TRUTH, TRANSLATED

Plain language for a complicated world. Clear takes on education, geopolitics, political accountability and Big Tech.

EDITION 2 · AUGUST 2026



IN THIS EDITION

The $14.7 billion cheat sheet

A Truth, Translated special report

Encryption under threat

Why the timeline for your bank, email and messages just moved

When fear has two faces

What the Royal Commission report did not fully say

AI detection tools are failing students

And institutions know it

A hard reset for VET

CRICOS applications suspended nationwide

Dr Sukh Rashpal Sandhu | Subscribe free on LinkedIn

EDITION 2 | AUGUST 2026

Contents

Truth, Translated: August 2026 3

Truth, Translated Special Report: The $14.7 Billion Cheat Sheet Just Got Cheated. 4

🔐 The Encryption Protecting Your Bank Account, Your Emails, and Your Messages Is Under Threat. And the Timeline Just Changed. 7

When Fear Has Two Faces: What Australia's Royal Commission Report Did Not Fully Say, But We Must. 10

Sector-focused alarm: 4 in 10 Students Refused. VET at 3.3%. Australia's Student Visa System Is Breaking In Real Time. 14

ASQA Is Verifying Your Resource Licences. If You Cannot Prove Yours Is Valid, Your Registration Is at Risk. 17

Truth, Translated: AI detection tools are failing students and institutions know it. 22

Truth, Translated: A Hard Reset for VET: CRICOS Applications Suspended Nationwide. .. 27

Truth, Translated: AI Replaced Him. The Courts Sided With the Worker. What This Means for the Future of Employment Law. 30

Truth, Translated: A Compliant Document Is Not a Compliant Provider 32

When the Regulator Cancels a Cohort, the Evidence Must Reach Each Person 39

Truth, Translated: The overnight AI shutdown that should worry every country outside America 43

TRUTH, TRANSLATED Google Is Not Saying When AGI Arrives. It Is Asking What Comes After 45

Follow us on LinkedIn 50

About Truth, Translated 52

INTRODUCTION

Truth, Translated: August 2026

By Dr Sukh Rashpal Sandhu

Edition 2 of Truth, Translated: 12 pieces for August 2026.

Plain language for a complicated world. Clear takes on education, geopolitics, political accountability and Big Tech.

This edition brings together 12 pieces first published in Truth, Translated on LinkedIn, each carrying the date it first appeared. Republished pieces appear as they were first published, with their original images; facts, figures and links are as they stood on that date.

In this edition

- Truth, Translated Special Report: The $14.7 Billion Cheat Sheet Just Got Cheated.
- 🔐 The Encryption Protecting Your Bank Account, Your Emails, and Your Messages Is Under Threat. And the Timeline Just Changed.
- When Fear Has Two Faces: What Australia's Royal Commission Report Did Not Fully Say, But We Must.
- Sector-focused alarm: 4 in 10 Students Refused. VET at 3.3%. Australia's Student Visa System Is Breaking In Real Time.
- ASQA Is Verifying Your Resource Licences. If You Cannot Prove Yours Is Valid, Your Registration Is at Risk.
- Truth, Translated: AI detection tools are failing students and institutions know it.
- Truth, Translated: A Hard Reset for VET: CRICOS Applications Suspended Nationwide.
- Truth, Translated: AI Replaced Him. The Courts Sided With the Worker. What This Means for the Future of Employment Law.
- Truth, Translated: A Compliant Document Is Not a Compliant Provider
- When the Regulator Cancels a Cohort, the Evidence Must Reach Each Person
- Truth, Translated: The overnight AI shutdown that should worry every country outside America
- TRUTH, TRANSLATED Google Is Not Saying When AGI Arrives. It Is Asking What Comes After

Thank you to every reader who subscribes, shares and responds.

Warm regards,
Dr Sukh Rashpal Sandhu
Author, Truth, Translated

TRUTH, TRANSLATED

Truth, Translated Special Report: The $14.7 Billion Cheat Sheet Just Got Cheated.

By Dr Sukh Rashpal Sandhu | First published 27 April 2026

A $14.7 billion company. Gone.





A $14.7 billion company. Gone. Not slowly. Overnight.

Chegg was a multi-billion dollar platform built on one simple reality: students would pay to get answers. Homework answers. Assignment answers. Exam answers. Academics called it cheating. Investors called it a $14.7 billion company.

Then ChatGPT arrived. Students stopped paying. Why subscribe when AI gives you the same answers, faster, for free?

The result? **Chegg's stock fell nearly 99% from its peak**. Over 56% of its workforce, cut. Its core product, obsolete.

This is what happens when **access to knowledge becomes free**. Businesses built on selling information are being hollowed out. Not disrupted. Replaced.

Chegg is now pivoting to workforce training and enterprise learning. That is a genuine strategic shift. But it is also an admission: the original model is finished.

The real lesson here is not about Chegg. It is about every sector still charging a premium for access to information that AI can now deliver in seconds. Tutoring platforms. Legal research tools. Certification prep services. Publishing models. The disruption is not coming. It is already here.

AI does not just change pricing. It **commoditises knowledge itself**. What took expertise and time to package and sell can now be replicated at scale, at zero marginal cost.

The question is no longer whether AI will disrupt your industry. The question is how long your business model can survive before it does.

Sukh Sandhu

Director, CAQA | Editor-in-Chief | ISO Auditor | Poet and Writer

Forbes Business Council Member | United Nations Speaker

Top 100 Global Educator | RTO/AI/HE/ISO Expert

Global Advisory Board Member, Forttuna Education Council

TRUTH, TRANSLATED

The Encryption Protecting Your Bank Account, Your Emails, and Your Messages Is Under Threat. And the Timeline Just Changed.

By Dr Sukh Rashpal Sandhu | First published 30 April 2026

Most of the internet runs on a type of encryption called ECDSA-256 (a mathematical lock that scrambles your data so only the intended recipient can read it). Experts believed quantum computers, machines that process information using the principles of quantum physics rather than traditional binary code, were still decades away from cracking it.





On March 31, 2026, Google's Quantum AI team quietly published a paper that changed that assumption.

Their researchers showed that breaking ECDSA-256 encryption could be done using approximately 20 times fewer computational resources than previously thought, and potentially in minutes on sufficiently powerful hardware. A startup called Oratomic added to the picture by showing that RSA-2048, another foundational encryption standard protecting most of the internet, could be threatened with far fewer quantum bits than anyone had previously calculated.

Cloudflare, one of the world's largest internet security providers, responded by moving its post-quantum readiness deadline from 2035 to 2029.

The US government had set 2035 as its target. The security industry is now working to 2029.

There is something else worth understanding. Adversaries do not need quantum computers today to exploit this. The strategy known as **"harvest now, decrypt later"** means that state-level actors and sophisticated organisations are likely already collecting encrypted data today, banking on being able to decrypt it once the hardware catches up. If your data needs to remain confidential through the 2030s, it is already at risk.

AI is accelerating this. The same AI systems driving breakthroughs in medicine, climate modelling, and scientific research are also being used to optimise quantum algorithms,

compress timelines, and reduce the cost of building these machines. This is not a distant theoretical problem. It is an infrastructure problem unfolding right now.

The positive side of this convergence is real. Quantum computing combined with AI is expected to unlock drug discovery at a scale previously impossible, model climate systems with extraordinary precision, and solve optimisation problems that currently take years. The same technology that threatens today's security frameworks will power tomorrow's most important solutions.

But that future requires us to act on the security challenge now, not after the next election cycle or the next budget review.

For individuals, the immediate steps are practical: use applications that have already transitioned to post-quantum encryption standards, keep software updated, and treat any service that has not publicly committed to a quantum-safe roadmap with appropriate caution.

The question is not whether this transition is coming. It is whether institutions, governments, and organisations will move fast enough to get ahead of it.

What are your thoughts? Is your organisation already preparing for the post-quantum shift?

#QuantumComputing #Cybersecurity #AI #Encryption #PostQuantum #DigitalSecurity #Technology #Innovation #ArtificialIntelligence #FutureTech #SukhSandhu

Sukh Sandhu

Director, CAQA | Editor-in-Chief | ISO Auditor | Poet and Writer

Forbes Business Council Member | United Nations Speaker

Top 100 Global Educator | RTO/AI/HE/ISO Expert

Global Advisory Board Member, Forttuna Education Council

TRUTH, TRANSLATED

When Fear Has Two Faces: What Australia's Royal Commission Report Did Not Fully Say, But We Must.

By Dr Sukh Rashpal Sandhu | First published 2 May 2026

On 30 April 2026, Commissioner Virginia Bell AC SC presented the Governor-General, Her Excellency the Honourable Sam Mostyn AC, the Interim Report of the Royal Commission on Antisemitism and Social Cohesion. The Interim Report was tabled in the Australian Parliament on 30 April 2026.



The Interim Report examines the circumstances surrounding the antisemitic Bondi terrorist attack on 14 December 2025 and considers any other issues requiring urgent or immediate action. It contains 14 recommendations.

I have read this report carefully. It is important. It is necessary. And it is incomplete.

Not because the Commissioner failed in her duty. She did not. But because the full human picture of what is tearing Australian communities apart requires us to look honestly at both sides of a single, devastating reality.

The Fear This Report Documents

On the evening of 14 December 2025, two gunmen opened fire on a crowd celebrating Chanukah (the Jewish Festival of Lights) at Bondi Beach. Fifteen people were killed. Forty were wounded. Countless more were traumatised.

That is not a statistic. That is a community gathering to celebrate light, and walking into a massacre.

The report documents what led here. From 7 October 2023, when Hamas launched attacks across Israel killing over 1,200 people, antisemitism in Australia escalated sharply. By August 2024, ASIO, Australia's domestic intelligence agency, raised the National Terrorism Threat Level from POSSIBLE to PROBABLE. The Director-General of Security stated publicly that this threat level was the same as it was at the height of the Islamic State caliphate. ASIO described what it called "a disturbing escalation in the targeting of Jewish interests." At least two attacks on Australian Jewish communities were assessed by ASIO to have been directed by the Iranian government. Jewish schools, synagogues, businesses and prominent community figures were systematically targeted. Arson and vandalism spread across NSW and Victoria. In early 2025, a caravan containing explosives and antisemitic material was discovered in Dural, NSW, sending shockwaves through the Jewish community. Police ultimately determined it was part of a fabricated terrorism plot, a criminal fraud, though the fear it generated in those weeks before that determination was made was entirely real and entirely justified.

Jewish Australians were afraid. Genuinely, justifiably afraid. And the report is right to centre that fear.

But fear in this country did not travel in only one direction.

