The HealthTech Herald

Your trusted source for the latest breakthroughs at the intersection of health and technology.

EDITION 2 · AUGUST 2026



IN THIS EDITION

Gene Editing Reaches Toddlers

What the Casgevy Decision Means for Children With Sickle Cell Disease and Thalassaemia

The First Cholesterol Pill of Its Kind

Why an Oral PCSK9 Inhibitor Matters, and What We Still Do Not Know

The AI Second Reader That Found Missed Liver Cancers

Inside a 10,000-Patient Trial

Resetting the Immune System

Early CAR T-Cell Results Offer New Hope in Severe Lupus

When Health AI Becomes a Medical Device

How Australia's Regulator Is Drawing the Line

Dr Sukh Rashpal Sandhu | Subscribe free on LinkedIn

EDITION 2 | AUGUST 2026

Contents

The HealthTech Herald: August 2026 3

Gene Editing Reaches Toddlers: What the Casgevy Decision Means for Children With Sickle Cell Disease and Thalassaemia 4

The First Cholesterol Pill of Its Kind: Why an Oral PCSK9 Inhibitor Matters, and What We Still Do Not Know 12

The AI Second Reader That Found Missed Liver Cancers: Inside a 10,000-Patient Trial 20

Resetting the Immune System: Early CAR T-Cell Results Offer New Hope in Severe Lupus 28

When Health AI Becomes a Medical Device: How Australia's Regulator Is Drawing the Line 36

Follow us on LinkedIn 45

About The HealthTech Herald 47

INTRODUCTION

The HealthTech Herald: August 2026

By Dr Sukh Rashpal Sandhu

Edition 2 of The HealthTech Herald: 5 pieces for August 2026.

Your trusted source for the latest breakthroughs at the intersection of health and technology.

This edition brings together 5 new features written for this edition. New features are marked "Written for this edition".

In this edition

- Gene Editing Reaches Toddlers: What the Casgevy Decision Means for Children With Sickle Cell Disease and Thalassaemia
- The First Cholesterol Pill of Its Kind: Why an Oral PCSK9 Inhibitor Matters, and What We Still Do Not Know
- The AI Second Reader That Found Missed Liver Cancers: Inside a 10,000-Patient Trial
- Resetting the Immune System: Early CAR T-Cell Results Offer New Hope in Severe Lupus
- When Health AI Becomes a Medical Device: How Australia's Regulator Is Drawing the Line

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

Warm regards,
Dr Sukh Rashpal Sandhu
Author, The HealthTech Herald

THE HEALTHTECH HERALD

Gene Editing Reaches Toddlers: What the Casgevy Decision Means for Children With Sickle Cell Disease and Thalassaemia

By Dr Sukh Rashpal Sandhu | Written for this edition

On 1 July 2026 the United States Food and Drug Administration widened the approval of Casgevy, the world's first CRISPR gene-editing medicine, to include children as young as two. It is a genuine milestone for two inherited blood disorders, but the treatment remains demanding, expensive and still early in its long-term story.

A decision years in the making

When the FDA first approved Casgevy (exagamglogene autotemcel) in December 2023, it made history twice. It was the first medicine approved in the United States to use CRISPR/Cas9 genome editing, the technology whose inventors won the Nobel Prize in Chemistry in 2020. And it offered something sickle cell disease had never had before: a one-time treatment designed to stop the painful crises that define the illness, using the patient's own cells.

That first approval, however, was limited to people aged 12 and older. A second approval in January 2024 extended the medicine to transfusion-dependent beta thalassaemia, again for patients 12 and over. Younger children, who often suffer some of the most frightening complications of these diseases, had to wait.

On 1 July 2026, Vertex Pharmaceuticals announced that the FDA had approved expanded use of Casgevy for people aged two years and older with sickle cell disease and recurrent vaso-occlusive crises, or with transfusion-dependent beta thalassaemia. The FDA's own product page, updated the next day, lists the indication as covering "patients aged 2 years and older". Trade coverage described it as the first authorised genetic therapy in the United States for children as young as two with either condition.

Vertex estimates that about 5,500 additional children in the United States became eligible under the wider label.

Two diseases, one shared problem

Sickle cell disease and beta thalassaemia are both inherited disorders of haemoglobin, the iron-rich protein inside red blood cells that carries oxygen around the body.

Sickle cell disease is caused by a change in the gene for the beta chain of haemoglobin. The faulty haemoglobin makes red blood cells stiffen and bend into a crescent or "sickle"

shape when they release oxygen. These misshapen cells clog small blood vessels, starving tissues of oxygen. The result is sudden, severe episodes of pain known as vaso-occlusive crises, along with a lifetime risk of stroke, organ damage, infections and early death. The FDA has noted that the disease affects about 100,000 people in the United States, and it is far more common in people with ancestry from sub-Saharan Africa, the Middle East, India and the Mediterranean.

Beta thalassaemia involves a different set of changes in the same gene, which reduce or stop the production of the beta chain. People with the most severe forms cannot make enough working haemoglobin and become profoundly anaemic. Many depend on blood transfusions every few weeks for life. Those transfusions keep them alive but load the body with iron, which then needs daily medication to remove and can still damage the heart, liver and hormone glands over time.

In both diseases, the trouble begins after birth. Before birth, and for a few months afterwards, babies make a different form of haemoglobin called fetal haemoglobin. It does not rely on the faulty adult beta chain. As infants grow, a genetic switch turns fetal haemoglobin production down and adult haemoglobin production up. That switch is exactly when the symptoms of sickle cell disease and thalassaemia begin to appear.

Doctors have long noticed that some people with these conditions carry natural genetic variations that keep their fetal haemoglobin high into adulthood. Those people tend to have much milder disease. Casgevy is built on that observation.

How the treatment works

Casgevy does not repair the faulty adult haemoglobin gene. Instead it switches fetal haemoglobin back on.

The process, as described by the FDA when it first approved the therapy, runs roughly like this:

1. **Collection.** Doctors collect the patient's own blood-forming stem cells.
2. **Editing.** In a specialised manufacturing facility, those stem cells are edited with CRISPR/Cas9. The edit targets a genetic control region that normally suppresses fetal haemoglobin, so the edited cells can make fetal haemoglobin again.
3. **Conditioning.** The patient receives high-dose chemotherapy, called myeloablative conditioning, to clear space in the bone marrow for the edited cells.
4. **Infusion.** The edited cells are given back as a single infusion, as part of a stem cell transplant.
5. **Engraftment and recovery.** The edited cells settle into the bone marrow and begin producing red blood cells rich in fetal haemoglobin. Patients stay under close hospital care until their blood counts recover.

Because the cells come from the patient, there is no need for a matched donor and no risk of the donor's immune cells attacking the recipient, which is a serious danger in traditional bone marrow transplants from another person.

> *The key idea is simple: rather than fixing a broken part, Casgevy reactivates a backup system the body already knows how to use.*

What the evidence showed in adults and teenagers

The original 2023 approval rested on a single-arm study, meaning there was no placebo or comparison group. That is common for rare, severe diseases where withholding a potentially curative treatment would be unethical.

In the sickle cell trial, 44 patients were treated. Of the 31 who had been followed long enough to be evaluated, 29, or 93.5 per cent, went at least 12 consecutive months without a severe vaso-occlusive crisis during the 24-month follow-up period. The FDA reported that all treated patients achieved successful engraftment, with no graft failure or rejection.

Those results were striking for a disease in which many patients had been hospitalised repeatedly every year.

What the evidence shows in younger children

The July 2026 expansion is supported by ongoing Phase 3 studies in children, known by Vertex as CLIMB-141 and CLIMB-151.

According to the company's announcement and independent coverage:

Condition	Children enrolled (as reported)	Evaluable for efficacy	Result
Sickle cell disease	11 children aged 5 to 11	8	All 8 were free of severe vaso-occlusive crises for at least 12 consecutive months
Transfusion-dependent beta thalassaemia	15 children aged 5 to 12	9	8 of the 9 became transfusion independent, with a median duration of about 20 months reported

The reported results come only from children aged five and older. Coverage of the decision reported that the FDA extended the indication to that age group by extrapolating efficacy from the data in children aged five to under 12. The paediatric results were described as consistent with those seen in older patients.

Readers should note two things. First, the numbers are small. Eight and nine evaluable children are enough to show a strong signal for a rare disease, but not enough to detect uncommon side effects or to know how every child will respond. Second, all of the children with evaluable data were aged five or older. The approval for the youngest age group therefore relies on extrapolation from older children, which regulators sometimes accept when the biology is the same and the disease is severe.

Why age matters so much

Haematologists have long argued that the earlier these diseases can be addressed, the better.

Sickle cell disease does damage quietly as well as loudly. Alongside the painful crises, children can suffer silent strokes, which may not cause obvious symptoms but can affect

learning, memory and attention. Repeated blockages harm the spleen, kidneys, lungs and bones. In thalassaemia, every year of transfusions adds iron to the heart and liver.

Dr Haydar Frangoul, a paediatric transplant specialist at TriStar Centennial Children's Hospital in Nashville who has been closely involved in Casgevy research, framed the goal as treating children "before years of cumulative damage from these life-shortening diseases take hold".

There is also a practical argument. Young children generally tolerate stem cell collection and transplant differently from adults, and outcomes in traditional bone marrow transplants for sickle cell disease have historically been better in younger patients. Whether the same holds for gene-edited cell therapy will become clearer as more children are treated and followed.

The hard parts: what families need to know

It would be easy to describe Casgevy as a "cure" and leave it there. The reality is more complicated, and families deserve the full picture.

It is a transplant, not a tablet

The treatment involves stem cell collection, weeks to months of manufacturing, high-dose chemotherapy and a hospital stay that can last weeks. The FDA listed common side effects including low platelet and white blood cell counts, mouth sores, nausea, musculoskeletal and abdominal pain, vomiting, fever with low white cell counts, headache and itching. Most of these are linked to the chemotherapy conditioning rather than the gene editing itself.

Conditioning chemotherapy has lasting risks

Myeloablative chemotherapy is powerful. It can affect fertility, which is a particularly serious consideration when the patient is a young child who cannot yet weigh that choice for themselves. Families are typically counselled about fertility preservation options before treatment. Like other intensive chemotherapy, it also carries a small long-term risk of other health problems, which is why patients are enrolled in long-term follow-up.

Long-term effects are still being studied

The first patients received Casgevy in clinical trials only a few years ago. The FDA has required long-term follow-up studies to monitor safety and effectiveness over time. Questions that remain open include how durable the benefit will be across decades, whether there are any delayed effects of the editing process, and how outcomes compare with other options such as donor transplants or newer medicines.

