident response preparedness.

911. While IHiS' existing security incident response framework and IR-SOP envisage proper and prompt incident detection, investigation and reporting, the evidence led reveals that many of the security incidents involved in the Cyber Attack went undetected, were poorly investigated or were unreported. Had early detection, proper investigation and timely reporting occurred, the unauthorised access to, and exfiltration of, patient data from the SCM database could likely have been prevented.

912. A proactive response is key to mitigating damage and facilitating recovery efforts. It is thus imperative that the incident response process is sharpened through the measures recommended in this section.

41.1 Incident response plans must be tested with regular frequency before a real incident occurs

913. To ensure that response plans are effective, they must be tested. Plans must not only be used in real-world incidents – they must be tested with regular frequency before a real incident occurs.

41.1.1 Testing of incident response plans is critical

914. Testing is critical because it provides an opportunity to reveal weaknesses and omissions that ought not to be discovered only after a breach already has occurred. Planning can only go so far, and while organisations can strive to create comprehensive incident response plans, failure to test such plans until a real event occurs, may result in the realisation (too late) that the plans fail at the first step because they are unworkable, or did not adequately consider real-world constraints or difficulties. The failure to frequently test an incident response plan could result in increased response time, confusion amongst the responders, and at its worst, a failure to even respond to a serious security incident.

915. Organisations, in particular IHiS, must ensure that training and building familiarity with incident response plans is ongoing. Training should be continuous and not limited to a one-time event. Continuous mechanisms must be in place for ensuring that reporting triggers and reporting procedures are known, understood, and complied with. This should be led by CEO, IHiS. At the same time, SingHealth and MOHH are to have oversight of this, as the system owner and the holding company respectively.

916. All relevant parties should be drilled on the response plan, with exercises and simulations carried out regularly. The creation of an incident response plan must not be viewed as a one-time exercise. It is an ongoing process, and refinements to the plan must be made when drills demonstrate the need for the plans to be modified. Ensuring that plans are reviewed and amended on an ongoing basis will allow incorrect information regarding tools and people to be updated, and for reviewing of response measures that do not work, or are out of order. This is consistent with Vivek's expert testimony that "*a plan that is a Word document that is filed somewhere, or a PDF that is filed somewhere does not help*" and that the plan should be kept current and effective by constantly updating it after every incident and after every TTX. For example, simulation exercises can prevent confusion by engaging with all the key stakeholders to set clear expectations, contributing to the completeness and clarity of post-breach actions and responsibilities. Gen. Alexander's evidence was explicit that,

"*[p]roper training and a solid exercise program would have ensured personnel knew and understood their roles and responsibilities in helping to prevent the Cyber Attack on SingHealth*." Gen. Alexander's conclusion was buttressed by other experts. Dr Lim testified that exercises are vital to training, help to strengthen SOPs, and ensure that staff do not just "*go through the motions*". Vivek also recommended that *everyone* participate in exercises that simulate real-life scenarios, so that they are able to respond in the event of an incident.

917. Evidence has been led that IHiS was involved in three TTXes from 2016 to 2018. However, a review of the list of participants for the exercises conducted in 2016 and 2017 reveals that only staff from CSG, SMD, and Cluster management including Cluster CIOs and Cluster ISOs, attended the exercises. It is telling that line operational IT staff were not involved, even though they would, in many instances, be the first responders involved in identifying, detecting and responding to a cyber attack. Responding to a security breach involves more than the people in charge of cybersecurity. As stated by Dr Lim, cybersecurity involves *all* staff in an organisation, because the impact of a cyber attack affects the *whole* organisation. Technical staff are usually the first to spring into action following an incident as they seek to identify the problem, assess damage and start remediation. It is therefore essential for them to be involved in exercises, for the response plans and procedures to be ingrained in them.

918. Running real-world drills beyond tabletop is also a good way to test an incident response plan. In the context of IHiS, this would have allowed the SIRT and senior management to go through the full process of responding to and managing an attack. For example, a third-party vendor can be hired to oversee running the drill, to avoid internal bias, and provide a report that can be used for later assessment.

919. Testing incident response processes should also involve senior management and even members of the board of directors. This is a basic requirement of corporate risk management. Senior executives and board members should be prepared to respond to major crises caused by cyber attacks,

and this level of preparation would be best achieved by participation in simulation exercises.

41.1.2 Employees must be made aware of the procedures in place for reporting security incidents

920. People have a key role to play in an effective cybersecurity strategy, with many of the most basic attacks being avoidable if existing policies and procedures are followed. There should be a clear and established procedure for reporting a security incident. Sufficient attention must therefore be placed on ensuring that employees are aware of response plans that have been put in place.

921. A more fundamental problem emerged during the Cyber Attack – many IHiS employees who first witnessed signs of the attack were not even aware of any response plan for dealing with a security incident. As regards the Cyber Attack, the first responders who demonstrated initiative, like Sze Chun, Lum, and Katherine (all IT staff) stated that they were completely unaware of *any* security incident reporting procedure and hence had no guidance or training on how to collectively respond to the incidents before them. This is an obvious area for improvement.

922. As explained above, IHiS' incident reporting processes as regards Cluster CII systems are covered in two documents – SIRF and IR-SOP.

923. The above documents appear to be focused on reporting by the *Cluster*. The reporting lines in the documents begin with the Cluster ISO and GCIO. As acknowledged by Director CSG Kim Chuan, there is no written protocol for how IHiS staff, who discover an IT security incident affecting a Cluster's IT system, are to escalate the matter internally within IHiS, or to determine when and how to inform the Cluster ISO and/or GCIO.

924. It is also clear that many front-line IT staff were not even aware of the above documents, including:

(a) Sze Chun;

(b) Katherine;

(c) Lum;

(d) Steven;

(e) Henry; and

(f) Chan Chee Choong.

925. There is no clarity on whom staff ought to raise any potential security incidents to. Director CSG, Kim Chuan's position is that staff should inform their 'boss' or the SMD. On the other hand, GCIO Benedict has emphasised that speed of reporting matters more than the chain of reporting, and maintained a presence in a TigerConnect chat group containing staff from the delivery group, whom he expected to raise IT issues directly to him. IHiS CEO Bruce stated that in addition to the GCIO, the SMD Lead, Hann Kwang, should also be kept informed of IT security incidents, even though Hann Kwang does not appear in any documented reporting flow.

926. Further, even within the SMD team for SingHealth, processes were inconsistent and unclear. During the response to the Cyber Attack, Benjamin was reporting his observations to various individuals including both Wee and Ernest through multiple modes, including TigerConnect, Whatsapp, email, and in person, and it was unclear who had the responsibility for reporting upwards. This lack of consistency had been flagged several times during earlier TTXes. During the 2016 TTX, the external conductors had found that the members of the SIRT were not familiar with the written incident response procedures. A TTX in 2018

conducted by another external consultant showed that SIRT members did not follow the steps defined in the IR-SOP and SIRF when responding to incidents.

927. Lack of awareness of the organisation's response plan can hamper timely reporting, or even result in non-reporting. Although speed of reporting is important, it is also important to have a clearly-defined and well-communicated reporting flow, so that uncertainty and confusion is reduced and reporting is encouraged. It is also critically important that staff are rigorously tested on their understanding of the plans, and actually *follow* the plans when an incident occurs.

928. Vivek's testimony is that it is important for exercises to have "*realistic contours*" which bring out the "*pressure points*" for participants. The Committee agrees with Vivek's testimony. The Solicitor-General suggested that one novel way of educating staff about IT security would be Gamification. Benefits of Gamification include improved motivation and increased engagement. Games allow for role-playing as both attackers and defenders, and challenges participants to make quick, high-impact decisions, which help them to understand which activities can make the biggest difference during a cyber attack. This can be explored, and should not only involve technical staff, but should also include senior management of an organisation, and can be complemented by other initiatives such as red teaming exercises.

929. Organisations, in particular IHiS, must engage every employee in data security by using positive reinforcement to reward good behaviour, instead of the more conventional approach of identifying negative behaviour and reporting that behaviour to management. This should be led by CEO, IHiS with oversight by the chief executives of SingHealth and MOHH.

41.2 Pre-defined modes of communication must be used during incident response

930. Communication and coordination between members of the CERT, and between the CERT, SIRT and management, is critical.

931. During the Cyber Attack, as observed by Vivek, communication within the CERT was "*ad-hoc using various means such as TigerConnect chat, WhatsApp, emails, Excel sheets, PPT and other undocumented discussions*". In Vivek's expert opinion, this lack of formal coordinated communication impacted the investigation in more ways than one – critical information was not captured properly, captured in a fragmented manner, or was not shared with, or communicated clearly to, the relevant individuals. For example, there were various occasions in June and July 2018 when Benjamin had shared his *ad hoc* observations on the incidents in the SingHealth network with Ernest and Wee *via* Powerpoint slides, but both Ernest and Wee had difficulty understanding the significance of the information Benjamin was sharing.

932. In the absence of a coordinated system for communication, it proved to be a major challenge to find, coordinate and communicate with the key parties involved in responding to the incident. Vivek also observed that "*[i]mportant action items were not tracked and followed up on*", and cited the following particular examples:

(a) The user account for Workstation A had been identified as an account involved in suspicious activity as early as January 2018 but no action was taken on this finding and it was not tracked to closure. In fact, the user account for Workstation A later played a significant part in the Cyber Attack in June 2018, when it was used to access Citrix Server 4 from workstation VM 2; and

(b) There was no follow-up on other instances of access to a foreign IP address logged in the PHI 1's firewall logs in January 2018. This

proved to be significant, as this IP address belonged to a malicious C2 server that was later used during the Cyber Attack.

933. In Vivek's expert opinion, it is possible that investigation and proper follow-up on the above activities would have offered the CERT a chance to hunt the attacker before he did further damage during the Cyber Attack.

934. Accordingly, a formal method of communication should be established by IHiS led by the CEO, in the form of a centralised communication dashboard. This central dashboard would display all the details of the current state of investigations, allowing all members of the incident response team to keep abreast of developments and retrieve the information necessary to perform their roles. This would provide a more coordinated means of communication and would serve to document all communications, and limit the disruption and confusion arising from constant messaging across multiple platforms. Multiple streams of communication across different channels could otherwise overwhelm individuals and lead to missed messages or conflicting information.

935. For example, there was no centralised way for members of the CERT to ascertain whether items were being followed up on. In January 2018, Benjamin had already discovered that there were instances of callbacks to a suspicious IP address from PHI 1 and SGH. He arranged for this IP address to be blocked from PHI 1's network, but not from the SGH network. Benjamin sent an email to Ernest and his other colleagues from SMD, but did not follow-up and was not personally aware if anyone had blocked the suspicious IP address from the SingHealth network. In fact, no one did. A centralised communication dashboard can also help in managing, tracking and segregating information and updates relating to multiple concurrent investigations that may be ongoing.

41.3 Correct balance must be struck between containment, remediation and eradication, and the need to monitor an attacker and preserve critical evidence

936. In responding to an incident, it is crucial that responders, in this case, IHiS' CERT/SIRT, strike the correct balance between attempting to stop the observable signs of attack, and preserving evidence such that it is possible to track the movements of the attacker and monitor its activities. In this case, the responders erred too much on the side of containment and eradication, resulting not only in the loss of opportunities to detect the full extent of the attacker's presence in the network but also in the loss of important pieces of evidence.

937. Vivek has correctly highlighted a number of missteps by the CERT:

(a) CERT resorted to reformatting several systems infected with malware (*e.g.* PHI 1 Workstion in January 2018). While at some point these systems needed to be reformatted, doing so in a hurry can seriously hamper the investigation as it leads to loss of potentially valuable forensic evidence. A better practice would have been to quarantine (*i.e.* isolate) the system on the network without turning off the power, so that the infected systems could be studied further (*e.g.* to identify C2 servers with which the workstation was communicating).

(b) CERT also resorted to shutting down systems that were exhibiting suspicious behaviour (*e.g.* Citrix Server 1, Workstation B, PHI 1 Workstation). While this may seem to be a natural thing to do, doing so could seriously hamper the investigation as it leads to loss of potentially valuable forensic evidence. Again, a better practice would have been to quarantine the system on the network without turning off the power, for further study.

(c) CERT resorted to blocking IP addresses that were identified as malicious (*e.g.* IP address range associated with workstation VM

2; blocking communications with a foreign IP address for PHI 1's network). While at some point these IP addresses must be blocked, doing so in hurry can hamper the investigation as it indicates to the attacker that its presence has been discovered, and attackers usually respond by moving their communications over to another IP address or URL that has not yet been flagged as malicious. A better practice would have been to first study the network traffic for signs of any active data exfiltration. If data is found to have been exfiltrated, then the IP address should be blocked. Otherwise, it should be actively monitored to learn more about the attacker's behaviour and presence in the network.

(d) CERT and other responders resorted to resetting several passwords during the investigation (*e.g.* the L.A. account, the D.A. account, the A.A. account). While at some point these passwords must be reset, doing so in hurry can hamper the investigation as it indicates to the attacker that its presence has been discovered, and attackers usually respond by using other accounts that have not yet been flagged as compromised. A better practice would have been to put the compromised passwords on active monitoring and use them to learn more about the attacker's behaviour and presence within the network.

938. In Vivek's expert opinion, a CERT team (even one formed only six months prior) should not be susceptible to the above missteps. Hence, the response plan (for example, the IR-SOP on the security incident response methodology) should be improved by setting out rules cautioning against the missteps identified above and other similar examples. This is also the expert opinion of Vivek. In addition, the response plan must also be made available to *all* IT staff, as they are potentially first responders (as was the case in the Cyber Attack). It cannot be confined just to the IT security personnel. All staff should be aware of what they should, and should not, do in a security situation, to ensure that the appropriate balance is struck between stopping the attack and gathering evidence.

41.4 Information and data necessary to investigate an incident must be readily available.

939. A lack of information in the early stages of the incident response process has negative knock-on effects for the entire duration of the incident response. Responders will struggle to assess the impact of the attack, contain the damage, and escalate to management. As regards the Cyber Attack, investigations were hampered by the SMD team for SingHealth's inability to promptly obtain accurate information and data. This led to delays which proved to be significant. Two examples were observed by Vivek:

(a) The CERT had to physically visit affected sites to obtain forensic images of the compromised workstations. This slowed down investigations considerably as the team would have to first locate, then subsequently arrange to visit and seize, the machines. Workstation C took five days to be located and was picked up only on 18 June 2018. Such delay would have given the attacker valuable time to penetrate deeper into the system.

(b) The CERT did not have direct access to logs. Again, this created delay that could have been exploited by the attacker to penetrate deeper into the system.

940. Specifically, in relation to the two issues above, CERT should have direct access to the logs; and asset management should be reviewed to accurately reflect the location of assets, so that action can be taken immediately at the desk side, if necessary. These issues should be addressed by the CERT working closely with IT staff, particularly of the Delivery Group, to understand what data sources they have, what data they are capable of producing, and how the data can be managed and accessed when needed, during an investigation into a security incident. Engaging the staff who manage the various systems, and evaluating the asset management system will help in uncovering the full range of potential data sources.

941. Accordingly, the CERT should identify events that serve as a sign or signal of an attack (*e.g.* failed logins, deletion of logs, communication to unusual IP addresses *etc*.) that could provide contextual information about an incident, and establish processes for recording, aggregating, and making sense of such data points. The crucial point is that individual events and pieces of evidence must be meticulously recorded, and aggregated in a single place, so that responders are easily able to look at the cumulative mass of evidence to determine if an attack is taking place. This can best be accomplished by the establishment of a single, consolidated ASOC.

41.5 An Advanced Security Operation Centre or Cyber Defence Centre should be established to improve the ability to detect and respond to intrusions

942. The traditional prevention-dominant approach to cybersecurity, which focuses on defending the perimeter, has failed to prevent intrusions. The reality is that no network is impenetrable. Prevention is crucial – organisations cannot lose sight of it as the primary goal. However, a new proactive approach to security is needed to enhance capabilities to detect threats that will inevitably slip through the perimeter defences.

41.5.1 Importance of a proactive defence strategy

943. It is therefore critical to move to a detection-oriented strategy to defend against cyber attacks. It is not possible to control when a security incident happens, whereas it is possible to control one's response to the incident. The strategy must be one of prioritising efforts that enhance visibility, allow early detection and enable a proactive response through monitoring, analytics and prompt detection. The best defence is a good offence – responding early and aggressively can deter attackers from penetrating further into the network and realising their ultimate objectives. Vivek gave the example of a bank that had been breached and successfully responded aggressively:

> Were they breached? Yes. Were they technically compromised? Yes. But did anyone know about them, no, because there was no impact. The impact was contained.

944. IHiS' current security detection capability rests largely on its outsourced managed security service (**"MSS"**) provider. A MSS provider is an IT service provider that provides an organisation with cybersecurity monitoring and management of various security systems, which may include antivirus and anti-malware, intrusion detection systems, intrusion prevention systems, firewalls *etc*. Alerts raised would be communicated to IHiS security team who would then have to evaluate the alerts for significance (*e.g.* signs of attack), before acting upon them. Thus, while the MSS provider is responsible for receiving alerts, ultimately, assessments of the seriousness of the alerts and consequent remedial actions are squarely within the remit of IHiS' security staff.

41.5.2 Overview of an Advanced Security Operations Centre

945. The better way of integrating both alerts and responses is to have an ASOC. An ASOC would consolidate the people, processes, and technologies necessary to monitor and respond to potential security incidents in a single place, facilitating detection, containment, and remediation of IT threats. An ASOC should be designed to monitor applications and network activity for unusual signs, then analyse those signs to determine whether an attack is in progress. If it is determined that an attack is taking place, the ASOC (also called a Cyber Defence Centre ("**CDC**"), where it incorporates incident response functions) can then coordinate investigations, reporting, and remediation efforts.

946. In Gen. Alexander's expert opinion, an ASOC is an especially important organisational measure to be put in place, to support the CISO. Vivek observed that an ASOC would be a better option than having outsourced MSS, as MSS providers are often limited to superficial reporting of alerts as they do not have full access to an organisation's systems. In contrast, an ASOC would have full access. This is key to responding effectively to an attack. As Vivek said:

> What is important is to respond without wasting any time and respond with full force, all hands on the deck, where everybody is focused on figuring out what's going on, understanding what the attackers will do next if they had this level of access, understanding what sort of steps can be taken to prevent them from spreading further, contain them within a certain environment, and then figure out the remediation steps which requires a response with full force.

947. Technology. The ASOC must be equipped with the right tools to do its job. This includes a suite of technology that provide the right level of visibility over the organisation's operations, commensurate with its security posture. The suite of technologies may need to be updated periodically, as cyber attack vectors evolve. Some examples include:

(a) Security information and event management ("**SIEM**") solutions;

(b) Intrustion Detection System ("**IDS**")/Intrusion Prevention System ("**IPS**") solutions;

(c) Threat and vulnerability management tools;

(d) Filtering technologies;

(e) Data loss prevention tools;

(f) Traffic/packet inspection solutions;

(g) Data analytics platforms;

(h) Reporting technologies; and

(i) Forensic tools.

