ALSID
ACADEMY

Episode 4

1.5 MILLION LESSONS FROM THE INFAMOUS SINGHEALTH HACK

EPISODE 4 1.5 MILLION LESSONS FROM THE INFAMOUS SINGHEALTH HACK

On January 10th, 2019, the Committee of Inquiry (COI) published its public report detailing the now-infamous SingHealth hack incident.

This 453-page report is already a difficult read for most security practitioners and, as long and exhaustive as it looks, the most determined readers amongst us might still feel unsatisfied by its lack of details on some of the attack's critical steps.

We provide here, in layman's words, a description of the incident. We also try to bridge the information gaps to provide a clearer description of the attackers' moves. While these deductions are not proven, cyber-criminals' modus operandi are predictable enough to consider them probable.

We will follow up on this article in the coming days to provide an analysis of what organizations can do to prevent this disaster from happening in the first place. And, in the unfortunate case that it may be too late, we will also discuss the most common options with regard to remediation steps and discuss whether they seem appropriate or not.

WHAT HAPPENED

While the aftermath of the attack is still uncertain, its sheer size only is sufficient to bring quite a little anxiety to Singapore's citizens: **1.5 million personal records, plus the prescription details of 160,000 patients**, including those of Prime Minister Lee Hsien Loong, have leaked to an unknown criminal.

As for attribution, the COI seems convinced perpetrators are state-sponsored operators. As a matter of fact, some of the attack techniques used, as well as the persistent approach utilized, clearly set the attackers out of your usual low-tech hit-and-run hackers' crowd.

Considering potential targets, and albeit the fact that millions or records have effectively leaked, it seems probable that this operation targeted specifically the private medical records of Prime Minister Lee Hsien Loong. It should be noted that the Prime Minister's Office publicly released in 2015 a statement on PM Lee Hsien Loong being diagnosed with a prostate cancer. Whether the attackers were looking for embarrassing details to be used as leverage is unknown but cannot be excluded.

After their initial reconnaissance, on which we naturally have no information, the operatives sent booby-trapped emails to SingHealth employees. The trap consisted in a **malicious attachment capable of exploiting a known Outlook vulnerability. When opened, the attachment dropped a rogue program on the victims' systems, matter-of-factly providing full control to the attackers over the targeted computers.**

There is not much to be said on this initial approach. While all organizations constantly do their best to detect those patient-zero infections, this will remain a cat-and-mouse game for the foreseeable future and SingHealth is not to be blamed for suffering such a breach, particularly considering the nature of its opponent.

Hardly harmful, this initial touchdown aimed at providing the attackers with the ability to explore SingHealth's IT infrastructure from within. And as with the vast majority of sophisticated attacks, the actual IT target of the hackers was SingHealth's Active Directory infrastructure. Owning your victim's AD systematically signs the end of the game: with full control over all IT resources, there is not much an attacker cannot access.

- **Using Active Directory as a transport for destructive malwares.** Destructive malware is not rocket science. Highly-sophisticated payloads such as Stuxnet are the exceptions, while today's consumer-level ransomwares are good enough to do the destruction job effectively. The only challenge in those attacks is distribution: getting these malwares installed on a sufficiently large number of endpoints so that recovery at scale becomes unrealistic. In this regard, exploiting Active Directory weaknesses is the only practical option for hackers to move laterally within the infrastructure. Every large-scale, infrastructure-wide attack that has crippled production capabilities in recent years has had an Active Directory exploit at its core.

In the SingHealth case, it is not entirely clear **how the Active Directory was compromised.** Though considering common industry practices, and details provided in the report, we can safely assume they followed either one of the following courses of actions:

- They used the current user's credentials (supposedly unprivileged) to propagate but did not get elevated privileges from it. As the report suggests though, they might have used weak passwords or clear-text passwords found in network shares to directly authenticate into more-privileged accounts. Which is of course in itself quite a vulnerability. Re-using those accounts, they might have ended up on a computer holding the credentials of a privileged user such as a domain administrator, thus gaining effective control over the whole Active Directory.

The rest of the attack simply consisted of **using legitimate accounts to query databases and scout resources,** like you would on your own systems, until they found what they were looking for.

> **"Owning your victim's AD systematically signs the end of the game: with full control over all IT resources, there is not much an attacker cannot access."**

FROM ACTIVE DIRECTORY, WITH LOVE

This attack sheds yet again a somber light on the state of insecurity of Active Directory infrastructures. There is no doubt that the AD was the primary IT target of the attackers.

In this respect, the SingHealth incident is only the latest sorry example of a long list of operations that became truly successful when and only when they gained access on their victim's Active Directory: Aurora, Target, Sony, Carbanak, NotPetya, the list goes on.