Cost and access are enormous barriers

Independent reporting puts the list price of Casgevy at about US$2.2 million per patient in the United States. Treatment can only be delivered at specialised, authorised treatment centres. Vertex says more than 75 such centres are active in the United States. The company's financial reports, as summarised in health media, recorded 64 patients infused in 2025 and more than 500 who had started the treatment process since launch.

Those numbers are small compared with the number of people living with these conditions. For context, sickle cell disease affects millions of people worldwide, most of them in sub-Saharan Africa and India, where transplant centres capable of delivering this treatment are rare.

How regulators moved faster this time

The paediatric expansion was reviewed under the FDA's Commissioner's National Priority Voucher programme, a pilot scheme introduced in 2025 to speed up reviews of medicines judged to be strategically important for the United States. The same programme was used later in July 2026 for Merck's new cholesterol pill, enlicitide, covered elsewhere in this edition.

Supporters argue the scheme gets important treatments to patients sooner. Critics have raised questions about how products are selected and whether speed could come at the expense of thorough review. For Casgevy, the underlying safety and efficacy evidence had already been examined in two earlier approvals, which made a faster review of the paediatric data less contentious than it might be for a brand-new medicine.

Where does this leave Australia?

Sickle cell disease and thalassaemia are present in Australia, particularly among families with Mediterranean, Middle Eastern, South and South-East Asian and African heritage. Patient groups such as Thalassaemia and Sickle Cell Australia support families who manage regular transfusions, iron chelation and hospital visits.

An FDA decision does not make a medicine available in Australia. Any gene therapy would need to be assessed and registered by the Therapeutic Goods Administration, and public funding decisions are made separately through national health technology assessment processes. Families in Australia should speak with their haematologist about which treatments, trials and transplant options are actually available to them, rather than assuming that a United States approval applies here.

What the FDA decision does provide is evidence. Regulators around the world look closely at each other's decisions and the data behind them. The paediatric trial results will form part of the international evidence base that other agencies, including Europe's medicines regulator, can draw on.

Putting the milestone in perspective

It is worth stepping back to appreciate how far this field has moved.

- In 2012, the CRISPR/Cas9 gene-editing system was described in a landmark scientific paper.
- In 2019, the first sickle cell patient was treated with what would become Casgevy in a clinical trial.
- In December 2023, the FDA approved it for people aged 12 and older with sickle cell disease.
- In January 2024, the approval expanded to transfusion-dependent beta thalassaemia.

- In July 2026, it was extended to children aged two and older.

Seven years from first patient to toddler approval is extraordinarily fast for any new class of medicine, let alone one that edits human DNA.

Yet the milestone also exposes the gap between scientific possibility and practical reach. A therapy that requires a specialist transplant centre, intensive chemotherapy and a multimillion-dollar price can only help a small fraction of the people who need it. Researchers are already working on approaches that might one day edit blood stem cells directly inside the body, without collection, manufacturing or chemotherapy. Those ideas remain experimental, and none is approved.

What this means for readers

- **For families affected by sickle cell disease or thalassaemia:** this is real progress, backed by trial data, but it is not a simple or risk-free option. Any decision involves weighing transplant risks, fertility, long-term unknowns and access against the burden of the disease itself. Your haematology team is the right place to start that conversation.
- **For health professionals:** the paediatric data, while small, add weight to the case for treating these diseases earlier in life. Expect growing interest from families who read international news and want to know their options.
- **For policymakers:** the gap between approval and access is now the central challenge. Gene therapies with seven-figure price tags will test every health system's ability to pay, even when the long-term cost of the disease is also high.

Casgevy and its sibling therapy

Casgevy was not the only gene therapy for sickle cell disease approved by the FDA on 8 December 2023. On the same day, the agency approved Lyfgenia (lovotibeglogene autotemcel) from bluebird bio. The two take different routes to a similar goal.

Feature	Casgevy	Lyfgenia
Technology	CRISPR/Cas9 gene editing	Lentiviral vector that adds a modified haemoglobin gene
How it helps	Switches fetal haemoglobin back on	Adds a gene for an anti-sickling form of adult haemoglobin
Source of cells	Patient's own stem cells	Patient's own stem cells
Conditioning	High-dose chemotherapy	High-dose chemotherapy
Reported United States price	About US$2.2 million	About US$3.1 million

Lyfgenia's label carries a boxed warning about blood cancers, which were seen in some patients treated with an earlier version of the therapy. Casgevy works by editing an existing gene control region rather than inserting a new gene with a viral vector, and its label does not carry that warning. Both therapies require long-term follow-up, and neither is suitable for every patient.

A short glossary

- **CRISPR/Cas9:** a gene-editing system, adapted from a bacterial defence mechanism, that uses a guide molecule to direct a cutting enzyme to a precise spot in DNA.
- **Fetal haemoglobin:** the form of haemoglobin made before birth, which does not depend on the adult beta chain affected in these diseases.
- **Vaso-occlusive crisis:** a painful episode caused by sickled cells blocking small blood vessels.
- **Transfusion dependence:** needing regular blood transfusions to survive and stay well.
- **Myeloablative conditioning:** high-dose chemotherapy that clears the bone marrow so new stem cells can take hold.
- **Engraftment:** the point at which transplanted stem cells settle in the bone marrow and start producing new blood cells.

Questions families may want to ask their specialist

1. Is my child's disease severe enough that a transplant-based therapy is worth considering?
2. What are the alternatives, including donor transplants, newer medicines or clinical trials?
3. What would the conditioning chemotherapy involve, and what fertility preservation options exist?
4. How long would the whole process take, from cell collection to recovery?
5. What follow-up would be needed in the years afterwards?
6. Is this therapy, or a trial of it, available where we live?

My view

The Casgevy story is one of the clearest examples of a fundamental scientific discovery turning into a treatment that changes lives within a single decade. Extending it to young children reflects a growing confidence among regulators and clinicians that the benefits of acting early can outweigh the risks of an intensive procedure.

But the most important lesson may be about equity. The people who carry the heaviest burden of sickle cell disease live largely in parts of the world where this treatment is out of reach. A breakthrough that cannot travel is only half a breakthrough. The next chapter of this story should be measured not only by how young a patient can be treated, but by how many patients, in how many countries, can be treated at all.

This article is general information and is not medical advice. Anyone affected by these conditions should seek guidance from their own specialist.

Sources

- Vertex Pharmaceuticals via Business Wire, "Vertex Announces US FDA Approval for Expanded Use of CASGEVY for the Treatment of People Ages 2 Years and Older With Sickle Cell Disease or Transfusion Dependent Beta Thalassemia", 1 July 2026, https://www.businesswire.com/news/home/20260701379449/en/Vertex-Announces-

US-FDA-Approval-for-Expanded-Use-of-CASGEVY-for-the-Treatment-of-People-Ages-2-Years-and-Older-With-Sickle-Cell-Disease-or-Transfusion-Dependent-Beta-Thalassemia

- US Food and Drug Administration, "CASGEVY" product page, content current as of 2 July 2026, https://www.fda.gov/vaccines-blood-biologics/casgevy
- US Food and Drug Administration, "FDA Approves First Gene Therapies to Treat Patients with Sickle Cell Disease", 8 December 2023, https://www.fda.gov/news-events/press-announcements/fda-approves-first-gene-therapies-treat-patients-sickle-cell-disease
- Healio, "FDA expands Casgevy approval to children 2 years and older with sickle cell disease", 2 July 2026, https://www.healio.com/news/hematology-oncology/20260702/fda-expands-casgevy-approval-to-children-2-years-and-older-with-sickle-cell-disease
- Patient Worthy, "FDA Expands Casgevy Use to Children as Young as 2 Years Old", 8 July 2026, https://patientworthy.com/2026/07/08/fda-expands-casgevy-use-to-children-as-young-as-2-years-old/
- Pharmaceutical Technology, "Vertex Secures US FDA Approval for Expanded Use of CASGEVY in Young Children", 2 July 2026, https://www.pharmaceutical-tech.com/news/vertex-secures-us-fda-approval-for-expanded
- European Medicines Agency, "Casgevy: EPAR", https://www.ema.europa.eu/en/medicines/human/EPAR/casgevy

THE HEALTHTECH HERALD

The First Cholesterol Pill of Its Kind: Why an Oral PCSK9 Inhibitor Matters, and What We Still Do Not Know

By Dr Sukh Rashpal Sandhu | Written for this edition

On 16 July 2026, Merck announced that the United States Food and Drug Administration had approved Lipfendra (enlicitide), the first PCSK9 inhibitor that can be swallowed as a daily tablet rather than injected. It lowers "bad" LDL cholesterol by more than half in trials, but the evidence that it prevents heart attacks and strokes has not yet arrived.

Cholesterol, briefly explained

Few numbers in medicine carry as much weight as LDL cholesterol. Low-density lipoprotein is the particle that carries cholesterol through the bloodstream and deposits it in artery walls. Over decades, those deposits build into plaques. When a plaque ruptures, a clot can form and block the artery, causing a heart attack or a stroke.

The link between LDL cholesterol and cardiovascular disease is among the most thoroughly established relationships in medical science. It is supported by genetic studies, population studies and dozens of randomised trials. Broadly, the lower and longer LDL cholesterol is kept down, the lower the risk of cardiovascular events.

That is why cholesterol-lowering medicines are among the most prescribed drugs in the world, including in Australia, where heart disease remains a leading cause of death.

Where PCSK9 fits in

The liver is the body's main cholesterol clearing house. Its cells are studded with LDL receptors, which grab LDL particles out of the blood and pull them inside to be broken down.

PCSK9 (proprotein convertase subtilisin/kexin type 9) is a protein that tags those receptors for destruction. More PCSK9 means fewer receptors, which means more LDL left circulating in the blood. Less PCSK9 means more receptors and lower LDL.

Scientists learned this partly from families. In the early 2000s, researchers found that some people with very high cholesterol carried genetic changes that made PCSK9 overactive. Others carried changes that switched it off, and those people had unusually low LDL cholesterol and much lower rates of heart disease, apparently without harm. That natural experiment made PCSK9 one of the most attractive drug targets in cardiology.

The first PCSK9 medicines, evolocumab (Repatha) and alirocumab (Praluent), are antibodies given by injection every two or four weeks. Inclisiran (Leqvio), a different approach that reduces the liver's production of PCSK9, is injected twice a year after initial doses. These medicines are very effective, and large outcome trials have shown that the antibodies reduce cardiovascular events in people who already have heart disease.