948. The SIEM solution chosen is particularly important, as ASOCs are most often organised around the SIEM, which aggregates and correlates data from the various tools employed by the ASOC onto a single platform, which then provides a comprehensive overview to security staff at a single glance.

949. People. An ASOC requires the right manpower to function well. The staff of the ASOC should be a mix of experienced security professionals and IT staff. The IT staff provide a solid understanding of the organisation's IT infrastructure, and are usually trained in computer engineering, network engineering, or computer science and may have credentials such as CISSP[87] or GIAC. [88] The security personnel can help to bring fresh perspectives based on their experience. Working together, the ASOC staff should be able to analyse large quantities of data and intuitively recognise the need for further investigation when it arises.

950. Processes. The ASOC needs to have well-defined processes that facilitate consistent operations and repeatable outcomes. The ASOC needs to be stable and functional at all times, as it is the heart of an organisation's security architecture. At the same time, the processes must be wide and flexible enough to accommodate possible incident scenarios and provide detailed guidance for response. Examples of incidents include:

(a) Phishing;

(b) Malware infections;

(c) Bring your own device-related incidents;

---

[87] Certified Information Systems Security Professional ("**CISSP**") is an independent information security certification granted by the International Information System Security Certification Consortium, also known as ISC.

[88] Global Information Assurance Certification ("**GIAC"**) is an information security certification entity that provides a set of vendor-neutral computer security certifications linked to the training courses provided by the SANS Institute.

(d) Website defacement; and

(e) Denial of service attacks.

951. Further, IHiS should consider designing the ASOC to integrate monitoring and incident response systems with emerging technologies even further upstream in the detection process, such as threat intelligence and security analytics.

41.5.3 Features of an ASOC

952. The key features of an ASOC are:

(a) Visibility of threats;

(b) Ability to detect sophisticated, targeted, persistent or previously unknown threats;

(c) Ability to process alerts, to analyse and understand them;

(d) Ability to respond to attacks, if a network is impacted;

(e) Preparation for the inevitable successful attacks that will impact their networks in future;

(f) Ability to discover and mitigate vulnerabilities before they are exploited by others; and

(g) Workflows, processes and teamwork.

953. Increased visibility. A well-designed and implemented ASOC thus maximises existing security investments by linking individual technological components (such as those mentioned in paragraph 947 (pg 326) above) in a manner that extends the benefits these systems provide. This allows analysts a full view of data from multiple sources within the network and its systems.

954. Increased capability for correlation and analysis of data. An ASOC security analyst must have the right tools to identify and analyse an attack. The ASOC pulls together information from multiple sources, including endpoints, gateways, or networked devices, to determine what is important. Without an ASOC, a security analyst would have to go through the laborious process of checking multiple sources of input. For example, during his investigation into the incidents taking place in January 2018, Benjamin had to check the antivirus programs, which led him to check PHI 1's IPS, and finally to look at firewall and proxy logs. He also had to ask the MSS service provider to continue monitoring traffic to the suspicious IP addresses, as the MSS were outsourced to the service provider.

955. Manually checking multiple sources of input is both time-consuming and prone to error. Important sources of input may be missed. Further, the disorganised nature of information gathering means that larger patterns of suspicious conduct might not be recognised. The better option is for the ASOC to utilise advanced behaviour-based analytics to determine if the pattern of activities across the entire network indicates a legitimate human user, an innocuous automated process, or malicious activity. This shifts the paradigm from log-based, post-incident security to more proactive intelligence-driven security.

956. Full lifecycle management of incidents. The key point is that an ASOC should cover the entire lifecycle of an incident, all the way from initial detection through response and resumption of normal operations. This includes 24 by 7 monitoring, coordination of response teams and processes, and containment and remediation activities, all under one roof to improve response time and reduce confusion.

957. IHiS is currently exploring the option for transitioning the current MSS to an ASOC. The proposal is for the ASOC to have proactive defence capabilities, including active Threat Hunting. Leong Seng has said that this ASOC will combine people, processes, and technology to better manage IHiS' overall security defences. Essentially, a good ASOC would pull together all the strands

mentioned in this recommendation, and act as a focal point for convergence of the organisation's incident response processes. This is a step in the right direction and is to be encouraged.

42 RECOMMENDATION #7: PARTNERSHIPS BETWEEN INDUSTRY AND GOVERNMENT TO ACHIEVE A HIGHER LEVEL OF COLLECTIVE CYBERSECURITY

#VIGILANCE #DETECTION #GOVERNANCE

958. A common thread running through the expert evidence is that the occurrence of a cyber attack is inevitable. The ever-increasing scale and sophistication of APTs means that APTs can and will find ways to breach networks and systems. Singapore is particularly vulnerable to cyber attacks for two reasons. First, the attack surface is large, as our society is highly connected and digitalised. Second, large amounts of sensitive data reside in our servers, as Singapore is a business, financial, and healthcare hub, making us a high-value target for cyber threat actors.

959. Recommendations #1 to #6 are essential. Recommendation #7 builds on these recommendations to bring our cyber defences to a new and higher level.

960. The high degree of interconnectivity and the potential risks at the national level make it imperative that there is collective security over our systems. As Gen. Alexander stated, collective security is key, and the government must be involved in cyber defence, especially against APTs. Dr Lim has also noted that many countries have classified cyber attacks targeting CII as a matter of national security. The Singapore government has recognised this, and the establishment of CSA was a public commitment to strengthen cybersecurity as a whole.

961. We will discuss four aspects of collective security here:

(a) Sharing of threat intelligence;

(b) Partnering with Internet Service Providers ("**ISPs**");

(c) Defence beyond borders; and

(d) Using a network to defend a network.

962. In considering the need for collective security, it is apposite to highlight Gen. Alexander's observations:

> "No individual company can stand alone against nation-state threat actors. Even if one company has strong defenses, a state aggressor will patiently probe the business' entire ecosystem, or even the entire business sector seeking a point of vulnerability – and there will be one. Network visibility and automated information sharing between companies, sectors, and governments are necessary to provide a comprehensive defense. Combining the capabilities of the public and private sectors is essential. … Governments possess a monopoly on the use of force and public/private collaboration is necessary to strike back using the full spectrum of governmental power. A solid collective defense foundation will allow high-speed, automated requests for government support."

42.1 Threat intelligence sharing should be enhanced

963. All the experts recognised that enterprises can and should purchase threat intelligence from commercial companies. This is a recommendation that should be adopted. Commercially available threat intelligence is at a basic level, and includes information on common threats across the world and the mitigation that can be done in response to these threats.

964. Apart from this basic level of threat intelligence, there are other sources of threat intelligence:

(a) Intelligence generated by CSA from their investigations with their investigative partners;

(b) Intelligence generated by each enterprise from their investigations and prevention and detection tools;

(c) Classified information provided by commercial companies to their trusted partners; and

(d) Classified information provided by security partners in other countries.

42.1.1 Intelligence generated by CSA from their investigations with their investigative partners

965. CSA operates an intelligence centre which analyses intelligence generated from its investigations. Where CSA is involved in containment and investigation, it will concurrently share threat intelligence from such investigations with all CII sectors so that protective and precautionary measures can be taken.

966. The threat intelligence is proactively shared in the form of actionable items, *i.e.* by providing malware indicators or specific instructions. CE, CSA's evidence is that *actionable* intelligence is important in order to let the enterprises know what steps to take. Dan's evidence is that CII operators have different levels of maturity and not all CII operators will be able to analyse the intelligence and translate it into useful technical information that they can pass to their IT departments for action. Actionable intelligence is thus required, so that CII operators can consume the intelligence for immediate use. CSA thus informs the CII of the potential threats they need to look out for in particular systems or applications, and how they should mitigate against the threats.

967. CSA has a few modalities of sharing threat intelligence:

(a) Alerts or advisories are sent to CII operators. In 2017, 80 alerts or advisories were sent. Where one sector is subject to a cyber attack, CSA shares actionable intelligence to enable CII sectors to level up across the board to prevent other sectors from being similarly attacked.

(b) Spot reports and intelligence summaries are sent to CII operators. These cover cyber attacks in other countries, so that CII operators can learn from what has happened in other countries, and take the necessary remediation or protection measures within their own systems. In 2017, around 20 spot reports and intelligence summaries were shared.

(c) Curated intelligence specific to a sector is sent to the particular sector. CSA may then work with that sector to ensure the necessary follow-up action is carried out.

(d) CSA conducts presentations on the threat landscape at meetings with CISOs and management in CII sectors.

968. CSA's distillation of threat intelligence into actionable items for CII sectors is a sensible approach. It has the twin benefits of (i) analysis by CSA of the nature of the threat; and (ii) clear directions to CII sectors of how they can take steps to mitigate the threat. This is crucial, because raw threat intelligence alone cannot form the basis of a detection program, and there must be some set of event data to which the threat intelligence is applied.[89]

969. To illustrate how CSA shares threat intelligence, its actions after the Cyber Attack are highlighted below:

(a) Concurrent to CSA's containment and investigation efforts, CSA provided intelligence and situation awareness to all the other CII sectors.

(b) CSA instructed the CII sectors to scan for newly discovered IOCs that would be indicative of the same attacker being present in their

[89] Michael Collins, *Network Security through Data Analysis*, (O'Reilly Media, Inc., 2nd Ed, 2017) at p329.

networks, and advised on possible measures to mitigate a similar incident.

(c) CSA called up other users of the SCM database to explain the vulnerabilities observed in the SCM database and to ask them to take immediate measures to protect themselves.

(d) CSA organised a briefing for relevant stakeholders of all CII sectors and recommended that they review their protection and management of large databases.

(e) Following the public announcement of the Cyber Attack, CSA directed that CII sectors adopt heightened measures, in anticipation of potential opportunistic attacks on sensitive systems.

(f) CSA published two advisories on protection and precautionary measures: (i) a technical advisory for companies on measures to protect their systems and customers' personal data; and (ii) to encourage members of the public to take personal precautionary measures against scams that could arise from the theft of the personal data that had been lost in the Cyber Attack.

970. In our view, it is critical for the government, through CSA, to continue to ensure sharing of threat intelligence across the CII sectors (in line with its information management process). As Vivek noted, the attackers have the ability to move across the whole fabric of systems, the defenders must thus have visibility across the same range of systems, in order to provide an adequate defence.

42.1.2 Intelligence generated by each enterprise from their investigations and prevention and detection tools

971. There should be sharing of threat intelligence within each sector and across sectors. This is valuable where the sectors are faced with like threats, or use similar systems and thus have similar vulnerabilities.

972. There should also be sharing of threat intelligence from the sectors to the government. Gen. Alexander has opined that if the cyber attack is meant to destroy a country's infrastructure, the government must have a role. The government has to have the ability to see the cyber attack in time, in order to have a role in defence that goes beyond incident response. Where enterprises encounter suspicious behaviour indicative of a cyber attack, we recommend that they share this information with CSA. Where the suspicious behaviour meets the threshold for reporting under the relevant reporting frameworks, the information will have to be shared with the CSA as soon as possible, or at the latest, in line with the timelines for reporting under the frameworks. Even where the suspicious behaviour may not meet the threshold for reporting, enterprises should exercise judgment on whether their observations should be shared with CSA nonetheless, to enable in-depth analysis and, if necessary, broader dissemination across the CII sectors.

42.1.3 Classified information provided by commercial companies to their trusted partners

973. Commercial companies which offer threat intelligence feeds may separately engage in a deeper analysis of the intelligence and further generate classified intelligence based on this analysis. Such analysis is not available commercially, but may be shared with trusted partners. CE, CSA's evidence is that CSA is a trusted partner of some of these commercial companies, and receives classified threat intelligence from them.

974. CSA will then distil this threat intelligence into actionable intelligence and share it with CII operators (see paragraphs 966 (pg 333) and 967 (pg 333) above).

42.1.4 Classified information provided by security partners in other countries

975. Gen. Alexander's view is that the sharing of threat intelligence is a good area for allies to work together, and should be driven by sharing of threat indicators across governments. Dr Lim's evidence is that there are a lot of collaborations at the country level, and that most countries are prepared to share threat intelligence where it does not target a specific sensitive area.

976. CE, CSA's evidence is that Singapore has memoranda of understanding with several countries to facilitate cooperation in sharing threat intelligence. Such sharing enables a broader view of threats and threat actors.

977. CSA will again distil this threat intelligence into actionable intelligence and share it with CII operators.

42.2 Partnerships with ISPs should be strengthened

978. Dr Lim gave evidence that ISP analytics with national network and DNS data is a valuable tool in Singapore's multi-layered cyber defence capabilities. This capability allows real-time streaming of data where anomalous or malicious activities can be identified. It also goes further to forewarn of imminent threats. This is a capability that should be further studied and developed.

42.3 Defence beyond borders – cross-border and cross-sector partnerships should be strengthened

979. We have covered the sharing of threat intelligence between countries above. In addition to this, partnership between countries can take place on a wider basis, including sharing of best practices, and response to cyber attacks.

980. In addition to government-to-government sharing, sharing can also take place between CII sectors and enterprises both within Singapore and from other jurisdictions. We recommend the continuation of sector-level sharing of best

practices, and that the CII sectors consider whether to establish forums for such sharing on a regular basis.

42.4 Using a network to defend a network should be explored

981. An advanced technique or strategy that may be considered for the longer-term is using a network to defend a network. This can be done by establishing a behavioural analytic capability, with an expert system and hunt platform, to provide network speed information to a collective cybersecurity platform for the sector and to the government. Gen. Alexander called this his most important recommendation, as the ability to see across companies and sectors allows the elimination of threats that are invisible to any one company or sector.

982. Behavioural analytics make collective defence a possibility, and produces a wealth of events that can be shared in a collective defence strategy at network speed. As noted by Gen. Alexander, if the government has the opportunity to see data gathered by enterprises from behavioural analytics, the government will be able to map such data onto classified intelligence, and can inform the sectors of which behaviours they need to focus on, and what remedial action to undertake as a priority. Gen. Alexander opined that "*[i]t is ironic that the network and associated devices have become the biggest technological advances of our time, yet we don't use a network to defend a network*".

983. As behavioural analytics is an advanced technique or strategy, it may not be necessary to immediately implement this recommendation. Even in the United States of America, behavioural analytics are still in the initial pioneering phase in their healthcare sector. CE, CSA's evidence is that cybersecurity is still a nascent area within Singapore's ecosystem. As such, before enterprises adopt more advanced new technology or methods, enterprises should first ensure that they get the basics right in the short-term. Even as Singapore works on getting the basics of cybersecurity right in the short-term, the sectors should continue to monitor developments in behavioural analytics and other advanced technology. This could involve dialogues with vendors of commercially available products to gain a better understanding of the products, test them, and discover their

limitations. As cybersecurity is an evolving and dynamic area, the CII sectors should continually educate themselves on the latest technology, so that they can be implemented at the appropriate time, without undue delay.

984. While behavioural analytics is more suited as a long-term recommendation, there is an aspect of collective security that can be implemented in the medium-term. To enable governments and companies to learn how to fight in cyberspace as a cohesive whole, there should be promulgation of a common doctrine, system interoperability, information sharing, regular exercises, and trust. A common doctrine of cybersecurity[90] may include (a) goals (*e.g.* the level of cybersecurity sought and the acceptable risks, costs, and trade-offs); and (b) means (*e.g.* protect, detect, respond, and recover). System interoperability will enable sharing and ready use of information securely and effectively. Information sharing may include sharing of threat intelligence and best practices, as we have elaborated on earlier in this section. We have elaborated on the need for regular exercises in in the context of improving incident response processes above.

985. Recommendation #7 will bring our cybersecurity posture to a higher level. Although it is the last of the Priority Recommendations, it is not the least important. CSA and relevant agencies should study this recommendation and consider how to implement measures to better achieve collective security, sharing of threat intelligence and networked defence.

[90] For more information on the scope of the doctrine for cybersecurity, see Deirdre K. Mulligan and Fred B. Schneider, "Doctrine for Cybersecurity".

43 RECOMMENDATION #8: IT SECURITY RISK ASSESSMENTS AND AUDIT PROCESSES MUST BE TREATED SERIOUSLY AND CARRIED OUT REGULARLY

#PREVENTION #VIGILANCE #GOVERNANCE

986. IT security risk assessments and audits are important for ascertaining gaps in an organisation's policies, processes and procedures, and must be treated seriously and carried out regularly, with findings followed up on religiously.

43.1 Risk assessments must be conducted at critical junctures

987. While the HITSPS does provide for the conduct of IT security risk assessments, the policy is not adequate, and worse, there were gaps in IHiS staffs' conduct of the risk assessments. We will elaborate on this with reference to our recommendations as follows.

43.1.1 IT security risk assessments must be conducted on CII and mission-critical systems annually and upon specified events

988. The HITSPS requires that an IT security risk assessment be done for all mission-critical IT systems, before they are commissioned and during the system design phase; and maintained whenever there are major changes. Under section 15(1)(*b*) of the Cybersecurity Act, however, CII owners are required to conduct cybersecurity risk assessments on CII at least once a year, and this risk assessment is to include each CII asset in the CII system. Under section 15(2) of the Cybersecurity Act, the CII owner must furnish a copy of the cybersecurity risk assessment report to the Commissioner not later than 30 days after completion of the risk assessment.

989. Accordingly, we recommend that IHiS must re-formulate its policy to require the conduct of cybersecurity risk assessments on CII and mission-critical systems at critical junctures:

(a) at least annually;

(b) in respect of new systems, during the design of the solution and before commissioning; and

(c) whenever there are major changes to the systems.

43.1.2 A written cybersecurity risk management framework must be established

990. The HITSPS does not set out a proper cybersecurity risk management framework. The CCoP requires CII owners to establish a written cybersecurity risk management framework, which shall include:

(a) roles and responsibilities in managing cybersecurity risk, including reporting lines and accountabilities;

(b) identification and prioritisation of CII assets;

(c) organisation's cybersecurity risk appetite, as well as thresholds or limits for residual risk;

(d) cybersecurity risk assessment methodology; and

(e) treatment and monitoring of cybersecurity risk.

991. We recommend that a comprehensive written cybersecurity risk management framework covering at least the above areas should be established. We elaborate on our recommendations in respect of some of these areas.

43.1.3 Risks must be thoughtfully identified and prioritised during each assessment

992. The HITSPS sets out an IT security risk assessment form containing pre-populated threats/risks. The SingHealth Cluster ISO Wee used the same template in 2016 and 2017 to conduct the risk assessment for the SCM system. No thought was given as to whether the same set of threats/risks listed in the template were applicable (at all, or year on year).

993. We recommend that IHiS/SingHealth should pro-actively identify the applicable threats/risks for each relevant system at each assessment. Risk assessment forms should not come hard-coded with a set of pre-populated threats/risks, such that the same template of fixed threats/risks are reviewed year on year without further thought. In particular, given what IHiS/SingHealth now know about the attacker's *modus operandi* in the Cyber Attack, and given that the healthcare sector may be subject to other APT attacks in future, the threat/risk areas pertaining to each relevant system should be re-looked and identified taking into account the new knowledge gained. As Vivek said, "*the way I recommend risk management be done is you apply your controls to the attackers' modus operandi and see where you have gaps*".

