OFFICE OF THE INFORMATION & PRIVACY COMMISSIONER FOR BRITISH COLUMBIA

CHECK AGAINST DELIVERY

SPEECH TO
BCNET CYBERSECURITY SYMPOSIUM 2026
November 24, 2026

Michael Harvey
Information and Privacy Commissioner for British Columbia

It's an honour to be here on the territories of the xʷməθkwəy̓əm (Musqueam), Skwxwú7mesh (Squamish), and Səl̓ílwətaʔ/Selilwitulh (Tsleil-Waututh) First Nations.

I am privileged to be able to live in this beautiful province with my family, surrounded by nature which reminds me to always reflect on how that comes with both privilege and responsibility, to the land itself, to the people that have lived here, and to the people that live here now.

As an Officer of the British Columbia Legislature, I also acknowledge that I am privileged to work with people across many traditional Indigenous territories, covering all regions of our province. This privilege comes with a responsibility to seek every opportunity to advance goals of truth and reconciliation as I conduct my legislative mandates and, as a Canadian, as I live my life.

Thank you to Dennis (Silva), Antonio (Brandao), Armi (Pamintuan), and the rest of the BCNET team for inviting me to join you here today.

Being here with you today takes me back to my own time in a provincial government executive role and how every October, during Cybersecurity Awareness Month, someone from our IT team would sit all of us executives down for a briefing on the latest threats to our systems.

I joked then that it was fitting that they held these briefings around Halloween because they were well calibrated to scare the wits out of every executive in the room.

That was some years ago, and those threats have only grown more complex and pernicious since.

This year it seems that November is the scariest month. Over the past few weeks. there have been an escalating number of breaches and other concerning incidents. Some of these are clearly cyberattacks from bad actors, apparently using AI to increase the potency of their attacks, and automate them. In other instances AI agents have escaped their training sandboxes and gotten up to no good. In yet other instances, AI has been used, helpfully, to identify critical security bugs in software, but sometimes this creates a race against time between security experts trying to patch the bugs and bad actors wishing to exploit them.

Coal Harbour, just a couple blocks away, seems awful calm for someone raised in Newfoundland on the shores of the North Atlantic. But the sight of water here inspires maritime cliches in a guy like me: we had better batten down the hatches and secure the rigging because the storm isn't coming. It's here and it's only going to get worse.

In recent years, in recent months, in recent *weeks*, the education sector has been the target of several high-profile cyber attacks.

Earlier this month, EducationPlannerBC publicly reported a cybersecurity incident. That came after the Canvas/Instructure breach last spring, and those follow other prominent breaches, including the 2024 PowerSchool breach, the Blackbaud breach in 2020, and other attacks at higher education institutions, as well as school districts.

Cybercriminals target the education sector because of the wealth of personal information institutions hold and its potential for exploitation – everything from academic and financial records to identification details, even health and counselling information.

What makes these high-profile attacks even more troubling is that so much of that information belongs to young people, just starting out in their lives.

The repercussions of a privacy breach leading to identity theft or their information being sold on the dark web could haunt them well into adulthood.

That's where I'd like to centre my remarks today. When we're talking about data, *we're really talking about people* – that protecting data means *protecting the people who entrusted us with it*.

For many of you, that means the identity documents and financial details of a registering student; the personnel records of staff and faculty; the information of research participants who shared their stories because they trusted the promise that they would be protected.

Indeed, trust is the prerequisite for all of these interactions and that trust is precious because of just how central personal information is to our lives.

In the Information Age, our personal information isn't about us, *it is us* – it's what we choose to share about our thoughts and feelings. It's the places we're going, the things we're buying – even our mental and physical health.

As artificial intelligence and other emerging technologies reshape our world, it's important to remember the value of our personal information, and what it represents – *who* it represents.

It represents us.

When we lose control over our information, we lose control over something precious and unique to each of us – what gives us personhood, identity and agency in the Information Age.

That surrender diminishes us as individuals. We're not only surrendering our individual privacy; we're losing our autonomy.

That has implications not only for the individual, but for the collective – for who we are as a province and a country.

Five years ago, conversations about data sovereignty were mostly about residency – where our data is stored.

Then we began to talk about the data in transmission – where our data went as it was processed.

Then we looked at implications of foreign laws on Canadian soil.

When we recognize that there is no longer any meaningful distinction between our online and offline selves, old demarcations fall away, and the conversation turns to more foundational issues.

How we treat our data is how we treat our people.

