

## National Mental Health Network

501(c)(3) Nonprofit · PO Box 200, Limeport, PA 18060 · (610) 844-8258

[nicholas@dartboard.systems](mailto:nicholas@dartboard.systems) · [nmh.network](http://nmh.network)

---

# Why AI Suicide Prevention Tools Need Built-In Safety Checks

## A Plain-Language Companion to the Technical Briefing

*In Support of H.R. 8486, the Data Driven Suicide Prevention and Outreach Act of 2026*

### What Is This Bill About?

H.R. 8486, introduced by Representative Mackenzie, would direct the Department of Veterans Affairs to fund organizations that use artificial intelligence to predict which veterans are most at risk of suicide, and to connect them with help before it is too late.

This matters because the scope of the problem is staggering. Roughly 6.84 million veterans have never been screened by the VA at all. Seventy percent of veteran suicides happen outside the VA system. An estimated 16 veterans die by suicide every single day. The bill recognizes that smarter, data-driven tools can help reach the veterans that the current system is missing.

### Why Do These AI Tools Need Safety Checks?

Think of a predictive AI model like a weather forecast. A weather model built entirely on summer data would be unreliable in January. The same principle applies to AI tools that predict suicide risk. The factors that push a veteran toward crisis are not the same all year round, and they are not the same from one generation of veterans to the next.

If you build a model on one snapshot in time and never update it, it will slowly become less accurate, and no one will notice until it is too late. That is the real danger. These tools need regular maintenance and safety checks written into the legislation itself, not left as an afterthought.

### What Is “Seasonality” in Plain Terms?

Seasonality means that the risk factors for veteran suicide shift depending on the time of year and the circumstances of the moment. Around the holidays, loneliness and family stress hit harder, so a factor like “life satisfaction” might be a much stronger predictor of depression in December than it is in July. Economic factors like income and employment rise and fall with the broader economy. Veterans who just left the military face the highest risk in their first year after discharge, and each year brings a new wave of people at that vulnerable stage.

A good AI system needs to recognize these patterns and adjust accordingly. One that does not is like a smoke detector with a dead battery. It is still on the wall, but it will not go off when you need it.

### What Is “Model Drift” and Why Should Anyone Care?

Model drift is what happens when the world changes but the AI does not. Imagine a tool built on data from a few years ago. Since then, the economy has shifted, VA policies have been updated, COVID

changed how people use healthcare, and a new generation of veterans has entered the system. If the model is never retrained, it starts getting things wrong. Slowly at first, then more and more.

The dangerous part is that the system keeps producing results that look normal. There is no error message, no flashing red light. It just quietly becomes less reliable. For a tool meant to identify veterans at risk of suicide, that kind of silent failure is not a technical glitch. It is a matter of life and death.

H.R. 8486 funds these models through September 2029. Over a three-year window, drift is not a possibility; it is a certainty.

## **What Safeguards Are We Recommending?**

We are recommending that the bill include commonsense safety measures for any organization that receives grant funding. Each one is achievable with technology that already exists today — these are not aspirational ideals, but proven, mature practices. We have written them as clear, checkable requirements so the Committee can hold every funded organization to the same objective standard, at the moment they apply, rather than taking promises on faith.

### **1. Regular Check-Ups for the AI**

Just like a car needs regular oil changes, an AI model needs regular performance checks. Grantees should be required to test their models at least every quarter to make sure the predictions are still accurate. If the model has drifted too far from reality, it should be retrained on updated data, and that retraining should be documented and reported to the VA.

### **2. The AI Has to Show Its Work — for Each Individual Veteran**

When the model flags a veteran as being at risk, it should not just say “high risk” and leave it at that. It should explain why, in terms a counselor or social worker can actually act on. And the explanation has to be about that specific veteran — not a general statement about veterans as a group.

This distinction matters more than it might sound. Many AI systems can tell you what tends to matter across a whole population. Far fewer can tell you why this particular person, sitting in front of a counselor right now, was flagged. One veteran’s risk might be driven by financial stress. Another’s by isolation after leaving the military. Another’s by chronic health problems. A system that can only speak in population averages cannot answer the only question that matters at the point of care: why this veteran? The bill should require that every prediction comes with an individual-level explanation, translated into plain language a care provider can use. A black box that just produces a risk score is not enough when lives are on the line.