Active Directory infrastructures remain the nexus point of everything that electronically-matters in organizations but are still dangerously ignored by security operators. In their defense, the size, complexity and volatility of a given AD makes it a singular security challenge. Still, it remains very concerning to witness how our industry underestimates the risks it incurs on global security.

In our next article, we will provide an analysis of what organizations can do to prevent this type disaster in the first place, notably regarding the protection and monitoring of their AD infrastructure. We will constructively criticize the most common security practices for Active Directory and will review potential remediation approaches for those organizations that are in the unfortunate position of having to recover from such an incident.

> **"Active Directory infrastructures remain the nexus point of everything that electronically-matters in organizations but are still dangerously ignored by security operators."**

DELVING DEEPER

A deep read through the 453-page report released by the Committee of Inquiry (COI) on January 10th allows for a safe assumption: **the attack indeed reached a turning point when it succeeded compromising SingHealth's Active Directory (AD).** But what exactly went wrong with this critical infrastructure?

6 MONTHS EXPLOITING KNOWN ACTIVE DIRECTORY VULNERABILITIES

You read it right: for a whopping 6 months - from December 2017 to June 2018 -, SingHealth's tormentors have been raiding their victim's Active Directory unchallenged. That's more than is necessary to achieve lateral movement on the organization's machines and to reach a critical mass of compromised systems. Beyond this point, probabilities work for you: it is very likely that you'll end up on a machine that either holds AD admin credentials or has direct access to the data or resources you were looking for.

There are 2 clear takeaways here:

- the audit methods used in anticipation of those threats did not translate into adequate hardenings of the infrastructure, and 2.
- the detection mechanisms SingHealth had in place were not sufficient.

AUDITS WITHOUT REMEDIATIONS ARE A WASTE OF TIME AND MONEY

On the first point, COI's report shows that SingHealth routinely performed technical audits that covered its Active Directory. Unfortunately, the recommended remediations weren't fully applied. Sadly, this observation is not a SingHealth exclusive: audits tend to become an administrative requirement, a box that needs checking, rather than an actual security improvement tool.

If a given organization does not enforce, as a process, its audits' recommended remediations, then it boils down to the particular willingness of its technical staff to implement them... which usually ends up into no remediation applied at all. Not judging those individuals here: manually mingling into Active Directory is not a risk-free nor a quick task, and it often comes in competition with the gazillion other duties that fall on these professionals.

As for SingHealth, there are a few examples of weaknesses that were explicitly reported on by auditors but which weren't followed-up on, such as trivial passwords and lenient access controls. Whether those following weaknesses were reported on by auditors is unknown, but, considering they are obvious, it seems safe to assume they didn't go unnoticed. Anyhow, the resulting non-remediations have had the same sorry consequences on SingHealth's Active Directory integrity.

- The Citrix servers that held access to the attackers targeted resources were in an Organizational Units (OU) that was blocked for Group Policy Object (GPO) propagation. As a result, the central password policy for AD accounts wasn't applied, allowing attackers to exploit weak passwords.
- Beyond those AD accounts, local admin accounts also seem to have been unmanaged, while they should have. Using Microsoft's Local Administrator Password Solution (LAPS) could have prevented hackers to exploit those.
- Finally, the report reveals the existence of a dormant privileged (service) account. Privileged and dormant: there go two words we don't like to see next to each other! However, hackers do: this account was exploited by attackers to access the data repositories they were targeting. Following standard AD security best practices would have required the deactivation of this account, therefore closing a pathway to SingHealth's critical assets.

Although these previous items look like evident flaws, Active Directory remains undoubtedly a complex, moving infrastructure that's made of thousands of different entities. Looking for weaknesses in this gigantic hay stack is a time consuming and error prone process (as it happens, we know of a **robust tool that can help** but that's another story).

Beyond the SingHealth incident, let's have a quick look at other weaknesses that are commonly exploited by attackers.

"Audits tend to become an administrative requirement, a box that needs checking, rather than an actual security improvement tool."

ACTIVE DIRECTORY: A FECUND PROVIDER OF VULNERABILITIES

Here, «vulnerabilities» is to be understood in its broader sense. While **software vulnerabilities exist**, exploiting misconfigurations or gaping legit holes is more leisurely and, therefore, way more common.