In the Information Age, data sovereignty *is* sovereignty.

Data sovereignty is about more than where data lives – it is about who controls it, under what laws, and with what safeguards for rights, transparency, accountability, and democratic values.

It's not about sealing our data behind our borders. Data sovereignty is not isolation. It is about Canada's agency and the decisions we make about what matters most.

This requires a life-cycle approach.

Institutions of higher education have a critical role here in equipping people with digital skills, nurturing homegrown talent and building Canadian industries that can shape – not just navigate – the Information Age.

This gives us a risk-based lens for decisions. At each stage we can ask: What is important to us? What is sensitive enough to demand tighter controls? What is off limits?

We have choices.

We have agency.

We are not bit players in a pre-written script of technological inevitability.

That's because behind every artificial intelligence application – indeed, behind any collection of personal information by any type of tech you can name – are people, the choices they make, and the values that inform those choices.

When we consider it from that perspective, what we have now is not inevitability, but opportunity – to chart our course toward a future where we benefit from technology as individuals and a society.

Where we choose to deploy and guide its use in a way that aligns with our values, with what we want for ourselves and for generations to come.

British Columbia has a tremendous resource economy, and that is not going away, but our future is also as an international innovation-driven economy, and that economy will depend on *trust* in how data is handled.

I'm excited by that possibility and believe that we can build a BC where people don't have to trade their privacy for progress –where innovation is strengthened and made more sustainable over the long term because of the trust that comes with respecting people's rights.

That's why promoting and facilitating trusted innovation is one of the core priorities in my office's strategic plan.

I'm a techno-optimist, but I'm also a realist, and in my role, I'm well aware of how things can and often do go wrong, and of the challenges all of you face in protecting the information entrusted to you.

Today, my aim is not to scare you but to encourage vigilance in protecting personal information, and provide my thoughts on today's cybersecurity threat landscape from the perspective of personal information.

I'll look at what my office is seeing in breach reports and what our public sector privacy law says about relationships with vendors, which are often a key factor in breaches.

I'll also look at breaches from the perspective of the executive suite and, finally, the looming spectre of quantum computing and what that means for privacy protection.

I'll finish with questions I'd encourage leaders in this room to take back to their teams.

…

Personal information has always been at the centre of a wide range of cyberattacks.

In many cases, AI makes these types of attacks more straightforward to execute and more far-reaching in their potential consequences.

Take spear-phishing, for example.

Prior to AI, successful spear-phishing campaigns required significant work and skill to pull off.

Attackers needed to research their targets, learn about them and their work relationships, and target their attacks so specifically and convincingly that their targets – most often executives – believed what they were seeing, and did what was requested of them.

That required time, effort, and skill.

AI removes those barriers.

First, the research.

AI makes it faster and easier to scrape personal information from much broader sources at scale. No longer do attacks need to be focused on executives to be appealing to attackers – any employee with a public profile is a viable target.

Second, the message itself.

The awkward or bizarre requests many of us remember getting from spear-phishing campaigns in the past have now given way to far more believable attempts with AI.

AI can produce messages that appear to be from someone in a position of trust or authority, and it's no longer just email. An attacker can pair a message with a cloned voice call or a deepfake video to impersonate a supervisor, with each appearing to corroborate the other.

Post-secondary institutions are a prime target for this type of attack given the amount of publicly available information out there: faculty profiles, staff directories, research bios, conference programs.

Third, the volume. An attacker only needs to succeed once and AI improves their odds of success by letting them make far more attempts, far faster. It's a numbers game.

As advanced as these methods have become, one of the most effective safeguards against them is human judgement – particularly the ability to recognize patterns, see something is off about the communication, and take action.

Those are valuable skills that require regular training to develop – cybersecurity training programs that looks at the many possible dimensions of these threats and train people to recognize them.

And the threats AI poses to privacy aren't only external – these tools can be part of breaches from *within* organizations. For example:

- An employee might disclose personal information while using an unauthorized chatbot at work.
- AI tools that *are* approved for work might not have been sufficiently scrutinized for privacy and security and might be sharing data with unauthorized third parties. I'll discuss how you can guard against that later in my remarks.
- AI tools may scrape or reuse publicly accessible information online in unexpected and unauthorized ways, and likewise, any data entered into an AI tool could be used in unforeseen ways when there aren't adequate safeguards in place.
- It can be difficult or impossible to correct or delete information that has been collected and processed by an AI too, which further undermines privacy rights.