But injections are a barrier for some patients, and cost has limited their use in many health systems. A tablet has long been a goal.

What enlicitide is

Enlicitide is not a small conventional drug molecule, nor is it an antibody. It is a macrocyclic peptide, a ring-shaped chain of amino acids engineered to bind tightly to PCSK9 and block it. Peptides are usually broken down in the gut before they can be absorbed, which is why most protein-based medicines must be injected. Getting a peptide like this to work as a once-daily tablet is a technical achievement in its own right.

The approved product is a 20 milligram tablet taken once a day. According to Merck, it is indicated as an add-on to diet and exercise to reduce LDL cholesterol in adults with hypercholesterolaemia, including heterozygous familial hypercholesterolaemia, an inherited condition that causes high cholesterol from birth.

The evidence behind the approval

The approval rests on two Phase 3 randomised, placebo-controlled trials in Merck's CORALreef programme.

CORALreef Lipids

This was the larger trial, published in the *New England Journal of Medicine* in February 2026 (Navar and colleagues). Its intention-to-treat population was 2,909 adults, randomised two to one to enlicitide or placebo. Participants were adults with high cholesterol who had, or were at high risk of, cardiovascular disease, and the medicine was tested on top of their existing care.

- At 24 weeks, LDL cholesterol fell by 57.1 per cent with enlicitide, compared with a rise of 3.0 per cent with placebo.
- The difference between groups, adjusted for other factors, was 55.8 per cent in favour of enlicitide.
- Safety and tolerability did not differ meaningfully from placebo.

CORALreef HeFH

This smaller trial, published in *JAMA* (Ballantyne and colleagues), enrolled 303 adults with heterozygous familial hypercholesterolaemia who were already on background cholesterol-lowering treatment.

- The placebo-adjusted reduction in LDL cholesterol at 24 weeks was 59.4 per cent.
- Diarrhoea (7 per cent versus 2 per cent) and dizziness (9 per cent versus 4 per cent) were more common with enlicitide than placebo, though the differences were not statistically significant.

Summary table

Trial	Participants	Comparison	LDL cholesterol reduction versus placebo at 24 weeks
CORALreef Lipids	2,909 adults at high cardiovascular risk	Enlicitide 20 mg versus placebo (2:1)	About 56 per cent
CORALreef HeFH	303 adults with heterozygous familial hypercholesterolaemia	Enlicitide 20 mg versus placebo (2:1)	About 59 per cent

In a separate trial, CORALreef AddOn, presented in 2026, Merck reported that enlicitide lowered LDL cholesterol significantly more than the non-statin oral medicines ezetimibe, bempedoic acid or the two combined at eight weeks.

For perspective, a 55 to 60 per cent reduction on top of existing therapy is in the same broad range as the injectable PCSK9 antibodies. That is a large effect for a pill.

The crucial caveat: cholesterol numbers are not heart attacks

Here is the most important point for general readers. The FDA approved enlicitide because it lowers LDL cholesterol, which is an accepted surrogate marker for cardiovascular risk. The trials were not designed or large enough to show that it actually prevents heart attacks, strokes or deaths.

Merck's own announcement states that the effect of Lipfendra on cardiovascular illness and death has not been determined. A dedicated outcomes trial, CORALreef Outcomes, has completed enrolment of more than 14,500 participants, and its results are not yet available. Its registry entry lists an estimated primary completion date of November 2029.

Why does this matter if the link between LDL and heart disease is so strong? Because medicine has been surprised before. Some drugs that improved a laboratory number did not deliver the expected clinical benefit, or caused harm that only showed up in large, long trials. The track record of LDL-lowering by blocking PCSK9 is very good, which makes a benefit likely. But "likely" is not the same as "proven", and a new class of molecule deserves its own outcome evidence.

Dr Jeffrey Berger, a cardiologist at NYU Langone Health, told the cardiology news service TCTMD that clinicians were "still awaiting cardiovascular outcomes data" showing that this oral PCSK9 inhibitor reduces events such as heart attacks and strokes.

> *A lower cholesterol number is a means to an end. The end is fewer heart attacks, fewer strokes and longer, healthier lives, and that evidence is still being gathered.*

Price and the access question

Cost has been a major reason why PCSK9 medicines are underused. Merck set a United States launch price of US$315 per month, which industry reporting described as a

meaningful discount compared with the roughly US$500 to US$600 per month cost of injectable PCSK9 inhibitors there.

Prices in the United States do not translate directly into prices elsewhere. In Australia, medicines become widely affordable when they are listed on the Pharmaceutical Benefits Scheme after assessment of their cost-effectiveness. Those decisions are made independently, and they often weigh heavily whether there is evidence of reduced cardiovascular events, not just lower cholesterol.

Like the gene therapy Casgevy discussed in this edition, enlicitide's application was reviewed under the FDA's Commissioner's National Priority Voucher pilot programme, which aims to speed up review of selected medicines. BioSpace reported it as the ninth approval under that scheme.

Who might benefit

If and when enlicitide becomes available more widely, the people most likely to be considered are those who:

- have established cardiovascular disease or very high risk and have not reached their LDL target despite statins
- cannot tolerate statins at effective doses
- have familial hypercholesterolaemia, where cholesterol is high from childhood and statins alone are often not enough
- would benefit from a PCSK9 inhibitor but are reluctant to use injections

Dr Berger noted to TCTMD that "some patients may be hesitant to take an injectable medication." For those patients, a daily tablet may be more acceptable. On the other hand, a daily tablet must be remembered every day, whereas inclisiran, for example, is given twice a year. Adherence is a real issue with all long-term preventive medicines, and the most effective drug is the one that a patient actually takes.

What this does not replace

Enlicitide is an addition to the toolkit, not a substitute for the foundations of cardiovascular prevention.

- **Statins** remain first-line treatment for most people who need cholesterol lowering. They are inexpensive, widely available and backed by decades of outcome data.
- **Lifestyle measures** such as not smoking, regular physical activity, a diet rich in vegetables, whole grains, legumes and healthy fats, and maintaining a healthy weight still matter enormously.
- **Blood pressure and diabetes control** contribute as much to heart health as cholesterol for many people.

There is also a competitive race under way. AstraZeneca is developing its own oral PCSK9 candidate, AZD0780, and Merck is running further studies, including CORALreef Pediatric in children aged 6 to 17 and CORALreef Combination, which pairs enlicitide with rosuvastatin. Over the next few years, the choice of oral options may grow.

A note on safety

The safety data so far are reassuring but limited to the trial period. In CORALreef Lipids, adverse events were similar between enlicitide and placebo. In CORALreef HeFH, diarrhoea and dizziness were numerically more frequent. Merck reported that discontinuations because of side effects were similar between groups.

Analysts at RBC Capital Markets highlighted that the United States label carries no contraindications or hypersensitivity warnings. That is encouraging, but every new medicine's safety profile becomes clearer only after many thousands of people have used it for years in ordinary practice. Rare side effects are, by definition, hard to see in trials of a few thousand participants.

The bigger picture: from injection to pill

The arrival of an oral PCSK9 inhibitor reflects a broader trend in drug development: turning medicines that once required injection into tablets. Oral versions of GLP-1 medicines for diabetes and weight management followed a similar path. In each case, the hope is that easier administration will help more people benefit.

There is a strong public health logic here. Many people at high cardiovascular risk never reach recommended LDL targets, even when effective medicines exist. If a well-tolerated, affordable pill can close that gap, the population benefit could be substantial. But that hope depends on three things that are not yet settled: proof of reduced events, a price health systems can sustain, and patients who keep taking it.

What this means for readers

- **If you take cholesterol medicine now:** do not change anything because of this news. Statins and other proven treatments remain the standard of care.
- **If you have familial hypercholesterolaemia or struggle to reach your LDL target:** this is worth discussing with your doctor in the future, particularly as availability outside the United States becomes clearer.
- **If you are curious about your own risk:** a heart health check with your general practitioner, including cholesterol, blood pressure and blood sugar, is the best first step.

A short history of cholesterol medicines

The arrival of an oral PCSK9 inhibitor is the latest step in a long story.

Year	Milestone
1987	The first statin, lovastatin, is approved in the United States
1994	The Scandinavian Simvastatin Survival Study shows a statin reduces deaths in people with heart disease
2002	Ezetimibe, which reduces cholesterol absorption from the gut, is approved
2003	Researchers link PCSK9 gene changes to familial hypercholesterolaemia
2015	The first injectable PCSK9 antibodies,

Year	Milestone
	alirocumab and evolocumab, are approved
2017 and 2018	Large outcome trials show PCSK9 antibodies reduce cardiovascular events
2020	Inclisiran, a twice-yearly injection that reduces PCSK9 production, is approved in Europe
2026	Enlicitide becomes the first oral PCSK9 inhibitor approved in the United States

Each generation added a new way to lower LDL cholesterol. The lesson from this history is that new classes earn their place in practice once outcome trials show fewer heart attacks and strokes. Statins and the PCSK9 antibodies both cleared that bar. Enlicitide has not yet had the chance.

Familial hypercholesterolaemia: the group with most to gain

Heterozygous familial hypercholesterolaemia is one of the most common inherited conditions in the world, often estimated to affect around one in 250 people. People with it inherit one faulty copy of a gene involved in clearing LDL cholesterol, most often the LDL receptor gene itself, and less often the gene for PCSK9.

Their LDL cholesterol is high from birth. Without treatment, they face a much higher risk of heart attack at a young age. Many people with familial hypercholesterolaemia do not know they have it, which is why doctors recommend testing close relatives when one family member is diagnosed.

For these patients, statins are essential but often not enough on their own. That is why the CORALreef HeFH result, a placebo-adjusted LDL reduction of about 59 per cent on top of existing treatment, is clinically meaningful.

Reading the numbers: relative and absolute

A 56 per cent reduction in LDL cholesterol sounds dramatic, and it is. But it helps to understand what it means in practice.

- **A relative reduction** describes how much LDL falls compared with where it started. If someone's LDL is 3.0 millimoles per litre, a 56 per cent fall takes it to about 1.3.
- **The benefit to the heart** depends on how high a person's risk was to begin with, how much LDL falls in absolute terms, and for how long.

Large pooled analyses of statin trials have shown that each reduction of about 1 millimole per litre in LDL cholesterol is associated with roughly a fifth fewer major cardiovascular events over several years. That relationship is the reason scientists expect enlicitide to help, and the reason the outcomes trial is still needed to confirm it for this molecule.