994. Vivek also proposed re-thinking the prevalent practice of using asset classification to prioritise risk. He explained that organisations have to operate within a budget, and that requires prioritising investments based on the risk so as to maximise the benefits derived from the budget. Many factors are considered while assessing risk, and asset classification or asset value is one of them. Most classification models are quite simplistic, in that they mostly ignore the effect of network connectivity between systems. As a result, several systems, and especially endpoints, get classified as low priority assets and consequently receive lesser degree of controls coverage including preventive and detective controls. Attackers know this very well, and exploit it using a simple and highly effective *modus operandi* involving penetrating lower priority assets which receive less coverage for defensive, preventive and detective controls. Thereafter, attackers would perform lateral movement and privilege escalation. It becomes

very difficult to distinguish the attacker's activity and even if they are discovered, it becomes extremely difficult to take meaningful action to contain them without breaking the business. Using asset classification to prioritise risk is a systemic weakness.

995. In line with the weakness identified by Vivek, we would caution against a fixed practice of prioritising cybersecurity risks according to asset classification. Instead, we recommend that, similar to the identification of risks, the prioritisation of risks also be carried out proactively and thoughtfully.

43.1.4 A clear process and methodology for cybersecurity risk assessment, and treatment and monitoring of cybersecurity risk should be established, and staff must be trained on the same

996. Wee explained the procedure which he followed for completing the 2016 and 2017 risk assessment forms as follows:

(a) His role was to initiate the annual risk assessment process for CII, and he would use the "IT Security Risk Assessment Form" template in the HITSPS. He would make an initial assessment of the risks and fill up the form. He would then submit the draft form to the Infrastructure and Application groups in IHiS' Delivery Group for review. Once they completed their reviews, he would send the form to GCIO Benedict to review. After GCIO Benedict reviewed the form, Wee would present it to the SingHealth IT Steering Committee (a management-level committee).

(b) According to Benedict, the form was sent to him for his "*reference and information, but [his] approval of the completed risk assessment is not required*". If new technical controls were required in response to the risks identified, Wee would coordinate with the relevant teams in the IHiS Delivery Group to ensure they provided and implemented the necessary measures.

997. The way in which the risk assessments were conducted is unsatisfactory. There was no clear ownership over the identification and assessment of risks and risk controls. While Wee was put in charge of the process, the technical knowledge of the system being assessed resided within the IHiS Delivery Group, over whom Wee exercised no control or oversight. There was no SingHealth management line of sight over the process either, although the SCM system belonged to SingHealth. This resulted in cursory risk assessments, as well as stark errors in the completion of the risk assessment forms. For example, in the 2016 risk assessment, in respect of item 9 concerning threats of malicious software being introduced by the developer programmer, it was stated that the existing risk control included bi-annual vulnerability assessment and annual penetration testing and code review – which was, as accepted by Kim Chuan, clearly wrong, because there was no penetration testing or code review of the SCM application. This mistake was repeated in the 2017 risk assessment form.

998. It is also unclear if anyone was tracking the risk assessments. Under the Processes for Management of Critical Information Infrastructures (CII) Systems in Health Sector ("**PMCII**") policy, CSG was supposed to be tracking the risk assessments of the CII in the healthcare sector, but CSG did not track the completion of the proposed action plans from the 2016 risk assessment, although Kim Chuan has stated that CSG is in the process of doing so for the 2017 risk assessment.

999. We recommend that IHiS/SingHealth set out a clear process and methodology for cybersecurity risk assessment, which should include:

(a) How to identify the threats/risks that the system is subject to, and who is in charge of such identification. For example, Kim Chuan has stated that for the conduct of risk assessments moving forward, there should be a look-back and identification of issues raised in internal audit reports or in other penetration test reports, which should then be taken into account when assessing the risk of a particular threat;

(b) How to identify the controls that are in place to address the risks, and who is in charge of such identification;

(c) How to assess the likelihood of the risk occurring, and who is in charge of such identification. For example, Kim Chuan has stated that efforts have begun to ensure IHiS staff are aligned on the understanding and assessment of risks, so as to reduce the element of subjectivity in risk assessment;

(d) How to identify the additional controls that may be needed to address the residual risks, and who is in charge of such identification;

(e) Who is in charge of formulating the action plan to implement measures for additional controls;

(f) How the action plan shall be tracked, by whom and when; and

(g) Who in management shall review and have oversight of the risk assessment process.

1000. We further recommend that once the process and methodology are established, there should be proper dissemination of the same to the relevant staff, who should also attend training to familiarise themselves with the process and what implementing it entails. Indeed, Kim Chuan testified that CSG was conducting workshops for Cluster security officers, SMD and the Delivery Group to train them on risk assessments. The workshops would harmonise the assessment of cybersecurity risks and effectiveness of controls by Cluster ISOs and GCIOs.

43.1.5 A policy should be established for a comprehensive risk register to be maintained and updated after every risk assessment

1001. The CCoP requires CII owners to maintain a list of all cybersecurity risks identified, by way of a risk register in respect of each CII. CII owners shall ensure all identified cybersecurity risks listed are monitored regularly with a view to ensuring that the organisation's thresholds or limits for risks are not breached. The risk register shall be updated after every cybersecurity risk assessment. A risk register shall document the following:

(a) Date the risk is identified;

(b) Description of the risk;

(c) Likelihood of occurrence;

(d) Severity of the occurrence;

(e) Risk treatment;

(f) Risk owner;

(g) Status of risk treatment; and

(h) Residual risk, which is defined in the CCoP as "*the risk exposure after risk mitigating controls are considered or applied*".

1002. While the HITSPS provides for an IT security risk register, the policy is inadequate, when compared against the requirements under the CCoP.

1003. We recommend that a policy be put in place that establishes:

(a) the requirement for a comprehensive risk register, documenting the items set out in the CCoP, in respect of each CII and mission-critical system on which a risk assessment is done;

(b) the requirement for the risk register to be updated after every cybersecurity risk assessment;

(c) the person(s) in charge of maintaining the risk register; and

(d) a protocol for surfacing of the risk register to senior management at regular intervals.

43.1.6 Senior management should be responsible for and clearly articulate the organisation's risk appetite

1004. The CCoP defines "risk appetite" as "*the amount and type of risk that an organisation is willing to take in order to meet their strategic objectives and it is often taken as a forward looking view of risk acceptance*". "Risk acceptance" means "*the informed decision to knowingly take a particular risk*".

1005. In our recommendations on the adoption of an enhanced security structure and readiness, we had explained the experts' view that cybersecurity should be treated as a risk management issue and not merely a technical issue, and that senior management had to have oversight of risks. For example:

(a) Gen. Alexander recommended that risks need to be elevated to the CEO level, and not stop at the CIO (who may have a conflict between the two missions of ensuring IT operations and IT security) and end up taking risks that the CEO is not aware of. It is very important to have the CEO or management at equivalent level know and discuss the risks. There must be sufficient senior management oversight of risks.

(b) Dr Lim recommended that cybersecurity risks should be treated as part of the enterprise risk management and should be regularly updated every quarter at the enterprise risk management meeting. This is to ensure that cybersecurity risk is given the necessary

attention and resources are directed and prioritised by the senior management within the organisation.

(c) Relatedly, and as we have also explained, Dr Lim also expressed the view that the senior management making decisions on risks would need to be equipped with technical expertise/competency to appreciate and manage the risks.

1006. In line with these recommendations, it would follow that it is for senior management to *articulate* the organisation's risk appetite, and we recommend that a clear cybersecurity risk appetite statement be drawn up and regularly reviewed and updated by senior management.

43.2 Audit action items must be remediated

43.2.1 Regular audits on CII systems must be conducted by an independent third party in line with the CCoP requirements and upon specified events

1007. Under the CCoP, CII owners shall carry out an *independent* cybersecurity audit of CII at least once every two years or at such higher frequency as may be directed by the Commissioner.

1008. The CCoP provides that the scope of the audit shall include:

(a) All CII owned by the CII owner; and

(b) Compliance with the Cybersecurity Act and the CCoP, and any applicable codes of practice, codes of standards of performance and directions that the Commissioner may have issued.

1009. A CII owner shall submit the audit report to the *Commissioner* within 30 days after the completion of the audit as required under section 15(2) of the Cybersecurity Act.

1010. Dr Lim also recommended that SingHealth and IHiS should conduct regular security reviews and audits to validate the security measures that have been put in place to protect the database systems. This is especially important when systems are being upgraded, maintained and serviced, as well as when there are changes in the system configuration. The audit should be done by an independent third party who has no preconceived opinions on the security or configuration of the system.

1011. We recommend that in respect of CII systems:

(a) Audits be conducted at the intervals specified by the Commissioner and under the CCoP;

(b) Audits also be conducted when the CII systems are being upgraded, maintained and serviced, and when there are changes in the system configuration; and

(c) Such audits be conducted by independent third parties who had no input into the design or operation of the system in question.

43.2.2 Periodic audits on other IT systems should be conducted in line with Audit Committee requirements

1012. Under the HITSPS, GIA shall conduct independent audits of PHIs' IT systems periodically to evaluate and test the adequacy of, and the compliance to prevailing IT security policies and standards. The HITSPS provides that the frequency and scope of audits shall be directed by the Audit Committees of the respective institutions. This policy remains acceptable in respect of other non-CII systems, although IHiS/SingHealth should consider whether there are non-CII but important systems which should be audited more regularly or frequently.

43.2.3 A written protocol for the remediation of audit findings must be established

1013. There were serious gaps in the way audit findings are remediated by IHiS, as evidenced by the problems with the remediation action plans arising from the FY16 IT security audit on H-Cloud which included network penetration testing. These problems can be summarised as follows:

(a) There was no consideration of whether audit findings applied across Clusters and remediation steps should also been taken across Clusters.

(b) At IHiS staff-level, remediation was stated to be done when it was not actually done or not done thoroughly. No verification was conducted by line management.

(c) There were misunderstandings with GIA on what the remediation measures were to entail, and when they were supposed to be completed.

(d) CSA found that several of these vulnerabilities were present during the Cyber Attack, and could have been exploited by the attacker.

1014. Witnesses from IHiS' senior and line management who testified before the Committee acknowledged these failings, and put forward suggestions on how the audit remediation process could be improved.

1015. Lum, as the supervisor of the staff who had failed to take steps to comply with the requirements under the audit remediation plan for password complexity and administrator credential issues, stated that he would ensure that such important tasks were verified personally by him or a designated person in future such that the audit findings would be properly addressed and closed.

1016. Leong Seng testified that:

(a) Since April 2018, IHiS has set up a centralised audit liaison team to pool all audit issues from all audit reports from all Clusters. The reports are maintained in a shared platform with all audit issues being tracked. The GIA has access to this shared platform so that everyone is looking at one common list of audit issues. There will be a service management team (inside the Delivery Group) to handle audit management and be the single point to do the overall tracking of the response to the audit issues.

(b) For specific audit findings, the Infrastructure team of the respective Cluster to which the finding related would come up with a remediation plan and deadline. That team has to execute the plan accordingly. The other Cluster Infrastructure teams (in respect of which the audit finding was not specifically made) would plan measures as well if the finding is relevant to their Cluster.

(c) The Infrastructure Services group is organised in a matrix manner, with a horizontal Cluster Infrastructure Lead, and vertical Tower Leads for specific domain competency areas of system management, security management, end-user and network. The Tower Lead would ensure that issues surfaced by an audit on any one Cluster is propagated to the other Clusters. The Tower Lead would ensure harmonisation and standardisation of the way the Clusters remediated and put in place measures. A similar structure applied to the Applications group. The Tower Leads would drive the efforts to remediate issues within their respective competency areas in a standardised manner, but the specific remediation plans and plan timings would be planned by the horizontal Cluster Leads.

1017. At senior management level, based on Benedict's evidence, there appeared to be processes in place for surfacing audit findings and escalating problems with remediation. We note that in the case of the GIA's FY16 audit on

H-Cloud, the problem was that line management did not verify staff's purported remediation actions, such that line management did not know that there were remediation issues to surface to senior management, until the GIA's escalation in mid-2018. Benedict explained how senior management was generally involved in the audit process, as follows:

(a) Results of internal audit reports were distributed to Cluster Audit Committee, Cluster senior management, and IHiS senior management for CII and non-CII audits. For CII audits, the reports were also sent to CSG for monitoring of the follow-up action.

(b) If, as GCIO, Benedict found that remediation was not being done *per* the stated timelines, he would escalate the matter to the Director of the Delivery Group (Leong Seng), and if necessary, the IHiS CEO, Bruce, and at the same time keep SingHealth management apprised of the potential delay.

1018. Benedict suggested that regular updates by the Delivery Group on the status of audit items should be provided at CIO forums, for CIO to update Cluster management, with urgent matters highlighted.

1019. The GIA's IT audit head, Thng Chiok Meng, suggested that the GIA's verification of audit remediation action items could be done on half-yearly basis for staggered batches of audit findings, instead of only being reviewed by the GIA in the next financial year.

1020. Separately, Dr Lim recommended that it should be an independent party who should confirm implementation of the remediation according to the audit recommendation. The Committee agrees with this.

1021. Having considered all the evidence, we recommend that a written protocol for the remediation of audit findings should be established, which should set out, minimally the following requirements:

(a) A process for surfacing audit findings and the status of audit findings, at regular intervals, to IHiS' Audit and Risk Committee and CII owners (in this case, SingHealth).

(b) A clear remediation plan by the Delivery Group for each audit finding must be drawn up that:

(i) details the actions which the issue owner will take to address all non-compliance; and

(ii) sets out the timeline(s) for implementing the actions stated.

(c) Clear communication and agreement on the remediation plan by the Delivery Group with the auditor.

(d) A system to be put in place for verification, at IHiS' line management level of the implementation, of remediation plans.

(e) A system for centralised tracking of the status of audit findings, and the propagation of remediation plans across Clusters, where relevant.

(f) A process for escalation to IHiS' Audit and Risk Committee and CII owners in the event of problems with the implementation of remediation plans.

(g) Verification by the GIA of audit remediation action items to commence within six months of the audit findings instead of only being reviewed by the GIA in the next financial year.

1022. There has to be a policy of zero tolerance towards false or incorrect reporting of remediation of audit findings.

44 RECOMMENDATION #9: ENHANCED SAFEGUARDS MUST BE PUT IN PLACE TO PROTECT ELECTRONIC MEDICAL RECORDS

#PREVENTION; #DETECTION; #GOVERNANCE

1023. EMRs undoubtedly present many benefits. They improve patient care, and coordination of care, through enhanced access to patients' medical information by all members of the healthcare team. The platform chosen for SingHealth to store EMRs was the SCM. The SCM operates like a dashboard, holding information such as patient records, diagnostic data, and medical history. This is very sensitive information. As the Cyber Attack has demonstrated, it is critical to protect the security and confidentiality of such medical records.

1024. The Cyber Attack aside, other recent cyber attacks have seen data breaches grow in size, number, and scope. Whether attacks are against telecommunications, financial services, entertainment, or healthcare institutions, data in respect of *millions* of users has been compromised. The attackers are no longer going after just credit card information. Attackers are after personally identifiable information ("**PII**").

1025. Breaches involving PII and patient data are particularly hazardous to both individuals and organisations. Harm to the individual may include tampering with medical records, identity theft, embarrassment, or blackmail. Harm to the organisation may include a loss of public trust, legal liability, or remediation costs.

1026. Protecting the perimeter proved insufficient against the attacker in this case, and in any event, the threat to EMRs may come from malicious insiders. It is recommended that, network security aside, data-centric security measures must be implemented to:

(a) Ensure the confidentiality[91] and integrity[92] of medical records;

(b) Protect against any reasonably anticipated threats or hazards to the security or integrity of such information; and

(c) Protect against any reasonably anticipated use or disclosure of such information.[93]

44.1 A clear policy on measures to secure the confidentiality, integrity and accountability of electronic medical records must be formulated

1027. The HITSPS is silent on the issue of measures (generally) to protect the confidentiality, integrity and accountability of EMR.[94] The HITSPS relates *only* to a narrow subset of "sensitive information" and even then, provides very little detail on control measures for "sensitive information".

1028. Given the importance and sensitivity of the PII contained in EMR, it is important to have a comprehensive policy document that applies to the protection of EMR. This policy must document and make clear the measures that are in place to protect the EMR. We elaborate on some key measures that should be addressed in the policy, in the following sections.

44.1.1 Role-based access for front-end users

1029. The policy should provide for limits on access, and provide screening controls so that only authorised staff can access patient data. Role-based access

---

[91] Confidentiality means the property that data or information is not made available or disclosed to unauthorised persons or processes.

[92] Integrity means the property that data or information have not been altered or destroyed in an unauthorised manner.

[93] 45 CFR (US) § 164.306: Security Standards: General rules, of the U.S. Health Insurance Portability and Accountability Act of 1996 (HIPAA), that provides data privacy and security provisions for safeguarding medical information.

[94] COI investigations did not uncover any other policy document covering this issue either.

control helps to restrict EMRs to users who are made members of a certain role according to their responsibilities (*e.g.* doctor, nurse, clinician *etc*) or corporate position. Role-based access is already in place, but the classes of persons to whom access is granted, the extent of the access granted, should be reviewed as part of the wider post-Cyber Attack review. The Committee notes SingHealth's perspective that the "*implementation of IT projects is meant to serve, support and improve patient care, and that an appropriate balance will have to be struck when assessing the feasibility of IT projects*".

1030. The policy must establish clear access controls including:

(a) Role-based security that restricts access to information based on pre-established categories of patients, duties and documents based on specific job requirements of the user; and

(b) Tagging of sensitive data with status indicators that enable restriction of identified patients and encounters to only those with permissions to access such data.

1031. In short, the policy should follow the principle of *least access* – that is, staff should have access only to the resources they need to perform their daily tasks, and no more. Access to confidential data should be on a strict, need-to-know basis. Further, there should be no *general* access to patient data – staff should only be able to access the data when they need it for a specific purpose, and the scope of the data accessed should be tightly controlled to include only data essential to the completion of the task.

44.1.2 Database-level access by administrators, developers and support team

1032. Security measures should not only be geared towards external attackers – there is a real risk of patient data being compromised by insiders too. We recommend that the need for administrators, developers and support team to access patient data be reviewed. IHiS should aim for the least number of people possible to have access to the database. To the maximum extent possible,

administrators, developers and support team should not be able to view actual patient data. Currently, IHiS staff such as database administrators are able to access medical records. The only control is that any access by such personnel is logged for audit purposes. This is insufficient, because it does not *stop* access, and by definition, the logs would only be useful to show that access had already taken place.

1033. Administrators should have only the bare minimum privileges they need to do their job, and only during periods while they need access. The policy should adopt best practices for database security:

(a) Unused accounts must be deleted.

(b) Shared accounts should be prohibited – While administrators may find sharing passwords convenient, doing so makes proper database security and accountability almost impossible.