This is why we emphasize thorough privacy and security vetting of any tools that your organization plans to use, and ongoing internal education, so that your staff know which tools are authorized, and how they can use them without compromising people's rights.

I'd like to fundamental and valuable privacy principle here and that's data minimization – minimizing the amount of personal information you collect, use, and disclose.

Scraping is easier when there's something to scrape. Sharing is easier when there's something to share. Personal information doesn't have to be as readily available as it is.

Consider what information *needs* to be out there and limit where you can.

Think of how the information you're sharing could be used in a potential AI-powered cyberattack, or how it could be inadvertently used and shared with unauthorized parties when using unvetted AI tools.

…

Just how much of a factor AI is in the privacy breaches reported to my office is unclear as most of the reports we receive don't specify or speculate on whether the technology may have been a factor in a breach.

Public bodies are required under our *Freedom of Information and Protection of Privacy Act*, or FIPPA, to report privacy breaches to my office that may result in significant harm – such as identity theft, financial and reputational damage, and other harms.

That requirement gives us a clearer picture of areas where all public bodies, including higher education institutions, are vulnerable, and cyberattacks are clearly a part of that picture.

We had 3 breaches involving cyber-attacks at post-secondary institutions during the 2025/2026 reporting year, of a total of 33 breaches overall.

Since April 1st , we've already had 7 cyberattack breach notifications out of a total of 23 breach reports.

Not all threats come from the outside – many of these breaches may start with an email sent to the wrong person, information shared with someone who shouldn't have it, or shared with an unauthorized piece of software.

While we can't put a number on how many breaches have involved AI, we know, as I've mentioned that it is being used to find vulnerabilities, make social engineering easier, and that the unauthorized use of these tools creates a significant new risk of breaches from within organizations.

…

As the education technology sector continues to advance, education systems are using third party platforms more and more. This trend creates both opportunities and risks for the security of the personal information of post-secondary employees, students and others.

Handing data over to a vendor does not mean handing over responsibility. Under FIPPA, public bodies remain responsible for protecting personal information under their control, including where that information is handled by a service provider.

Some of you may be familiar with the investigations into the PowerSchool breach by the Ontario and Alberta privacy commissioners. In this incident, an unauthorized party accessed and stole personal information that was in the vendor's *custody* but under the *control* of various school boards. The investigations' findings and my office's experience offer some practical insights into protecting privacy in these circumstances.

When advertising their services, vendors will often make claims to be compliant with privacy laws in places where they sell their services – for example, in Europe, the US or at the federal level.

But compliance with other laws is insufficient – it's FIPPA and PIPA that matter here.

So wen public bodies come looking for advice when onboarding third party vendors, my policy team recommends a few key requirements.

For service provider contracts:

- First, that providers provide prompt notification of any breach or suspected breach of systems that impact personal information under the public body's control.
- Second, an explicit prohibition on secondary uses. Some providers will try to de-identify information and use it for their own purposes, such as product development. That would not be compliant under FIPPA.
- Third, measures to make sure that the vendor's involvement, such as system maintenance or upgrades, doesn't lead to unauthorized collection, use or disclosure.
    - We often recommend that public bodies have their own encryption key, and they decrypt the personal information vendors store on their behalf only for authorized and pre-determined purposes.
    - Ideally, an employee from a public body would oversee the vendor's access to personal information in real time. However, depending on the needs of the public body and the services provided by the vendor, this may not be reasonable, or even possible.

Beyond the contract, we also expect to see reasonable security arrangements under section 30 of FIPPA, including:

- **Multi-factor authentication**, including two factor authentication, for *all* access to the vendor tool solution to mitigate common attack methods such as keylogging and credential stuffing.

- **Immutable audit logs**. These are essential because they provide a tamper-proof record of events, ensuring personal information cannot be accessed or altered by individuals looking to cover their tracks.
- **A records retention and destruction schedule**. FIPPA rarely supports indefinite retention of personal information, as it significantly increases the risk of unauthorized access to or disclosure of information that the public body no longer requires. Beyond the need to keep information used when making decisions about an individual for at least one year, FIPPA does not include set timelines. However, our office recommends keeping personal information only for as long as necessary to fulfill its original purposes.

You may also wish to ask if the vendor has cybersecurity insurance.

Having it is no guarantee of security, but insurers on high and increasing levels of safeguards, so it would be a good signal. In fact, if you don't have cyber insurance yourself, it might be worth considering.

As your institutions are trying to modernize with new technology, a robust privacy management program becomes essential.

