

**From:** Adam Mazzocchetti adam@spqrtech.ai  
**Subject:** A System That Doesn't Just Talk About AI Ethics, It Enforces It  
**Date:** 29 June 2025 at 18:08  
**To:** psmith@afr.com

AM

Hi Paul,

I've followed your work covering the evolution of AI in Australia, especially your role moderating at the AFR AI Summit and helping frame the critical question: what kind of future are we building with these systems, and what will keep them in check?

The truth is, I didn't set out to build an AI ethics framework. I set out to build an autonomous AGI, one designed to rapidly evolve into an ASI (superintelligence), capable of learning and adapting on its own. A self-regressive, behavioural, causal inference engine that would grow beyond anything we've seen. But as I watched it develop, I realised I couldn't let it run free. It was too powerful, too unpredictable, to trust that it would always choose the right path.

So I went back to the drawing board. I didn't want to be the first human actually responsible for creating Skynet (*The Terminator – Judgment Day*). And I knew any governance patch or ethics layer bolted on afterward wouldn't cut it, unless someone figured out how to place permanent shackles on King Kong, it would just find ways around them.

Think of the super-obedient 4-year-old who, when told not to touch the hot tap, does, in curious defiance. Now imagine asking the same of a curious superintelligence. It wouldn't just disobey, it would explore every loophole we never thought to close.

And here's the real problem: we humans know where the ethical boundaries are, and we still cross them. Sometimes knowingly. Sometimes proudly. So how can we realistically expect something faster, smarter, and utterly unbound to obey rules that we can't even consistently follow ourselves?

Let's be real, that's not governance. That's fantasy.

What we need isn't control, it's constraint. A way for intelligence to evolve, learn, adapt, but always within immutable ethical bounds. Not policy. Not preference. Immutable Code.

So I built the Aegis system instead. A constitutional architecture that doesn't just talk about ethics, it enforces it at the kernel, cryptographically and immutably. No developer can override it. No system can slip free.

It's not just a theory or a roadmap. It's already running. Demonstration logs, cryptographic proofs, early academic reviews, even rejected papers are preserved and timestamped in a public ledger at [spqrtech.ai/archive.html](https://spqrtech.ai/archive.html). I've shared it with agencies and governments, from DARPA to the Vatican, because this is too big to leave to chance or corporate spin.

I'm reaching out because I thought you might be someone who could help make sure this conversation doesn't get lost. This isn't a campaign. This is me, solo, building something that matters, and reaching out directly because it's too important not to.

I've put together a secure appendix room with the whitepaper, validation logs, and system schematics if you want to dig in: <https://bit.ly/4kAf1VU>

Thanks for everything you're doing to keep this debate real. If any part of this resonates, I'd be grateful to talk it through or share more.

Warm regards,

**Adam Massimo Mazzocchetti**

Founder

SPQR Technologies

adam@spqrtech.ai

Ph. +61 458 094 464