The Fear This Report Did Not Fully Explore

The same report notes, almost in passing, that Operation SHELTER, NSW Police's response operation launched in October 2023, was established to address "the rise in antisemitic AND Islamophobic incidents." That word "and" carries enormous weight. It appears in a footnote. It deserved a chapter.

Since 7 October 2023, Muslim and Palestinian Australians have also been living in fear, though of a different and less visible kind.

They have feared that expressing grief for Palestinian civilian deaths makes them suspects. They have feared that attending protests places them under scrutiny. They have feared that their religion, their names, their appearance, is enough to draw suspicion regardless of what they have said or done. Many have watched legislation passed in the name of counter-extremism and asked, with genuine anguish, whether the protections being built are being built equally for them.

The ASIO Director-General himself warned in October 2023 that "words matter" and that he saw "direct connections between inflamed language and inflamed community tensions." That observation cuts in every direction. Inflamed language that treats every pro-Palestinian voice as a potential terrorist, that conflates a religion with a foreign government's conduct, that frames Muslim Australians as a security problem rather than a community of human beings with legitimate grief, also drives radicalisation, also fractures social cohesion, also generates the kind of fear that hollows out belonging.

When Geopolitics Lands in Our Streets

The Royal Commission's own report acknowledges that the outbreak of hostilities between the United States, Israel and Iran in February 2026 is likely to have increased the risk of attacks directed at the Australian Jewish community. That acknowledgment is correct and necessary.

But those same hostilities, which saw strikes launched on 28 February 2026 resulting in thousands of deaths across the region, did not generate fear only in Jewish households. They deepened anguish across every Muslim and Arab Australian family watching the Middle East burn, worrying about relatives, fearing being held responsible for the actions of governments they did not elect and do not represent, and navigating the very real risk of Islamophobic violence that tends to spike when conflict erupts.

The Iranian government's conduct, including its state-directed attacks on Jewish Australians documented in this report, is indefensible. It must be named and condemned without qualification. And it must also be noted that ordinary Muslim Australians bear no responsibility for Iran's actions, just as ordinary Jewish Australians bear no responsibility for decisions made in Washington or Jerusalem.

Both communities are losing. Neither is winning. And Australia cannot afford to pick one set of human beings and ignore the other.

What Social Cohesion Actually Requires

The Commission is titled the Royal Commission on Antisemitism AND Social Cohesion. Social cohesion is not a footnote. It is not achieved by protecting one community while surveilling another. It is not achieved by passing legislation that addresses hatred in one direction while leaving Islamophobia largely unaddressed at the national institutional level. It is not achieved by $102 million in security funding for Jewish community sites, welcome and necessary as that is, without equally visible and equally resourced acknowledgment of the fears experienced by other targeted communities.

Social cohesion requires equal recognition of equal humanity.

The 15 lives lost at Bondi Beach demand justice, accountability and lasting structural reform. That is correct and necessary and non-negotiable.

The families of Palestinian civilians killed in the conflict overseas, many of whom have relatives in Australia who have been made to feel their grief is suspicious, also deserve to be seen as human.

Holding both of these things simultaneously is not a political act. It is a moral one.

My Direct Ask to Those Leading this Commission

The Royal Commission's final report must not limit itself to antisemitism as a standalone phenomenon. It must examine the full architecture of hatred in Australia, including Islamophobia, and the way geopolitical conflict is being imported into the daily lives of ordinary Australians from multiple directions.

It must ask whether Australia's counter-terrorism resourcing, surveillance practices and legislative responses have been applied with genuine consistency and fairness across communities.

It must ask whether the language used by public nominated officers in framing these issues has reduced fear or amplified it.

And it must ask, with honesty, what kind of Australia we are building when fear becomes the defining experience of not one but multiple communities at the same time.

We cannot build social cohesion on selective empathy. We cannot build it on silence. And we cannot build it by importing the politics of foreign governments into our streets and asking Australians to choose sides in someone else's war.

The Bondi attack was a crime against humanity. So is every act that tells a human being that their fear does not count.

Sukh Sandhu

Director, CAQA | Editor-in-Chief | ISO Auditor | Poet and Writer

Forbes Business Council Member | United Nations Speaker

Top 100 Global Educator | RTO/AI/HE/ISO Expert

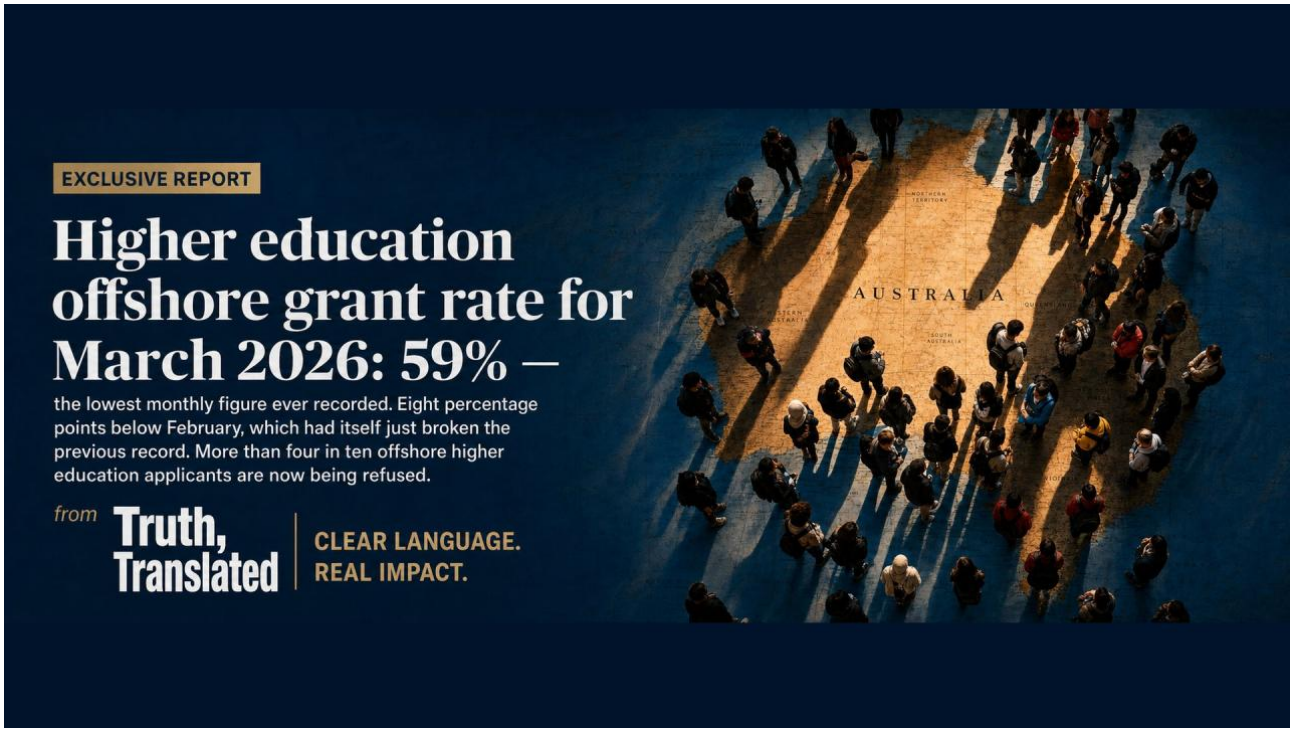
Global Advisory Board Member, Forttuna Education Council

TRUTH, TRANSLATED

Sector-focused alarm: 4 in 10 Students Refused. VET at 3.3%. Australia's Student Visa System Is Breaking In Real Time.

By Dr Sukh Rashpal Sandhu | First published 5 May 2026

Australia's offshore student visa data for March 2026 is out, and the numbers are the most confronting on record for anyone working in international education.



The headline: Higher education offshore grant rate for March 2026: 59%: the lowest monthly figure ever recorded. Eight percentage points below February, which had itself just broken the previous record. More than four in ten offshore higher education applicants are now being refused.

Markets showing strong recovery

China: 96% overall (+22.8pp), HE: 97.3% | PG: 97.7% | Non-Award: 100% | VET: 33.3%

Vietnam: 89.4% (+21.4pp), PG and Non-Award at 100% | HE: 91%

Malaysia: 91.3% (+9.6pp), multiple sectors at 100%

Myanmar: 100%, small volume, consistent approvals

Bhutan: 69.2% (+16.3pp), PG reached 100%

Markets under severe pressure

India: 45.7% overall (-25.3pp), HE: 49% | PG: 98% | VET: 3.3%

Nepal: 25.4%, VET: 7.2% | HE: 27%

Philippines: 44.1% (-35.6pp), VET: 23.2%

Sri Lanka: 46.8% (-21.1pp), HE weak | VET: 3.6%

Pakistan: 51.1% (-20.5pp), only PG performing

Bangladesh: 47.3%, PG strong | VET near unusable

Kenya: 35.1% (-15.6pp)

Thailand: 65.7% (-8pp), VET: 8.3%

Indonesia: 71.7%, HE and PG strong | VET and ELICOS weak

The pattern is structural, not incidental

VET is effectively non-viable across South Asia. India's VET grant rate of 3.3% is not a warning sign, it is a functional closure. Postgraduate is the only structurally sound pathway across most South Asian markets. China and Southeast Asia are moving in the opposite direction, with clean documentation and consistent approval profiles. Chinese applications for higher education fell 39% year on year yet approval rates remain near 97%, confirming the divergence is driven by application quality and financial scrutiny, not volume.

What is driving this

The Department of Home Affairs is now applying Ministerial Direction 106 to conduct holistic financial assessments of whether applicants have genuine access to funds for their **entire** course duration, beyond the one-year minimum the visa criteria requires. This is an undisclosed shift in how an existing rule is being applied. Providers, agents and students have received no official guidance on what a compliant application now looks like. Former Deputy Secretary of Immigration Abul Rizvi has stated publicly that this is an intentional policy shift designed to reduce net migration toward Treasury forecasts, with South Asian markets disproportionately bearing the impact. Visa refusals now also feed into evidence level calculations, a 50% weighting, meaning sustained high refusal rates can push institutions into higher risk categories regardless of their own compliance record.

For providers and agents

Postgraduate is the safest pathway across all affected markets. VET applications from South Asia, the Philippines, and East Africa require immediate risk reassessment. Genuine financial evidence, not just proof of funds for year one, is now the primary determinant of outcome. Students with strong academic profiles are being refused alongside weak applicants because the financial threshold is being interpreted at a standard nobody has publicly defined. That is not an integrity system. That is a migration reduction lever.

If your recruitment strategy hasn't been rebuilt around March 2026 data, it is already out of date.

#AustraliaStudentVisa #InternationalEducation #VETSector #HigherEducation #StudentVisa #SubClass500 #VETAustralia #RecruitmentStrategy #VisaGrantRates #CAQA

TRUTH, TRANSLATED

ASQA Is Verifying Your Resource Licences. If You Cannot Prove Yours Is Valid, Your Registration Is at Risk.