I encourage you to check out our guidance on Accountable Privacy Management in BC's Public Sector to build those programs. I understand that the BCNET team will share a document containing a list of these resources and key takeaways from my speech this morning after today's session.

A strong privacy management program means that privacy will shape the design of any digital transformation project rather than just being a checkbox before launch.

…

Robust privacy management programs and thorough privacy impact assessments will reduce your risk of a privacy breach – they won't eliminate it entirely.

Breaches happen.

There's an axiom in our circles that there are two types of organizations: those that have experienced a breach, and those that will.

I've also heard it said that the best way to prepare for a cyberattack is to have already been through one.

If experience is the best way to prepare, the second best is to have a plan, and to know that there is help out there if you need it.

I've been on both sides of a privacy breach, so I know what one looks like from the executive suite, and there are some commonalities that bear discussing.

First, it's important to understand that all the high emotions you feel in the wake of a breach – the chaos and fear – are valid and to be expected.

You're faced with both how to get systems back online again, and how to deal with bad actors to avoid or at least mitigate the potential damage to the people who entrusted their information to you.

Then there are operations people on the ground operating in a cloud of confusion of their own because they have no idea what's happening at higher levels.

And it's amid this fog of war that communications folks are charged with getting clear, coherent information out to people as fast as they can.

It's in situations like this that the silos in your organization become most starkly apparent, and where they do the most damage.

Some level of disturbance is normal and to be expected in the wake of a breach.

What you can do to minimize that disturbance is *plan*.

Executives can be proactive in their planning by ensuring their incident response plans are up to date and that they have physical copies of them, in case they can't access digital systems.

They also should check with their CISOs and CIOs about red-teaming, where simulated cyberattacks are carried out to test vulnerabilities.

We would advise executives to have a confirmation codeword for extremely sensitive and urgent decisions that is shared through non-digital means, so that threat actors are unlikely to know it.

We also advise executives to make sure their networks are properly segmented, and that they have proper logging, detection and auditing capabilities.

If a breach does occur, my office recommends four steps:

1) Contain the breach: Stop the unauthorized access right away.
2) Evaluate the risks: What information was affected, how many people were impacted, and whether there's a real risk of significant harm.
3) Notification: Under FIPPA, public bodies must notify my office and affected individuals of breaches where there is a risk of significant harm without unreasonable delay.
4) Prevent reoccurrence: Once the immediate crisis is managed, find the cause, fix vulnerabilities, update policies and train your staff.

We have detailed guidance on our website on these steps, including checklists for public bodies that have experienced a breach.

No two breaches are the same, so the specifics of the response to them will vary –

First, you might consider hiring a "breach coach," who can assess the situation and advise on what specific technical experts you might need to hire.

For example, they may advise you to engage a specialist in ransomware negotiations, or they may provide that service themselves, and may advise on which forensic experts to engage, whether as part of the immediate incident response or afterwards as part of an investigation.

Include the information of one or several breach coaches in your breach response plan so that you have it on hand when you need it.

You may want to hire a forensic cyber security expert, and you may consider offering ID theft protection to people.

Our website has two detailed guidance documents and an infographic about identity theft, including steps both affected organizations and individuals can take.

Notifying our office in the wake of a serious breach is mandatory under FIPPA and voluntary under PIPA – it's smart no matter what.

Under FIPPA, the law says that notification of the commissioner must occur without unreasonable delay.

You may interpret that to mean that we should be notified only when you have sufficient information about what is going on to notify us. And it's not that there isn't some wisdom in that; but it is common, particularly if the breach is significant, for organizations to reach out to us early, informally, with formal notification to come later.

This is the practice that the government itself follows, and I would recommend it to you too.

Public bodies and organizations often see my office's role strictly in terms of enforcement.

However, very few of the breach reports that come to us result in us launching an investigation.

In the wake of a breach, our top priority is not to find fault and penalize, it's to help, to support an organization through a difficult time, and mitigate the damage to people and their personal information.

If you're dealing with a breach, don't get stuck in an internal huddle, reach out to us and make use of our expertise.

We're here to help.

…

Earlier, I said my intent wasn't to scare anyone, but to promote vigilance.

I would be remiss, though, if I didn't follow the standard horror movie template and save the scariest big baddie in the realm of cybersecurity for the end of my remarks – quantum computing.

I won't go into a technical discussion, but for our purposes, I think it suffices to say that quantum computers use the principles of quantum physics to process data.

They have the potential to do so magnitudes of degrees faster than today's computers.