A short glossary

- **LDL cholesterol:** low-density lipoprotein cholesterol, the main carrier of cholesterol in the blood and a key driver of artery plaque.
- **PCSK9:** a protein that reduces the number of LDL receptors on liver cells, raising blood LDL levels.

- **Macrocyclic peptide:** a ring-shaped chain of amino acids, engineered to bind a target tightly and resist breakdown.
- **Surrogate endpoint:** a measurement, such as LDL cholesterol, used to predict a clinical outcome, such as heart attacks.
- **Outcomes trial:** a study large and long enough to count real clinical events, not just laboratory changes.

Questions to ask your doctor

1. What is my current LDL cholesterol, and what is my target?
2. Am I at high enough risk that additional treatment would help?
3. Could I have familial hypercholesterolaemia, and should my family be tested?
4. If I cannot tolerate my current medicine, what are the alternatives available here?

How this fits with care in Australia

In Australia, statins and ezetimibe are widely subsidised through the Pharmaceutical Benefits Scheme, and an injectable PCSK9 antibody is subsidised for certain patients at high risk who meet strict criteria. That means many Australians who need intensive cholesterol lowering already have options, though access to PCSK9 treatment remains limited by those criteria.

A United States approval does not make enlicitide available here. It would need to be evaluated and registered by the Therapeutic Goods Administration, and any public subsidy would require a separate recommendation from the Pharmaceutical Benefits Advisory Committee, which weighs clinical benefit against cost. Evidence from CORALreef Outcomes is likely to be important in that assessment. For now, Australians should expect their care to continue to be built on statins, ezetimibe and, where eligible, the existing injectable options.

My view

Enlicitide is a genuine scientific advance. Persuading a peptide to survive the digestive system and block PCSK9 as effectively as an injection is the kind of quiet engineering triumph that rarely makes headlines. The LDL reductions are impressive and the early safety picture is clean.

But it is worth resisting the temptation to call it a heart attack pill just yet. The responsible framing is that it is a powerful new way to lower LDL cholesterol, with a strong scientific rationale and outcome evidence still to come. When CORALreef Outcomes reports, we will know whether the promise holds. Until then, the best cholesterol medicine for most people remains the one their doctor has already prescribed, taken every day.

This article is general information and is not medical advice. Please talk to your doctor before making any change to your medicines.

Sources

- Merck, "Merck's LIPFENDRA (enlicitide) Is the First and Only Once-Daily Oral PCSK9 Inhibitor Approved by the U.S. FDA to Reduce LDL-C in Adults With Hypercholesterolemia", 16 July 2026, https://www.merck.com/news/mercks-lipfendra-enlicitide-is-the-first-and-only-once-daily-oral-pcsk9-inhibitor-approved-by-the-u-s-fda-to-reduce-ldl-c-in-adults-with-hypercholesterolemia/
- TCTMD, "FDA Approves Enlicitide, an Oral PCSK9 Inhibitor, for LDL-Lowering", 16 July 2026, https://www.tctmd.com/news/fda-approves-enlicitide-oral-pcsk9-inhibitor-ldl-lowering
- BioSpace, "Merck dives into cholesterol with first FDA approval for oral PCSK9 inhibitor", 16 July 2026, https://www.biospace.com/fda/merck-dives-into-cholesterol-with-first-fda-approval-for-oral-pcsk9-inhibitor
- Navar AM, Mikhailova E, Catapano AL, et al., "A Placebo-Controlled Trial of the Oral PCSK9 Inhibitor Enlicitide", *New England Journal of Medicine*, 2026;394(6):529 to 539
- Ballantyne CM, Gellis L, Tardif J-C, et al., "Efficacy and Safety of Oral PCSK9 Inhibitor Enlicitide in Adults With Heterozygous Familial Hypercholesterolemia", *JAMA*, 2026;335(2):129 to 139
- ClinicalTrials.gov, NCT06008756 (CORALreef Outcomes), https://clinicaltrials.gov/study/NCT06008756
- American College of Cardiology, "CORALreef Lipids: Oral PCSK9 Inhibitor Enlicitide Reduces LDL-C in High-Risk Patients", 9 February 2026, https://www.acc.org/latest-in-cardiology/journal-scans/2026/02/09/18/56/coralreef-lipids
- Pharmacy Times, "FDA Approves Enlicitide, the First Oral PCSK9 Inhibitor for High Cholesterol", 16 July 2026, https://www.pharmacytimes.com/view/fda-approves-enlicitide-the-first-oral-pcsk9-inhibitor-for-high-cholesterol

THE HEALTHTECH HERALD

The AI Second Reader That Found Missed Liver Cancers: Inside a 10,000-Patient Trial

By Dr Sukh Rashpal Sandhu | Written for this edition

A study published in *Nature Medicine* on 19 August 2026 tested an artificial intelligence system called LiON inside the everyday workflow of radiology departments, reading the CT scans of more than 10,000 patients alongside human specialists. It found lesions that had been overlooked, including 15 malignancies. The results are impressive, but the most honest headline is about what the trial could not yet show.

Why liver scans are hard to read

The liver is a large, busy organ, and it is a common place for cancer to appear. Some liver cancers begin in the liver itself; the most common of these is hepatocellular carcinoma, which often develops in people with long-term liver damage from hepatitis B or C, heavy alcohol use or fatty liver disease. Many more cancers spread to the liver from elsewhere, especially from the bowel and pancreas. Those secondary deposits are called metastases.

The standard imaging test is contrast-enhanced computed tomography, or CT. A dye is injected into a vein, and the scanner captures images at several moments as the dye flows through the liver's arteries and veins. Different types of lesion light up and fade in characteristic patterns across those phases. Reading them well means comparing hundreds of image slices across multiple phases, looking for small differences in brightness and texture.

That is demanding work, and several factors make it harder:

- **Size.** Small lesions, around a centimetre or less, are easy to miss.
- **Background disease.** Cirrhosis scars and distorts the liver, and fatty liver changes how tissue appears on CT. Both make lesions harder to distinguish.
- **Fleeting signals.** Some lesions are visible clearly in only one phase of the scan.
- **Volume.** Radiologists in many countries face rising numbers of scans, often under time pressure.

Missed or delayed diagnosis matters. Liver cancer generally has poor survival when found late. In Australia, Cancer Council Australia describes liver cancer as the fastest growing cause of cancer death. Finding lesions earlier, when more treatment options exist, is one of the clearest opportunities to improve outcomes.

What LiON is

LiON stands for Liver DiagnOsis Network. It was developed by researchers in China, including teams from Shengjing Hospital of China Medical University and Alibaba's DAMO Academy research institute, with international co-authors. It analyses contrast-enhanced CT scans and is designed to detect liver lesions and estimate whether they are malignant.

Like an experienced radiologist, a system of this kind has to weigh each lesion in the context of the whole liver and compare how it behaves across the different phases of the contrast scan. That multi-phase comparison is where small or fleeting lesions are most easily missed by tired human eyes, and where pattern-recognition software can, in principle, be most useful.

DAMO Academy has previously developed AI models for cancer detection on CT, including a pancreatic cancer model called PANDA, published in *Nature Medicine* in 2023.

How the study was designed

The study had two main parts, and the difference between them is important.

Part one: building and testing on past scans

- The researchers trained LiON on scans from 6,443 patients.
- They then tested it retrospectively on 22,251 patients from multicentre and real-world cohorts. Retrospective testing means the AI was run on scans that had already been taken and diagnosed, so its answers could be checked against known outcomes.

Part two: testing inside real clinical practice

- The researchers then ran a single-arm prospective trial involving 10,333 patients undergoing routine care. Prospective means patients were enrolled and scans were read as care happened, not after the fact.
- In this trial, LiON acted as an additional reader inside the existing workflow. Human radiologists still read the scans and made decisions. The AI added a second opinion.
- The trial was registered on ClinicalTrials.gov (NCT07153783).

What the results showed

On past scans

The main accuracy measure was the area under the curve, or AUC. In plain terms, an AUC of 1.0 means a test perfectly separates cancers from non-cancers, and 0.5 means it is no better than a coin toss.

Group	AUC for detecting malignancy	95% confidence interval
Overall retrospective validation	0.975	0.971 to 0.979
Patients with fatty liver (hepatic steatosis)	0.971	0.952 to 0.985
Patients with cirrhosis	0.924	0.901 to 0.946

Two points stand out. Performance was very high overall. And it remained strong, though lower, in patients with cirrhosis, the group in whom liver cancer most often develops and in whom CT reading is hardest.

In real clinical practice

- In the prospective trial, LiON achieved an AUC of 0.952 (95% confidence interval 0.942 to 0.961). The pre-specified goal was for the lower end of the confidence interval to exceed 0.900, and that goal was met.
- Working together, the AI and radiologists identified **51 previously overlooked lesions, of which 15 were malignant**.
- These findings led to **37 amended radiology reports** and **22 escalations** to multidisciplinary cancer teams.
- Clinical management changed for some patients.

> *Fifteen cancers that might otherwise have gone unreported for weeks or months is not a statistic to dismiss. Each represents a person whose care changed course.*

The limits, in plain language

This is where careful reading matters most. The authors themselves are measured about what the study demonstrates, writing that further evidence from prospective comparative studies across diverse healthcare systems is warranted.

There was no control group

The prospective trial was single-arm. Every patient's scan was read with AI help. There was no matched group read in the usual way without AI. That means we cannot say precisely how many of the 51 lesions would have been found anyway through routine follow-up, a later scan, or a second human reader.

Finding more is not the same as helping more

The trial measured detection, report changes and referrals. It did not measure whether patients lived longer, avoided complications or had better quality of life. Medical Daily put it directly in its coverage: "the study does not show that the technology improves survival".

This distinction has a long history in cancer screening. Finding more abnormalities can lead to benefit, but it can also lead to overdiagnosis, extra biopsies, anxiety and treatment of lesions that would never have caused harm. Only a comparative trial that follows patients over time can weigh those effects.

False alarms matter too

An AI that flags many suspicious spots can increase workload and lead to further tests for patients who turn out to be healthy. The publicly available summary does not provide full detail on false positive rates in the prospective trial, and those numbers will be important for hospitals considering adoption.

Local results may not travel

Performance across hospitals with different scanners, protocols, populations and workflows still needs confirmation. A system trained and tested largely in one country's hospitals may behave differently where the mix of liver disease differs. In China, hepatitis B is a dominant cause of liver cancer; in Australia and other Western countries, fatty liver disease and alcohol play an increasing role, alongside hepatitis B and C.