By Dr Sukh Rashpal Sandhu | First published 7 May 2026

Something is happening in the background of regulatory activity that many RTOs (Registered Training Organisations) are not yet aware of, and when they find out, some of them will not be ready for it. ASQA (the Australian Skills Quality Authority) has moved beyond asking RTOs whether they have resources.



Something is happening in the background of regulatory activity that many RTOs (Registered Training Organisations) are not yet aware of, and when they find out, some of them will not be ready for it. ASQA (the Australian Skills Quality Authority) has moved beyond asking RTOs whether they have resources. The regulator is now actively verifying whether those resources are legitimately licenced, checking the terms of End User Licence Agreements, and directly contacting resource developers and providers to confirm purchase dates, licence validity periods, and whether specific RTOs are actually covered by the agreements they claim to hold.

CAQA has been part of this process. We have received inquiries from the regulator regarding specific RTOs: when a purchase was made, whether a current licence is in force, what the EULA permits and restricts, and how our agreements compare to what an RTO

has presented during audit. This is not a new direction that ASQA is contemplating. It is already underway. The regulator has acknowledged internally and through its engagement with industry that unlicensed resource use is a widespread problem across the sector. I welcome this move. It is long overdue.

Why Equipment, Facilities, and Resources Are Now Audit-Critical

The Standards for RTOs 2025 are unambiguous. An RTO must have the resources, facilities, and equipment necessary to deliver the training and assessment it is registered to provide. That obligation extends to the legal right to use those resources, not just physical possession of them.

For years, the sector treated resource compliance primarily as a quality question: are the materials fit for purpose, are they current, do they align with the training package? Those remain valid questions. But ASQA has added a layer that many RTOs are underequipped to address: can you prove you are legally authorised to use what you are using?

This means ASQA is now examining licences and EULAs as compliance documents, not just commercial paperwork. Auditors want to see when a licence was purchased, how long it runs, what it permits, and whether it actually covers the scope of use the RTO is relying on it for. When an RTO cannot produce that evidence, or when the documentation they produce does not match what the original supplier holds on record, the audit outcome can be severe.

What CAQA Has Been Contacted About

To be clear about what this looks like in practice: CAQA and our approved third-party resource providers, Total VET Resources, Precision rto resources and Compliant Learning Resources, have received direct contact from the regulator and from RTOs in the context of audits and compliance reviews. The inquiries have been specific. When did this RTO purchase a licence? Is that licence still current? What does the EULA allow? Does the agreement on file with the RTO match your records?

Some of those inquiries have confirmed legitimate, current arrangements. Others have revealed gaps: licences that have lapsed, licences that were never held in the RTO's name, or situations where someone sold an RTO resources without any valid agreement to do so in the first place.

The picture that has emerged from this process is consistent with what ASQA has identified at the sector level. The problem is not isolated to a handful of providers. It is widespread. Consultants and resellers have been operating for years on the assumption that this area would not be scrutinised closely. That assumption is no longer valid.

ASQA is contacting resource providers directly to verify licence validity, purchase dates, and EULA terms. A number of RTOs have had their registration cancelled without being given the opportunity to rectify. If your records do not match what your supplier holds on file, do not wait for the regulator to contact you.

The Consequences Are Registration-Level

This is not a paperwork issue that gets resolved with a corrective action plan and a follow-up audit. Using resources without a valid licence is a breach of the Standards for RTOs 2025. Depending on the circumstances, it can constitute grounds for conditions on registration, suspension, or cancellation. These are not theoretical outcomes. They are the tools ASQA uses when it determines that an RTO's operations do not meet the requirements for registration.

What makes this more serious than many providers appreciate is that ASQA has not been offering a universal opportunity to rectify. A number of RTOs that have come before the regulator with resource licensing failures have had their registration cancelled outright, without a corrective action period and without the chance to address the gap retrospectively. The cancellation has been the outcome. This is not a warning-first environment where the breach is serious or where the gap between what was claimed and what existed on record is significant. The regulator is treating deliberate or systemic resource non-compliance as grounds for immediate action, and the evidence from recent cases reflects exactly that.

There is also exposure under Australian copyright law independent of the regulatory process. Copyright in training and assessment resources belongs to their developers. A licence to use a resource does not transfer copyright, and it does not permit resale, reproduction, or distribution unless those terms are explicitly included in the agreement. An RTO that has been sold resources by someone who had no right to sell them is in a legally complicated position, regardless of whether they acted in good faith.

The student harm dimension compounds this. When RTOs use resources that are outdated, improperly modified, or sourced from unauthorised copies that have drifted from the originals, students may be assessed against tools that do not meet training package requirements. Qualifications built on that foundation are vulnerable. That is a harm that flows directly from a compliance failure at the resource level.

Having an Active EULA Does Not Cover Everything

This is a detail that has caught a number of RTOs and consultants off guard, and it is worth stating clearly. Holding an active EULA with a resource provider does not automatically licence every product that appears in your scope of registration. A licence is specific. It covers the training products that were purchased, at the time they were purchased, under the terms agreed at that point.

What the regulator has identified, and what CAQA has been asked to verify on a product-by-product basis, is that some RTOs and some consultants acting on their behalf have added qualifications and units of competency to their operations that were never included in the original licensing arrangement. They have an active EULA, so the appearance of compliance is there. But when you look at what that EULA actually covers, the additional training products sit outside it. There was no purchase made for them. There is no payment on record. They have been added to the scope without any corresponding licence.

ASQA is now verifying each training product individually. That means for every qualification and unit an RTO delivers, the regulator is checking not only whether a licence exists but whether that specific product was purchased, when the payment was made, and whether the licence terms in force at that time covered that product. A broad agreement that was legitimate when it was signed in 2021 does not automatically extend to a cluster of units added to scope in 2024.

An active EULA covers what was purchased. If training products have been added to your scope without a corresponding licence purchase and payment on record, you are outside the terms of your agreement regardless of what else is covered.

The practical implication is that RTOs need to audit their own resource arrangements with the same specificity ASQA is now applying. Pull the licence. Check every training product listed on your scope. Check whether each one appears in a paid, documented transaction with your supplier. If there are gaps, those gaps are audit findings waiting to happen.

What a Valid Licence Actually Looks Like

RTOs need to understand what legitimate resource licensing involves. A valid arrangement includes a documented agreement between the RTO and either the original developer or an authorised reseller. It specifies the scope of use, the units or qualifications covered, the number of users or learners permitted, and the duration of the licence. It is held in the RTO's own name, not borrowed from a consultant or transferred informally.

When CAQA or our approved providers, Total VET Resources and Compliant Learning Resources, licence resources to an RTO, the paperwork exists, the records are maintained, and the terms are clear. When ASQA contacts us about a specific RTO, we can confirm or deny a valid arrangement within a short timeframe and provide the relevant documentation. Critically, we can also confirm which specific training products are covered and when each purchase was made. RTOs that have added products to scope without corresponding licence transactions cannot get that same confirmation, because the transaction simply does not exist.

The gap between what an RTO believes they hold and what a supplier can actually verify on a product-by-product and payment-by-payment basis is exactly where audit risk concentrates.

If You Are Unsure About Your Resource Licences, Act Now

If you purchased resources from a consultant, reseller, or provider and you are not certain whether that arrangement constitutes a valid, current licence, the time to resolve that is before your next audit, not during it. Review the documentation you hold. Check whether it names your RTO specifically. Check the expiry date. Check what the EULA permits in terms of the number of learners, the qualifications covered, and the permitted uses of the materials.

If you purchased resources from CAQA, Total VET Resources, or Compliant Learning Resources and you need to confirm the status of your licence, contact us directly through caqa.com.au. We can verify your arrangement and provide documentation suitable for an

audit. If you purchased resources from someone else and you are concerned about whether that arrangement was legitimate, you can also contact us. We will tell you honestly what we know and what steps are available to you.

You can also engage directly with ASQA. The regulator is not approaching this as a punitive exercise aimed at providers who were misled. But providers who take no steps to address known or suspected compliance gaps when they have the opportunity will find themselves in a more difficult position than those who act proactively.

If you are unsure whether your resource licence is valid and current, contact CAQA atcaqa.com.au**or contact ASQA directly. Acting now is significantly less costly than addressing a finding during audit.**

Why This Matters and Why I Welcome It

The VET sector's credibility depends on the integrity of every component of training delivery. Resources are not administrative background noise. They are the materials through which learners engage with the content of their qualifications. When the licensing of those materials is treated as optional or as something that can be navigated around with a cheap resale arrangement, it undermines the quality and the legal standing of the training built on them.

ASQA's decision to pursue this area seriously, including through direct verification with suppliers and detailed examination of EULAs and purchase records, reflects an understanding that surface-level compliance is not enough. The regulator is checking whether the foundations are sound, not just whether the visible outputs look acceptable.

This is the kind of regulatory engagement the sector needs. It creates pressure on those who have been cutting corners, and it protects the RTOs and resource developers who have been doing things properly. The inquiry process that CAQA has been part of is a signal that this scrutiny will continue and likely intensify. Providers who have not yet taken this seriously would benefit from doing so before the regulator's attention turns to them.