(c) Grant privileges to administrators, developers and support team only to the extent needed (read only *vs* insert/delete records, for example).

(d) Access by administrators, developers and support team must be controlled/restricted to only the tables to which they need access.

(e) A system for managing privileged accounts should be in place to provide authorised users with a temporary password with the privileges they require each time they need to access a database.

44.1.3 Logging policy and audit trails

1034. The EMR system must document and keep up-to-date logs and maintain an audit trail of authorised access to the system by users. This means it must record how medical records are accessed, by whom, what information was accessed, and when. That way, security personnel can quickly investigate if they suspect an insider was involved in a data breach. As shown in the Cyber Attack,

an external actor can also obtain credentials and masquerade as an authorised insider. Logging of access to the EMR from the front-end client can also therefore be essential to investigating unauthorised access by external attackers.

1035. This can be accomplished through the use of audit trails which allow organisations to precisely monitor who has accessed patient information by tracking all system activity, modifications, generating timestamps for entries, listing what was viewed, for how long, and by whom. Alerts can then be set to flag unusual activity.

1036. Although it appears that IHiS did have some policy for logging access, this was not reduced to writing. Audit trails were in place for access to the SCM medical records, and for sensitive records in particular. IHiS should rationalise which systems are subject to audit trails and reduce the policy to writing, so that it is clear and any gaps in coverage can be identified. Further, the policy should also detail what logs are kept, and how long they are kept.

44.1.4 Rate limiting

1037. Rate limiting refers to controlling the number of medical records that can be accessed by a user at one time. It appears that IHiS did have some sort of rate limiting policy in place, but it appears not to have been documented. Dr Chong testified that, when the SCM was initially procured, it was decided that if more than a certain set number of records were accessed at the same time, an alert would be sent to the IHiS security team, and the Cluster IT and Operations teams.

1038. Again, the existing policy should be reduced into writing, so that there is clarity about its requirements and scope. Any gaps can then be identified and addressed.

44.1.5 Tagging of sensitive data

1039. We cover the topic of sensitive data later in this Recommendation. For the moment, it suffices to say that the HITSPS is silent on this issue. The policy, again, should be formalised so that gaps can be identified and addressed. The written policy should also address the issues raised in the following sections.

44.2 Databases containing patient data must be monitored in real-time for suspicious activity

1040. Bulk queries during the Cyber Attack were not detected by any monitoring systems and came to light only by chance, when it was noticed by an alert employee (Sze Chun). Monitoring for such queries, which are indicative of unauthorised data harvesting, must be implemented at database-level.

1041. On 4 July 2018, Sze Chun noticed that an unusual query had been run. Sze Chun was aware that the SCM front-end application does not allow for bulk queries. Bulk queries in and of themselves would therefore have been suspicious. However, the bulk queries run from 27 June to 4 July 2018 had not been picked up because there was no mechanism in place to detect bulk queries to the SCM database.

1042. These queries were repeatedly run a few minutes apart over several days. Given the frequency of the attempts and the large number of records sought, it should have been clear that there was no legitimate reason for these queries.

1043. It is recommended that a system of database activity monitoring ("**DAM**") be implemented. DAM is the process of observing, identifying and reporting a database's activities in real-time. DAM tools help in detecting unusual and unauthorised, internal or external activities and will serve in the prevention and protection of sensitive data from intruders.

1044. DAM solutions possess the following capabilities:

(a) Monitoring of database. These tools audit database activity on a 24/7 basis in real-time. DAM monitors the activity of:

(i) Privileged users (including database administrators and system administrators), to ensure that data is not accessed or modified without authorisation;

(ii) Users, to check for unusual or malicious activity; and

(iii) User accounts, to check if the accounts are dormant or inactive.

(b) Attack prevention. DAM also helps to prevent attacks by:

(i) Providing alerts in real-time to notify security personnel of suspicious activity detected; and

(ii) Blocking attacks in real-time, based on recognition of known database exploits and unusual patterns of activity.

(c) Auditing for forensic investigations. DAM solutions are able to track the source of data leaks by recording the who, what, when, where and how of every query and identifying which records exactly have been exposed.

1045. Following the Cyber Attack, IHiS procured a DAM solution. This solution is capable of detecting anomalous database activity, like bulk queries, and can automatically trigger alerts or block the activity. IHiS is still testing it before rolling it out fully, as there are concerns about whether the implementation of the DAM solution will negatively affect the performance of the IT systems, either by causing lag or by triggering too many false positive alerts. Although there

appears to be no set timeline for the rolling out of the DAM solution at present, this is a positive step forward and should be encouraged.

44.3 End-user access to the electronic health records should be made more secure

1046. Although the attacker compromised the A.A. account in this case and was able to retrieve patient data in bulk by querying the database directly, there is also a significant risk of an attacker using stolen credentials to access the EMR *via* the front-end client, masquerading as a legitimate user, and carrying out targeted retrieval of medical records of specific pre-identified individuals. This would not trigger alerts tied to the volume of records retrieved.

1047. More rigorous authentication methods should therefore be considered. Because passwords are so vulnerable, requiring people to use at least two forms of authentication – *e.g.* a password and token – to access the EMR would appreciably enhance protection against unauthorised access. A multifactor authentication process would make it significantly harder for an attacker to impersonate a user, even if the primary password has been exposed. Experts Dr Lim, Gen. Alexander, Vivek and Richard all concur with the recommendation to implement two-factor authentication ("**2FA**").

1048. Gen. Alexander testified that 2FA has been successfully implemented in a number of health services in the USA, including Centura Health, UC Health, National Institute of Health, and Raleigh Regional Hub. Gen. Alexander also said that it is possible for 2FA solutions to be extremely quick, and to enable a one-time log in process, such that once logged in, medical personnel can carry on accessing the EMR while walking around the wards. Vivek has said that implementation of 2FA would not necessarily be too onerous, and could be accomplished by simply issuing smart ID cards to users, which is already done in the government context.

1049. IHiS and the PHIs have valid concerns that the implementation of 2FA will be burdensome and may slow down or otherwise negatively impact the provision of healthcare services. Vivek recognises that there may be challenges in patient care and/or other operational impact with the implementation of 2FA on corporate user accounts. 2FA is also not foolproof; there are vulnerabilities in 2FA platform itself which can be exploited, and it needs to be "*monitored with a hawk eye*".

1050. Nevertheless, given the importance of security and the effectiveness of 2FA as a security control, it should still be implemented where patient safety is not affected. For example, while the emergency room may not be an appropriate place for 2FA, 2FA might be implementable in normal wards. As Vivek and Gen. Alexander have noted, depending on the exact solution chosen, the disruption to existing workflows can be minimised to a large extent.

1051. IHiS and the PHIs should very carefully consider which roles must be exempted from the requirements of 2FA. Security cannot be sacrificed simply for the sake of expediency and convenience. Any exception to the normal 2FA policy creates a weakness that can be exploited. Vivek gave the example of a company where just 13 out of 45,000 users were not required to use 2FA, and an attacker managed to locate their identities and use their accounts to break into the system. The Cyber Attack cannot be viewed as a one-off. The number of breach incidents in healthcare continues to grow about 10 percent each year according to Symantec.[95] Taken together, it is clear that cyber attacks pose a clear and present danger to PHIs, and it would be foolhardy to forgo security simply for the sake of convenience.

1052. 2FA should thus be implemented for PHIs. The Committee notes the MOH family's concerns that the implementation of 2FA on corporate user accounts will pose patient safety issues. An independent study should be carried out on the jurisdictions that have successfully implemented 2FA for PHIs, to

[95] Symantec 2018 Internet Security Threat Report – Executive Summary for Healthcare Professionals.

learn how patient safety concerns were dealt with and disruption to provision of medical services minimised.

44.4 Measures should be considered to secure data-at-rest

1053. In the Cyber Attack, the attacker was able to view the full details of the medical records stored in the SCM database, once he had gained access. This was so as there were no measures in place to secure the data-at-rest in the database.

1054. Data-at-rest refers to information stored in databases in filesharing servers, in backup tapes *etc*, and generally includes any data that is not being transmitted through a network (which is known as data-in-motion).

1055. The amount of data that is being generated daily continues to increase exponentially. Given the rapid pace of development of cyber attacks, data-centric security measures must be deployed. These measures include safeguarding the data itself as it resides in repositories such as databases.

1056. In general, mechanisms to protect data involve coding data in such a way that access to the data is restricted. This process can generally be referred to as "masking"[96] and can occur at the central record repository. Techniques used to mask information in a patient's medical record include data encryption and tokenisation.

(a) Encrypting data-at-rest prevents unauthorised access by anyone who defeats normal system access controls. It alters the content of the data and stores it in encrypted form. This makes health data unreadable unless an individual has the necessary key or code to decrypt it. This would ensure that unauthorised individuals are not able to see the data in its original form. Dr Lim has recommended encrypting all data-at-rest, where possible, to protect against both

[96] Data masking is the process of hiding original data with random characters or data.

internal and external malicious actors. Dr James Yip ("**Dr Yip**"), MOH's Chief Data Advisor, also testified that it would be possible to encrypt patient databases, and provide tiered access to the decrypted data.

(b) Tokenisation [97] also prevents unauthorised access to selected columns[98] of data. Tokenisation can be used as an alternative to encryption on a column-by-column basis. Even if a database is compromised, tokenising PII (personally identifiable information, such as name and NRIC number) would effectively frustrate an attacker's ability to query for the medical records of specific individuals. Dr Lim testified that even if the data cannot be wholly encrypted, key information can at least be anonymised and hashed. Even bulk downloads of medical records would provide the attacker with no means of ascertaining who the individual records relate to. As the full medical record is not encrypted, there would be less performance-overhead related issues, as compared with encryption.

1057. It is acknowledged that encryption and tokenisation of data may have some impact on the operations of the PHIs, in terms of speed of access to patient records. However, such adverse impact should not be presumed without further study. As before, security should not be sacrificed merely for convenience, given the high-threat environment that exists today. Implementation needs to be carefully handled to minimise disruption to operations. An independent study should be conduct on the feasibility of implementing these measures in the EMR systems of the PHIs.

---

[97] Tokenisation is the process of substituting a sensitive data element with a non-sensitive equivalent, referred to as a token, that has no exploitable meaning or value. The token is a reference that maps back to the sensitive data through a tokenisation system.

[98] In a relational database, a column is a set of data values of a particular type (*e.g.* NRIC, Name *etc*.)

44.5 Controls must be put in place to better protect against the risk of data exfiltration

1058. In many cases, victims of cyber attacks are not aware that the sensitive data is leaving their systems because their data outflows are not monitored. The movement of data across network boundaries must be carefully scrutinised to minimise its exposure to attackers.

1059. CSA's analysis of the network logs revealed that the main bulk of the traffic between SingHealth's network and a malicious IP address was from Workstation A between 27 June to 4 July 2018.

1060. This unusual network activity went undetected until after 10 July 2018. Typical use of workstations does not involve the uploading to the internet of anywhere near as large quantities of data, and constituted a clear red flag that could have been detected, had the right controls been in place at the time.

1061. The Committee accepts CSA's recommendation that a Data Loss Prevention ("**DLP**") solution should be implemented to prevent such occurrences in future. DLP solutions detect potential data breaches/data exfiltration transmissions and prevent them by monitoring, detecting and blocking sensitive data while in-motion. DLP helps to prevent end users from sending sensitive or critical information out of the corporate network.

1062. Alerts and/or blocking can be set based on either the volume of data being sent out, or the content. It is possible, for example, to prevent data from being transferred out of the network or even out of endpoints. DLP solutions typically have a degree of machine learning capability and are able to, in conjunction with the rules set manually, determine what constitutes unusual activity and block it, or trigger alerts to relevant personnel for a response. DLP solutions are already used widely among many enterprises.

1063. The Committee notes the MOH family's concerns about the effectiveness of DLP solutions in the healthcare context, where most parts of its IT network

relate to sensitive data; and that such solutions do not work well with encrypted data. The MOH family should study the feasibility of implementing data encryption and if such a solution is assessed to be not suitable, then the DLP solution should be implemented.

44.6 Access to sensitive data must be restricted at both the front-end and at the database-level

1064. There is no written policy relating to IHiS' treatment of access to and monitoring of sensitive EMR. The HITSPS is silent on this issue.

1065. Front-end controls. The SCM application supports the tagging of sensitive data within its system. For these tagged patients, only selected users are allowed access to the medical records. All instances of access to sensitive data are subject to logging and alerts.

1066. The current approach enforces security of sensitive records through a corrective mechanism: authorised persons have almost unrestricted access to the records, but there is a strict *ex post facto* audit process for inappropriate accesses. This process is purely retrospective, as it occurs after damage may have been incurred. Particularly when an attacker has stolen credentials and is masquerading as an authorised user, an *ex post facto* audit process would be ineffectual in preventing breaches of sensitive data. Requiring 2FA to access the EMR, and sensitive medical records in particular, would significantly reduce the risk of such records being compromised.

1067. Database-level controls. During the Cyber Attack, there were no database-level controls that would have restricted the querying of sensitive data using SQL commands. This was a significant omission in the security of the SCM database, and was exploited by the attacker, who ran multiple queries to retrieve medical records of PM Lee.

1068. Leong Seng testified that a DAM solution, which is currently being tested by IHiS, is capable of monitoring and blocking attempts to access specific

records in a database. As a priority, DAM must be implemented for those entries tagged as sensitive.

1069. Similarly, even if encryption and tokenisation cannot be applied to all databases wholesale for performance reasons, steps should nonetheless be taken to encrypt or tokenise sensitive data. This is because such data constitutes an obvious high-value target for attackers. All the measures we have proposed including encryption and tokenisation apply with particular urgency to such sensitive data.

45 RECOMMENDATION #10: DOMAIN CONTROLLERS MUST BE BETTER SECURED AGAINST ATTACK

#PREVENTION #VIGILANCE

1070. Protecting CII in a Windows network environment necessarily requires protection of other components of the network. Domain controllers in particular must be secured, as compromise of a domain controller can lead to extremely serious consequences for the entire network.

1071. Windows domain controllers host the Active Directory Domain Services ("**AD DS**") database, in addition to providing the services and data that allow for effective management of servers, workstations, users, and applications. If privileged access to a domain controller is obtained by a malicious user, he has full control over the entire Windows domain and servers. The malicious actor can then modify, corrupt, or destroy the AD DS database and, by extension, all of the systems and accounts that are managed by active directory.

1072. An external consultant observed the following during the penetration test conducted on the H-Cloud in FY16:

> Domain Admin has full control on the servers in the network domain of the organization including creating administrator accounts in any local servers. By default, a Domain Admin account holder has complete unrestricted access to all resources in the entire network.
> By gaining Domain Admin access in an organisation, the following damages could happen:
>
> - Install ransomware to lock down the data.
> - Access, tamper, destroy organizational IT resources.
> - Create any number of accounts and grant them admin access in the Active Directory, such as OUs, admin accounts/Groups, *etc.*
> - Place time-bombed malicious software on any domain-joined machine.

- Since it's a root access, he/she can turn off, disable, bypass any additional security measure that might be put in place to prevent he/she accessing other resources.
- Add new account into Domain Admin group and use those new accounts permanently without being detected (unless domain admin list is constantly being reviewed)

1073. IHiS was made aware of these risks by May 2017, when the FY16 GIA Audit Report was released. This was more than a year before the Cyber Attack. Nonetheless, these weaknesses were not adequately addressed, and the evidence points strongly to compromised domain controllers having played a key role in the Cyber Attack. Further, domain administrator accounts, like the D.A. account, had been compromised during the Cyber Attack.

45.1 The operating system for domain controllers must be more regularly updated to harden these servers against the risk of cyber attack.

1074. Ideally, the operating system ("**OS**") for all servers should be kept up to date. However, it is accepted that this may not be feasible given the sheer number of servers involved. Nonetheless, priority must be given to domain controllers when rolling out OS upgrades. Domain controllers play a critical role in a Windows network as they are the servers that function as a detailed map of the network and set the basic rules that determine which users are allowed access to which systems.

1075. The use of older OSes means that vulnerabilities that have been addressed in newer versions of the OS can still be exploited. Vivek testified that in general, newer versions of an OS are more secure than older versions, as they benefit from developments in secure coding practices. In today's IT world, ignoring the security risks inherent in outdated server infrastructure and operating systems is tantamount to disregarding the obvious. Failure to upgrade weakens the ability to respond to the changing cyber threat landscape, and results in the inability to

provide protection against new and emerging threats, which more up-to-date versions of the software may have addressed.

1076. Software infrastructure (and critical server infrastructure, in particular) must be modernised in order to attain adequate levels of protection – not just once as in the case of Y2K, but continuously. Continued reliance on older, and more easily compromised computer infrastructure running OS versions that cannot be patched to address critical vulnerabilities, creates an unacceptable level of risk where infrastructure supporting CII systems is concerned. Methods to 'hack' and compromise older systems are well documented and widely distributed through the internet, social media, and hacking forums. Continuing to use such OSes exposes the domain controllers to targeted exploits.

1077. It is acknowledged that upgrading is a time- and resource-intensive process. Resource constraints notwithstanding, the pace of upgrading is really a question of assessment of risk, *prioritisation*, and management buy-in. This makes it important that such issues are also raised to the attention of senior management, so that appropriate appreciation of risk can be made, and support given where needed to push through with upgrading. Given the severity of the risk involved, it is incumbent on IHiS management to make time and allocate the required resources to ensure that domain controller OSes are kept up to date.

45.2 The attack surface for domain controllers should be reduced by limiting login access

1078. During the Cyber Attack, the attacker accessed domain controllers from the SingHealth end-user zone using RDP. The fact that domain controllers were accessible *via* RDP unnecessarily increased the attack surface. In general, insufficient network segregation increases the surface that can be exploited by attackers, and correspondingly increases the risk level of the network.

1079. This problem should be addressed by prohibiting remote connections to the domain controllers *via* RDP and other remote management solutions. Access to domain controllers should be limited to dedicated workstations, which would

be made available on a needs-only basis for the performance of administrative tasks. These workstations should be isolated from the internet and have no email access, to further limit the attack surface. IHiS is considering implementing such measures, and going further to physically limit access to such dedicated workstations by placing them in secured server rooms. This is a positive move and should be encouraged.

45.3 Administrative access to domain controllers must require two-factor authentication

1080. Passwords alone are insufficient protection for domain controllers. Given the importance of domain controllers to the network, and the various ways in which passwords may be acquired by attackers, it is crucial that 2FA be implemented to protect the domain controllers against attackers who have already managed to obtain passwords. Experts Dr Lim, Gen. Alexander, Vivek and Richard all concur with the recommendation to implement 2FA for servers.

1081. With 2FA in place, any attacker would be prompted for a second factor during the authentication process. MOH family accepts this. This second factor would need to be provided in addition to the user's password for the attacker to successfully authenticate and gain access as that user. Since that second factor is based on something that the user possesses (either a device, an account, or token), this would offer a good level of protection against this type of attack where the password is compromised in some way.