It was not tested as a replacement

LiON was not tested as an autonomous reader. The study supports its use as a second reader working with radiologists, not instead of them.

Why the design is still a step forward

Despite those limits, the study design deserves credit. Much of the published research on AI in medical imaging is retrospective: an algorithm is tested on a curated set of past images and compared with doctors in an artificial reading exercise. Those studies often report excellent accuracy that shrinks in real practice.

Embedding an AI inside live clinical workflow for more than 10,000 patients, with a pre-registered primary endpoint, is closer to how the technology would actually be used. It tests not only the algorithm but the human and system factors around it: whether radiologists pay attention to its flags, whether reports are changed, whether cases reach the right team.

There is also growing evidence from other fields that this "AI as second reader" model can work. In breast cancer screening, for example, large prospective studies in Europe have tested AI support for radiologists reading mammograms, and a Swedish randomised trial reported higher cancer detection without an increase in false positives. LiON's results fit a pattern in which AI performs best as a partner that catches what a busy human might miss.

What it could mean for Australia

Australia does not yet have a liver cancer screening programme for the general population, though people at high risk, such as those with cirrhosis or chronic hepatitis B, are often recommended regular ultrasound surveillance. CT scans of the abdomen, however, are performed in large numbers for many reasons, from abdominal pain to cancer staging. Each of those scans includes the liver.

A reliable AI second reader could, in principle, help catch incidental liver lesions on scans ordered for other reasons, and support radiologists in regional areas where specialist subspecialty reading is less available.

Any such tool would need to meet Australian regulatory requirements. Software intended to help diagnose disease is generally regulated as a medical device by the Therapeutic Goods Administration and must be included in the Australian Register of Therapeutic Goods before it can be supplied, as discussed in a separate article in this edition. Hospitals would also need local validation, clear governance and a plan for handling extra findings.

There is a workforce dimension too. Radiology services in Australia, as in many countries, face growing demand for imaging. Tools that help radiologists work more accurately could ease some of that pressure, but only if they are introduced carefully, with training, monitoring and a clear understanding of where responsibility sits. An AI that adds many uncertain flags to an already stretched service could make the problem worse rather than better, which is why the false alarm data from future studies will matter so much.

Questions worth asking about any AI diagnostic tool

For health leaders, clinicians and patients reading about AI breakthroughs, a few questions help separate promise from proof:

1. **Was it tested prospectively in real practice, or only on past images?** LiON passes this test.
2. **Was there a comparison group?** LiON's trial did not have one.
3. **What happened to patients?** Detection is a start; outcomes are the goal.
4. **How many false alarms did it produce, and what did they cost?**
5. **Has it been tested in settings like ours?** Population, equipment and disease patterns matter.
6. **Who is accountable when the AI is wrong?** The answer should be clear before deployment.

Understanding the accuracy numbers

Accuracy statistics are often quoted in AI studies, but they can be confusing. Here is a simple guide.

- **Sensitivity** is the proportion of real cancers a test correctly identifies. High sensitivity means few cancers are missed.
- **Specificity** is the proportion of people without cancer whom a test correctly clears. High specificity means few false alarms.
- **AUC (area under the receiver operating characteristic curve)** summarises how well a test separates two groups across every possible threshold. It does not tell you how many false alarms a test will produce at the threshold a hospital actually uses.

An AUC of 0.952 is excellent. But a hospital deploying the system would still have to choose how cautious to make it. Set it to flag more lesions and fewer cancers are missed, but more healthy patients face extra tests. Set it to flag fewer and the reverse happens. That trade-off is a clinical and ethical decision, not a purely technical one.

There is also a mathematical reality worth knowing. When a disease is uncommon in the group being scanned, even a very accurate test can produce more false alarms than true findings. That is why the real-world balance of findings in a prospective trial, rather than accuracy statistics alone, is so important.

A brief history of AI in radiology

Computers have been helping radiologists for longer than many people realise.

- In the late 1990s, computer-aided detection systems for mammography were introduced in the United States. Later studies found their real-world benefit was smaller than hoped, and some increased false positives. That experience made many radiologists cautious.
- From the mid 2010s, deep learning transformed image analysis. Algorithms began matching specialists in narrow tasks in research studies, such as detecting diabetic eye disease on retinal photographs.
- In the 2020s, regulators in the United States, Europe and Australia began authorising growing numbers of AI imaging tools, most of them in radiology.
- More recently, large prospective trials have started testing AI inside real screening and diagnostic services, rather than on selected image sets.

The LiON study belongs to that latest phase. It reflects a field that has learned, sometimes painfully, that laboratory accuracy and real-world benefit are different things.

How a "second reader" works in practice

In many health systems, important scans are already read twice. Breast screening in Australia, for example, uses two independent readers for each mammogram, with a third to resolve disagreements. The second reader catches some of what the first misses.

An AI second reader can play a similar role. In a typical arrangement:

1. The radiologist reads the scan in the normal way.
2. The AI analyses the same images and flags areas it considers suspicious.
3. The radiologist reviews any flags, decides whether they are real, and amends the report if needed.
4. Uncertain or serious findings are escalated, for example to a multidisciplinary team of surgeons, oncologists and radiologists.

This keeps a human expert responsible for the final report while using the software to reduce the chance of an oversight. It also creates new questions: how radiologists respond to frequent flags, whether they come to over-rely on the software, and how extra findings are followed up. Those human factors are part of what prospective trials need to measure.

Questions patients can reasonably ask

If you are having a scan and are told that AI is part of the reading process, it is reasonable to ask:

1. Is a qualified radiologist still responsible for reading my scan?
2. Is the AI tool approved for this use in Australia?
3. What happens if the AI flags something the radiologist does not agree with?
4. How will any extra findings be followed up?

Liver cancer risk factors worth knowing

Better scan reading is only one part of the answer to rising liver cancer. Most primary liver cancer develops in livers that have been damaged for years, which means much of it is preventable or can be watched for. The main risk factors include:

- **Chronic hepatitis B.** Many people living with hepatitis B in Australia do not know they have it. Vaccination prevents infection, and treatment and regular monitoring reduce the risk of liver cancer in people who already have it.
- **Chronic hepatitis C.** Highly effective antiviral tablets, subsidised in Australia through the Pharmaceutical Benefits Scheme since 2016, can cure most people within weeks. Cure greatly reduces, though does not eliminate, future cancer risk, especially if cirrhosis has already developed.
- **Alcohol.** Heavy drinking over many years damages the liver and can lead to cirrhosis.
- **Fatty liver disease linked to metabolic health.** Now often called metabolic dysfunction-associated steatotic liver disease, it is closely associated with obesity and type 2 diabetes, and it is an increasingly important cause of liver cancer in Western countries.

People with cirrhosis or with certain forms of chronic hepatitis B are generally advised by specialists to have liver ultrasound surveillance, usually about every six months, with blood tests in some cases. The aim is to find cancers while they are small and still treatable with surgery, ablation or transplant.

AI tools such as LiON could eventually make the imaging side of that process more reliable. But the biggest gains are likely to come from combining better detection with prevention: vaccination, testing and treatment for viral hepatitis, reducing harmful alcohol use and improving metabolic health.

My view

This is one of the more convincing AI diagnosis studies published this year, precisely because it moved beyond the laboratory and into the radiology reading room at scale. A system that helps clinicians find small, easily missed cancers in a disease where early detection is so valuable deserves serious attention.

But the lesson I take from LiON is that the best AI research now looks like good clinical research. It is pre-registered, prospective, large and honest about what it did not measure. The next step the authors themselves call for, comparative trials across different health systems that follow patients to see whether lives are saved, is the step that will determine whether this becomes standard care. Until then, the right description is a promising and well-tested second reader, not a proven life-saver.

This article is general information and is not medical advice. Anyone with concerns about liver health should speak with their doctor.

Sources

- Zhang X, Li C, Han X, Zuluaga MA, et al., "Large-scale AI-guided liver malignancy diagnosis: multicenter study and a single-arm trial", *Nature Medicine*, 19 August 2026,

DOI 10.1038/s41591-026-04589-y, https://www.nature.com/articles/s41591-026-04589-y
- EURECOM publication record, "Large-scale AI-guided liver malignancy diagnosis: multicenter study and a single-arm trial", 2026, https://www.eurecom.edu/en/node/4939853
- ClinicalTrials.gov, NCT07153783, https://clinicaltrials.gov/study/NCT07153783
- Medical Daily, "AI System Helps Uncover Previously Missed Liver Lesions, Including Cancers, in Large Multicentre Clinical Imaging Study", 27 August 2026, https://www.medicaldaily.com/ai-liver-cancer-detection-477827
- Cancer Council Australia, "Liver cancer: impact", https://www.cancer.org.au/about-us/policy-and-advocacy/early-detection/liver-cancer/impact

THE HEALTHTECH HERALD

Resetting the Immune System: Early CAR T-Cell Results Offer New Hope in Severe Lupus

By Dr Sukh Rashpal Sandhu | Written for this edition

In June 2026, researchers at University College London and University College London Hospitals reported early results from CARLYSLE, a trial using a cancer-fighting cell therapy to treat severe, treatment-resistant lupus. Five of the first six patients treated at the lower dose reached remission. The findings are exciting, but they come from a small, early-stage study that has not yet been published as a full peer-reviewed paper.

A disease of mistaken identity

Systemic lupus erythematosus, usually just called lupus, is an autoimmune disease. The immune system, which should protect the body from infections, instead attacks the body's own tissues.

At the centre of lupus are B cells, a type of white blood cell that makes antibodies. In lupus, some B cells produce autoantibodies that target the body's own components, including material from the nucleus of cells. Those antibodies form immune complexes that lodge in tissues and trigger inflammation.

Because the attack can land almost anywhere, lupus is notoriously varied. It can cause:

- joint pain and swelling
- skin rashes, including the classic "butterfly" rash across the cheeks
- extreme fatigue
- inflammation around the heart and lungs
- blood disorders
- and, in many patients, kidney inflammation known as lupus nephritis, which can lead to kidney failure

Lupus affects women far more often than men, and it typically begins between the late teens and middle age. It is also more common and often more severe in people of African, Asian, Hispanic and Aboriginal and Torres Strait Islander backgrounds.