There are thousands of configuration 'atoms' that, if badly configured, can cascade into major infrastructure-level vulnerabilities. This profusion is a nightmare for security professionals and a blessing for their opponents. Let's go through some common exploits:

- Dangerous credentials exposure is the #1 threat that Active Directory infrastructures face. Due to some legacy administration practice which go back as far as to Windows 2000, accounts with complete access to their organization's most private resources (like user passwords) have proliferated. This situation dramatically increases the opportunities for an attacker to leverage process injection techniques (available as ready-to-use packages in various open-source tools such as Mimikatz).
- It's frequent, albeit unfortunate, to find privileged accounts running Kerberos services. Kerboeroasting is a widely known technic that, in this context, allows for stealthily extracting those accounts' credentials through offline brute-force attacks.
- Omitting to restrict delegation on sensitive accounts allows for their impersonation. This Kerberos delegation exploit is as simple as running a command line but remains a major hit in many large infrastructures.
- Dangerous access control on GPO consists in, for the attacker, modifying a legit GPO to inject malicious commands that allow him or her to get control over privileged accounts and servers.

This list can go on seemingly endlessly, and attackers are not short of options when it comes to exploiting bad configurations, and new clever ways emerge regularly.

HOW SINGHEALTH DETECTED INTRUDERS IN THEIR AD

Well, there's a catch here: they didn't.

Imagine: we just went through a dozen of reasons that prove it's tough to manually harden an Active Directory, while this can be planned and architectured. Now, what about monitoring those thousands 'atoms', continuously, in real time, seeking weird-looking **modifications and unexpected behaviors**? This is just humanly impossible.

As a matter of fact, SingHealth's incident detection came from observing the last near-death symptoms of the attack: the eventual queries ran against the medical database. Admittedly, later is better than never. However, let's agree that detecting an attack after it's burned its way through your infrastructure and has exfiltrated your most sensitive assets is not a win.

In this regard, the lack of AD-centered **detection technology** had a destructive effect: instead of running emergency mitigation tactics to contain the attackers' movements and avoid that they reach their targets, SingHealth was condemned to a post-heist reaction aiming, at best, at banning their opponents out.

By the way, did they truly got banned? Let's take a closer look at the remediation tactics SingHealth implemented.