46 RECOMMENDATION #11: A ROBUST PATCH MANAGEMENT PROCESS MUST BE IMPLEMENTED TO ADDRESS SECURITY VULNERABILITIES

#PREVENTION #VIGILANCE #GOVERNANCE

1082. The initial entry to SingHealth's network was likely by way of a phishing email containing malicious code. The attacker was able to compromise Workstation A that was running Microsoft Outlook (**"Outlook"**), which was vulnerable to a publicly available hacking tool. The attacker then used the tool to drop malware onto Workstation A, which was subsequently used to escalate the attack. CSA assessed Workstation A to have been a key pivoting point in the overall scheme of the attack.

1083. In fact, a patch[99] for Outlook, that would have rendered the hacking tool ineffective, had been made available by Microsoft in late-2017. However, this patch was not installed on workstation A as at 1 December 2017, when the malicious code was executed. The failure to patch in a timely fashion essentially led to the success of this phase of the attack. This constituted a missed opportunity for IHiS which, if addressed, would have stopped or significantly arrested the progress of the attack.

1084. To avoid attacks through known issues or vulnerabilities, systems should be fully up to date with the latest security patches. A robust security patch management process must be implemented as a critical component in maintaining the security of SingHealth IT systems. Patching is of critical importance in a networked environment. Patches do not only ensure the security of individual devices, but also that of the network as a whole. This is because the security of a network is only as strong as its weakest link – it only takes one unpatched device for an attacker to get into a network, and from there, to move laterally through the network towards his objective. As such, a failure to patch

[99] A patch is a piece of code that can be applied to a software program after it has been installed.

has ramifications that extend well beyond the security of any individual unpatched device.

46.1 A clear policy on patch management must be formulated and implemented

1085. Patch management is the process of identifying, acquiring, installing, and verifying patches for software and systems. Patches correct security and functionality problems in software and firmware. From a security perspective, patches serve to mitigate software vulnerabilities. Applying patches to eliminate these vulnerabilities can significantly reduce the risk of exploitation.

1086. A *detailed* policy must be formulated to put in place a rigorous and timely patching regime. In doing so, reference may be made to best practice documents such as the NIST Guide to Enterprise Patch Management Technologies[100] and technical reference papers issued by other jurisdictions. The Solicitor-General referred to the Government of the Hong Kong Special Administrative Region's ("**HKSAR**") paper on Patch Management.[101]

1087. Organisations should have clear and stringent patch management timelines, and adhere to these timelines to ensure that security patches for IT systems are tested and implemented in a timely manner. This will minimise the window of opportunity in which attackers can exploit system vulnerabilities to perform malicious activities. The written policy should make clear that patch management is not merely operational in nature but is integral to a defence-in-depth [102] strategy, where patching represents one layer of a multi-layered

[100] NIST.SP.800-40r3.

[101] This paper provides a core set of principles and methods that can be used as a reference in putting together an effective patch management programme.

[102] The fundamental principle behind defence-in-depth is that no single security product is foolproof and that an organisation should be required to have several layers of security in place. This was also discussed earlier under Recommendation #1.

approach to security countermeasures. Importantly, any proposed deviation from policy, should be brought to the attention of senior management, so that a decision can be made at the right level, after weighing all relevant considerations.

1088. Core elements that should form part of the patch management policy are detailed in the following sections.

46.1.1 Maintenance of an organisational-level software inventory[103]

1089. The policy should require that an accurate inventory be maintained of all software packages, along with version numbers of those software packages. This inventory would help administrators better monitor and identify vulnerabilities and patches that are applicable across the organisation.

46.1.2 Vulnerability identification and patch acquisition

1090. The policy should require administrators to refer to a number of information resources in order to monitor vulnerabilities and patches that may be applicable to the installed software systems. As each type of resource has its own specialised area, administrators need to be able to refer to more than one source for accurate and timely information on new vulnerabilities and patch releases. Common resources include product vendor websites and third-party security advisory websites (run by CERTs and security vendors). There is no evidence that any such pro-active monitoring is currently carried out by IHiS, beyond rolling-out patches made available by product vendors for the various software systems.

46.1.3 Patching timelines

1091. Software security patches which fix security vulnerabilities and other bugs for software installed on SingHealth and IHiS issued endpoint devices (*e.g.* operating system software, application software) are applied to on a specific

[103] Government of HKSAR's paper on Patch Management at p4.

posting cycle, with emergency or critical security patches, such as the WannaCry patch, applied as soon as possible, outside the patch cycle.

1092. A specific posting cycle in deploying patches is essentially patch bundling, which has a downside – it lengthens the time from when a patch becomes available to the time the vulnerability is fixed on the unpatched systems. If an attacker exploits the same vulnerability before the patch is installed, the delayed patching is clearly detrimental. The attacker effectively has a longer window of opportunity to exploit the vulnerability because of the delay in installing the patch. This is all the more so as the release of a patch may provide attackers with the information that they need to exploit the corresponding vulnerability (*e.g.* reverse engineer the vulnerability from the patch), meaning that a newly released patch might need to be applied immediately to avoid the vulnerability it is designed to address from being exploited.

1093. It is imperative that timelines for deploying patches are actually adhered to on the ground. The policy *cannot* simply be a theoretical framework. The importance of timely patch management cannot be overstated. The draft HITSPS Version 4.0 prescribes a two-week timeline for implementation of patches for endpoint machines. If this had been in place in 2017, Workstation A would have been patched against the publicly available hacking tool well before the malicious code was executed in December 2017.

46.1.4 Risk assessment and prioritisation

1094. The policy should acknowledge limited resources, which make it unfeasible to roll-out all patches immediately, and address the fact that administrators will need to prioritise the deployment of new patches, by performing a risk assessment to determine which systems, and which software, should be patched first.

1095. In general, this prioritisation should be based on the following criteria:

(a) **Threat** – A threat is any potential direct danger to information systems, or software that is exposed to a higher degree of risk (*e.g.* by virtue of its exposure to the internet). Examples of systems facing high threat levels are web servers, email servers and applications, and servers containing sensitive information. Special focus must be placed on patching of email applications, as email attacks are now the most common vector for initial intrusions into systems.[104] Indeed, in this case, CSA's hypothesis is that the initial infection originated from a phishing email.

(b) **Vulnerability** – A vulnerability signifies the absence of, or a weakness in, a safeguard which could be exploited by an attacker. It could be outdated software which is less secure *etc*.

(c) **Criticality** – This is a measure of how important or valuable a system is to operations. For example, database servers and network infrastructure would be considered more critical to operations.

1096. Systems facing more threats, or that are more vulnerable, or are mission-critical should be accorded a higher priority in the patch management process. MOH family's view is that patching should be carried out comprehensively for all assets connected to the network, in a manner which poses the least cybersecurity exposure.[105] Should a patch be assessed to be less urgent or critical, steps should be taken to mitigate any exposure before the patch is deployed. In general:

---

[104] SANS Institute, "Securing Against the Most Common Vectors of Cyber Attacks", SANS Institute Reading Room, August 2017.

[105] This section maps to CIS Control 3 "Continuous Vulnerability Management" and CIS Control 8 "Malware Defences". Comprehensive patching of all assets connected to the network greatly mitigates the risk associated with unpatched machines.

(a) Patches addressing security concerns should take priority over patches dealing more with functionality;

(b) Patching of software which provides a broader attack surface (*e.g.* with a connection to the internet) should take priority over patching of software with a more limited attack surface (*e.g.* software used internally to a network);

(c) The *type* of security vulnerability should also play a part in determining the priority with which patches are applied. For example, patches addressing vulnerabilities related to remote code execution on internet systems, like email applications, should be higher priority; and

(d) Similarly, the type of *application* should also be a criterion for priority. For example, email applications should be given a higher priority for patching, as email is the most common attack vector.

46.1.5 Patch testing

1097. IHiS, in practice, carries out patch testing before the patches are deployed. The practice notwithstanding, the patch management policy should be explicit in addressing this issue.

1098. Patch testing is vital to ascertain whether or not a new patch will affect the normal operation of any existing software. Patch testing[106] should consist of the following:

---

[106] Vinod Mohan. (1 Aug 2013). <https://thwack.solarwinds.com/community/solarwinds-community/geek-speak/blog/2013/08/01/why-should-you-test-patches-before-deployment>

(a) Simulate test cases and check if the patches are getting deployed successfully on the target platform(s);

(b) Compare application performance before and after the patch deployment and check if there are any issues;

(c) Test if other applications running on the target environment are impacted by the patch update;

(d) Ensure that if the patch is successfully removed, no application or system issues will occur; and

(e) Incorporate patch testing as part of IT security risk assessment plan.

1099. There should be clear and stringent patch testing timelines, and a means to ensure that these timelines are adhered to.

1100. In addition to identifying any unintended problems, patches themselves should ensure that they have fully addressed the vulnerability in question or corrected the performance issue as intended.

1101. If it is not feasible to install the patch because, for example, testing results show that the patch will crash or seriously disrupt the production system, *alternate security controls* should be implemented and monitored for signs of the unpatched system being exploited.

1102. The Committee notes that MOH is committed to ensure that patches are effected in a timely way which minimises cybersecurity and operational risks.

46.2 The patch management process must provide for oversight with the reporting of appropriate metrics

1103. Once again, it must be highlighted that patch management cannot be a merely theoretical exercise. Processes must be in place to ensure that patch management policies are understood and complied with. In this regard, it is important for there to be a system for the recording of patch management metrics, and regular checking of said metrics to ensure that patch management policies are effective.

1104. It is almost impossible to set appropriate patching objectives and check if said objectives have been achieved without using a set of appropriate metrics. The metrics will also offer a wealth of information to security staff, and allow them to communicate more meaningfully with management and others about the status of the organisation's patch management policies. The status of an organisation's patch management must be measured using objective metrics, and cannot be left to subjective and unreliable judgements about the efficacy of implementation efforts.

1105. IHiS should undertake a comprehensive review and determine what metrics would be meaningful and feasible to track and regularly analyse. At a basic level, the following metrics with clear timelines should be considered:[107]

(a) Number of machines scanned;

(b) Number of machines not scanned;

(c) Number of patches found; and

(d) Number of patches not found.

---

[107] SANS Institute, "Patch Management and the Need for Metrics", SANS Institute Reading Room, July 2004.

1106. Furthermore, the collected metrics and analyses thereof should be subject to regular management oversight. IHiS should review and determine which body would be the most appropriate to have oversight of this function. The policy should set out explicitly what the lines of reporting are, who has responsibility for reporting, and how regularly reports on metrics should be issued. It is suggested that there be two concurrent lines of reporting to:

(a) Director, Delivery Group – This is to ensure oversight of the personnel managing the systems and applications, as they should be the ones with the primary responsibility to ensure that patches are applied; and

(b) Lead, SMD – This is to ensure oversight from a security perspective, so that there can be heightened security monitoring even as systems and applications are waiting to be patched, and also so that generally, vulnerabilities and lapses can be picked up and addressed by staff with a dedicated cybersecurity portfolio.

47 RECOMMENDATON #12: A SOFTWARE UPGRADE POLICY WITH FOCUS ON SECURITY MUST BE IMPLEMENTED TO INCREASE CYBER RESILIENCE

#PREVENTION #VIGILANCE #GOVERNANCE

1107. A software[108] upgrade is a newer or better version of the software, in order to bring the system up to date, which typically offers a significant change or major improvement over the current version. OS upgrades in particular can make significant changes to a system in functionality, security, user interface *etc* over the previous version.

1108. In CSA's assessment, outdated software was a contributing factor to the Cyber Attack. For instance, there was a vulnerability in Microsoft Outlook which was exploitable by a publicly available hacking tool, which allowed the attacker to install malware on compromised workstations. Microsoft Outlook is part of the software package, Microsoft Office. As at August 2017, when the initial infection took place, only a few workstations in SingHealth were running an updated version of Microsoft Office, while the majority were still running the vulnerable version of Microsoft Office.

1109. Vivek testified that in general, newer versions of software are more secure than older versions, as they benefit from developments in secure coding practices. IHiS and CII operators in general must actively update their software so that outdated and unsupported software, which significantly increase exposure to security risks, are replaced on a timely basis.

1110. Upgrading software allows systems to benefit from additional protections and ensures that systems have the latest security solutions to help limit the cyber

[108] The term "software" as used in this recommendation refers both to systems software and application software. Systems software includes the programs that are dedicated to managing the computer itself, such as the operating system. Application software includes programs that are used to complete tasks, such as creating documents, spreadsheets, and publications, doing online research, sending email *etc*.

threat. Malicious parties are continually innovating, devising new ways of attacking systems, and in response, the IT security industry has to find ways of reducing or eliminating this threat. However, systems can only benefit from the latest security tools and if the software is kept up to date.

47.1 A detailed policy on software upgrading must be formulated and implemented

1111. IHiS' policy documentation, and the HITSPS in particular, are silent on the issue of software upgrades. The *draft* version of HITSPS Version 4.0 provided in IHiS' evidence *also* omits any mention of a software upgrade policy.

1112. A detailed policy must be formulated to make clear that security is an important consideration when determining if and when software should be upgraded, and how such upgrades should be prioritised.

1113. We set out the core elements that should form part of the software upgrade policy in the following sections.

47.1.1 Maintenance of an organisational-level software inventory

1114. The policy should require that an accurate inventory be maintained of all commercial off-the-shelf software packages in use by the Clusters, along with version numbers of those software packages. This inventory would help administrators better monitor and identify which endpoints require software upgrades to be rolled out, when the decision is made to do so.

47.1.2 Planning process for upgrades

1115. The policy should provide for a continuous planning process that does not only involve operational IT staff. For example, a cross-functional team can be formed, comprising all stakeholders – users of the software from Clusters; IT operational staff; IT security staff; and IHiS/Cluster management. When a new version of the software is released, the team can map the new functions to the current system and the business processes that are affected to determine whether

the updates are worth incorporating; and must *further*, give adequate attention to and place due emphasis on security improvements in the upgrades. In essence, the team can make a holistic and comprehensive assessment of the implications of the upgrade on their respective areas of expertise, and then make a combined recommendation to Cluster management as to how and when the upgrade should be adopted.

47.1.3 Identification of upgrades significant to security

1116. There must be an identifiable individual or appointment holder taking current responsibility for every piece of software deployed, from a security standpoint. Those responsible for software must monitor relevant sources of information which may alert them to a need to act in relation to new security vulnerabilities.

1117. The policy should require the individual to closely review all software releases, and to critically assess whether there are security improvements which are significant. An upgrade significant to security would, for example, be one where a known vulnerability (*i.e.* one that has been publicised or one that has been exploited in a cyber attack) has been fixed. The functional improvements in the upgrade aside, a risk-assessment based approach should be adopted in assessing the software release. Staff involved in managing software must have experience, training or qualification commensurate with the importance of the software and risk levels involved. Staff involved must be aware of, and proactive in managing, information security-related risks associated with the software.

1118. There is no evidence that any such proactive security assessment of new software releases is currently carried out by IHiS.

47.1.4 Risk assessment and prioritisation

1119. The policy should acknowledge limited resources, which make it unfeasible to purchase and install all upgrades immediately, and address the fact that administrators will need to prioritise the deployment of new upgrades, by performing a risk assessment to determine which software should be upgraded first.

1120. In general, this prioritisation should be based on the following criteria:

(a) **Threat** – A threat is any potential direct danger to information systems, or software that is exposed to a higher degree of risk (*e.g.* by virtue of its exposure to the internet). Special focus must be placed on upgrading of email applications, as email attacks are now the most common vector for initial intrusions into systems.

(b) **Vulnerability** – A vulnerability signifies the absence of, or a weakness in, a safeguard which could be exploited by an attacker.

(c) **Criticality** – This is a measure of how important or valuable a system is to business operations. For example, OSes on domain controllers would be considered more critical to network security.

1121. Software facing more threats, or that are more vulnerable, or are mission-critical should be accorded a higher priority in the upgrade process. The outcome of this exercise would be to rate each individual piece of software in terms of priority (High, Medium or Low) for upgrades, as assessed from a security standpoint. For example, email software, which by its very nature involves communication with the internet, and which is a common attack vector for malicious exploitation, should be rated as High Risk. A standalone software system which is not connected to the Local Area Network would be rated as Low Risk.

47.1.5 Upgrade timelines

1122. The general (unwritten) policy of upgrading software at EOL should be reviewed. IHiS' approach was to consider factors like cost, user needs, proximity to EOL, and compatibility with existing environment, when deciding whether to upgrade software. Security was not one of the main considerations. This strategy may have been acceptable six or seven years ago. It is not today. This approach emphasises cost and operational ease over defence, and is at best naïve in the current cybersecurity environment.

1123. The longer a piece of software has been released, the longer malicious actors and security researchers will have to find vulnerabilities that can be exploited. Often, the exploitable entry points for commercial software are made public by researchers so that users can be made aware, and so that software companies can take steps to fix them. Software companies then release a new version of their software to address these security issues. Many cybercriminals track information about vulnerabilities. Once they find a new vulnerability, the criminals work as quickly as possible to develop an "exploit" to abuse the vulnerability. Using such an exploit, hackers can then target victims who have yet to update their software.

1124. Ideally, all information technology software applications should generally reflect the most recent version of the application software that is properly vendor-supported. Where this is not possible, as a rule of thumb, the installed version of the software should be no more than two versions behind the most recent commercially available version. This is because the longer software goes without upgrades, the longer the list of accumulated exploits to which it may be vulnerable.

1125. Separately, Vivek testified that recent years have seen software evolving with an increased focus on the escalating cybersecurity threats, and versions released recently are far more secure than those deployed eight, nine or ten years ago, when the constant threat did not exist.

1126. CII operators that need to maintain a high security posture cannot afford to ignore the dormant security vulnerabilities that lie waiting to be exploited in outdated software.

1127. No timeline can be fixed as to how quickly an upgrade should be installed, after it is released, as there are considerations such as availability of budget, size of the installed base that needs to be upgraded, and the length of downtime or disruption to operations. However, any enterprise-wide security plan that ignores planned upgrades to software is incomplete. Routine and regular software upgrades are an essential element in every security and risk mitigation plan, and a well thought-out upgrade strategy is a critical component of overall IT security. Upgrading software to make one's network more secure is not just a defensive strategy – it is a proactive one that protects one's business and provides necessary stability to one's network.

47.2 An appropriate governance structure must be put in place to ensure that the software upgrade policy is adhered to

1128. It bears repeating that the software upgrade policy, like all other written policies, cannot simply be treated as a theoretical exercise. It must be implemented and diligently enforced. As such, IHiS and Cluster management must put in place an appropriate governance structure to:

(a) Ensure that the software upgrade policy is adhered to;

(b) Ensure that security considerations are given due weight in decisions regarding software upgrades; and

(c) Ensure that any decision to forgo and upgrade or deviate from the upgrade policy is properly considered and documented.