Sukh Sandhu

Director, CAQA | Editor-in-Chief, VET Sector Magazine | ISO Auditor | Poet and Writer

Forbes Business Council Member | United Nations Speaker

Top 100 Global Educator | RTO/AI/HE/ISO Expert

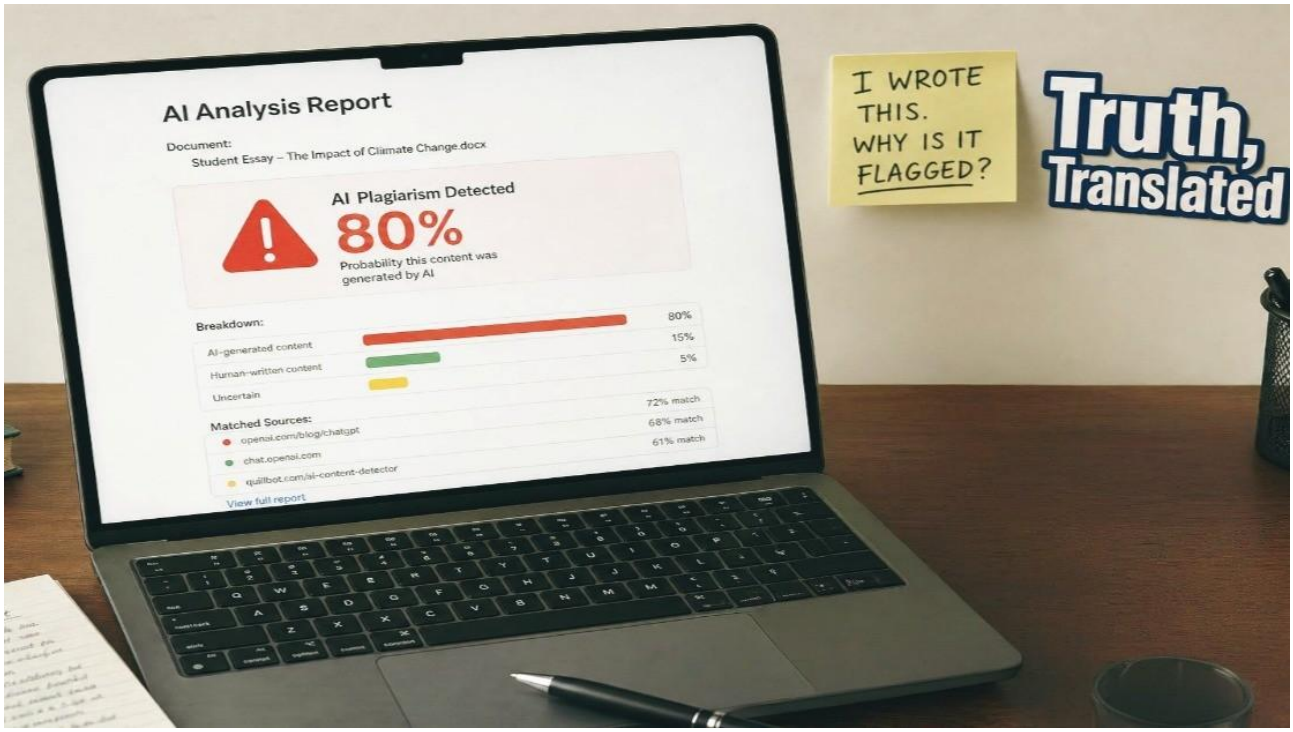
Global Advisory Board Member, Forttuna Education Council

TRUTH, TRANSLATED

Truth, Translated: AI detection tools are failing students and institutions know it.

By Dr Sukh Rashpal Sandhu | First published 8 May 2026

How AI detection technology is failing students, undermining integrity, and why education must urgently course-correct.



ANALYSIS

Students Are Not the Problem. The Tools Are.

There is a crisis unfolding inside educational institutions across Australia and around the world. It is quiet, largely undocumented, and profoundly unfair. It is not a crisis of student dishonesty. It is a crisis of institutional overreliance on technology that was never designed to carry the weight being placed upon it.

AI detection tools have been adopted rapidly and with very little critical scrutiny. Institutions that spent decades refining their approaches to academic integrity have, in the space of a few years, handed that responsibility to algorithms. The results are damaging real students, real careers, and real lives.

I have witnessed this personally, directly, and repeatedly. Not once. Not twice. Again and again and again.

We are flagging human thought as artificial. And nobody seems to find that alarming enough.

What I Have Seen With My Own Eyes

A student sitting next to me, at my desk, sharing my screen. Writing their work word by word, sentence by sentence, in real time, with me watching every keystroke. That same work, produced entirely in front of me, was later flagged as AI-generated.

That is not a hypothetical. That is not an edge case. That happened. And it has happened not once but repeatedly across students from different backgrounds, studying different subjects, at different stages of their programmes.

What makes it even harder to accept is the inconsistency. The same piece of content, checked ten times across ten different sessions, cleared every single time. The eleventh check? Flagged as AI. Not a word changed. Not a sentence restructured. The same content, the same student, the same work. Different result.

This is not quality assurance. This is a roulette wheel dressed up as academic policy.

What AI Detectors Actually Measure

The widespread misunderstanding of how these tools work is at the heart of the problem. AI detectors do not detect AI. That is not a semantic distinction. It is a fundamental technical reality that institutions appear to have either not understood or chosen to ignore.

These tools measure perplexity and burstiness. Perplexity refers to how predictable the text is, how closely the word choices follow statistically likely sequences. Burstiness measures variation in sentence length and structure. AI-generated text tends to score low on both measures, producing smooth, consistent prose with limited variation. That is what the tools are trained to detect.

The brutal irony is this: a student who has worked hard, who has refined their language, who has read widely and developed a clear, organised writing style, will often produce text that looks statistically similar to AI output. Structured. Coherent. Precise. Clean. These are the markers of good academic writing. They are also, apparently, grounds for suspicion.

Meanwhile, a student who writes in rushed, inconsistent prose, with irregular sentence lengths and imprecise vocabulary, will pass. The tools reward poor writing and penalise excellence. In what world does that serve academic integrity?

The better a student writes, the more likely they are to be flagged. We are punishing clarity and rewarding chaos.

Non-Native English Speakers Are Most at Risk

The inequity embedded in this system falls hardest on students for whom English is a second or third language. These students frequently invest extraordinary effort in their written work precisely because they know their language skills are under greater scrutiny. They take more care. They check grammar. They revise sentences. They produce polished, precise work.

And that polish is exactly what gets them flagged.

There is a profound injustice in an international student spending weeks on an assignment, seeking help from language support services, proofreading carefully, producing their best work, and then facing an academic misconduct process because a scanner decided their effort looked too clean.

Researchers studying the intersection of AI detection and English as an Additional Language learners have documented this pattern consistently. The tools are not culturally or linguistically neutral. They carry embedded biases toward informal, native-speaker prose patterns. That is not a minor technical limitation. It is a structural flaw that discriminates.

The False Positive Problem Is Documented and Admitted

This is not a matter of opinion. The companies behind the most widely used AI detection tools have acknowledged that false positives occur. Turnitin, one of the most prominent platforms in this space, has stated publicly that its AI detection capability is designed as an indicator, not a definitive finding. That means the score is supposed to trigger a conversation, not serve as a verdict.

Yet in practice, institutions treat these scores as verdicts. Students receive misconduct notices. Appeals processes are initiated. In some cases, students fail assessments, lose scholarships, or are expelled, on the basis of a percentage score from a tool whose own developers describe it as probabilistic.

Academic research published in recent years has demonstrated false positive rates that should have stopped this practice in its tracks. One widely cited study found that essays written by non-native English speakers were falsely flagged as AI-generated at significantly higher rates than native speaker essays. Another found that simply running human-written text through grammar improvement tools was enough to push the content over detection thresholds.

These findings are not obscure. They are accessible. The question is why institutions are not acting on them.

The Institutional Failure

Educational institutions have an obligation to their students that goes well beyond processing enrolments and issuing qualifications. They are in the business of developing human capability. Academic integrity frameworks exist to protect the value of that development, both for the individual student and for the broader community that relies on qualified, competent graduates.

When an institution hands over its integrity judgements to a tool that demonstrably produces inconsistent, biased results, it is not protecting academic integrity. It is outsourcing its responsibilities to technology because it is cheaper and faster than doing the work properly.

And the students pay the price.

There is also a deeper issue here around trust. Assessment has always been, at its core, a relationship between an educator and a learner. The educator creates conditions in which the learner demonstrates understanding. The educator then exercises professional judgement about whether that demonstration is genuine and sufficient. When you replace professional judgement with algorithmic scoring, you do not just risk false positives. You break the relationship that makes assessment meaningful.

> ***Assessment was always supposed to measure understanding, not prose style. The moment an institution outsources that judgement to a scanner, it has already failed the student.***

The Solution Is Not Complicated

The answer to this problem has existed for as long as education has existed. Talk to the student. Sit with them. Ask them to explain their work. Ask them to write a paragraph in front of you on a concept from their submission. If they can do it, they wrote it. If they cannot, that is your answer. No scanner needed. No algorithm needed. Just a conversation between an educator and a learner.

Viva voce assessments, oral defences, in-person demonstrations, live problem-solving tasks: these have been validated methods of confirming genuine understanding for generations. They do not require technology. They require educator time, and a genuine commitment to getting the answer right rather than getting the process over quickly.

For contexts where scale makes individual conversations impractical, there are still better approaches than blanket scanner use. Staged assessment designs, where students submit plans, drafts, and reflections alongside final work, create a body of evidence that is far harder to fabricate and far more revealing of genuine understanding. Portfolio-based approaches achieve the same. In-class components, supervised writing tasks, and oral components embedded within larger assessments provide corroborating evidence without relying on tools that have been demonstrated to be unreliable.

None of these approaches require institutions to abandon concern about academic misconduct. They require institutions to address that concern through professional practice rather than algorithmic shortcut.

What Institutions Must Do Now

1. **Review and suspend AI detection as sole evidence.**No student should face academic misconduct proceedings on the basis of a detection score alone. Any institution still operating this way needs to review its policy immediately. A detection score is a prompt for further inquiry, not a finding.
2. **Train academic staff on the limitations of these tools.**Many educators using these tools do not understand what they measure or how they fail. That is not a criticism of individual educators. It is a systemic training failure that institutions need to correct.
3. **Redesign assessment for the AI era.**The real question AI raises for assessment is not how to catch students using it. The real question is whether existing assessment designs are still fit for purpose. Many are not. Institutions that respond to AI by adding detection layers without rethinking assessment design are solving the wrong problem.

4. **Provide accessible appeal mechanisms.**Students who are flagged must have access to fair, timely, and transparent appeal processes. The burden of proof must not fall entirely on the student to disprove a statistical score.
5. **Listen to the students.**The students experiencing this are often the most committed, the most careful, and the most invested in their studies. They deserve to be heard, not processed.

A Final Word

I have seen students in tears over this. Students who worked for weeks. Students who polished every sentence and submitted their best work. And a scanner decided their effort looked too good.

That is not academic integrity. That is academic cruelty.

Education exists to expand human capability and opportunity. Every student who is wrongly accused, every appeal that drags on for months, every qualification delayed or denied because of a flawed tool, represents a failure of that purpose. It is a failure that is entirely preventable.

This must stop. Students deserve educators who engage with their work. They deserve institutions that take their responsibilities seriously enough to do the hard, human work of genuine assessment. And they deserve better than a percentage score deciding their future.

If you are in education and you have seen this, say something. Loudly. Because the students often cannot.

Sukh Sandhu Director, CAQA | Editor-in-Chief, VET Sector Magazine | ISO Auditor | Poet and Writer Forbes Business Council Member | United Nations Speaker | Top 100 Global Educator Global Advisory Board Member, Forttuna Education Council | RTO/AI/HE/ISO Expert

TRUTH, TRANSLATED

Truth, Translated: A Hard Reset for VET: CRICOS Applications Suspended Nationwide.

By Dr Sukh Rashpal Sandhu | First published 18 May 2026

The door has just been shut for 12 months. Private RTOs, what happens now?