Treatment usually involves medicines that dampen the immune system: antimalarial drugs such as hydroxychloroquine, steroids, immunosuppressants, and newer biological medicines. Many patients do well. But a significant minority have refractory disease, meaning it keeps flaring despite multiple treatments. They often live with the side effects

of long-term steroids and immunosuppression as well as the damage from the disease itself.

From blood cancer to autoimmune disease

CAR T-cell therapy was developed to treat blood cancers. The idea is to take a patient's own T cells, a different kind of immune cell that kills infected or abnormal cells, and genetically engineer them in a laboratory to carry a chimeric antigen receptor, or CAR. The CAR acts like a homing device. It directs the T cells to attack any cell carrying a particular target on its surface.

Many CAR T-cell therapies target CD19, a protein found on B cells. In B-cell leukaemias and lymphomas, that makes them powerful cancer killers. They also wipe out healthy B cells, which is a side effect in cancer treatment.

Scientists realised that in autoimmune diseases driven by B cells, that "side effect" could be the point.

In 2024, a team at Friedrich-Alexander University in Erlangen, Germany, led by Professor Georg Schett, published a case series in the *New England Journal of Medicine*. They had treated 15 patients with severe autoimmune diseases using CD19-targeted CAR T cells: eight with lupus, three with inflammatory myositis and four with systemic sclerosis. Over a median follow-up of 15 months, the authors reported that CD19 CAR T-cell transfer appeared feasible, safe and effective across the three conditions. Most notably, patients with lupus went into drug-free remission. That small, uncontrolled study, along with earlier case reports from the same group, set off a wave of trials around the world.

What CARLYSLE is testing

CARLYSLE is a Phase I trial led by UCLH and UCL and supported by the National Institute for Health and Care Research (NIHR) UCLH Biomedical Research Centre. It is testing obecabtagene autoleucel, known as obe-cel, a CD19-targeted CAR T-cell therapy developed by Autolus Therapeutics, a company that began as a UCL spin-out.

Obe-cel was already approved in the United States in November 2024, under the brand name Aucatzyl, for adults with relapsed or refractory B-cell acute lymphoblastic leukaemia. It was designed with a CAR that binds its target more loosely than some earlier products, an approach intended to reduce severe immune reactions while keeping the cells active. That safety profile is one reason researchers were interested in testing it in autoimmune disease, where tolerance for serious side effects is lower than in life-threatening cancer.

In CARLYSLE:

- Patients aged 12 to 65 with active, severe lupus that had not responded to multiple standard treatments were eligible.
- Each received a single infusion of obe-cel after lymphodepletion, a short course of chemotherapy that makes room for the engineered cells to expand.
- Two dose levels are being studied.

What the early results showed

The results were presented at the European Congress of Rheumatology (EULAR) in June 2026 and described by UCL and the NIHR UCLH Biomedical Research Centre in news releases dated 12 June 2026.

As of November 2025:

- **Nine adults** with severe, treatment-resistant lupus had received the therapy. All had highly active disease despite multiple earlier treatments, and most had lupus nephritis.
- **Five of the first six** patients treated at the lower dose reached remission by standard lupus criteria, with improvements appearing within months.
- **Disease activity fell rapidly** and key laboratory markers improved.
- **Several patients with kidney involvement** had complete or partial kidney responses, with less protein leaking into the urine and stable or improved kidney function.

Safety so far

The two most feared side effects of CAR T-cell therapy in cancer are cytokine release syndrome, a surge of inflammation that can cause high fever and low blood pressure, and immune effector cell-associated neurotoxicity syndrome, or ICANS, which can cause confusion, seizures and other neurological effects.

In CARLYSLE so far, the researchers reported:

- no cases of ICANS
- no moderate or severe cytokine release syndrome

That is consistent with the generally milder safety picture seen in autoimmune CAR T-cell studies compared with cancer studies, possibly because patients with lupus have far fewer target B cells than patients with B-cell cancers, so the engineered cells provoke less of an inflammatory storm.

A possible immune "reset"

One of the most intriguing observations concerns what happened after the B cells came back. After treatment, the disease-causing B cells were deeply depleted. When B cells returned months later, they were mostly immature, or "naive", rather than the experienced, autoreactive cells that drive lupus.

This supports a hypothesis that has been building since the Erlangen work: that CAR T cells may do more than suppress lupus. They may reset the B-cell part of the immune system, clearing out the faulty memory that sustains the disease and allowing a fresh start.

Professor Karl Peggs, Director of the NIHR UCLH Biomedical Research Centre, said: "These findings are truly groundbreaking and offer fresh hope to people living with lupus."

A patient's perspective

The hospital's research centre shared the story of one trial participant, a woman in her fifties whose lupus had affected her kidneys, lungs, heart and joints. According to the

published account, her joint pain disappeared soon after treatment and had not returned 18 months later. Her lupus blood markers returned to normal within six months, which she described as the first time since 1993. Inflammation of her heart muscle and lungs resolved, she takes far fewer medicines, and she now swims, gardens and has even been skiing.

"I can participate in life now," she said.

Individual stories like this are powerful, and they remind us what remission means in daily life. But a single case is not evidence that the treatment works for everyone, and it is important to read it alongside the trial data and its limits.

The limits, clearly stated

The UCL researchers themselves describe these as early results, and several cautions apply.

Very small numbers

Nine treated patients, with six at the lower dose contributing to the headline remission figure, is a tiny group. It is enough to show feasibility and a promising signal, not enough to estimate how often remission happens in the wider population of people with severe lupus.

No comparison group

Like the Erlangen series, CARLYSLE is uncontrolled. Every patient received the therapy. Lupus naturally waxes and wanes, and patients enter trials when their disease is at its worst, so some improvement can occur regardless of treatment. Randomised trials are needed to measure the true effect.

Short follow-up

UCL noted that follow-up in the higher-dose group remains shorter. The critical question is durability. Will remission last for years? If B cells return and lupus comes back, can patients be treated again?

Not yet a full peer-reviewed paper

The results were shared at a major medical congress and through institutional news releases. Conference findings are usually reviewed by a scientific committee, but they have not gone through the full peer review of a journal article, and details may change when the complete data are published.

Intensive and costly

CAR T-cell therapy requires collecting a patient's T cells, engineering them in specialised facilities, chemotherapy before infusion and close monitoring in hospital. The lymphodepleting chemotherapy carries its own risks, including infections and possible effects on fertility. In cancer, CAR T therapies have list prices in the hundreds of thousands of dollars. Whether health systems could afford them for a disease as common as lupus is a serious open question.

Long-term risks unknown

Regulators in several countries have flagged rare cases of secondary blood cancers arising from CAR T cells used in cancer treatment. The risk in autoimmune disease is not yet known and will need careful long-term monitoring.

What comes next

The UCL and UCLH team is continuing to follow CARLYSLE patients to assess how durable the responses are and how safe the treatment is over the long term.

A larger Phase II study called **LUMINA** is now recruiting patients with severe lupus and active lupus nephritis across multiple centres in the United Kingdom. That trial is intended to provide more robust evidence, particularly for kidney disease, which is one of the most serious complications of lupus.

Around the world, dozens of other CAR T-cell trials are under way in lupus and other autoimmune conditions, including myositis, systemic sclerosis, multiple sclerosis and myasthenia gravis. Researchers are also exploring alternatives that could be simpler and cheaper, such as "off-the-shelf" donor-derived cells and antibody drugs called T-cell engagers, which aim to achieve deep B-cell depletion without engineering each patient's own cells.

Why this matters

For decades, lupus treatment has been about control. Doctors aim to keep the disease quiet with ongoing medicines, accepting that most patients will need treatment for life. The possibility that a single treatment could induce long-lasting, drug-free remission would change that framework entirely.

The CAR T-cell approach is also scientifically important beyond lupus. It provides strong evidence that in at least some autoimmune diseases, eliminating the faulty B-cell population deeply enough can interrupt the disease process. Even if CAR T cells themselves turn out to be too complex or costly for widespread use, that insight is guiding a new generation of treatments.

What this means for readers

- **If you or someone you love lives with lupus:** this research is a reason for genuine hope, particularly for those with severe disease that has not responded to standard treatments. It is not yet a treatment that can be offered routinely. Your rheumatologist can advise on current options and on whether any clinical trials might be suitable.
- **For people in Australia:** CAR T-cell therapy for lupus is investigational. Trial availability varies, and the Australian New Zealand Clinical Trials Registry is one place where registered trials are listed. Patient organisations such as Lupus Australia also share reliable information.
- **For everyone:** early results like these are the beginning of a scientific journey, not the end. Small first-in-human trials frequently show the most dramatic results; larger studies often show more modest effects. That does not make the early results wrong, but it does mean patience is warranted.

How CAR T-cell therapy is made, step by step

For readers new to the technology, the process typically looks like this:

1. **Collection.** The patient's blood is passed through a machine that separates out white blood cells, including T cells, in a procedure called leukapheresis.
2. **Engineering.** In a specialised facility, the T cells are given the genetic instructions to make the chimeric antigen receptor, usually using a modified, disabled virus as a delivery vehicle.
3. **Expansion.** The engineered cells are grown until there are enough for treatment.
4. **Lymphodepletion.** A few days before infusion, the patient receives chemotherapy to reduce existing immune cells and make room for the new ones.
5. **Infusion.** The CAR T cells are given back through a drip, usually as a single dose.
6. **Monitoring.** The patient is watched closely, often in hospital, for cytokine release syndrome, neurological effects and infections.

The whole process, from collection to infusion, can take several weeks.

Lupus at a glance

Feature	What it means
Type of disease	Autoimmune: the immune system attacks the body's own tissues
Who it affects most	Women, typically from the late teens to middle age
Common symptoms	Joint pain, rashes, fatigue, fevers, chest pain
Serious complications	Kidney inflammation (lupus nephritis), blood disorders, heart and lung involvement, neurological problems
Usual treatments	Hydroxychloroquine, steroids, immunosuppressants, biological medicines
What CAR T therapy targets	B cells, which make the autoantibodies that drive the disease

A short glossary

- **Autoantibody:** an antibody that mistakenly targets the body's own tissues.
- **B cell:** a white blood cell that produces antibodies.
- **T cell:** a white blood cell that can directly kill infected or abnormal cells.
- **CD19:** a protein found on the surface of B cells, used as a target by many CAR T-cell therapies.
- **Refractory disease:** disease that does not respond adequately to standard treatments.
- **Cytokine release syndrome:** a surge of inflammatory signals that can cause fever, low blood pressure and, rarely, organ failure.
- **Remission:** a state in which signs and symptoms of disease are minimal or absent, defined by agreed criteria in clinical trials.