1129. At present, it does not appear that any such structure is in place. For example, a security deviation form approved by Benedict for the postponement of certain software upgrades does not appear to have been escalated to anyone

else in Cluster management. Given the importance of ensuring security for software, it is suggested that there be a dual reporting structure to both IHiS and Cluster management, including:

(a) Lead, SMD (IHiS) – to ensure that security considerations are given adequate weight;

(b) Cluster Infrastructure Lead (IHiS) – to ensure that the upgrades are appropriate in the current environment;

(c) GCIO (Cluster) – for operational concerns; and

(d) Dy GCEO (Cluster) (or equivalent) – to ensure that Cluster management is apprised of and agrees to any upgrades or to forgo said upgrades. Alternatively, this role can be filled by a dedicated Cluster CISO, as we have proposed in Recommendation #1.

48 RECOMMENDATION #13: AN INTERNET ACCESS STRATEGY THAT MINIMISES EXPOSURE TO EXTERNAL THREATS SHOULD BE IMPLEMENTED

#PREVENTION #VIGILANCE

1130. Temporary Internet Surfing Separation (**"ISS"**) was implemented in SingHealth from 20 July 2018, and in the other two clusters from 22 July 2018. In this section, we will consider the expert opinions on (a) whether ISS should be lifted, and if so, (b) what are the alternatives to ISS, and (c) whether additional mitigating controls are required.

1131. The appropriate internet access strategy is an issue of risk management. It requires consideration of resources, demands, infrastructure constraints, and operational imperatives. It is thus a decision that should be undertaken by the healthcare sector, weighing the full range of considerations. MOH has not come to an official position on the appropriate internet access strategy, and has formed a horizontal committee to look into this issue, and weigh the balance between cybersecurity risks, patient safety, and cost.

1132. While this is an issue for the healthcare sector's ultimate decision, there are guiding principles that the healthcare sector should apply in determining its internet access strategy. First and foremost, we caution that the operational need for internet usage should not be conflated with the need for internet usage on the *same device* which has access to internal networks and databases containing confidential information (including the EMR). If the internet can be accessed on a separate device and/or *via* separate networks, the costs or operational drawbacks from an efficiency perspective of doing so must be balanced against the security gains. Second, while we accept that patient safety must be the predominant concern for the healthcare sector, it would be apposite for the healthcare sector to also bear in mind that inasmuch as patient safety relates to *treatment,* it also entails protection of patients' confidential and sensitive medical information and records. As recognised by the Minister for Health (Mr Gan Kim Yong), patient wellbeing "*includes safeguarding the confidentiality of patient*

data".[109] To the extent that patient confidence in the confidentiality of patient data is undermined, the ability of the healthcare sector to engage and provide healthcare to patients is correspondingly reduced. The trade-off may thus not be a simplistic one between binary objectives of patient safety versus cybersecurity, as both objectives may be twinned or interdependent. Third, as is already being considered by the healthcare sector, it would be sensible to make distinctions between different internet use-cases in the healthcare sector to determine where internet usage for work is really needed, where workarounds can be implemented, and what mitigating measures should be put in place where internet connectivity is permitted – a careful balance of all these considerations is needed for the healthcare sector to arrive at an optimal tiered internet access strategy.

1133. In the latter regard, the experts have recommended a tiered internet access strategy as follows:

(a) Where devices or databases do not need to be connected to the internet, they should not be connected. This recommendation should be implemented, as it will reduce the attack surface.

(b) ISS for all endpoints. While devices connected to internal networks and databases are isolated from the internet, internet usage for operational needs can be carried out on separate internet-surfing devices. Depending on user needs, this separate device can be either the user's personal device, or enterprise-procured devices.

(c) If ISS is unsuitable:

(i) where only one-way communication is required, there should be a unidirectional gate (*e.g.* data diodes) to prevent data leakage; and

[109] *Singapore Parliamentary Debates, Official Report* (6 August 2018) vol 94 (Ministerial Statement).

(ii) where two-way communication is required, internet isolation technology ("**IIT**") should be employed. IIT creates a secure remote infrastructure, so that information from the internet is transmitted to the endpoints in a "flattened" state, *i.e.* with all its macros removed. The technology combines Firewall and Content Threat Removal technology that cleanses all traffic that passes through it. If IIT is implemented, mitigating measures should be implemented simultaneously to address the residual risk.

48.1 Healthcare Sector's pre-Cyber Attack internet access strategy

1134. As early as 2015, CSA had made a security observation that hospitals, including those in the SingHealth cluster, had endpoints which could access both the internet and the EMR concurrently, and this could lead to medical records being accessed by unauthorised personnel *via* the internet. CSA thus made two recommendations: (a) to use a thin client to access the internet (understood by IHiS to mean IIT through a virtual browser ("**VB**") or remote browser ("**RB**")); or (b) to disable internet access from hospital terminals (understood by IHiS to mean ISS).

1135. At that time, ISS was not considered feasible, as internet usage enabled the PHIs' core operations, including patient care, education and research, and administration and operations. Hence, as consideration went into which staff did not really need the internet for work, IIT in the form of VB was concurrently studied as an alternative solution. After studying the VB solution, IHiS recommended the use of VB over ISS, as it would be less disruptive. IHiS trialled a proof of concept of a RB product and concluded that it would be an effective solution. In June 2017, CSA gave the conditional go-ahead for the RB solution, provided that mitigating controls were put in place to address the residual risks, and that Senior Management in the healthcare sector accepted these risks.

1136. Hence, by June 2017 the healthcare sector had already determined that internet access would be removed for staff that did not require the internet for

work and that internet access (save for certain high-risk sites) *via* IIT in the form of the RB solution would be provided to staff that required the internet for work. However, the status quo as at the time of the Cyber Attack was not acceptable – (i) internet access removal for staff that did not require the internet had only been implemented in some PHIs; (ii) there had been no firm decision on which staff really needed the internet for work; (iii) there had been no consideration of whether the RB solution should be deployed in the same or different device from which clinicians accessed the EMR, and the concomitant risks of either option; and (iv) in any event, the RB solution had not yet been rolled-out.

1137. While ISS was not the preferred solution for the Health sector pre-Cyber Attack, we note CE, CSA's evidence that public acceptance of cybersecurity measures changes over time, and in particular, after cyber attacks happen. The internet access strategy should thus be considered afresh, in the light of the Cyber Attack.

48.2 Benefits and drawbacks of Internet Surfing Separation

48.2.1 Benefits

1138. ISS prevents an attacker from gaining direct access into the CII systems that are providing essential services, and prevents any attacker that may remain in the network to callback (*i.e.* establish connections out) and steal further data.

1139. In the case of the Cyber Attack, despite a suite of containment measures, it was discovered on 19 July 2018 that the attacker was trying to re-enter the network. ISS was necessary to contain the threat and prevent further compromise. The implementation of ISS achieved this aim, and no further suspicious activity was detected in SingHealth's network thereafter. This safeguarded key public healthcare IT systems and confidential patient data. The benefit of ISS as a cybersecurity measure is thus clear.

48.2.2 Drawbacks

1140. MOH representative Dr Yip gave evidence that ISS has not been implemented across the public healthcare sector in any other country. There is thus no case study on the long-term effects of ISS on the public healthcare sector.

1141. Dr Yip also testified that temporary ISS created challenges in the provision of patient care, lowered efficiency in frontline and backend operations, created impediments in meeting reporting requirements, and resulted in constraints on research, education and innovation activities, such as:

(a) Back-end administration and hospital operations:

(i) The servicing and maintenance of hospital equipment has been affected, as remote troubleshooting and pushdown of software updates is no longer possible. Significant manpower and scheduling is required to perform troubleshooting and updates. There is also the risk of patches/updates being missed, which could in turn affect the quality of patient care and pose cybersecurity concerns.

(ii) Payment processing[110], as well as processes essential for business operations (*e.g.* procurement, payroll, staff claims processing).

(b) Reporting requirements:

(i) National registries cannot be directly accessed, and staff have to transcribe the information onto internet-surfing devices, and causes delay.

[110] With temporary ISS, payments *via* EZLINK are no longer available. Patients have to use alternative modes of payment.

(c) Teaching, research and education, and innovation:

(i) Temporary ISS has made it more challenging to conduct literature research for statistical and epidemiological support. It has also impaired teaching opportunities. Due to the proliferation of personal and research organisation laptops, there is also a new challenge of how to safeguard and monitor the data being stored on such laptops.

1142. In the short period of time after temporary ISS was implemented, workarounds have been deployed for all the work streams affected by ISS, and there was little impact on patient care in the short term. In particular, a key workaround was providing sufficient separate devices for internet-surfing. While this was no doubt resource-intensive for the public health institutions, it is now a sunk cost. Another key workaround relies on human effort, where staff have to be more meticulous in: (a) using internet-surfing devices to do their "last mile checks" before prescribing treatment; and (b) transcribing data onto/from internet-surfing devices; and (c) servicing and maintaining hospital equipment.

1143. On the implementation of ISS after the Cyber Attack, Prof. Ivy in her evidence said:

> "There has been some loss of productivity. People are working longer and harder. People are using their mobile phones and their own devices to do some of the work that they need to do, but there has been relatively little noise, I would say, about it. Even though I think there is hope that we will review it at some point, and, certainly, we work with MOH to look at this, but, you know, staff have taken it in their stride, because I think that the horror of our patients' data having been breached is an unacceptable risk at this point for us to even consider just opening to internet again. So we certainly hope the virtual browser platform, other solutions will come to play in the future. But I would say, at this point of time, there are

> inconveniences, workarounds, but nothing too major. I think we've continued to be able to deliver the care to patients that we need to."

1144. In the long run, Dr Yip's evidence is that there is nothing in MOH's healthcare transformation strategy that will be affected by ISS, provided that sufficient time, money, and effort are expanded to find workarounds.

1145. The real issue is how optimal these workarounds are. Dr Yip testified that the workarounds have come at the price of increased time and costs, loss in productivity and new risks, and in the long-term, may have adverse impacts, including manpower constraints and lower staff morale. We recognise these challenges, and note that the healthcare sector will have to balance this challenge against the cybersecurity risks.

48.3 Benefits and drawbacks of internet isolation technology

48.3.1 Benefits

1146. The experts were of the view that if the internet is required for operational purposes, IIT, such as VB or RB should be implemented. IIT isolates and executes all internet content in a secure browser located in sandboxes instead of the host machine, which eliminates the risk malware infection on the organisation's workstations and network. Risks of phishing are contained as phishing sites are prevented from delivering malware and harvesting private information. Further, even if the IIT platform is compromised, it can be easily restored to its last known proper configuration, which will prevent malware from spreading further and can also be used for intelligence gathering.

1147. When IHiS did a proof of concept of a RB solution, it found that this would be effective and viable as a secure internet access platform. RB is a purpose-built solution for organisations to securely access the internet using the concept of virtualisation.

1148. Dr Yip's evidence is that if VB or RB was implemented instead of ISS, this would go a very long way in helping clinicians do their work, depending on how the solution is deployed. Dr Yip noted that there were several permutations to how the solution is operationalised, for instance:

(a) VB or RB could be deployed in either the same or different device from which clinicians access the EMR.

(b) The content allowed in VB or RB has to be calibrated.

48.3.2 Drawbacks

1149. IIT is arguably less secure than ISS. CSA's view is that that while the remote browser solution does mitigate some of the risks of internet surfing, there are still risks that ISS mitigates that the RB solution does not. Whether there are any residual risks and what these risks are will depend on how the product implements the solutions. If VB or RB is implemented, there will need to be careful consideration as to what product is chosen, and how to calibrate the particular product.

48.3.3 Mitigating controls to address the residual risks

1150. As explained above, ISS prevents an attacker from gaining direct access into the CII systems that are providing essential services – it provides a high degree of security. At the same time, the evidence of MOH representative Dr Yip highlights the potential drawbacks – increased time and costs, lost productivity and new risks. If ultimately, the considered decision taken is to implement VB or RB instead of ISS, the healthcare sector must ensure that the residual risks of not implementing ISS are adequately addressed by strong mitigating controls. One mitigating control that was put in place before the Cyber Attack was internet-whitelisting. Another mitigating control, the ATP solution, was in the process of being deployed before the Cyber Attack. The containment measures implemented by IHiS after the Cyber Attack may also go some way to address the residual risks. These should be augmented with the other recommendations listed in this Part which the healthcare sector should carefully study.

1151. The Committee notes that IHiS is conducting a trial of a VB solution with a select group of hospital users to ensure smooth usability with minimal disruption to hospital operations. This trial will allow IHiS to study how VB can be deployed effectively, how functions within the healthcare setting where internet access is integral can be identified, and how to minimise impact and disruption on other systems.

49 RECOMMENDATION #14: INCIDENT RESPONSE PLANS MUST MORE CLEARLY STATE WHEN AND HOW A SECURITY INCIDENT IS TO BE REPORTED

#VIGILANCE #DETECTION #RESPONSE

1152. Employees should be trained on how to respond to security incidents so that they know what to do when an attack occurs. Without an incident response plan, it will be difficult to minimise the damage of a security breach as employees will be left to their own devices. Precious time can be lost trying to figure out what actions to take. Some malware infections spread at lightning speed as was seen in May 2017 with the WannaCry ransomware outbreak, where infections crossed borders and hopped between continents in a matter of hours.

49.1 An incident response plan for IHiS staff must be formulated for security incidents relating to Cluster systems and assets

1153. As mentioned before, IHiS' incident reporting processes are set out in the following documents:

(a) SIRF – translates the requirements of the NCIRF into the context of PHIs; and

(b) IR-SOP – cluster-level standard operating procedure for responding to security incidents.

1154. The SIRF is meant primarily for a sector-to-CII level, and it is for the Cluster CIOs and their IT leads to develop lower level processes to comply with its requirements.

1155. In relation to the IR-SOP, the reporting lines in the document begin with the Cluster ISO and GCIO, but there is no established procedure for reporting a security incident *to* the Cluster ISO or GCIO. There was no written protocol for how IHiS staff were to escalate a matter internally or determine when to report

to the Cluster ISO or GCIO. This was not covered in either of the documents mentioned above. As such, key front line personnel like Katherine, Lum, and Sze Chun, were unaware of who to report to when suspicious indicators were observed. This resulted in confusion and consequent delays in response. While a process for Cyber Incident Security Response was developed for IHiS staff, this relates to IHiS company systems, rather than to Cluster systems. In addition, no incident reporting process was developed for SingHealth officers.

1156. IHiS is not alone in this regard. A 2018 study[111] found that found that 77 percent of organisations surveyed did not have a formal security incident response plan. Almost half of the organisations indicated that their plan was either informal and *ad hoc*, or non-existent.

49.1.1 The need for an incident response plan

1157. An effective incident response plan is critical for all levels of employees, with specific plans in place for Cluster staff and IHiS employees. This is essential because it is not a matter of *if* a cyber attack will happen; it's a matter of *when*. As CE, CSA said:

> "[W]e need to assume the mindset that it is a matter of *when*, not *if*, our systems are breached. There is no such thing as "100% cybersecurity", and defending our cyberspace will be a ceaseless battle."

1158. The lack of an incident response plan increases the likelihood of security incidents going undetected and unreported. Even where an incident is detected, the lack of a clear and well-thought out response plan would result in confusion and fragmentation of response. This would give the attacker valuable time in

[111] The 2018 study was by the Ponemon Institute, which conducts independent research on consumer trust, privacy, data protection and emerging data security technologies.

which to penetrate further into the system, and hamper the efficacy of the response.

1159. Even the process of formulating an incident response plan for IHiS staff could prove valuable. Leveraging the Cyber Attack, initial planning for the incident response plan will reveal gaps in communication, policy, technical capability, roles and responsibilities that require urgent attention at the organisational level.

49.1.2 Contents of an effective incident response plan

1160. Broadly, an effective incident response ("**IR**") plan should provide a well-defined, organised approach to handling both suspected and confirmed security incidents. The IR plan must cover:

(a) Processes for identifying whether an attack is in progress (including common signs of an attack, and should specify that the signs must be considered cumulatively rather than in isolation);

(b) How employees should respond to an attack;

(c) Steps to be taken to mitigate the effect of the attack;

(d) How and when employees should report an attack (or signs of an attack);

(e) To whom the report should be made;

(f) The means by which the report should be made; and

(g) How employees should document their observations and actions.

1161. In drafting the IR plan, guidance can be sought from relevant documents produced by standards bodies, such as, the NIST Computer Security Incident Handling Guide.[112]

1162. In essence, the IR plan must address the immediate questions that would come to an employee during the course of an attack. It must make clear:

(a) Who is in charge of the response process;

(b) Who should be alerted; and

(c) Who can be approached for help.

1163. The IR plan must have a special focus on the reporting responsibilities of line staff. As emphasised by Dr Lim, cybersecurity involves *all* staff in an organisation, because the impact of a cyber attack affects the *whole* organisation. As demonstrated by the facts of the Cyber Attack, it is the line staff, like Sze Chun, that will often be the first responders. Line staff must be encouraged to take the initiative and report proactively. As CE, CSA said:

> "Staff should have a clear and common understanding of the incident reporting framework, the relevant reporting structures and processes, and what measures must be immediately taken in the event of a cybersecurity incident. New staff should be on-boarded in a timely manner, and regular refresher training should be conducted to ensure compliance with these SOPs."

1164. The IR plan for line staff should be augmented with playbooks (focusing on step-by-step directions) that act as helpful manuals for more specific threat situations. This is especially important for the line staff, whose normal functions do not involve security and incident reporting. As Vivek testified, effective

[112] NIST.SP.800-61 Revision 2.

communication of potential threat situations should be by way of reference to real-world examples that are easy to absorb: the message must be by way of narrative, and not simply an abstract concept.

1165. When deciding on what playbooks should be developed, IHiS should consider the types of incidents that are likely to occur, based on an understanding and evaluation of the relevant risks. The existing IHiS playbooks as at June/July 2018 were geared more towards conventional attacks, including ransomware and website defacement. There was no APT playbook. However, IHiS clearly already had some visibility in this area, as the Cybersecurity Threat Assessment for the healthcare Sector, presented by Kim Chuan to the IHiS Audit and Risk Committee on 5 June 2018, does specifically highlight APTs as a threat to PHIs. Furthermore, the cybersecurity exercises conducted in March 2017 and March 2018 featured APTs as one of the threat scenarios. IHiS should continue to proactively monitor the evolving threat landscape and craft playbooks accordingly. These playbooks should be forward-looking, and should not simply cover areas of past significance.

49.2 The incident response plan must clearly state that an attempt to compromise a system is a reportable security incident

1166. It is absolutely crucial that an unambiguous and easily understandable definition of the term "security incident" is adopted *uniformly* across all security documents (including the general IR, and any IR SOPs for security personnel). This is to prevent confusion and to facilitate ease of reporting. It must also be made crystal clear that suspicious *attempts* to access IT systems are reportable security incidents.

1167. A key stumbling block in the case of the Cyber Attack was that different personnel held different view as to the definition of a "security incident", and consequently, reporting was delayed:

(a) Benjamin knew that the definition of a "security incident" included attempts to compromise a CII.