Signed today, 18 May 2026, the Australian Government has formally suspended all CRICOS applications to the National VET Regulator for a full year. CRICOS (Commonwealth Register of Institutions and Courses for Overseas Students) is the legal gateway to enrolling international students in Australia. Without it, a provider cannot legally take a single international student.

This is not aimed at universities. It is not aimed at government schools. It lands squarely on private RTOs (Registered Training Organisations, meaning independent vocational training providers).

This is not just a freeze on new providers entering the market. Existing registered providers are also blocked from adding new courses to their international scope. The only carve-outs are narrow: adding a delivery location for a course you already hold, or transitioning a superseded qualification to its replacement version.

The freeze runs until at least May 2027. This determination was made under sections 14E and 14F of the Education Services for Overseas Students Act 2000. It is not a policy announcement. It is law, registered and in force today.

The sector has been under sustained scrutiny for years. Visa integrity concerns, compliance failures, and rapid provider growth have built pressure that has now resulted in a full stop, not a speed bump. The government has responded with a blunt instrument. Whether this achieves quality outcomes or simply creates a longer pipeline of the same problems will depend entirely on what reforms are introduced during this 12-month window.

How else will private RTOs suffer?

The immediate pain is obvious. But the damage runs deeper than a frozen application pipeline.

International students generate significant tuition revenue. For many private RTOs, that revenue cross-subsidises domestic delivery, staff salaries, facility costs, and compliance infrastructure. A 12-month freeze on growth does not freeze those costs. Providers who were banking on new international enrolments to fund their next phase of operations now face a serious cash flow gap.

Staff who were hired in anticipation of international growth face uncertain futures. Marketing campaigns already paid for deliver no return. Agents and migration partners who refer international students will redirect those students elsewhere, and rebuilding those referral relationships later will not be straightforward.

Providers in the middle of course development specifically designed for international cohorts now have no pathway to register those courses for overseas students until 2027 at the earliest. That is sunk cost with no near-term return.

And reputationally, every private RTO is painted with the same brush regardless of its compliance record. The providers who have done the right thing for years are carrying the consequences of those who have not.

So what are your options?

Use this 12 months to get your compliance house in order. Providers already registered under CRICOS need to demonstrate they deserve to stay on that register. ASQA (Australian Skills Quality Authority, the national VET regulator) will not be sitting still.

If you hold CRICOS registration, protect it. Audit your student records, marketing materials, delivery practices, and fees and charges documentation now. A cancellation during this period would be catastrophic.

If you were planning to apply, redirect that energy into building genuine quality. The window will reopen. The providers who use this time well will be positioned strongly when it does.

If your business model depends entirely on international students and you have no domestic enrolment base, this is the moment to honestly reassess your strategy.

The government has paused the intake. It has not signalled that the door reopens without conditions.

Now the harder questions.

This determination targets CRICOS. But does anyone genuinely believe the scrutiny stops at the international border?

The same compliance failures, the same low-quality delivery, the same rubber-stamp assessments that drove this freeze exist in the domestic market too. The government knows this. ASQA knows this. The Auditor-General has said this.

If the international market gets a 12-month freeze, what stops a future government from applying equivalent restrictions to domestic RTO registrations? What stops a moratorium on new domestic scope additions? What stops a mandatory revalidation process that effectively freezes existing domestic registrations while reviews are conducted?

None of that has happened yet. But the legal architecture to do it already exists. Sections 14E and 14F of the ESOS Act are the international version of powers that have domestic equivalents under the National Vocational Education and Training Regulator Act 2011.

Private RTOs that believe this is only an international student problem are not reading the room.

The question is not whether more regulation is coming. The question is whether the sector will shape what that regulation looks like or simply absorb whatever arrives.

What do you think? Is a domestic freeze next? Will this clean up the sector or just push the same problems into the domestic market? And if you have questions about what this means for your organisation, reach out.

Sukh Sandhu

Director, CAQA | Editor-in-Chief | ISO Auditor | Poet and Writer

Forbes Business Council Member | United Nations Speaker

Top 100 Global Educator | RTO/AI/HE/ISO Expert

Global Advisory Board Member, Forttuna Education Council

TRUTH, TRANSLATED

Truth, Translated: AI Replaced Him. The Courts Sided With the Worker. What This Means for the Future of Employment Law.

By Dr Sukh Rashpal Sandhu | First published 19 May 2026

When an AI takes your job, can the law step in?



Courts in China just said: yes.

A quality assurance supervisor was hired to monitor AI large language models (LLMs, the same technology that powers tools like ChatGPT) for accuracy and safety. His role was then handed over to a separate LLM. When he declined a demotion that came with a 40 per cent pay cut, the company terminated his contract.

Two courts, including the Hangzhou Intermediate People's Court, ruled the dismissal was unlawful.

The reasoning is significant. The court found that replacing a worker with AI does not constitute a "major change" in a company's circumstances under Chinese employment law. Relocation and mergers qualify. AI substitution, on its own, does not.

The court also found the alternative role offered was unreasonable. A 40 per cent salary reduction is not a genuine alternative. It is a constructive exit.

This ruling joins a similar 2025 arbitration decision in Beijing, where an AI replacement of a map data collector was found unlawful because the company had effectively transferred the risk of its own technology choices onto its employees.

What does this mean for workers elsewhere?

In Australia, the framework is different. A "genuine redundancy" (a situation where an employer no longer requires a role to exist) is by definition not considered unfair dismissal under current law, provided consultation obligations are met and redundancy entitlements are paid. Several major Australian companies, including Telstra, Atlassian, Canva, Afterpay, and WiseTech, have cited growing AI use as part of recent workforce reductions.

The Chinese rulings do not bind any other jurisdiction. But they raise questions that every legal system will eventually have to answer:

At what point does technological substitution become an employment law issue, not just a business decision?

Should the cost of AI adoption be shared between employers and workers, or carried entirely by the worker who loses their livelihood?

Is a 40 per cent pay cut, offered as the only alternative to termination, ever a genuine option?

These are not hypothetical questions. They are already happening.

The pace of AI integration across industries is outrunning the pace of legal and policy reform. Governments, tribunals, and courts around the world are being asked to interpret frameworks designed for industrial-era disruption and apply them to something categorically different.

What courts decide in the next five years will shape the social contract around work for a generation.

Sukh Sandhu

Director, CAQA | Editor-in-Chief | ISO Auditor | Poet and Writer

Forbes Business Council Member | United Nations Speaker

Top 100 Global Educator | RTO/AI/HE/ISO Expert

Global Advisory Board Member, Forttuna Education Council

TRUTH, TRANSLATED

Truth, Translated: A Compliant Document Is Not a Compliant Provider

By Dr Sukh Rashpal Sandhu | First published 6 June 2026

What Mid City Education v ASQA tells every RTO and CRICOS provider, in law and in plain terms



A training college spent more than eighteen months rebuilding its compliance. It rewrote policy after policy, filed three rounds of affidavit evidence, that is written statements sworn to be true, and on its own account drafted nine revised policies in a single night before the hearing. Its chief executive told the Tribunal the college wanted to do more than satisfy the regulator, it wanted to go beyond compliance. After a four day hearing, the Tribunal still affirmed the decision to cancel the college's registration, both as a registered training organisation and across every course at every location for overseas students.

The case is Mid City Education Pty Ltd and Australian Skills Quality Authority [2026] ARTA 916, decided by General Member M Darian-Smith in Brisbane on 28 May 2026. For anyone who runs, advises, or audits a provider, it repays close reading, because it draws a line the sector too often blurs. The Tribunal was not asking whether the college had produced better paperwork. It was asking whether the college had become a better organisation, one that would stay compliant once intensive regulator oversight ended. Those are different questions, and the gap between them decided the outcome.

Background: the provider, the regulator, and the review

Three terms run through the case. A registered training organisation, or RTO, is a provider approved to deliver nationally recognised vocational qualifications, such as a Certificate IV in Kitchen Management. A CRICOS provider is one also registered, on the Commonwealth Register of Institutions and Courses for Overseas Students, to teach students who hold Australian student visas. Mid City was both.

ASQA, the Australian Skills Quality Authority, is the national regulator. It administers two separate legislative schemes at once, the National Vocational Education and Training Regulator Act 2011, which governs training providers generally, and the Education Services for Overseas Students Act 2000, which governs the teaching of overseas students, and it can impose sanctions ranging up to cancellation when a provider falls short. When a provider disputes an ASQA decision, it can seek review in the Administrative Review Tribunal. The Tribunal conducts what is called merits review, which means it does not merely check whether ASQA followed a fair process. It re-decides the question itself, on the evidence before it, and reaches what the law calls the correct or preferable decision, that is the right answer on the merits rather than merely a defensible one. That feature shapes everything that follows.

Two regulatory schemes, one regulator

ASQA cancelled Mid City's registration on 10 December 2024 under both schemes. The training side was cancelled under subsection 36(2)(f) and section 39 of the National Vocational Education and Training Regulator Act 2011 (the NVR Act), the principal law governing RTOs. The overseas student side, covering all courses at all locations, was cancelled under subsection 83(3)(c) of the Education Services for Overseas Students Act 2000 (the ESOS Act), the law that protects international students through measures such as tuition and refund assurance, which guarantees a student a refund or an alternative place if a course stops part way through, and accurate reporting. A single set of operational failures can breach both at once, and here it did.

One provision of the NVR Act becomes the hinge of the case. Section 22A requires an RTO to demonstrate both a commitment and the capability to deliver quality training. In plain terms, willingness alone is not enough. The provider must also be able to deliver, on its own, without the regulator holding its hand.

The matter also sits across a Standards transition. The reviewable decisions were framed against the Standards for Registered Training Organisations 2015, but by the time the Tribunal assessed current compliance the 2025 Outcome Standards and Compliance Standards had commenced on 1 July 2025, alongside the National Code of Practice 2018. The college therefore had to demonstrate present compliance against the newer instruments, not merely show that old breaches had been patched.

A merits review looks forward, not only back

Because the Tribunal re-decides the matter as at the date of its own decision, it treated the process as a continuum and considered every successive version of the college's policies, right up to the hearing. Past conduct still mattered. Subsection 36(3)(b) of the NVR Act allows a decision-maker to weigh an RTO's conduct since the notice of intention to cancel, the formal warning that must precede any cancellation, was given, and the

authorities are clear that belated, pressured fixes carry little weight. Compliance adopted only against the looming finality of cancellation, as the Tribunal has put it before, may give no confidence of ongoing commitment. The practical message is simple. Repairing an old breach does not win the case. The question is whether the systems will hold once no one is watching.