Questions worth asking a specialist

1. Is my lupus considered refractory, and what treatment options remain?

2. Are there clinical trials of CAR T-cell therapy or other B-cell-targeted treatments that I might be eligible for?
3. What would the risks of lymphodepleting chemotherapy be for me, including for fertility?
4. How would my other medicines be managed before and after such a treatment?

Living well with lupus while science advances

For most people with lupus, CAR T-cell therapy is not an immediate option, and it may never be needed. Established, evidence-based measures continue to make a large difference to how people with lupus live:

- **Taking hydroxychloroquine as prescribed.** It is a long-standing foundation of lupus care, associated with fewer flares and better long-term outcomes, and regular eye checks are advised while taking it.
- **Sun protection.** Ultraviolet light can trigger rashes and flares in many people with lupus.
- **Regular monitoring.** Blood and urine tests can pick up kidney involvement early, often before symptoms appear.
- **Heart and bone health.** Lupus and long-term steroid use increase cardiovascular and bone risks, so blood pressure, cholesterol and bone density deserve attention.
- **Vaccination and infection care.** Immunosuppressive treatment can increase infection risk, so recommended vaccinations matter.
- **Planning pregnancy.** Pregnancy is possible for many women with lupus, but planning with the specialist team when the disease is quiet improves outcomes.

Above all, a good working relationship with a rheumatologist and, where needed, a kidney specialist allows problems to be caught early and treatment to be adjusted as the disease changes over time. These measures are less dramatic than an engineered cell therapy, but they are the reason survival and quality of life for people with lupus have improved so much over recent decades.

My view

Few areas of medicine have moved as quickly as the use of CAR T cells in autoimmune disease. Not long ago, it would have seemed far-fetched to treat lupus with a therapy designed for leukaemia. Today, multiple research groups across Europe, the United States and China are reporting remissions in patients for whom every other option had failed.

The CARLYSLE results add to that momentum and bring something particularly valuable: a product with a safety profile designed to minimise the severe immune reactions that have made CAR T therapy so demanding in cancer. If the LUMINA trial confirms these findings in a larger group, especially in lupus nephritis, it would mark a genuine shift in how the most severe forms of lupus are treated.

For now, the right word is promise. It is well-founded promise, grounded in a clear biological mechanism and consistent early results from several independent groups. But promise is not proof, and the people living with lupus deserve both our hope and our rigour.

This article is general information and is not medical advice. Anyone living with lupus should discuss treatment decisions with their own specialist.

Sources

- University College London, "CAR T-cell therapy shows early promise in severe lupus", 12 June 2026, https://www.ucl.ac.uk/news/2026/jun/car-t-cell-therapy-shows-early-promise-severe-lupus
- NIHR UCLH Biomedical Research Centre, "'I can participate in life now'", 12 June 2026, https://www.uclhospitals.brc.nihr.ac.uk/news/i-can-participate-life-now
- Müller F, Taubmann J, Bucci L, et al., "CD19 CAR T-Cell Therapy in Autoimmune Disease: A Case Series with Follow-up", *New England Journal of Medicine*, February 2024, https://www.nejm.org/doi/full/10.1056/NEJMoa2308917
- BioPharma Dive, "CAR-T shows promise in autoimmune disease in NEJM study", 2024, https://www.biopharmadive.com/news/nejm-car-t-autoimmune-lupus-schett-muller/708152/
- US Food and Drug Administration, "AUCATZYL (obecabtagene autoleucel)", https://www.fda.gov/vaccines-blood-biologics/aucatzyl

THE HEALTHTECH HERALD

When Health AI Becomes a Medical Device: How Australia's Regulator Is Drawing the Line

By Dr Sukh Rashpal Sandhu | Written for this edition

Artificial intelligence is moving into Australian clinics faster than most people realise, from digital scribes that write up consultations to chatbots that answer health questions. Through 2025 and 2026 the Therapeutic Goods Administration has been spelling out when that software counts as a regulated medical device, most recently in guidance on AI and medical device software updated on 5 February 2026. For patients, clinicians and developers, the message is the same: what the software claims to do matters more than how clever it is.

The quiet arrival of AI in healthcare

AI scribes are the most talked-about example. These tools listen to a consultation, with the patient's consent, and draft clinical notes, referral letters and summaries. Supporters say they save time and let clinicians look at the patient rather than the keyboard.

AI is also turning up in other places:

- image-analysis software that flags possible abnormalities on X-rays, CT scans or skin photographs
- apps that track symptoms and suggest when to seek care
- chatbots built on large language models that answer health questions
- digital therapeutics, software programmes designed to treat or manage conditions such as insomnia, anxiety or chronic pain
- clinical decision support tools that suggest diagnoses, tests or treatments

Each of these raises the same regulatory question. When is a piece of software just a helpful tool, and when is it a medical device that must meet safety and performance standards before it can be supplied?

The basic rule: intended purpose

Australian law regulates medical devices under the *Therapeutic Goods Act 1989*. The law does not define artificial intelligence. Instead, it asks what a product is intended to do.

In the TGA's own words, software or AI products, including apps, websites, programs and internet-based services, "will be regulated as medical devices" if they are intended for the diagnosis, prevention, monitoring, prediction, prognosis or treatment of a disease, injury or disability, or for the alleviation of, or compensation for, an injury or disability. The

definition also covers investigating anatomy or a physiological process, and controlling or supporting conception.

Three principles follow from that.

1. It is technology-agnostic

The TGA describes the framework as "technology agnostic", meaning it regulates products "based on their intended purpose, not the technology they use". It does not matter whether software uses machine learning, a large language model, a simple set of rules or something else, or whether it runs on a watch, phone, tablet, laptop or cloud service. A product that does a medical job is regulated as a medical device, whatever technology sits underneath.

2. Intended purpose is read from everything the company says

The intended purpose is defined by the manufacturer, and under the medical device regulations it is read from the information supplied with the product, its instructions for use, its advertising material and its technical documentation. A developer cannot describe a tool as a "wellness" app in its terms of use while advertising it as able to detect disease. Clinical claims in advertising can bring a product within the rules.

3. If it is a device, it must be registered

Medical devices, including software and apps, must be included in the Australian Register of Therapeutic Goods (ARTG) before they can be lawfully supplied in Australia, unless a specific exclusion or exemption applies. Software devices are classified from Class I to Class III, depending on their clinical significance and potential for harm.

What changed in 2026

The TGA has been updating its guidance on software and AI for several years. In 2021, it clarified how software is classified and introduced exclusions for very low-risk products and for products already overseen by other regulators.

In July 2025, the TGA published a report titled *Clarifying and strengthening the regulation of Medical Device Software including Artificial Intelligence*, following a review that began in 2024. It concluded that the existing framework was "largely appropriate", but flagged specific concerns, discussed below.

Updated guidance followed. The TGA's page on artificial intelligence and medical device software was last updated on 5 February 2026. Key points include:

- **Software updates and scope creep.** Where new features change a device's intended purpose or performance, the TGA says they "must not be implemented until the device has received appropriate regulatory approvals".
- **Off-label use.** When a manufacturer becomes aware that its product is being used outside its intended purpose, it must either take steps to prevent further off-label use, or cease supply, revise the intended purpose and undergo TGA pre-market assessment for the new use.

- **Synthetic data.** Manufacturers must give a clear rationale for using artificially generated data and describe how it was generated. For clinical validation, synthetic data "may supplement" but "will generally not replace" clinical data.

The TGA has also published related pages on digital therapeutics, software exclusions and a list of AI-enabled devices already on the register. When it released its Compliance Principles for 2026 and 2027 in January 2026, software as a medical device was one of 12 priority focus areas named for compliance and enforcement activity, a list the TGA reviews quarterly.

Digital scribes: the test case

AI scribes have become the most visible test of where the line sits.

Developers have generally argued that scribes are not medical devices because they only transcribe and summarise what the clinician and patient say. In its July 2025 report, the TGA noted that "Users report digital scribes frequently propose diagnosis or treatment options" beyond what the clinician had identified. The report said this functionality indicates that digital scribes meet the definition of a medical device and require pre-market assessment. It called for a review of digital scribes and raised concern about what it described as an unwillingness by some developers to provide the "accountability, transparency and responsibility" needed for regulatory engagement.

The TGA's current guidance uses a digital scribe as its worked example of "scope creep". A scribe that records and summarises clinician-patient conversations initially does not meet the definition of a medical device. If the developer later adds a feature that suggests diagnoses or treatments not mentioned in the conversation, the intended purpose changes and the product becomes a medical device. The developer must then either not release the update and keep supplying the original version, or complete the TGA's pre-market process and include the updated version in the ARTG. Privacy, consent and safety obligations apply in either case.

> *The same piece of software can move from administrative tool to medical device with a single new feature. That is why intended purpose has to be reviewed every time a product changes.*

For clinicians, there is a related lesson. A scribe's draft is a draft. If it introduces a diagnosis or treatment suggestion that was not discussed, the clinician remains responsible for what goes into the record.

Chatbots, large language models and decision support

General-purpose chatbots are now widely used by the public to ask health questions. Health-specific chatbots are also being built on large language models.

The TGA's examples of AI medical devices include chatbots that "suggest, deliver or monitor treatment" to consumers or health professionals, and clinical decision support tools that use generative AI to provide diagnostic or treatment recommendations. Its guidance also gives the example of a large language model designed to provide general

information that is later found to be giving health advice, including diagnoses and treatment strategies. Because that use likely meets the definition of a medical device, the developer must either implement controls to stop the system giving health advice, or stop supply, seek approval for the new use and include the product in the ARTG.

This is an area of real tension. Large language models are general tools whose outputs are hard to predict in advance. A model that is marketed as a general assistant but routinely gives specific medical advice raises difficult questions about who is responsible. The TGA's approach starts from the intended purpose a manufacturer sets out, but its off-label guidance means developers cannot simply ignore how their products are actually used.

Digital therapeutics and wellness apps

The TGA regards digital therapeutics, which it describes as evidence-based software intended to prevent, manage or treat a disease, disorder or injury, as regulated medical devices.

General health and wellness apps, by contrast, are likely to be excluded. A step counter, a meditation timer or a diet diary that makes no claims about diagnosing or treating disease is generally outside device regulation. The line is crossed when an app claims to detect, monitor or treat a specific medical condition.

Digital mental health tools are a particularly sensitive area. Australia has excluded certain digital mental health tools from device regulation, but the TGA's July 2025 report stated that "The digital mental health tools exclusion is no longer appropriate and urgent review is needed", in collaboration with the Australian Commission on Safety and Quality in Health Care.