(b) The SIRM, Ernest, failed to appreciate that the definition of a "security incident" *included* attempts to access a CII. Ernest claimed that only successful attacks that had been 100% confirmed to possess malicious intent would be reportable.

(c) The Cluster ISO, Wee, understood that attempts would constitute a reportable security incident, but did not apply this definition consistently when it came to the crunch – although Wee knew that someone had been trying to access the SCM database, he did not report it as he was waiting for confirmation.

1168. Ernest and Wee's misinterpretation of the definition of a security incident was at odds with the understanding possessed by IHiS management. Bruce, Kim Chuan and Benedict all expected that attempts to access a CII would be escalated and reported. The author of the IR-SOP, Hann Kwang, also never intended that there be a requirement for an incident to be "confirmed" for it to be considered reportable.

1169. The definition of a security incident is currently found in the SIRF and the IR-SOP, and the Committee has found that there are ambiguities in the language used in these documents. Any ambiguity in the definition of security incidents should be addressed going forward. Language can be adapted from other comparable security documents. For example, the US Code of Federal Regulations, which in relation to the "Security Standards for the Protection of Electronic Protected Health Information", applicable to information systems that come under the purview of the U.S. Department of Health and Human Services, uses the following definition:

> "Security incident means the attempted or successful unauthorized access, use, disclosure, modification, or destruction of information or interference with system operations in an information system."[113]
> (emphasis added)

1170. Another example is the definition of a computer security incident in the NIST Computer Security Incident Handling Guide:[114]

> "A computer security incident is a violation or imminent threat of violation of computer security policies, acceptable use policies, or standard security practices."[115]
> (emphasis added)

1171. In other words, the definition must unambiguously and clearly state the both *attempts* and *successful* attacks are to be reported. Further, staff can be encouraged to err on the side of over-reporting. Bruce said that IHiS has now implemented two policies:

(a) For all staff to keep their reporting officers informed if the incident is not resolved within 24 hours; and

(b) To inform supervisors even of incidents that turn out not to be security-related.

1172. These are steps in the right direction. The IR plan should emphasise that, where staff are unclear on the definition or on how to apply the definition to the current situation, they should seek guidance and report the incident so that it can be properly assessed.

---

[113] 45 CFR (US) § 164.304: Definitions.
[114] NIST.SP.800-61 Revision 2.
[115] *Ibid* at [2.1].

49.3 The incident response plan must include wide-ranging examples of security incidents, and the corresponding indicators of attack

1173. There are many different ways to tell if a system has been or is being compromised, but unless employees are able to detect, alert, and respond to these indicators in real-time, the ability to stop a cyber-attack in its tracks will be very limited.

1174. Incident response plans should therefore include a wide variety of examples of possible security incidents. There are many types of security incidents that may require reporting. These incidents should be described broadly, and accompanied by detailed descriptions of corresponding indicators of attacks of that type. Some examples of the types of security incidents to be highlighted to employees include:

(a) Breach of personal information;

(b) Denial of service/Distributed denial of service;

(c) Unauthorised port scanning;

(d) Firewall breach;

(e) Virus outbreak;

(f) Computer account(s) accessed by an unauthorised person;

(g) Compromise of credentials resulting from malware infection, phishing attack, or improper disclosure of password(s) to an unauthorised person;

(h) Device(s) infected with ransomware; and

(i) Unauthorised modification of content or data (unexplained or unauthorised code changes, compromised/defaced website, *etc*.)

1175. It is important that the IR plan emphasises that *context* is crucial to understanding whether a cyber attack is taking place. For example, a single "ping" (a utility used to determine whether a specific Internet Protocol ("**IP**") address, or host, exists or is accessible) on the network initiated from an external source may require minimal, if any, response. No mitigating actions may be necessary since no harmful effects were caused by the incident. However, a suspicious pattern of "pings" on the communications network initiated from an external source or a specific malicious security incident would require a more detailed response, mitigation steps, and more detailed documentation of the incident and outcome. Again, it must be highlighted that employees must look at the indicators *cumulatively*, and not in isolation, to determine if an attack is in progress.

1176. There should be a particular focus on familiarising staff with APTs, as the signature feature of an APT attack is its propensity to remain under the radar, exploiting weaknesses in the ability of employees to detect and respond to subtle signs of attack. The Committee was informed that IHiS is adding a playbook for APT. IHiS should consider automating the playbook as an online knowledge retention tool for the purpose of guiding frontline responders. The plan should also familiarise staff with indicators of attack. Some suggested indicators are in the following sections.

49.3.1 Suspicious Privileged Account Activity

1177. As was seen from the Cyber Attack, should an attacker gain access to a user account on the network, the attacker will often seek to elevate the account's privileges, or use it to gain access to a different account with higher privileges. Staff need to be told to watch out for out-of-hours account usage, and account activity which is out of character for that particular user, *etc*.

49.3.2 Suspicious Outbound Traffic

1178. Staff need to have regard to the traffic that goes out of the network. In particular, staff need to be aware that hackers often make use of C2 servers to enable and maintain threat persistence. Staff should be able to spot and report any unusual patterns of outbound network traffic.

49.3.3 Anomalous login failure

1179. Staff need to be informed that signs of repeated failed logins to an account, or attempting to log in to an account that no longer exists, are clear signs that someone is up to no good.

49.3.4 Spikes in Database Activity

1180. Staff should closely monitor any spikes in database activity, as that could be an indicator that the database has been compromised.

49.3.5 Anomalous registry changes

1181. Staff should be made aware that one of the ways APTs are able to establish persistence and remain covert is by making changes to the system registry. Staff should be informed that should they become aware of registry settings deviating from its typical state, they should report the matter to minimise the potential damage caused by the attack.

49.3.6 Unusual port usage

1182. Staff should be sensitised to the fact that attackers will often use obscure port numbers in order to circumvent firewalls. Record must be kept of which ports are being used legitimately, and for what purpose. Should a port be used that is not in the 'whitelist', staff must be informed to report the matter immediately.

49.3.7 Suspicious File and Folder Activity

1183. Staff should be alerted that activity such as suspicious file or folder creation, modification or deletion, may be indicative of an ongoing attack. Large amounts of data in the wrong place should also be reported.

1184. Ultimately, the IR plan must be as comprehensive and practical as possible. It should be user-friendly and easy to absorb. Nonetheless, there will inevitably be situations which cannot be provided for in advance. It is therefore important for the IR plan to also promote a culture of proactive and early reporting – if in doubt, it is far better for employees to report and seek help than to stay silent.

50 RECOMMENDATION #15: COMPETENCE OF COMPUTER SECURITY INCIDENT RESPONSE PERSONNEL MUST BE SIGNIFICANTLY IMPROVED

#RESPONSE #PEOPLE DEVELOPMENT

1185. While IHiS does appear to have some in-house capability for dealing with cyber threats, the evidence shows that insufficient emphasis was placed on ensuring that security personnel were adequately trained and equipped to perform their functions effectively and competently. Although the IR-SOP does provide for a Security Incident Response Team ("**SIRT**"), a Computer Emergency Response Team ("**CERT**"), and a Security Incident Response Manager ("**SIRM**"), the reality was that the CERT was almost untrained, poorly equipped, and badly led, as the SIRM was unsure of his role and functions. This section elaborates on how these shortcomings should be addressed. The key point is that security personnel must be taken seriously, and cannot simply be left to languish in obscurity without adequate training and support, both managerial and material.

50.1 The Computer Emergency Response Team must be well trained to more effectively respond to security incidents

1186. When computer security incidents occur, it is critical for an organisation to have an effective way to respond. Organisations which are adequately resourced establish in-house CERTs[116], who act as first-responders to security incidents, when the need arises. Failure of these teams to quickly and effectively respond to security incidents can have far-reaching effects.

[116] CERTs are also sometimes called "Computer Security Incident Response Teams" ("**CSIRTs**").

1187. Composition of the CERT. The SingHealth CERT was formed in March 2018 and comprised three people:

(a) Benjamin

(b) Zac; and

(c) Azzlan.

1188. Out of the three members of the CERT, only Benjamin had been with IHiS for a significant period of time – Zac and Azzlan only joined IHiS in April and February 2018 respectively. The only training conducted for the CERT was a half-day course conducted by an external consultant on the use of forensic software. Benjamin had gone for one incident response course ("*Hacker Tools, Techniques and Incident Handling*" by SANS Institute), but had not otherwise received any formal incident response training. Zac and Azzlan did not receive any formal training for their roles. Furthermore, there was no reporting hierarchy within the CERT, and there were no proper procedures for assigning cases to members of the CERT.

1189. Deficiencies in CERT training. Vivek observed that the following deficiencies with the CERT's training contributed to IHiS' failure to mount a proper response to the Cyber Attack:

(a) The team was provided training on how to use certain tools. However, this was only a half-day training. These tools are very complex and advanced, and half a day is not enough to understand even the basic features of one of the two tools. Therefore, it is impossible that the CERT could have been adequately trained to use these two tools.

(b) The CERT was not trained to understand the tactics, techniques, and procedures ("**TTPs**") of advanced attacks. Benjamin had previously never encountered malware that called back to a remote C2 server – a rather common TTP used by advanced attackers.

1190. Recommendations for improvement of CERT training. Vivek's expert recommendation was that the CERT should be expanded and trained to detect and respond to advanced attacker activity. Practical goals for training should be set and the CERT should be provided with access to experts they can tap on as needed. Members of the CERT should possess a comprehensive understanding of attack methods, vulnerabilities, and the impact of attacks on IT systems and networks.

1191. Training should focus on building the competencies required in members of a CERT. To build a CERT with capable incident handlers, one needs individuals with:

(a) Relevant technical knowledge and expertise;

(b) The ability to recognise indicators of attack, collect forensic evidence, perform analysis and arrive at reasoned conclusions; and

(c) The ability to communicate effectively within the team and with others across the organisation.

1192. The Skills Framework for Infocomm Technology[117] provides a useful reference on the skills and domain knowledge that CERT members should consider acquiring or upgrading. In addition, CERT members ought to be sent

[117] This framework has identified cybersecurity as an emerging trend which requires skills such as cyber forensics, cyber incident management and cyber risk management.

for proper training conducted by reputable training providers like the SANS institute.

1193. The overall skill-set required of computer security incident responders can include:[118]

(a) Personal skills -

(i) Communication (written and oral);

(ii) Presentation;

(iii) Diplomacy;

(iv) Ability to follow policies and procedures;

(v) Team skills;

(vi) Integrity;

(vii) Knowing one's limits;

(viii) Coping with stress;

(ix) Problem solving; and

(x) Time management.

(b) Technical skills –

[118] Software Engineering Institute, Carnegie Mellon University, "What skills are needed when staffing your CSIRT?".

(i) Security principles;

(ii) Security vulnerabilities/weaknesses;

(iii) The internet;

(iv) Risks;

(v) Network protocols;

(vi) Network applications and services;

(vii) Network security issues;

(viii) Host/system security issues;

(ix) Malicious code; and

(x) Programming skills.

(c) Incident handling skills –

(i) Local team policies and procedures;

(ii) Understanding/identifying intruder techniques;

(iii) Communicating with sites;

(iv) Incident analysis; and

(v) Maintenance of incident records.

1194. Reinforcement and real-world application of training. The war against cyber attackers is unpredictable and rapidly evolving. A well-prepared CERT would be a powerful weapon in the arsenal of any defender. However, training

alone would not be adequate to ensure that the CERT fulfils its potential as the guardian of an organisation's IT assets. It is crucial that training and theoretical knowledge is made real and ingrained in the CERT members through the conduct of regular practical exercises. As CE, CSA said:

> "CSA recommends that IHiS should conduct a thorough review of their processes to ensure that there are no gaps, followed by a thorough and systematic training process to ensure that all staff have internalised these processes, and know exactly what steps to take in the event of a cybersecurity incident"

1195. In this regard, TTXes (Table Top Exercises) and drills are key to ensuring that CERT members familiarise themselves with incident response plans and processes. Repeated execution of these plans and processes will lead to increased efficiency, and reduce the chances of confusion and hesitation causing a delayed response in the event of a real attack. Furthermore, these exercises will inevitably expose weaknesses in the plans and processes that can then be addressed. The importance of practice and practical application cannot be overstated. As Gen. Alexander said:

> "[P]roper training and a solid exercise program would have ensured personnel knew and understood their roles and responsibilities in helping to prevent the cyber attack on SingHealth. Personnel involved in detecting and mitigating this attack would have benefited from an individual and collective training program."

1196. The conduct of exercises will also help to alert management to the natural abilities (or inabilities) of CERT members. This will allow for an assessment of the initiative of individual team members, and also for the evaluation of potential bottle necks in the incident response process due to the failings of individual officers (*e.g.* the SIRM).

50.2 The Computer Emergency Response Team must be better equipped with the necessary hardware and software

1197. A CERT requires the necessary tools in order to effectively perform its role and functions; without such tools, it would be inherently difficult to get anything done. It is highly recommended for the CERT to have software and hardware that can be readily utilised during an incident; this can range from anti-malware to laptops with packet sniffers, as well as incident response checklists *etc.*

1198. The CERT was not provided with equipment necessary for proper forensic investigations. As at June 2018, the *total* capacity of the CERT's office laptops was 500 GB. Benjamin therefore had to install forensic tools on his personal laptop to carry out forensic investigations of the PHI 1 Workstation and Workstation C. This proved to be a significant bottleneck.

1199. Deficiencies in CERT equipping. Vivek observed that the CERT was not well equipped to respond to security incidents:

(a) One program used by the CERT only allows individual systems to be imaged and does not scale well for an enterprise scale incident where multiple systems may be infected and may need investigation. Another program is quite complex and requires a lot of training and expertise to be used effectively. Therefore, such tools do not lend themselves well to a security incident wherein attackers are advanced and able to spread across a wide cross-section of the network.

(b) The team only had one computer to carry out forensic investigations. Therefore, even to process evidence using the traditional dead-disk forensics approach required a painfully long amount of time.

(c) The team did not have access to EDR software that would have allowed rapid isolation and containment of the infected systems and enabled rapid collection of forensic evidence from multiple systems at the same time. Vivek emphasised that use of an EDR could have cut response time down from almost one month to a single day.

(d) The team also did not have the tools and software needed to analyse malware and reverse engineer it to identify the malware's capabilities.

(e) The team did not have any case management software to log all the investigative updates and track the progress.

1200. Recommendations for improvement of CERT equipping. It goes without saying that the CERT must be provided with both the hardware and software, as mentioned above, to do its job properly. Furthermore, the tools provided to the CERT must be organised properly, to ensure that they are available for use at a moment's notice. Organisations may see their incident investigation and remediation processes experience unexpected delays, or even grind to a halt, if the tools teams rely on to unearth information about affected systems and people are inadequate, mismanaged or misused.

1201. IHiS should maintain an inventory of tools in a centralised location; team members should be trained across the entire tool set on an ongoing basis. Finally, tools should be regularly assessed to determine if they can address the most current threats.

50.3 A competent and qualified Security Incident Response Manager who understands and can execute the required roles and responsibilities must be appointed

1202. It is tempting to think of cyber defence primarily as a technical challenge, but the actions of people also play a critical part in success or failure of incident

response[119]. Defending and responding to cyber attacks requires the right people, who act responsibly and in the best interests of the organisation. Vulnerabilities in human assets can be just as dangerous as those in information systems.

1203. Role of SIRT and SIRM. Security incidents would be investigated by the SIRT, led by the SIRM. The SingHealth SIRM was Ernest. The CERT reports to the SIRM. In addition, the Infrastructure Services Lead, and the Application Services Lead, also report to the SIRM. It is the SIRM's responsibility to coordinate these inputs and report to the Cluster ISO. It is then the Cluster ISO's responsibility to escalate the issue to the GCIO. The structure of the SIRT is highlighted in the diagram below.

Figure 15: SIRT Reporting Structure

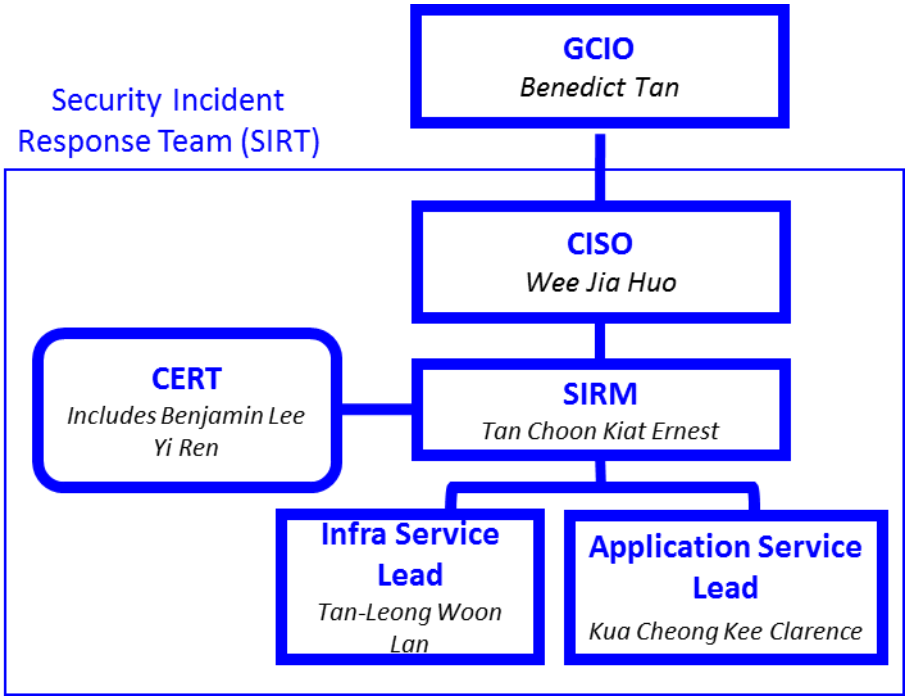

[119] CIS Controls Version 7 at p43.

1204. The SIRM's responsibilities include:

(a) Leading and coordinating activities during incident response;

(b) Managing technical activities during the incident response;

(c) Assigning responsibilities;

(d) Ensuring compliance with the incident handling procedures and guidelines in the IR-SOP;

(e) Receiving incident response alerts about security incidents;

(f) Managing the incident response process from the discovery, to assessment, remediation and resolution stages;

(g) Report to the Cluster ISO; and

(h) Developing IT security incident handling and response policies and processes.

1205. Deficiencies observed in the SIRM. Given the responsibilities and accountability needed to execute the incident response plan, the right SIRM must be in place. The SIRM must be empowered, competent, and possess the right skills sets for the job. Ernest *woefully* failed to meet these criteria. As Vivek observed:

> "The most glaring failure in my opinion was with the role of SIRM. I have read the latest testimony where Ernest seems to indicate that he was not looking forward to the additional workload in the event that the incident got escalated. While this certainly may be a contributing factor and may explain some of his actions or lack thereof, I also believe that Ernest did not fully understand their responsibilities of the SIRM role and did not have the necessary competencies to effectively discharge his duties."