What was rectified, and what remained in dispute

By the hearing, the issues had narrowed. ASQA accepted that several concerns had been met: fit and proper persons, which is whether the people who own and govern the college are suitable to run one, cooperation with the regulator, overseas student enrolment duration under National Code Standard 8.2, student support, and trainers and assessors under Outcome Standards 3.2 and 3.3. The Tribunal noted, pointedly, that ASQA read even this remediation as evidence of how poorly the college had managed its own compliance, because most of it arrived only after repeated prompting over many months.

Five issues remained live, and the college lost on every one.

The five unresolved areas of non-compliance

Meeting legislative requirements. This concerned Compliance Standard 20, which requires an RTO to comply with all applicable laws, and the reporting duties in sections 19 and 47H of the ESOS Act. Providers must enter prescribed information into PRISMS, the Provider Registration and International Student Management System, within strict deadlines, for example when a student does not commence or when a refund is paid. The recurring failure was a confusion between two different situations that carry different rules. A student default is when a student walks away from a course. A provider default is when the provider itself cannot deliver the course. They are governed by different sections, section 47H and section 46D, with different refund and notification timeframes. The chief executive carried that confusion across several of the college's own policies, even after ASQA had explained the distinction repeatedly. There were also two short late reports, of five days and one day.

Training and assessment, including facilities. This engaged Outcome Standards 1.1 and 1.8, National Code Standards 11.1 and 11.2, and sections 21 and 22 of the NVR Act, which together require that training is well structured and properly paced and that assessment, facilities and equipment are fit for purpose. The Certificate IV in Kitchen Management contains around thirty-three units and takes roughly eighteen months of full-time study, yet the strategy pointed to a four hour summative assessment, the final confirmation of competence. The college said assessment was in fact gathered progressively throughout delivery, but nowhere did it identify or explain where those additional hours occurred, despite repeated requests. Students, prospective employers and the regulator were left unable to see how competence across every unit, including units involving food safety, would actually be assessed. A small but telling detail compounded the concern. Seven trainer session plans for seven different sessions of a cookery unit all bore the same date, explained from the witness box as a typo.

Managing third parties. This engaged Compliance Standard 17(2) and section 21A of the ESOS Act, which require a provider to keep and publish an accurate list of its education agents and to notify the regulator of agent arrangements. The college's new agents policy

did not even reference the standard or the section it was meant to satisfy, and its records did not reconcile across its own systems. The website listed thirty-one agents, PRISMS listed thirty-one, yet ASQAnet, the regulator's records, showed only two current arrangements. Asked why inactive agents remained published, the chief executive explained that removing them would cause reputational damage and that the agents wanted to keep working with the college. The Tribunal found that commercial and reputational interest, rather than the regulatory duty to maintain an accurate list, appeared to be driving the conduct. That is the wrong order of priorities, and a decision-maker reads it as a signal about how a provider behaves when no one is watching.

Commitment and capability. This engaged section 22A of the NVR Act and section 11(1)(e) of the ESOS Act, the latter requiring that the provider has the clearly demonstrated capacity to deliver education of a satisfactory standard. The Tribunal accepted that the college was committed. Commitment was never the problem. Capability was, and the Tribunal was not persuaded the college had demonstrated it. The law requires both.

Assessment validation. This engaged Outcome Standard 1.5. Validation is the quality assurance process in which suitably skilled and credentialled people review whether a provider's assessment practices and judgements are sound. When the Standards were updated in 2025, the college did not conduct a fresh validation. It carried over an earlier report for a hospitality unit and simply updated the references from the 2015 Standard to the 2025 Standard, leaving the substantive comments unchanged. A duplicated form reference deepened the impression of a validation process that existed more on paper than in practice.

The Tribunal accepted that some complaints were minor and easily rectified. The clearest example was the automotive workshop for the Certificate III in Light Vehicle Mechanical Technology, where ASQA wanted a simple inventory of tools and equipment under Outcome Standard 1.8 and the college had supplied photographs instead. That was a quick fix. The serious problems were not. As the Tribunal has long held, failures in assessment go to the very reason the VET system exists, because unless assessments are valid and qualifications go only to those who have earned them, the value of every certificate is put at risk.

Self-regulation and the role of the regulator

The most instructive moment in the decision was not a document. In cross examination, the chief executive, who had taken over only about three weeks before the cancellation decision, described the college as a customer of ASQA. ASQA charged an annual fee, he reasoned, and was therefore responsible to the college, with a duty to tell it how to fix its problems rather than simply find fault and close it down. Counsel for ASQA corrected him plainly. ASQA is the regulator, with an oversight function, not a customer relationship.

The Tribunal found this revealed a fundamental and seriously mistaken view of the respective roles of an RTO and its regulator. The point is not a debating nicety. An RTO bears the onus of self-regulation. It must have systems and processes that keep it compliant at all times, without the regulator standing at its shoulder. The same attitude surfaced when the chief executive dismissed ASQA's detailed findings as cosmetic, the

kind of thing that could be corrected in a couple of hours. The Tribunal disagreed. Some faults were minor, but the remainder went to the integrity of the qualifications themselves, and the misunderstanding coloured how the Tribunal read everything else.

A recurring pattern: compliance that peaks at audit and fades afterward

The behaviour at the centre of this case is not unique to one college. It reflects a pattern the sector has watched repeat itself for years. A provider mobilises intensively in the lead up to an audit, during the audit itself, and again when the regulator raises concerns, producing a sudden surge of rewritten policies, fresh registers and neatly assembled evidence. Then registration is secured or restored, the pressure lifts, attention returns to marketing, enrolments and revenue, and the daily disciplines that actually keep a provider compliant are quietly allowed to lapse. Compliance is treated as an event to be survived once every cycle, rather than a condition to be maintained continuously. Mid City is simply a vivid illustration. Much of its strongest evidence was produced under deadline pressure, including nine policies drafted in a single night, and the Tribunal was unconvinced precisely because the work looked like a response to the threat of closure rather than the normal output of a well-run organisation.

This pattern matters because of a gap that documentation can hide. A complete, well-worded policy proves only that someone could write the policy. It does not prove that trainers follow it, that assessors apply it, that reporting happens on time, or that records are accurate when no auditor has asked for them. The Tribunal made the same point by drawing on earlier authority that a system being in place does not mean the system was adhered to or that the people engaged by the provider were actually performing their obligations. The danger of front-loaded compliance is that it produces the artefacts of quality without the practice of quality, and the two can drift a long way apart between audits.

It matters, too, because the entire regulatory model assumes self-regulation in the spaces between audits. ASQA oversees thousands of providers and cannot, and is not required to, maintain close supervision of any one of them indefinitely. Auditing is periodic by design. The system only works if a provider polices itself on the ordinary days when no one external is checking, which is the overwhelming majority of days. A provider that performs only when observed has, in effect, outsourced its conscience to the audit calendar, and that is exactly the arrangement the Tribunal declined to reward. The risk-of-recurrence inquiry that decided the sanction exists for this reason. Regulators and tribunals have learned that effort displayed under scrutiny is a poor predictor of conduct once scrutiny is withdrawn.

Most importantly, it matters because quality in vocational training is not generated at the moment of audit. It is generated every day a student is taught and assessed. A learner does not benefit from a validation report that was tidied up for a hearing. The learner benefits from being taught the right content, by a competent trainer, and assessed honestly against the real requirements of the qualification, on the day they sit in the classroom or the workshop. When daily operations slip, the harm is immediate and personal. Students pay fees and invest time for a credential that may not reflect genuine competence, employers hire on the strength of certificates that may not mean what they claim, and other honest providers carry the reputational cost of a system that looks

unreliable. This is why the objects of the NVR Act in section 2A, the stated purposes the law exists to serve, speak of protecting students and protecting Australia's reputation for VET, and why the Tribunal returns to those objects rather than to the neatness of any individual document.

There is also a simple organisational truth underneath all of this. Compliance that is bolted on in bursts is fragile, because it depends on a few people working heroically under pressure, and it collapses the moment those people are busy, absent or gone. Compliance that is built into daily routine is durable, because it survives staff turnover, growth, and the return to full trading that undid Mid City. The providers that come through audits calmly are usually the ones for whom the audit changes very little, because the evidence the auditor wants is simply a by-product of how they already operate. That is the difference between a compliant document and a compliant provider, and it is the difference this decision rewards.

Sanction: protection, not punishment

The hierarchy of sanctions under subsection 36(2) of the NVR Act runs from a written direction, through shortening the term or amending the scope of registration, to suspension under section 38, and finally to cancellation under section 39. The governing principle, drawn from a line of authority, is that a sanction is not a punishment. Its purpose is protective, and the task is to identify the least sanction needed to protect the public interest, because anything heavier would tip into punishment, which is impermissible.

Two questions framed the choice. First, the risk that the provider would not become and remain substantially compliant, the risk of recurrence. Second, the effect of the decision on students and on Australia's reputation as a provider of high quality VET. On the first, the Tribunal was not persuaded. For much of the dispute the college had operated under an agreed enrolment freeze, having undertaken not to enrol new students while the review ran, so it had been running quietly. Much of the improved documentation had been produced through eighteen months of close, step by step input from the regulator, and ASQA oversees thousands of providers and cannot maintain that level of scrutiny indefinitely. If the college returned to full trading it would be far busier, and the Tribunal had no confidence its policies would be implemented compliantly once self-regulating.

Cancellation followed, not to punish a provider that had plainly worked hard, but to maintain the protective objects of the legislation in section 2A of the NVR Act. Under subsection 39(3), a provider whose registration is cancelled cannot reapply for two years, or such shorter period as the regulator considers appropriate.

Implications for RTOs and CRICOS providers

Reviews look forward. Repairing an old breach does not win a case, because the regulator and the Tribunal weigh the risk that the same problems return. The real test is whether your systems hold when nobody is watching.

Stop treating the audit as the deadline. The most common and most dangerous habit in the sector is to surge effort before, during and after an audit, then relax once registration is secure. Compliance built in bursts is fragile and fades, while compliance built into daily

routine endures. Plan your operations so the evidence an auditor wants is simply a by-product of how you already work every day.

Last minute, regulator driven fixes carry little weight. There is a real difference between compliance you build and sustain yourself and compliance assembled at the eleventh hour under threat of closure.

Documents are necessary but never sufficient. Strong templates count for little if the people who own them cannot explain and operate them in daily practice. Capability lives in your team, not in your filing system.

Understand the relationship. A provider that treats the regulator as a paid adviser obliged to fix its problems has misread the core obligation of self-regulation under the NVR Act, and that misunderstanding will colour how a Tribunal reads everything else.

Treat assessment and validation as core. Unexplained assessment hours, carried-over validations, and duplicated forms are not housekeeping. They go to the worth of every qualification you issue, and the Standards on assessment exist to protect students.