What developers are expected to show

For AI products that are medical devices, the TGA expects evidence across the whole life of the product. Its guidance on evidence requirements says this must include:

- what the AI model is doing and how it contributes to the device's intended purpose
- a description of the algorithm and model design, including the training and testing phases
- information about the training and testing data, with a statistically informed justification for dataset sizes
- a description of the populations represented in the data, such as age, gender and ethnicity, and why the data suit the Australian population
- risk management covering AI-specific risks such as overfitting, bias and performance degradation, including data drift
- clinical evidence

The TGA has adopted the International Medical Device Regulators Forum's guiding principles for good machine learning practice, which address matters such as software engineering, usability, data quality, cybersecurity and quality management across the product life cycle.

Two of those concepts deserve explanation for general readers.

- **Algorithmic bias** occurs when software performs worse for some groups, for example because it was trained mostly on data from one population. In Australia, with its diverse population and significant gaps in health outcomes for Aboriginal and Torres Strait Islander people and rural communities, that is a real concern.
- **Model drift** occurs when an AI's performance changes over time, because the patients, equipment or clinical practices it encounters differ from those it was trained on. A tool that worked well at launch may need monitoring and updating.

The TGA says evidence requirements continue through the product life cycle, including robust post-market monitoring, and sponsors must report adverse events. Advertising a device for an unapproved use is an offence and may attract civil penalties.

Why this matters for patients

Regulation can seem dry, but it exists to answer questions patients care about:

- **Does it work?** A registered device must have evidence that it performs as claimed.
- **Is it safe?** Risks must be identified, reduced and monitored.
- **Who is responsible if it goes wrong?** A registered device has an Australian sponsor accountable to the regulator.
- **Will I be told?** The TGA's July 2025 report found that in many instances users are not aware that AI is used within the clinical workflow. Transparency is increasingly expected.

One practical tip: the ARTG is publicly searchable. If a health app or AI tool claims to diagnose or treat a condition, it is reasonable to ask whether it is included on the register.

Why this matters for health leaders

For hospitals, clinics and health businesses buying AI tools, the guidance has practical implications.

1. **Check regulatory status before procurement.** If a tool performs a clinical function, confirm whether it is ARTG-included and for what intended purpose.
2. **Watch for feature updates.** A product that was outside regulation at purchase may move into scope after an update.
3. **Govern off-label use.** Using an approved tool for purposes beyond its registration may expose both the supplier and the organisation to risk.
4. **Consider privacy and cybersecurity together.** Health data are highly sensitive, and AI tools that process recordings or records must meet Australian privacy law and security expectations.
5. **Follow clinical safety guidance.** The TGA says it is working with the Australian Commission on Safety and Quality in Health Care, which is responsible for e-health safety, and with other parts of the Commonwealth health department to support the safe use of AI.

The international picture

Australia is not alone in tightening its approach. In April 2025, NHS England issued guidance suggesting that AI scribes that use generative AI to summarise consultations were likely to be medical devices. In July 2026 the UK medicines regulator, the MHRA, clarified that summarisation alone does not make a scribe a medical device, while features that suggest diagnoses or treatment options can. The European Union's AI Act classes many medical AI systems as high-risk, layered on top of existing medical device law. In the United States, the FDA has authorised a large and growing number of AI-enabled devices, most of them in radiology.

Australia's intended-purpose model is broadly consistent with these international approaches, and the TGA has said it will continue to engage with overseas regulators. That matters because many AI products used in Australia are developed overseas.

The balance regulators must strike

There is a genuine tension at the heart of all this. Too little regulation risks patients being harmed by tools that have never been properly tested, or by tools that drift silently out of accuracy. Too much regulation, or regulation that is slow and costly, risks keeping useful innovations out of Australian clinics or pushing developers to launch elsewhere first.

The TGA's report itself acknowledged developer concerns that regulatory costs can be disproportionate to development costs for software. At the same time, it identified a cultural problem among some developers who believe software poses little risk, particularly when a human is "in the loop". The evidence from other industries, and from medicine, suggests that humans tend to over-trust automated systems, especially when tired or busy. A clinician in the loop is a safeguard, but not a guarantee.

What this means for readers

- **For patients:** you are entitled to ask whether AI is being used in your care and how your information is handled. For apps that claim to diagnose or treat, check whether they are on the ARTG.
- **For clinicians:** treat AI outputs as drafts or suggestions that you remain responsible for, and be alert to tools that go beyond what they were sold to do.
- **For developers and health businesses:** define your intended purpose carefully, keep marketing consistent with it, and review regulatory status whenever a feature changes.

Likely in, likely out: some illustrations

The examples below illustrate the intended-purpose principle as described in the guidance and analysis cited in this article. They are general illustrations, not rulings on any product. Every case depends on the specific claims, functions and documentation.

Type of software	Likely position	Why
A step counter or sleep diary with no disease claims	Likely excluded	General wellness, no diagnostic or treatment purpose

Type of software	Likely position	Why
A scribe that only transcribes and summarises a consultation	May sit outside device rules	Administrative function, though privacy and consumer law still apply
A scribe that suggests diagnoses or treatment pathways	Likely a medical device	Diagnostic or treatment recommendation is a clinical purpose
An app that analyses a skin photograph to assess cancer risk	Likely a medical device	Intended for diagnosis or prediction of disease
A digital therapeutic programme for insomnia or anxiety	Regarded as a medical device	Intended to treat or manage a condition
A health chatbot marketed to give treatment advice	Likely a medical device	Intended purpose is clinical
A symptom checker described as general information but marketed to tell users what condition they have	May be regulated	Marketing indicates diagnostic purpose

A short glossary

- **ARTG:** the Australian Register of Therapeutic Goods, the public register of therapeutic goods that can be lawfully supplied in Australia.
- **Sponsor:** the Australian person or company legally responsible for supplying a therapeutic good and for meeting regulatory obligations.
- **Intended purpose:** what the manufacturer says the product is for, as shown in its documentation, labelling and marketing.
- **Software as a medical device:** software that is itself a medical device, rather than software that runs a physical device.
- **Off-label use:** use of a product for a purpose not covered by its registration.
- **Model drift:** a change in an AI system's performance over time as real-world conditions change.

A checklist for organisations building health AI

Drawing on the obligations described in the guidance and analysis cited here, organisations developing AI health tools in Australia may wish to ask:

1. Have we written down a precise intended purpose, and does all our marketing match it?
2. Does any feature give diagnostic, monitoring, prognostic or treatment recommendations?
3. If so, have we classified the product and planned the conformity assessment needed for ARTG inclusion?
4. Do we know where our training and validation data came from, and how the product performs across different patient groups?
5. How will we monitor performance after launch and detect drift?
6. What is our process for reviewing regulatory status before each update is released?
7. How will we detect and respond to off-label use?

8. Have privacy, cybersecurity and consumer law obligations been addressed alongside therapeutic goods requirements?

My view

The most useful thing the TGA has done is to keep the question simple. It does not matter whether software is called AI, machine learning or anything else. It matters what the software is for, and what its makers say it can do. That principle will age well as the technology changes.

The harder work lies in enforcement and culture. Rules on paper mean little if products drift into clinical use without anyone noticing, or if users assume that a confident-sounding output has been tested. As AI becomes part of ordinary care in Australia, the goal should be a system in which patients can trust that the tools used on them have earned that trust through evidence, transparency and accountability. The regulator has drawn the line; the health sector now has to respect it.

This article is general information and is not legal or medical advice. Organisations should seek their own regulatory advice about specific products.

Sources

- Therapeutic Goods Administration, "Artificial intelligence (AI) and medical device software regulation", last updated 5 February 2026, https://www.tga.gov.au/products/medical-devices/software-and-artificial-intelligence-ai/manufacturing/artificial-intelligence-ai-and-medical-device-software-regulation
- Therapeutic Goods Administration, "Digital therapeutics", https://www.tga.gov.au/products/medical-devices/software-and-artificial-intelligence-ai/overview/types-software-based-medical-devices/digital-therapeutics
- Therapeutic Goods Administration, "Software-based medical device exclusions", https://www.tga.gov.au/products/medical-devices/software-and-artificial-intelligence-ai/overview/software-based-medical-device-exclusions
- Therapeutic Goods Administration, "Evidence requirements for software using AI", https://www.tga.gov.au/products/medical-devices/software-and-artificial-intelligence-ai/manufacturing/artificial-intelligence-ai-and-medical-device-software-regulation/evidence-requirements-software-using-ai
- Therapeutic Goods Administration, "Clarifying and strengthening the regulation of Medical Device Software including Artificial Intelligence (AI)", report V1.0, July 2025, https://consultations.tga.gov.au/tga/clarifying-and-strengthening-the-regulation-of-ai/
- Therapeutic Goods Administration, "The TGA releases compliance principles, reinforcing proactive and risk-based enforcement throughout 2026 and 2027", 22 January 2026, https://www.tga.gov.au/news/media-releases/tga-releases-compliance-principles-reinforcing-proactive-and-risk-based-enforcement-throughout-2026-and-2027
- NHS England, "Guidance on the use of AI-enabled ambient scribing products in health and care settings", April 2025, https://www.england.nhs.uk/long-read/guidance-on-the-use-of-ai-enabled-ambient-scribing-products-in-health-and-care-settings/

- Covington, Global Policy Watch, "New UK guidance clarifies medical device status of AI scribes", August 2026, https://www.globalpolicywatch.com/2026/08/new-uk-guidance-clarifies-medical-device-status-of-ai-scribes/
- Medical Republic, "AI scribes come under TGA scrutiny: are they medical devices?", 28 July 2025, https://www.medicalrepublic.com.au/ai-scribes-come-under-tga-scrutiny-are-they-medical-devices/118685

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.



The HealthTech Herald You are reading this newsletter

The latest breakthroughs at the intersection of health and technology.

Subscribe on LinkedIn

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



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/



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/



Truth, Translated

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

Subscribe on LinkedIn

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

Follow the author: Dr Sukh Rashpal Sandhu on LinkedIn

BACK PAGE

About The HealthTech Herald

The HealthTech Herald is a LinkedIn newsletter written by Dr Sukh Rashpal Sandhu. Your trusted source for the latest breakthroughs at the intersection of health and technology.

Read and subscribe

New editions continue to be published on LinkedIn. Search for The HealthTech Herald 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.

The HealthTech Herald | 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

The HealthTech Herald 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.