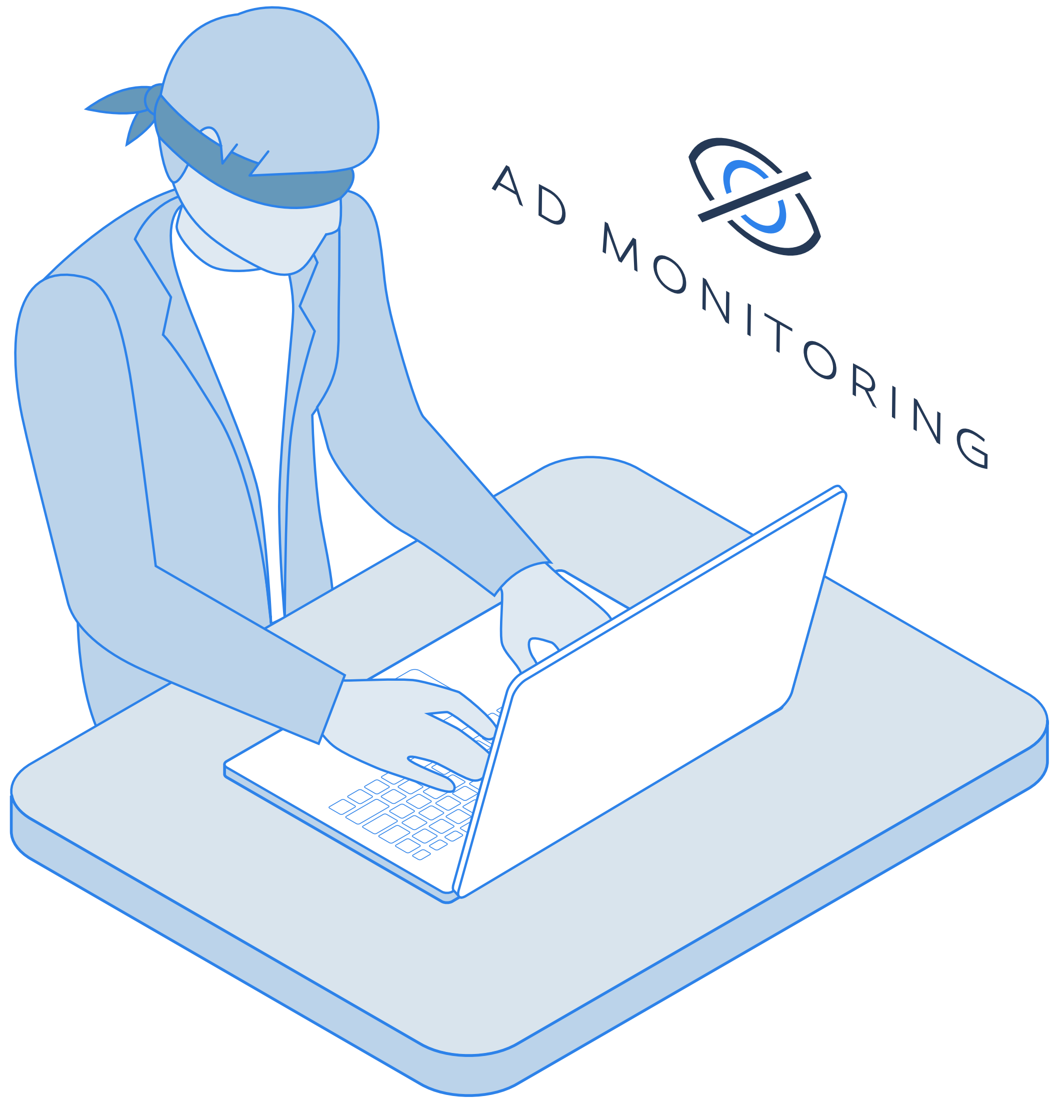

GET OUT!