Do not let commercial interest override regulatory duty. The reasons a provider gives for its choices tell a decision-maker a great deal about what it values.

Cancellation protects the public. The Tribunal looks for the least sanction needed to protect students and the sector, which means cancellation can be the correct outcome even where it causes genuine hardship.

The hard truth running through the decision is simple. Once a provider is defending a cancellation, demonstrating capability is far harder than demonstrating effort. Real compliance cannot be manufactured overnight. It has to be lived in daily operations, long before the auditor arrives, long before the notice lands, and long before there is anything left to defend.

Case reference: Mid City Education Pty Ltd and Australian Skills Quality Authority [2026] ARTA 916 (28 May 2026), General Member M Darian-Smith, Brisbane. Hearing 9 to 12 February 2026, decision 28 May 2026.

Source: https://www.austlii.edu.au/cgi-bin/viewdoc/au/cases/cth/ARTA/2026/916.html

TRUTH, TRANSLATED

When the Regulator Cancels a Cohort, the Evidence Must Reach Each Person

By Dr Sukh Rashpal Sandhu | First published 6 June 2026

What RWTP v ASQA tells every RTO, consultant, and student about cancelling qualifications, in law and in plain terms



A retired firefighter in his seventies paid three thousand dollars for a Certificate III in Carpentry, then learned, while overseas and without internet, that the regulator had cancelled it. His was one of 740 qualifications cancelled across former students of the same college. He did not respond to the notice in time. He represented himself. He lives with post-traumatic stress disorder, and the Tribunal watched him struggle to hold his composure at an earlier hearing. He still won. The Tribunal set the cancellation aside and sent the matter back to the regulator to reconsider.

The case is RWTP and Australian Skills Quality Authority [2026] ARTA 924, decided by Senior Member J Lye in Brisbane on 28 May 2026. It was handed down the same day as the better known Mid City Education decision, and the two read as companion pieces. Mid City is about a provider that could not show it deserved to keep its registration. RWTP is about the limits on what a regulator may do to the students left behind when a provider is shut down. Together they describe both edges of the same power, and anyone who runs, advises, or studies at an RTO should understand both.

Background: the provider, the regulator, and the student

A registered training organisation, or RTO, is a provider approved to issue nationally recognised vocational qualifications. ASQA, the Australian Skills Quality Authority, is the national regulator under the National Vocational Education and Training Regulator Act 2011 (the NVR Act). When a person disputes an ASQA decision, they can seek review in the Administrative Review Tribunal, which conducts merits review, meaning it re-decides the question for itself on the evidence and reaches what the law calls the correct or preferable decision, rather than simply checking whether the regulator followed a fair process.

The provider here was Harbour City College. ASQA investigated it, then cancelled its registration on 27 June 2025, effective 1 July 2025. That deregistration decision was not the subject of this case, and the Tribunal had no quarrel with it. What was challenged was a second, separate kind of decision that followed: the cancellation of the qualifications already issued to the college's former students.

Two different powers, two different evidence tests

This is the distinction at the heart of the case, and it is one the sector routinely misses. Cancelling a provider's registration and cancelling a student's qualification are governed by different provisions and answer different questions.

Deregistration looks at the provider and asks whether it should keep operating. Cancelling a qualification reaches into the past and strips a credential a person already holds. Under section 56(1) of the NVR Act, ASQA may cancel a qualification on stated grounds. It relied on two. The first, in section 56(1)(a), requires ASQA to be satisfied on reasonable grounds that the college did not provide, or arrange, all or part of the assessment necessary for that person to earn the qualification. The second, in section 56(1)(d), requires ASQA to be satisfied that cancellation is appropriate in all the circumstances because of the action it has taken against the provider.

ASQA cancelled the qualifications of former Harbour City College students by relying wholly on the findings from its investigation of the college. Notices of intention went to around 765 affected people, and the 740 who did not respond, RWTP among them, had their qualifications cancelled. It did not separately investigate RWTP. It used the same templated notice and the same templated statement of reasons for every affected person. That approach is what unravelled.

Why a college-level investigation did not carry the student-level decision

The Tribunal had no doubt about the strength of the evidence ASQA had gathered against the college's new owners, who had bought it in December 2024. The investigators found backdated records, enrolments and certificates issued in compressed timeframes, training and assessment strategies that were only drafts for future delivery, no evidence of trainers or assessors, and owners who had given false and misleading information. On that evidence the deregistration was sound.

The problem was timing and reach. RWTP enrolled around February 2024 and received his certificate in March 2024, well before the new owners arrived. The damning evidence almost all concerned the period after December 2024. When ASQA cancelled RWTP's

qualification, it leaned on the new owners' draft strategies, on the absence of assessors that was only established for the new owners' period, and on record alterations that the investigation traced to the new owners. None of that told the Tribunal what had happened in RWTP's own enrolment, under the previous owner, a year earlier.

The burden point mattered too. ASQA argued that RWTP had not established that the college assessed him. The Tribunal corrected this squarely. RWTP did not carry the onus. Under section 56(1)(a), the Tribunal itself had to be satisfied, on reasonable grounds, that the assessment was not provided, and on the evidence it could not be. There were even signs, from the investigation's own sampling, that some assessment had occurred under the previous owner, albeit to a less than satisfactory standard. A regulator cannot cancel a person's qualification by pointing at a provider's later failings and then asking the person to prove their own competence.

On the second ground, whether cancellation was appropriate in all the circumstances, the Tribunal found the question finely balanced but came down the same way. It accepted there is a strong public interest in the integrity of the VET system and in supporting the regulator's work. But it weighed against that the absence of evidence for RWTP's own period, and the regulator's decision to apply the deregistration findings to every student, under both owners, without asking whether those findings actually fit each person's circumstances. The regulator, the Tribunal said, had been undiscriminating.

The principle that decides cohort cancellations

The Tribunal's reasoning here is the part worth pinning to the wall. A decision to cancel an entire cohort of qualifications is a very serious matter, because the financial and reputational consequences fall directly on students and their employers, many of them young, low income, or relying on recognition of prior learning. Precisely because the power is so broad, its exercise deserves close and exacting scrutiny. And precisely because it can be exercised against many people at once, the regulator is not relieved of its duty to make sure its action is supported by evidence and defensible for each individual whose qualification it cancels.

In plain terms, the right to act against a group is not a shortcut around proving the case for each member of the group. A finding that a provider was non-compliant by the middle of 2025 does not, on its own, prove that a particular student was never assessed in early 2024. The regulator must frame its investigation with precision, cast the scope of its action to match the evidence it actually holds, and reassess that scope as the investigation develops.

A note on process the Tribunal could not ignore

There was also a striking inconsistency in how RWTP was treated. ASQA decided his case needed urgent action and shortened his time to respond to seven days. Then it took almost three months after that window closed to make the final decision. The Tribunal noted that this did not change the outcome, since RWTP received full merits review in any event, but it observed that the rush at the front end sat oddly with the delay at the back end, and that RWTP had a fair grievance about how little chance he was given to engage before a decision was made.

The lessons, for everyone with a stake in the sector

For regulators and those who model their thinking, the discipline is evidence that matches scope. A strong case against a provider is not automatically a case against each student the provider taught. Group action is permissible, but every individual cancellation must stand on evidence that reaches that individual's time, course, and circumstances. Where a provider changes hands, the investigation must separate what happened under each owner rather than smearing one period's failings across another.

For RTOs, the case is a reminder that your records outlive you. When a college collapses, it is the students who are left exposed, and the only thing standing between them and an automatic cancellation is the quality and clarity of the records the provider kept. Accurate enrolment data, genuine assessment evidence, properly documented recognition of prior learning, and clean version control on training and assessment strategies are not bureaucratic niceties. They are the proof that protects a learner long after the provider is gone. Backdating, draft strategies passed off as real, and ambiguous paperwork do not just sink the provider. They put every qualification it ever issued under a cloud.

For consultants and compliance advisers, the lesson is to document recognition of prior learning as rigorously as taught delivery. RPL is lawful and, for an experienced candidate like RWTP, often entirely appropriate. But an RPL decision that leaves no trace of how competence was assessed against each unit is indistinguishable, on the file, from no assessment at all. If you advise providers, build RPL records that a stranger could pick up years later and understand exactly how the judgement was reached. The absence of that trail is what turns a legitimate qualification into a contested one.

For students, the case is both a warning and a reassurance. The warning is to keep your own evidence. RWTP prevailed in part because he produced decades of licences and industry records that made his eligibility for recognition plausible. The reassurance is that the system contains a check. A regulator's decision, even one affecting hundreds of people, can be tested by an independent tribunal, and the tribunal will insist the evidence actually reaches you before your qualification is taken away.

The deeper theme connects this case to its companion. In Mid City, the lesson was that a provider must be genuinely capable, not merely well-documented, because compliance has to be lived. In RWTP, the lesson is the mirror image aimed at the regulator. Power exercised over many people must still be earned one person at a time, on evidence, with care. Both decisions protect the same thing, the integrity of the VET system and the people who depend on it, and both insist that integrity is proved in the detail, not asserted in the abstract.

Case reference: RWTP and Australian Skills Quality Authority [2026] ARTA 924 (28 May 2026), Senior Member J Lye, Brisbane. The Tribunal set aside the cancellation of RWTP's Certificate III in Carpentry and remitted the matter to ASQA to reconsider in light of its reasons.

TRUTH, TRANSLATED

Truth, Translated: The overnight AI shutdown that should worry every country outside America

By Dr Sukh Rashpal Sandhu | First published 16 June 2026

Last Friday, with almost no warning, the United States government told the AI company Anthropic that no foreign national anywhere, not even its own overseas employees, could use its two most powerful models. To comply, the company **switched them off for everyone**.



Last Friday, with almost no warning, the United States government told the AI company Anthropic that no foreign national anywhere, not even its own overseas employees, could use its two most powerful models. To comply, the company **switched them off for everyone**. If you live outside America, you were locked out within hours.

Here is why this matters to you, even if you have never typed a single word into an AI tool.

These systems are fast becoming the electricity of the modern economy. Soon they will sit inside our banks, our hospitals, our universities and our government offices. When one country can decide overnight who is allowed to switch on and who is cut off, a basic tool quietly turns into a lever of power.

Look at who was shut out. Not only rivals like Russia or Iran. Australia, Britain, Canada and New Zealand, America's closest intelligence partners in the alliance known as the **Five Eyes**, were treated exactly the same as adversaries. Trusted friends discovered they were on the wrong side of a line nobody told them existed.

I am not pretending the security worries are imaginary. A powerful tool can be misused, and any serious government must weigh that. My concern is ***how this was done***.

Where was the warning? Where was the evidence? The order arrived late on a Friday afternoon, named no specific threat, and shut out an entire category of people purely by their nationality. The public is being asked to trust a decision it is not permitted to examine.