1206. Ernest clearly failed to carry out any of the SIRM's responsibilities listed above, during the course of the Cyber Attack. Vivek further observed that time and again Ernest failed to not only apply his knowledge, but also failed to properly follow the IR-SOP:

(a) Ernest said that malware infections are not reportable if the malware is detected and cleaned without any network propagation. However, he seems to have done nothing to validate whether the "network propagation" indeed occurred. This was a failure to implement the IR-SOP.

(b) As one of the reasons to not report the January 2018 incident, Ernest said that "*suspected malware infection of a workstation are a very common occurrence*" even though Benjamin has indicated that he had never dealt with an advanced malware like this before. This was either a failure to correctly take stock of the incident, or a complete lack of understanding of malware infections.

(c) Per the IR-SOP, it is the responsibility of SIRM to lead and coordinate activities during an incident response. However, there was virtually no formal coordination happening between the different teams. The communication during the incident response was *ad-hoc* at best, and at worst counterproductive to the investigation because it wasted valuable time without making any real progress.

(d) Per the IR-SOP, the SIRM needs to report the incident up the command chain so a formal incident can be declared, and all available resources can be deployed/redeployed to respond to the incident. However, no formal incident was declared and therefore key experts and stakeholders kept operating in silos, which significantly hampered the incident response.

(e) The IR-SOP requires that post-incident reviews are conducted by the Cluster ISO. However, for these to happen, the incidents must be reported first. The incident in January 2018 was not reported up, and therefore never reviewed. A post-facto review, even if it was done later in February 2018, may have uncovered the need for taking additional action and may have helped prevent the incidents in June/July 2018.

1207. The IT Security team should be helmed by an individual who is motivated and interested in learning, as the field of information security is constantly evolving, and complacency leads to weakness.

1208. Detecting and effectively responding to incidents requires strong management processes, and managing an incident response team requires specific skills and knowledge. A background in information security management or security engineering would be ideal. The following competencies should be considered when filling the position of SIRM:

(a) Critical reasoning and analysis – The SIRM must be clear about the criteria to be applied from the various security policies and have the ability to apply those criteria to the situation presented to him;

(b) Gathering evidence – The SIRM must know what the relevant evidence is and how to preserve, collate, and analyse it;

(c) Problem-solving and creative thinking – The SIRM must be able to come up with solutions on the fly, to counter the cyber attackers; and

(d) Communication and leadership – Above all, the SIRM is the person responsible for managing the "boots on the ground", and must be a master communicator, ensuring that information flows in an orderly, efficient, and comprehensive manner to all the relevant individuals.

1209. SIRMs must be provided with the relevant training to shore up these competencies.

1210. To sum up, well-rounded security professionals who combine industry certifications with experience and education should be chosen. The best person for the job must be found, taking into account availability of candidates, cost, and the potential benefit to the entire organisation. This person must then be trained and developed to ensure that they reach their full potential, and can be an asset, rather than a liability, to the organisation.

51 RECOMMENDATION #16: A POST-BREACH INDEPENDENT FORENSIC REVIEW OF THE NETWORK, ALL ENDPOINTS, AND THE SCM SYSTEM SHOULD BE CONSIDERED

#VIGILANCE #GOVERNANCE

1211. An important post-breach action is that of ensuring that the threat is eradicated – *completely*. This means that all breach points must be identified and all attack traces/artefacts must be removed.[120] This includes malware, spyware or any other types of software. This exercise can be complex, lengthy and may require the work of outside experts.[121] Accordingly, IHiS should consider conducting an independent review of the SingHealth network, all endpoints and the SCM system.

1212. Over the course of the Inquiry, concerns were raised on whether the SingHealth network was "clean" post-breach. The concern is a real and urgent one because in the short-term, IHiS will be proceeding with a pilot deployment of their remote-browser solution for internet access at one PHI and in the long-term, the ISS temporarily in place now may be lifted. If the attacker is still in the network, it will spring to life when the system goes online.

1213. On whether the SingHealth network is "clean", CSA's evidence is that the network has been scanned and cleared of the malware or indicators of compromise ("**IOCs**") that were discovered through the course of investigation. CSA has pointed to the following measures:

(a) All Citrix servers have been reloaded with a clean image on 14 and 15 July 2018; and

---

[120] Alexander Ellrodt, *"If a Breach Happens – An Action Plan for Response and Damage Containment"* in *Managing Cybersecurity Risk* at p96.
[121] *Ibid.*

(b) IOCs discovered through CSA's forensic investigations were applied to enterprise-wide scanning, which has been completed.

1214. However, it is possible that not all the IOCs/signatures of the malwares, sleepers and backdoors potentially left in the network by the attacker have already been identified and scanned for. CSA's own evidence is that the attacker had established multiple footholds in the network.

1215. The Committee agrees with Vivek's expert recommendation that IHiS should perform a comprehensive forensic review of all workstations and servers to ensure that are no remnants of the attacker. In his opinion, it is not possible to ensure that the network and system is "100% clean", but he has explained what IHiS can do to ensure that the threat is eradicated as far as possible:

(a) Even if certain IOCs/signatures have been missed or lost, there will be enough information about the attacker to piece together the attack pattern and look deeper into the network to find remnants of the attack. It is also possible to use available intelligence on the attacker to look within the environment for artefacts left behind by the attacker.

(b) Expert investigators will be able to assist IHiS in such a forensic review, using EDR technology which can be licensed for the period of the forensic review.

(c) Such forensic reviews have been carried out by organisations who have been impacted by a significant cyber attack, and such a review can be done across a large number of endpoints, even in the hundreds of thousands.

1216. Gen. Alexander cautioned that we should *"[a]ssume compromised systems are forever compromised*", and that compromised systems should be replaced with new systems if financially feasible. This is an extreme position and we do not expect IHiS and SingHealth to overhaul its current systems. Instead,

IHiS should consider doing the next best thing – work with experts to ensure that no traces of the attacker are left behind.

1217. IHiS agrees with this recommendation and intends to engage an independent consultant to do this review, and eradicate any element of the malware tools used during the Cyber Attack.

52 CONCLUSION ON RECOMMENDATIONS

1218. The Committee's recommendations provide a comprehensive suite of measures that will enhance the capability of IHiS, SingHealth, and other organisations to deter, detect, respond to, and recover from IT security incidents. They range from basic cyber hygiene measures to more advanced measures which are better-suited after a certain level of cybersecurity maturity has been attained by the organisation.

1219. Implementation of the recommendations requires effective and agile leadership from senior management, and necessary adjustments to organisational culture, mindset, and structure. In this regard, the Committee is heartened to note that the MOH family is committed to learn from the Cyber Attack and will continually strengthen its systems against evolving cybersecurity threats. The Committee also notes that IHiS has already taken action following the Cyber Attack, accelerating three ongoing security projects, proposing six more measures, and considering an additional twelve measures (see **Annex B**).

1220. In the implementation of the Committee's recommendations and the measures from IHiS, appropriate oversight of the implementation process, and verification that the measures have been effectively and adequately implemented, is vital.

1221. In this regard, the Committee proposes that IHiS and SingHealth provide updates to the HITSC (being the healthcare sector's highest level platform for cybersecurity issues) every six months on the progress of the implementation of the Committee's recommendations and measures from IHiS, and for the HITSC to consult CSA should any issues arise regarding their implementation. MOHH has informed the Committee that the CSC "*stands ready to play a part in the process*". The HITSC is best placed to identify any such role for the CSC.

1222. The Committee also agrees with the Solicitor-General's recommendation that the GIA should conduct audit checks to verify that the Committee's recommendations and the measures from IHiS are implemented. These checks

should be performed shortly after the specified implementation dates. Copies of the GIA's audit reports should be furnished to the HITSC.

1223. IHiS and SingHealth should give priority to implementing the recommendations. This imperative applies equally to all organisations responsible for large databases of personal data. Cybersecurity threats are here to stay, and will increase in sophistication, intensity, and scale. Collectively, these organisations must do their part in protecting Singapore's cyberspace, and must be resolute in implementing these recommendations.

Annex A – The Members of the Committee

Chairman, Mr Richard Magnus

Mr Richard Magnus is a retired Senior (subsequently termed Chief) District Judge and is currently a member of the Public Service Commission. In addition to his vast experience as a judicial officer, Mr Magnus also led two previous Committees of Inquiry, the first in 1992, to investigate a fire incident at the Sembawang Shipyard involving the ship tanker 'M.T. Stolt Spur'; and the second in 2004, to investigate the cause of the incident at the MRT Circle Line worksite that led to the collapse of the Nicoll Highway. Mr Magnus is also Chairman of the Public Transport Council, the Political Films Consultative Committee, the Bioethics Advisory Committee, and the Ministry of Home Affairs' Remote Gambling Act Appeals Advisory Panel. He was awarded the Meritorious Service Medal in 2009, and the Public Service Star by the State in 2015.

Member, Mr Lee Fook Sun

Mr Lee Fook Sun is currently Chairman of Ensign InfoSecurity Private Limited. Mr Lee is also the Chairman of the Building and Construction Authority, and sits on the boards of the DSO National Laboratories, SMRT Corporation Limted, and Great Eastern Holdings Limited. He was previously with the ST Engineering Group for 17 years until he retired in 2017. He held various appointments such as President of Defence Business, Deputy CEO of Singapore Technologies Engineering Ltd, and President of Singapore Technologies Electronics Ltd. Prior to that, he served in the Ministry of Defence and the Singapore Armed Forces in various positions such as Director of Joint Intelligence Directorate, Director of Military Security Department, and Assistant Chief of General Staff (Logistics). He was also conferred an Honorary Fellowship from the ASEAN Federation of Engineering Organisations (AFEO).

Member, Mr T K Udairam

Mr T K Udairam is currently the Group Chief Operating Officer of Sheares Healthcare Management Private Limited. Mr Udairam has over 40 years of healthcare experience in Singapore with substantial experience in the operation and management of hospitals. Prior to joining Sheares Healthcare Management, he managed a regional healthcare system responsible for approximately 1.1 million people. He also served as the Chief Executive Officer of Changi General Hospital from February 2000 to May 2012. Mr Udairam was the chairperson of various IT committees in public healthcare, including the National TeleHealth Committee. He was also involved in the development and implementation of Medisave. Presently, he serves on the boards of the Tote Board and the Healthcare Information and Management Systems Society (HIMSS).

Member, Ms Cham Hui Fong

Ms Cham Hui Fong is currently the Assistant Secretary-General and Director, Industrial Relations, at the National Trades Union Congress. She also serves on the Board of the Central Provident Fund and is an Authority Member of the Civil Aviation Authority of Singapore. Ms Cham has served in many tripartite committees addressing issues such as wage restructuring and the employment and re-employment of mature workers, including the National Wage Council and the Tripartite Committee on Employability of Older Workers. Ms Cham served as a Nominated Member of Parliament from November 2006 to April 2009.

Annex B – Actions taken by IHiS following the Cyber Attack

Media Release

Update on Steps Taken to Strengthen Cybersecurity Across Public Healthcare System

Singapore, 1 November 2018 – With heightened cybersecurity threats following the SingHealth cyberattack, the Integrated Health Information Systems (IHiS) has taken several steps to strengthen cybersecurity measures across all public healthcare clusters and agencies. These measures help to improve our capacity to **prevent** cyberattacks, and also strengthen our ability to **detect** and **respond** should an intrusion take place on our critical systems.

Additional Security Measures Implemented Progressively

2. IHiS has expedited the planned implementation of Client Advanced Threat Protection (ATP). Client ATP goes beyond conventional defence against malwares by blocking threats based on the techniques used by advanced threat actors. As of 26 October 2018, ATP has been deployed in over 6,000 servers and over 60,000 endpoint devices such as PCs, laptops and others. Full deployment is expected to be completed by the end of the year.

3. Temporary Internet Surfing Separation[1] (ISS) was implemented across the public healthcare sector earlier as a precaution. Mitigating measures were introduced for 46 public healthcare institutions to minimise the operational disruptions caused by the ISS. IHiS is working with the Ministry of Health on a long term approach to the internet access strategy for public healthcare, as Internet connectivity plays an integral part in many aspects of healthcare delivery. On-the-ground implementation of ISS and alternative methods of protecting internal networks, such as the use of a virtual browser[2] solution, are being actively studied. A trial was carried out earlier to assess the technical feasibility of the virtual browser solution and test the compatibility with corporate applications. A pilot with a small group of users will be conducted to evaluate the user experience and further assess the security of the solution. The pilot is expected to complete by mid next year.

[1] Internet Surfing Separation is the practice of disallowing computers that are connected to the internal networks and systems from accessing the Internet. To access the Internet, staff will need to use separate terminals which are not connected to internal networks and systems.

[2] Virtual Browser is a solution where the content of websites that users visit is displayed and executed on an isolated and contained environment. As users only access reproduced content, it minimises the risks of their machine downloading and executing malicious files which may reside on the original sites.

4. IHiS has also identified and initiated 18 security measures that are being implemented progressively. These measures include:

a) Addressing Advanced Persistent Threats (APT) by sophisticated actors: Several measures are being initiated to improve our ability to detect indicators-of-compromise, record and monitor endpoints' system-level behaviours and events, detect advanced malwares and remove the threats, if any. Two-factor authentication will also be implemented for endpoint local administrators who manage end-user devices and installation of software.

An expanded suite of managed security services will be implemented via the Advanced Security Operations Centre including proactive threat hunting, threat intelligence, response services, and more. In threat hunting, for example, proactive and iterative searches are conducted to detect malicious or suspicious activities that may have evaded detection.

b) Addressing Vulnerabilities to Prevent Unauthorised Access to Public Healthcare Clusters' IT Networks: To further **prevent** the use of weak passwords, IHiS is enhancing the access management capability to manage complex passwords centrally, and automatically update and protect administrator accounts. More stringent restrictions will also be imposed on administrative access to servers within the network. The access management will be boosted with threat analytics to provide earlier **detection** of suspicious account activities by applying a combination of statistical modelling, machine learning, as well as behaviour analytics to identify unusual activities, and **respond faster** to threats.

To secure the network against unpatched equipment, the access control will be enhanced to allow only authorised devices that are patched with the updated anti-virus and anti-malware signatures to join the network.

c) Enhancing Security of the Allscripts Sunrise Clinical Manager (SCM): IHiS is enhancing the SCM infrastructure to strengthen security and reduce the risks for the SingHealth SCM database. Database activity monitoring for SCM (which processes an average of 42,000 queries per second) is already in place and is being enhanced with more comprehensive blocks and alerts on execution of bulk queries.

Reviewing Key Systems

5. A comprehensive review of cybersecurity safeguards for key systems including the Electronic Medical Record systems for all public healthcare clusters will be conducted as a precaution. The National Electronic Health Record system is also being reviewed and tested by GovTech and the Cyber Security Agency of Singapore, as well as by PwC, an independent IT consultant. This will ensure that these systems have adequate and appropriate cybersecurity measures to safeguard patient data.

Improving Organisational Processes and SOPs

6. Other than infrastructural and software enhancements, IHiS has also improved its organisational processes and standard operating procedures (SOPs) to reduce the risks and impact of human errors. For example, IHiS has instituted a requirement for suspicious IT incidents to be reported within 24 hours, even if initial investigations cannot determine that they are security incidents. Additional checklists will be progressively put in place to ensure compliance with the SOPs.

7. We have also stepped up staff engagement to heighten vigilance against potential threats. This includes increased alerts and reminders to staff, as well as planned roadshows and briefings on cybersecurity. Training for the security team will also be strengthened to enhance their ability to prevent, detect, and respond to advanced and evolving cyber threats. This includes understanding advanced hacker tools, techniques and exploits, in-depth intrusion detection and advanced digital forensics.

Conclusion

8. As cyber threats evolve and become increasingly advanced, sophisticated, and persistent, IHiS will continue to work with our public healthcare users, CSA, industry cybersecurity experts, and regulators to continually strengthen security measures for our public healthcare system.

About IHiS

9. IHiS was formed in 2008 for an integrated approach in the development and management of IT sytems in public healthcare. Today, IHiS supports the operations of 46 public healthcare institutions including acute hospitals, specialty centres, and polyclinics, as well as over 1,400 partners such as community hospitals, nursing homes, general practitioner clinics and voluntary welfare organisations. IHiS' objectives are closely aligned to the priorities of the Ministry of Health, to provide our citizens with quality healthcare that is accessible and affordable. *For more information on IHiS, pls refer to Annex A.*

Steps Taken to Strengthen Cybersecurity Across the Public Healthcare System

IHiS
Because IT's in Healthcare

With heightened cybersecurity threats, IHiS has taken active steps to strengthen cybersecurity across the public healthcare system. As cyber threats become increasingly advanced, sophisticated, and persistent, IHiS has identified and initiated further measures that are being implemented progressively.

01 PREVENT

Client Advanced Threat Protection
Block threats based on exploit techniques and sophiscated malwares used by advanced threat actors

Temporary Internet Surfing Separation
Machines connected to the internal network are not allowed to access the internet

Advanced Authentication
- Two-factor authentication for local administrators managing end-user devices and software installation
- Complex passwords managed centrally with automated updates and protection of accounts

Access Restriction
Enhanced access control to allow only authorised devices that are patched with updated anti-virus and anti-malware signatures to connect to the network

Virtual Browser
Pilot of solution where users can only access reproduced content on the web to minimise risks of downloading and executing malicious files which may reside on the original sites

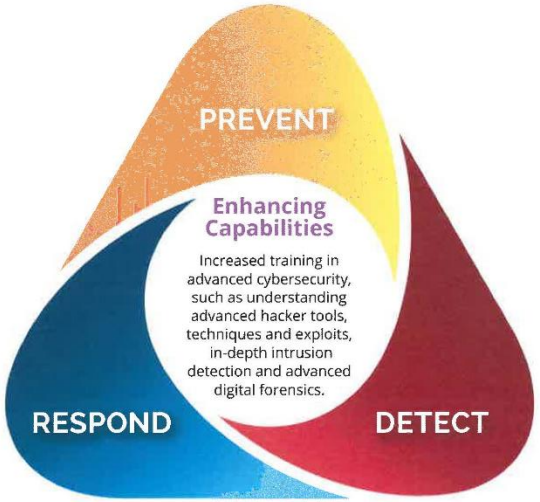

02 DETECT

Detection of Threats
Address advanced persistent threats by improving ability to detect indicators-of-compromise, record and monitor endpoints' system-level behaviours and events, detect advanced malwares and remove the threats, if any

Threat Hunting
Where proactive and iterative searches are conducted to detect malicious or suspicious activities

Threat Analytics
For earlier detection of suspicious account activities by applying a combination of statistical modelling, machine learning, as well as behaviour analytics to identify unusual activities

Database Activity Monitoring
Enhanced monitoring to include more comprehensive blocks and alerts on bulk queries

03 RESPOND

Advanced Security Operations Centre
Expanded suite of advanced cybersecurity services that include threat intelligence, response and remediation measures to abnormal activities and potential threats

Tightened Process
- Suspicious IT incidents to be reported within 24 hours, even if initial investigations cannot determine that they are security incidents
- Additional checklists to be put in place to ensure compliance with the SOPs