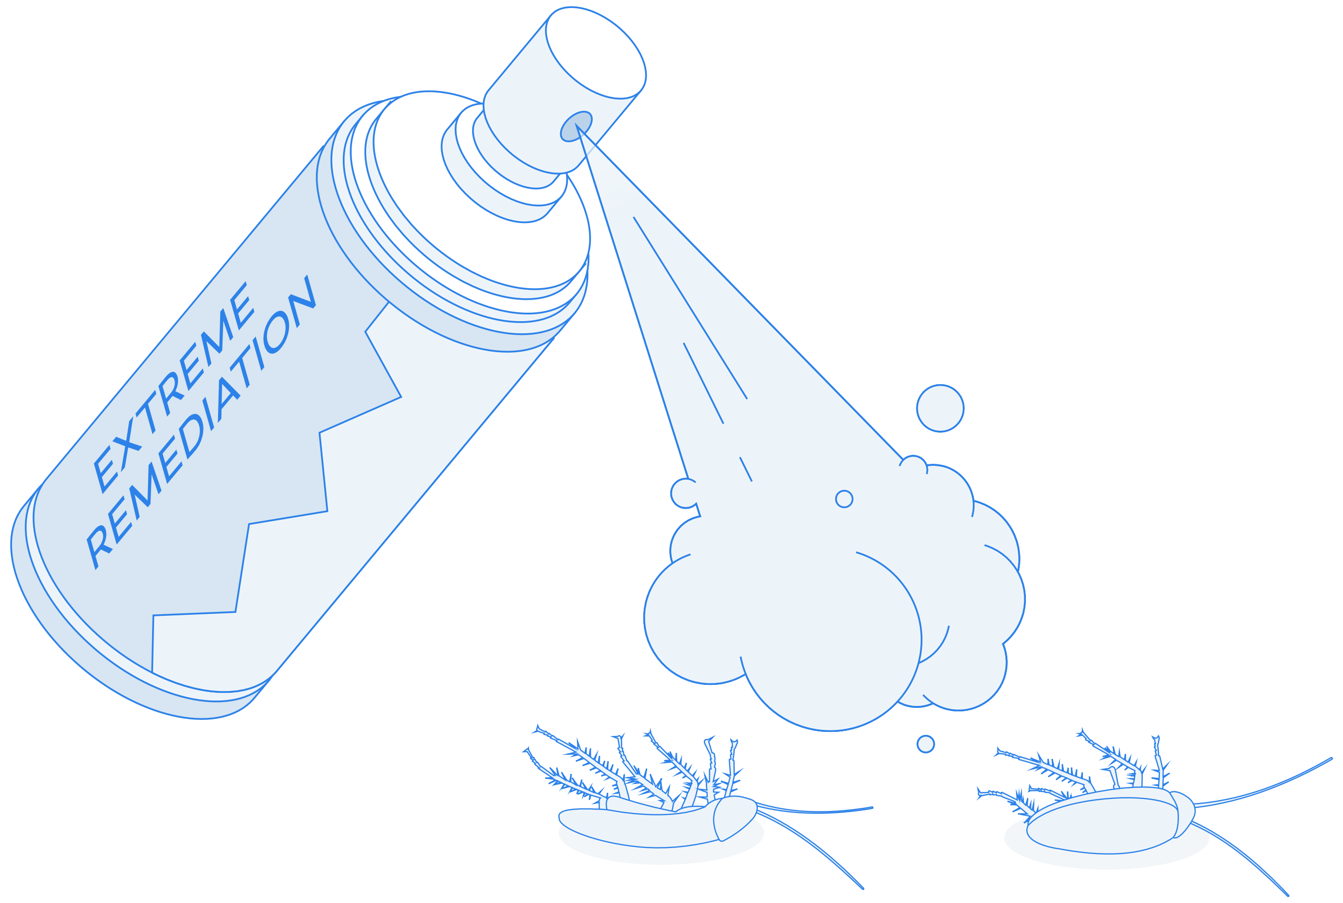

There is no denying that SingHealth's reaction after detection has been swift. It came in two waves: the first one, immediate, was orchestrated by SingHealth's teams, while the second stage was led by the Cyber Security Agency of Singapore (CSA).

Phase 1

Firstly, SingHealth's remediation team created a **new set of domain administrators and revoked the access of the former ones**. This is a must-do that's never sufficient but always necessary.

Then, incident responders enforced the deployment of a GPO preventing Domain Administrators from logging into servers. The intent here was to **avoid that a DA logs into a compromised system**, thus **exposing their new immaculate credentials**. It's anyhow a widely known best practice to limit privileged accounts to particular use cases.

Antiviruses were next set in motion to scan Domain Controllers. Although the effectiveness of those scans is debatable, it would have been unnecessarily risky not to do it.

Phase 2

At this point, CSA took over the rest of the incident response efforts.

Notably, they started with **2 consecutive resets of the KRBTGT account**. This KRBTGT account is a unique account used to encrypt the access key of every domain user. If an attacker owns it, it can easily create legit authentication material to impersonate any other account. Obviously, not what you want. By the way it's a best practice to renew the KRBTGT's password regularly (e.g., on a yearly basis), since it's not an automatic process.

Unfortunately, the operational steps to get a potential attacker out of KRBTGT's reach are not as trivial as it seems. Because the last 2 KRBTGT's passwords are valid (that's normal behavior), CSA performed 2 consecutive resets. The tricky part here is that you absolutely **must ensure with a 100% certainty** (110% is better) **that your attackers don't get the chance to compromise the KRBTGT account between those 2 consecutive resets**. In this instance, CSA performed those 2 rotations within a 24-hour time window. We have no insight into the measures they applied to make sure no breach happened during those 24.

Then CSA went into a **changing-passwords spree on accounts and applications**. Although operationally painful, this step is necessary to invalidate stolen credentials. Of course, this measure is only as good as the hardening implemented to block the offensive tactics used to steal those creds in the first place. Again, we have no insight on whether CSA hardened SingHealth's Active Directory as a pre-requisite to changing passwords at scale.

ONE LAST THING

Let's assume - and there's no reason not to - that all those remediations measures were conducted thoroughly and adequately. Is there any other way an attacker could have persisted in SingHealth's systems?

Guess what, there are. The COI purposefully omitted to publish some of the incidents' details in its report, so that there is no way to know whether those vulnerabilities were actioned upon or not. However, they do theoretically exist. Here are a couple of examples.

- It is possible to ensure persistency by adding malicious permissions on the adminSDHolder object. This object is used as a template for the permissions applied to newly-created privileged accounts. By tampering with it, an attacker can maintain permissions on every privileged account, even those created after the remediation mentioned above.
- Another classic consists in injecting a privileged Security IDentifier (SID… rings a bell? AISID maybe?) to the SidHistory attribute of an otherwise harmless-looking account. This would provide attackers with de facto elevated privileges, without the caveat of utilizing an account that's explicitly in a privileged group.
- SingHealth's threat actors could also have associated their own certificates with existing privileged users. Those certificates' validity does no perish when their related account's password is reset. And because they can be used as you would a password to authenticate, they represent a stealthy way to maintain control over your victim's accounts.

We don’t intend to be exhaustive here, but it should now be evident that Active Directory is complex machinery with hundreds of tracks, branches, and hidden pathways that lead to its core. After a successful attack, cleaning and sanitizing can be complex and time consuming, and regaining trust in this system is not a given.

It is not surprising that AD teams, everywhere and not only at SingHealth, struggle to audit, harden, and monitor it as a whole.

IT'S TIME TO CARE (ABOUT AD)

Though it's past time people care. We, as an industry, have spent literally billions on endpoint, data, and network protections. At a time when all major attacks have a strong AD component to them, we have no excuse for our collective carelessness. Whether it's auditing, **software**, or processes, there are options out there to make a difference. None of them is perfect but having none makes us part of the problem.

ALSID