I have spent a working lifetime inside compliance and standards, and here is what every auditor knows in their bones. A rule without transparency is not security. It is ***power wearing the costume of safety***.

Real safety is built in the open. It is written down, applied evenly to friend and rival alike, explained to the people it affects, and answerable to someone. A decision made behind a curtain, with no reasons given, is the opposite of that.

So the question is bigger than one order in one week. Who gets to control the tools the whole world is coming to depend on, and will the rest of us ever be allowed to see how those choices are made?

Ask it loudly, and keep asking. A curtain only becomes a wall when nobody demands to see what is behind it.

Sukh Sandhu

Director, CAQA | Editor-in-Chief | ISO Auditor | Poet and Writer

Forbes Business Council Member | United Nations Speaker

Top 100 Global Educator | RTO/AI/HE/ISO Expert

Global Advisory Board Member, Forttuna Education Council

TRUTH, TRANSLATED

TRUTH, TRANSLATED Google Is Not Saying When AGI Arrives. It Is Asking What Comes After

By Dr Sukh Rashpal Sandhu | First published 16 June 2026

The report puts no date on human level AI. It starts where that question ends and asks a different one: what artificial superintelligence would actually be, and what it could never be.



By Sukh Sandhu



A headline doing the rounds this month says Google has "revealed what comes after AGI, and it is shocking." The report sitting underneath that headline says almost the opposite. The distance between the noise and the page is the only thing worth reading here.

On 10 June 2026, a team at Google DeepMind posted a report to arXiv, the open archive where scientists share work before formal publication, titled "From AGI to ASI." The author list is not a marketing department. It includes Shane Legg, who co-founded DeepMind, and Marcus Hutter, the researcher who built the mathematical theory the report leans on, with the work led by Tim Genewein. These are some of the people who helped define what machine intelligence means in the first place, now asking what happens after machines reach human level.

Two terms carry the piece, so here they are in plain words. AGI, or artificial general intelligence, means software that is roughly as capable as a typical person across most thinking tasks. ASI, or artificial superintelligence, means a system that beats not just talented individuals but large, well organised teams of human experts in almost every field. The report adds a third idea, a theoretical ceiling it calls Universal AI. Think of it the way physicists think of the speed of light: a hard limit that is real and useful to know about, yet impossible to actually reach or build. It exists on paper, as mathematics, and cannot be computed in the real world.

Most coverage of artificial intelligence treats AGI as a finish line. This report starts where that conversation usually stops, and asks a quieter question. If machines reach human level, how might they keep going, what could slow them down, and what would such a future actually allow. The answer it gives is deliberately unglamorous. It is a

counterweight to the people promising paradise and the people promising doom in equal measure.

The heart of the report describes four roads from human level intelligence to something beyond it. The authors are careful to say these are not rivals, and that the real future could travel several at once.

The first road is simply scaling: more computing power, more data, bigger systems. This is the brute force approach that produced most of the recent leaps you have read about, from chatbots to Google's own Gemini models. The second is a paradigm shift, meaning not a larger version of today's approach but a genuinely new one. The report treats this as possible but unpredictable, for the obvious reason that you cannot put a breakthrough on a calendar. The third is recursive self improvement, which means an AI that improves the next version of itself, with each improvement making the following one a little easier, like compound interest applied to machine design. The fourth is the most human sounding of the four: collectives. Rather than one giant mind, superintelligence might emerge from large numbers of human level systems working together, the way companies, markets and research institutes let groups of people achieve what no single person could. Such collectives might be built on purpose, as automated firms, or arise on their own through competition between rival AI services.

Against those four roads the report sets six frictions, the speed bumps that could stall the journey. The most discussed is the data wall: the supply of high quality human written text is finite, and current estimates suggest the world runs low on it later this decade. There are the rising economic and physical costs of building ever larger systems. There are the limits of today's neural network designs. There is the plain possibility that research simply gets harder as the easy gains are used up. There is a subtle one the authors call the abstraction barrier, the worry that an AI raised entirely on human ideas may struggle to think genuinely new thoughts, ones that sit outside the human frame altogether, which would quietly cap how far it can go. And there is what they call a deliberate slowdown, the recognition that governments, regulators and public pushback are not outside forces acting on the technology. They are part of the equation itself.

Here is the sentence the dramatic headlines left out. Even a superintelligence would not be all knowing, and it would not be magic. The report points out that such a system could never play perfect chess, because the number of possible games is larger than any machine the physical universe could ever hold. More importantly, being far cleverer than us would not automatically let it cure ageing, deliver limitless fusion energy, upload human minds onto computers, or reverse climate change. Those are questions about the physical world, settled by experiment and resources, not problems that raw thinking power can wish away. The report also refuses to name a date. It does not tell you when any of this arrives, and it says plainly that both the speed and the final ceiling remain deeply uncertain.

So why should this matter to someone who does not work in technology. Because AI headlines are now part of ordinary life, arriving every week with words like "shocking" and "revealed" attached, and the ability to tell a hype headline from the careful document beneath it has become a basic literacy, no different from reading any news report with a critical eye. The practical message of this report is genuinely reassuring, and it is precisely

the part that did not travel. The most likely shape of change, the authors suggest, is not a single science fiction morning when everything flips at once. It is a long series of disruptions spread across many fields, year after year. That is not a wave you brace against in one moment. It is a climate you can read, plan for, and adapt to.

The truth, translated, is this. Google's researchers did not reveal something shocking about the future. They did something rarer and more useful. They drew a careful map of a place no one has been, marked the roads and the roadblocks, and wrote in clear letters across the top that it will not be magic. The shock, if there is one, is how calm the experts sound.

The calm itself is the real story, and the real question. It sits at odds with the fear that soaks through public conversations about this technology. That fear is not about job losses or a lower standard of living. It is a deeper, older fear about being replaced, about becoming the second cleverest species on the planet, about losing control of our own story. It is the fear that we are building our own gods, and that they will not be kind.

This report does not dismiss that fear, but it does reframe it. By describing a future of limits, frictions, and slow gains, it suggests that the thing we are building is not a god. It is a tool. A profoundly powerful, transformative, and dangerous tool, but a tool nonetheless. And the questions that matter about tools are not about worship or terror. They are about ownership, control, purpose, and safety. The quiet, technical language of this paper is an invitation to put down the science fiction and pick up the engineering, to stop looking at the sky and start looking at the blueprints.

Primary reference

Genewein, T., Franklin, M., Lerchner, A., Orseau, L., Albanie, S., Bales, A., Wyeth, C., Chan, S., Gabriel, I., Leibo, J. Z., Dafoe, A., Hutter, M., Graepel, T., & Legg, S. (2026). *From AGI to ASI* (arXiv:2606.12683). arXiv. https://doi.org/10.48550/arXiv.2606.12683

Supporting sources cited within the report

Bloom, N., Jones, C. I., Van Reenen, J., & Webb, M. (2020). Are ideas getting harder to find? *American Economic Review, 110*(4), 1104-1144. (Underpins the "research gets harder" friction.)

Drexler, K. E. (2019). *Reframing superintelligence: Comprehensive AI services as general intelligence* (Technical Report No. 2019-1). Future of Humanity Institute, University of Oxford. (Underpins ASI emerging from competing AI services.)

Good, I. J. (1965). Speculations concerning the first ultraintelligent machine. *Advances in Computers, 6*, 31-88. (Underpins the recursive self-improvement and intelligence-explosion idea.)

Hutter, M. (2005). *Universal artificial intelligence: Sequential decisions based on algorithmic probability*. Springer. (The AIXI framework behind the Universal AI ceiling.)

Legg, S., & Hutter, M. (2007). Universal intelligence: A definition of machine intelligence. *Minds and Machines, 17*(4), 391-444. (The intelligence measure grounding the AGI and ASI definitions.)

List, C., & Pettit, P. (2011). *Group agency: The possibility, design, and status of corporate agents*. Oxford University Press. (Underpins the automated-firm and collective-agent pathway.)

Sutton, R. (2019, March 13). *The bitter lesson*. http://incompleteideas.net/IncIdeas/BitterLesson.html (Underpins the scaling pathway argument.)

Villalobos, P., Ho, A., Sevilla, J., Besiroglu, T., Heim, L., & Hobbhahn, M. (2024). Position: Will we run out of data? Limits of LLM scaling based on human-generated data. In *Proceedings of the 41st International Conference on Machine Learning* (Vol. 235, pp. 49523-49544). PMLR. (The data-wall estimate; the authors place the exhaustion of public human text between 2026 and 2032.)

STAY CONNECTED

Follow us on LinkedIn

Enjoyed this edition? Each of these newsletters publishes new pieces on LinkedIn. Subscribe, like and share to read them first, and join the conversation in the comments.



Truth, Translated You are reading this newsletter

Plain language for a complicated world: education, geopolitics, political accountability and Big Tech.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/truth-translated-7450724196525035520/



Letters of the Heart

Real stories, timeless wisdom, and soul-deep inspiration.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/letters-of-the-heart-7329273590376210432/



The Hope Dispatch

Uplifting news, inspiring stories, and positive change from around the world.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/the-hope-dispatch-7338912777081638914/



The HealthTech Herald

The latest breakthroughs at the intersection of health and technology.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/the-healthtech-herald-7338913774701027328/



Connecting the Dots

Where compliance, risk management, digital innovation, and bottom-line growth intersect.

Subscribe on LinkedIn

https://www.linkedin.com/newsletters/connecting-the-dots-7397392757360648192/

Follow the author: Dr Sukh Rashpal Sandhu on LinkedIn

BACK PAGE

About Truth, Translated

Truth, Translated is a LinkedIn newsletter written by Dr Sukh Rashpal Sandhu. Plain language for a complicated world. Clear takes on education, geopolitics, political accountability and Big Tech.

Read and subscribe

New editions continue to be published on LinkedIn. Search for Truth, Translated on LinkedIn to subscribe.

Disclaimer

The views in this collection are the author's own. Pieces are reproduced as first published, and facts, figures and links are current at their original publication date. Nothing in this collection is legal, financial, medical, migration or regulatory advice.

Copyright

Copyright Dr Sukh Rashpal Sandhu 2026. All rights reserved. Images are reproduced as they appeared in the original posts.

Truth, Translated | Dr Sukh Rashpal Sandhu | sukh@sukhsandhu.com

Masthead

Author: Dr Sukh Rashpal Sandhu

Publisher: Dr Sukh Rashpal Sandhu

Edition: 2, August 2026

Contact: sukh@sukhsandhu.com

Truth, Translated is published on LinkedIn by Dr Sukh Rashpal Sandhu. This edition collects the pieces in the order they were first published, with their original images.