It's a troubling convergence of supercharged *software*, with AI accelerating the speed and volume of scraping attacks, and super-charged *hardware*, with quantum anticipated to rapidly accelerate processing speed.

Combined, that gives attackers unprecedented power to steal and exploit people's personal information.

Experts expect that quantum computing be powerful enough to break many of the encryption algorithms we rely on to protect data today.

If you experience a breach today, you may be confident that the stolen data was encrypted to today's highest standards and therefore protected.

Quantum undoes that assurance.

Attackers now can store encrypted data and wait until a time when quantum computing will allow them to break into it – a strategy called "harvest now, decrypt later."

Today's breaches aren't just a problem for today, they're a long-term liability.

When this will happen is a subject of debate, with some estimates saying around 10-15 years from now, others earlier. But we cannot wait that long to prepare. Preparation must start now.

The United States' National Institute of Standards and Technology (NIST) long ago started to make the transition towards post-quantum cryptography (PQC), and new quantum-resistant algorithms were standardized in 2024, with more on the way.

However, this migration takes time, with experts recommending that organizations start now and be fully migrated by 2029 to 2035.

The Government of Canada has a roadmap for the migration to post-quantum cryptography, which recommends that the migration of high-priority systems be complete by 2031, with all remaining systems migrated by 2035.

I believe much of our resilience in the face of the threats posed by quantum computing in this province will come from this room and the institutions you represent.

As you work to migrate your own systems, I believe the higher education sector can and should play a leadership role by training the specialists who will be essential for the post-quantum transition.

We'll need specialists to assist organizations to migrate their systems and applications to quantum-resistant states, and to contribute to the development of additional quantum-resistant algorithms.

Many of those specialists will come from your institutions.

…

The Canadian Centre for Cybersecurity and NIST in the US have published guidance on the migration to post-quantum cryptography. I encourage you all to review their

recommendations."However, I'd also like you to consider a time-tested principle that I touched on briefly earlier, and one that was as valid in the days of folders locked in filing cabinets as it is in the AI era and will be in the age of quantum computing and that's data minimization.

Data you never collected or no longer keep can't be harvested now and decrypted later.

People who entrust you with their information do so for a specific period and a specific purpose.

Keep to those purposes. Collect only the information you need, and keep it only as long as you need it.

This is not only a fundamental principle of good privacy; it's also a strategy that saves you money on having to store and protect unnecessary information and reduces your legal liability in the event of a breach.

Privacy-enhancing technologies are also important here.

As I mentioned, look for services that offer end-to-end encryption where only your organization holds the keys, not the vendor. That protects personal and sensitive data, even in the event of a data breach or vendor misuse.

Also worth considering is synthetic data. If your data can be synthesized, attackers won't get people's information.

Federating data, spreading it across many of places instead of having it in one central place, is also a useful strategy.

Making personal information private by default and redacting public information where possible also reduces what scrapers have to work with.

Here are some questions I'd ask you to take back to your teams:

1) Is our breach incident response plan current, does everyone on it know who to call, including the OIPC, and do you have physical copies of these plans on hand?
2) In the event of a breach, how would we verify information and messages coming from our leaders?
3) Do we have cyber-insurance? Should we consider getting it?
4) How can we break down silos within our organization to make sure we can talk to each other and reduce the "fog of war" in the event of a breach?
5) Are we building privacy into our programs, particularly when we bring new technology on board – how can we do so before launch?
6) What personal information could we stop collecting, delete, or replace with synthetic data?
7) Do we have a plan to migrate to quantum-resistant cryptography – and in the meantime, are we focusing on data minimization to reduce what might be exposed?

…

I'd like to pause before I close to recognize and thank you for the work you do at your institutions.

FIPPA requires you to protect the personal information in your custody or control by making "reasonable security arrangements."

My office is well aware of the complexity of that requirement in the face of the types of threats you're facing today and those on the horizon.

You're not alone in that work. You have each other through BCNET, and you also have my office – we're here to help, and we're in this together.

At the outset, I said that the success of BC's innovation economy will depend on whether people trust how their data is handled.

And that when we're talking about data – we're talking about people, about the first-year student who hands over their sensitive information on the basis of trust.

What I've discussed this morning is aimed at honouring that trust – protecting that data, that person, who they are, and who they can be.

Your institutions are uniquely placed to lead BC in this important work, to building an innovation economy where our rights our respected, our values reflected, and where technology serves people's needs and its use earns their trust.

Thank you. Now I understand we have some time for questions.