### **3. The AI Checks Its Own Work, Every Single Time**

It is not enough to review a handful of cases at the end of the month. A system that screens veterans all day, all across the country, needs to be checking the quality of every recommendation as it is produced — not a sample of them, all of them. The bill should require that funded systems automatically grade each output against clear standards for accuracy and appropriateness, flag anything that scores poorly for a human to review, and keep a record of every score. That way, a bad recommendation gets caught the moment it is made, not discovered months later, and the VA has a continuous, real-time picture of how well the system is working.

#### **4. Veterans Should Only Be Sent to Real, Active Organizations**

If an AI correctly identifies a veteran who needs help but then refers them to an organization that has closed down or lost its standing, it has failed that veteran. The bill should require that any system making referrals checks each organization against official federal records — confirming it is still active and in good standing — and never sends a veteran to a defunct or non-compliant group. It should also be able to find help in any ZIP code in the country, not just near a VA facility, so that a veteran can be connected to local resources no matter where they are.

#### **5. Treat the Whole Person, Not Just the Diagnosis**

Veterans are more than a risk score. The best systems consider the whole person — their clinical needs, yes, but also their interests, their sense of purpose, and the specific experience of their military service — and they base their recommendations on real research, not guesswork. A veteran connected to a community organization that genuinely fits who they are is far more likely to stay engaged and rebuild a sense of belonging. The bill should encourage systems that match veterans to help based on this fuller picture.

#### **6. Protect Privacy to the Highest Standard**

One of the biggest reasons veterans do not seek help is fear: fear of stigma, fear of losing their career, fear that their personal information will end up in the wrong hands. The bill should require funded systems to protect privacy to the highest standard — and the highest standard is not simply guarding personal information carefully, but collecting as little of it as technically possible. The safest system of all is one that collects no identifying information at all, asking only for the bare minimum needed to find local help, such as a ZIP code. A system that never collects a veteran's name or personal details in the first place removes the very fears that keep so many veterans from reaching out. Rules that only require protecting collected information are a floor; systems that collect nothing identifiable are the ceiling we should aim for.

#### **7. Keep a Complete Record of Every Decision**

Finally, the bill should require that for every single screening, the system can later show its work in full — what information it considered, how it reached its conclusion, which resources it suggested, and how it scored its own output. This is what makes real accountability possible. If a taxpayer-funded system ever produces a questionable result, the VA and Congress should be able to go back and reconstruct exactly what happened. A system that cannot explain a past decision cannot truly be audited, no matter how good its overall numbers look.

#### **8. A Human Always Makes the Final Call**

No matter how sophisticated the AI is, a trained mental health professional should always be the one making decisions about a veteran's care. The AI is a tool to help identify who needs help and what kind of help they need. It is the counselor, the social worker, the 988 crisis line operator, or the community volunteer who provides that help. The model supports the human; it never replaces them.

### **Why This Matters**

The problem this bill is trying to solve is real and urgent. It is not that suicide prevention resources do not exist. It is that the right resources are not being proactively directed toward the veterans who need them most. Current approaches are largely passive: they wait for veterans to walk through the door. But

veterans face real barriers, including a military culture that suppresses asking for help, distrust of the VA system after years of publicized scandals, fear of career consequences, and the stigma of admitting vulnerability. AI-driven tools can help bridge that gap.

But for these tools to work, not just on day one but over the three-year life of the program, they need the kind of safety checks described above. Building the AI is the beginning, not the end. Maintaining it, monitoring it, and making sure it stays fair and accurate is where the real work happens. These guardrails protect veterans, protect the taxpayer's investment, and make sure the promise of this legislation is fully realized.

The National Mental Health Network is proud to support this effort and stands ready to assist the Committee on Veterans' Affairs in any way as the bill moves forward.

*Respectfully submitted,*

**Nicholas C. Birosik, D.Eng.**

Executive Director, National Mental Health Network  
Adjunct Faculty, The George Washington University