Cyber-Security for Sovereign AI

Protecting National Infrastructure

Frank Meehan

Chief Strategy Officer and Co-Founder FrontierOne

FrontierOne builds highly **cyber-secure Sovereign AI factories** for Governments, Corporates and Financial Services.

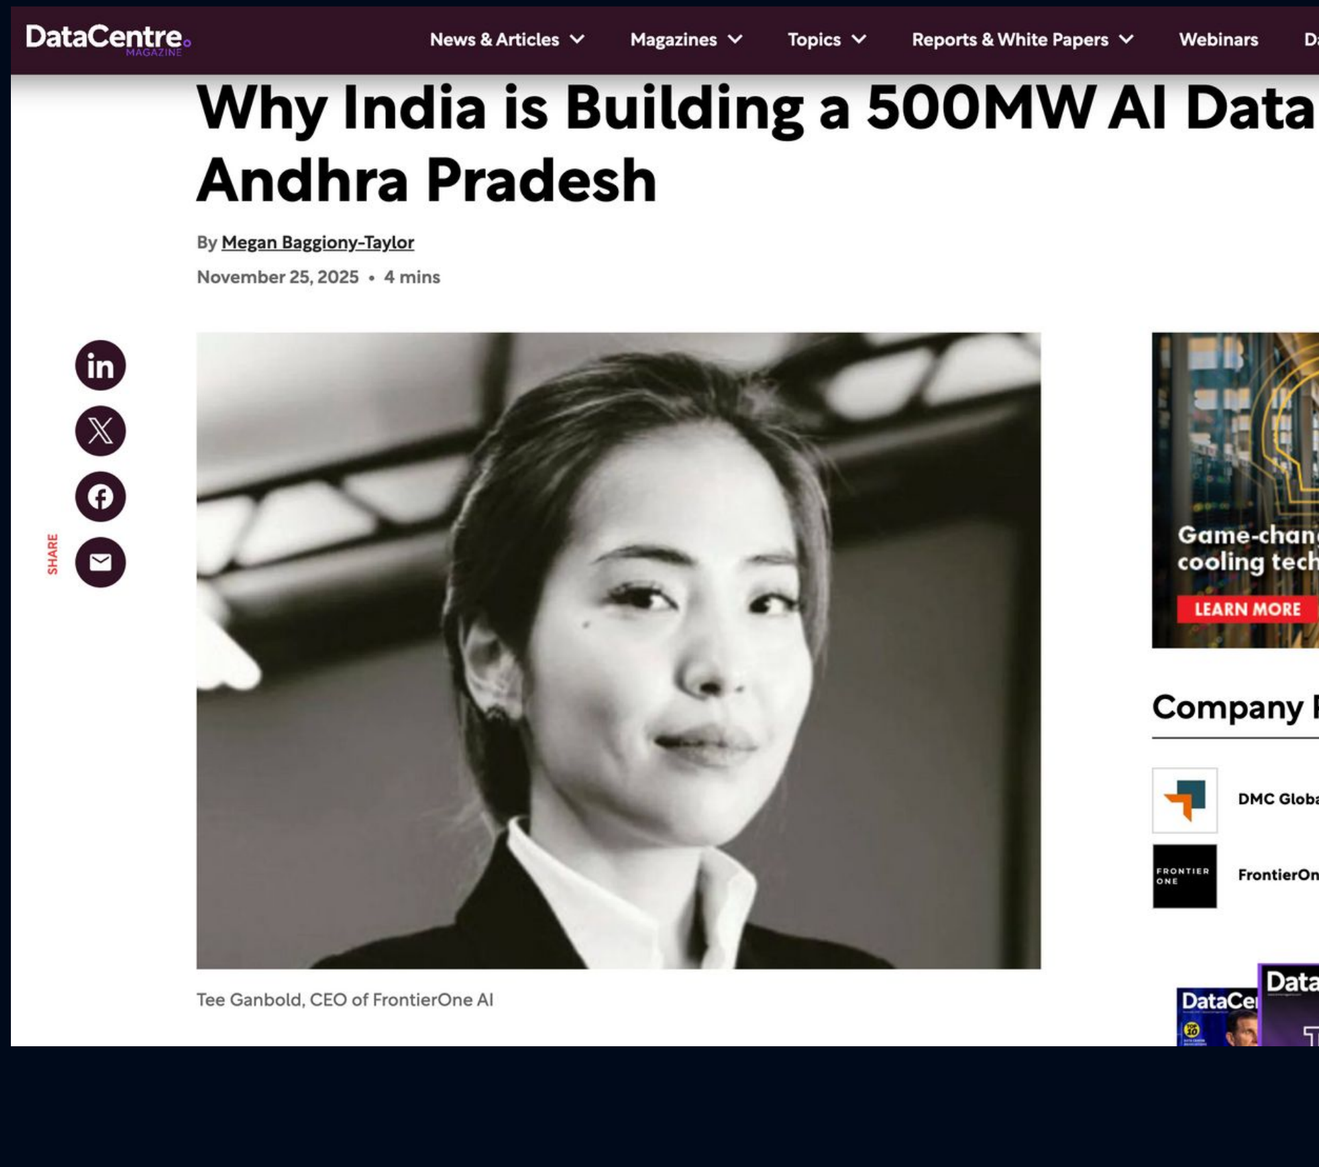

We secure our AI Factories from

“Concrete to Cloud” with our

ASGARD Cyber Operating

System .

Because AI is changing everything and

Cyber-Security is at the heart of it.

Like electricity in the 20th century, AI is now a utility embedded into core national functions.

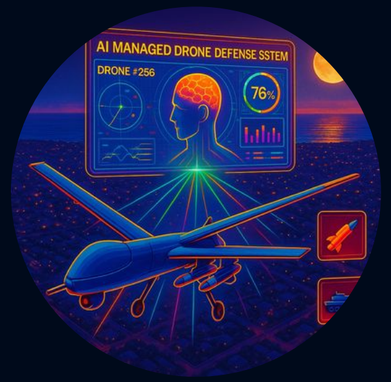

Defense

Energy

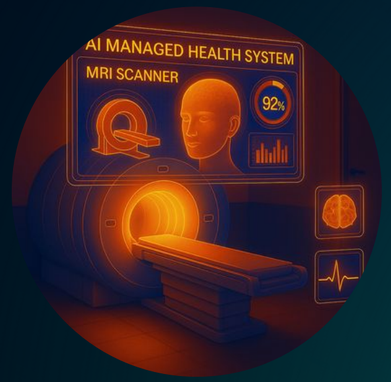

Health

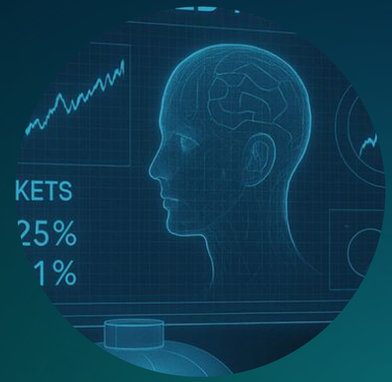

Finance

Everyone is talking about Sovereign AI

But what exactly is it, and who is it for?

Sovereign AI is primarily about **Trust and Security.**

Hyper-Secure

Hyper-Local.

Today, most national AI workloads look like this

- Models are trained in one country,
- Data is stored in another,
- And inference is served from a third, often on a public hyperscale cloud.

That means three big things are outside your control

- Data — where it lives, who can legally access it.
- Models — who can see, copy, or tamper with the weights.
- Alignment & governance — whose rules and values the AI is actually following.

How are AI systems attacked?

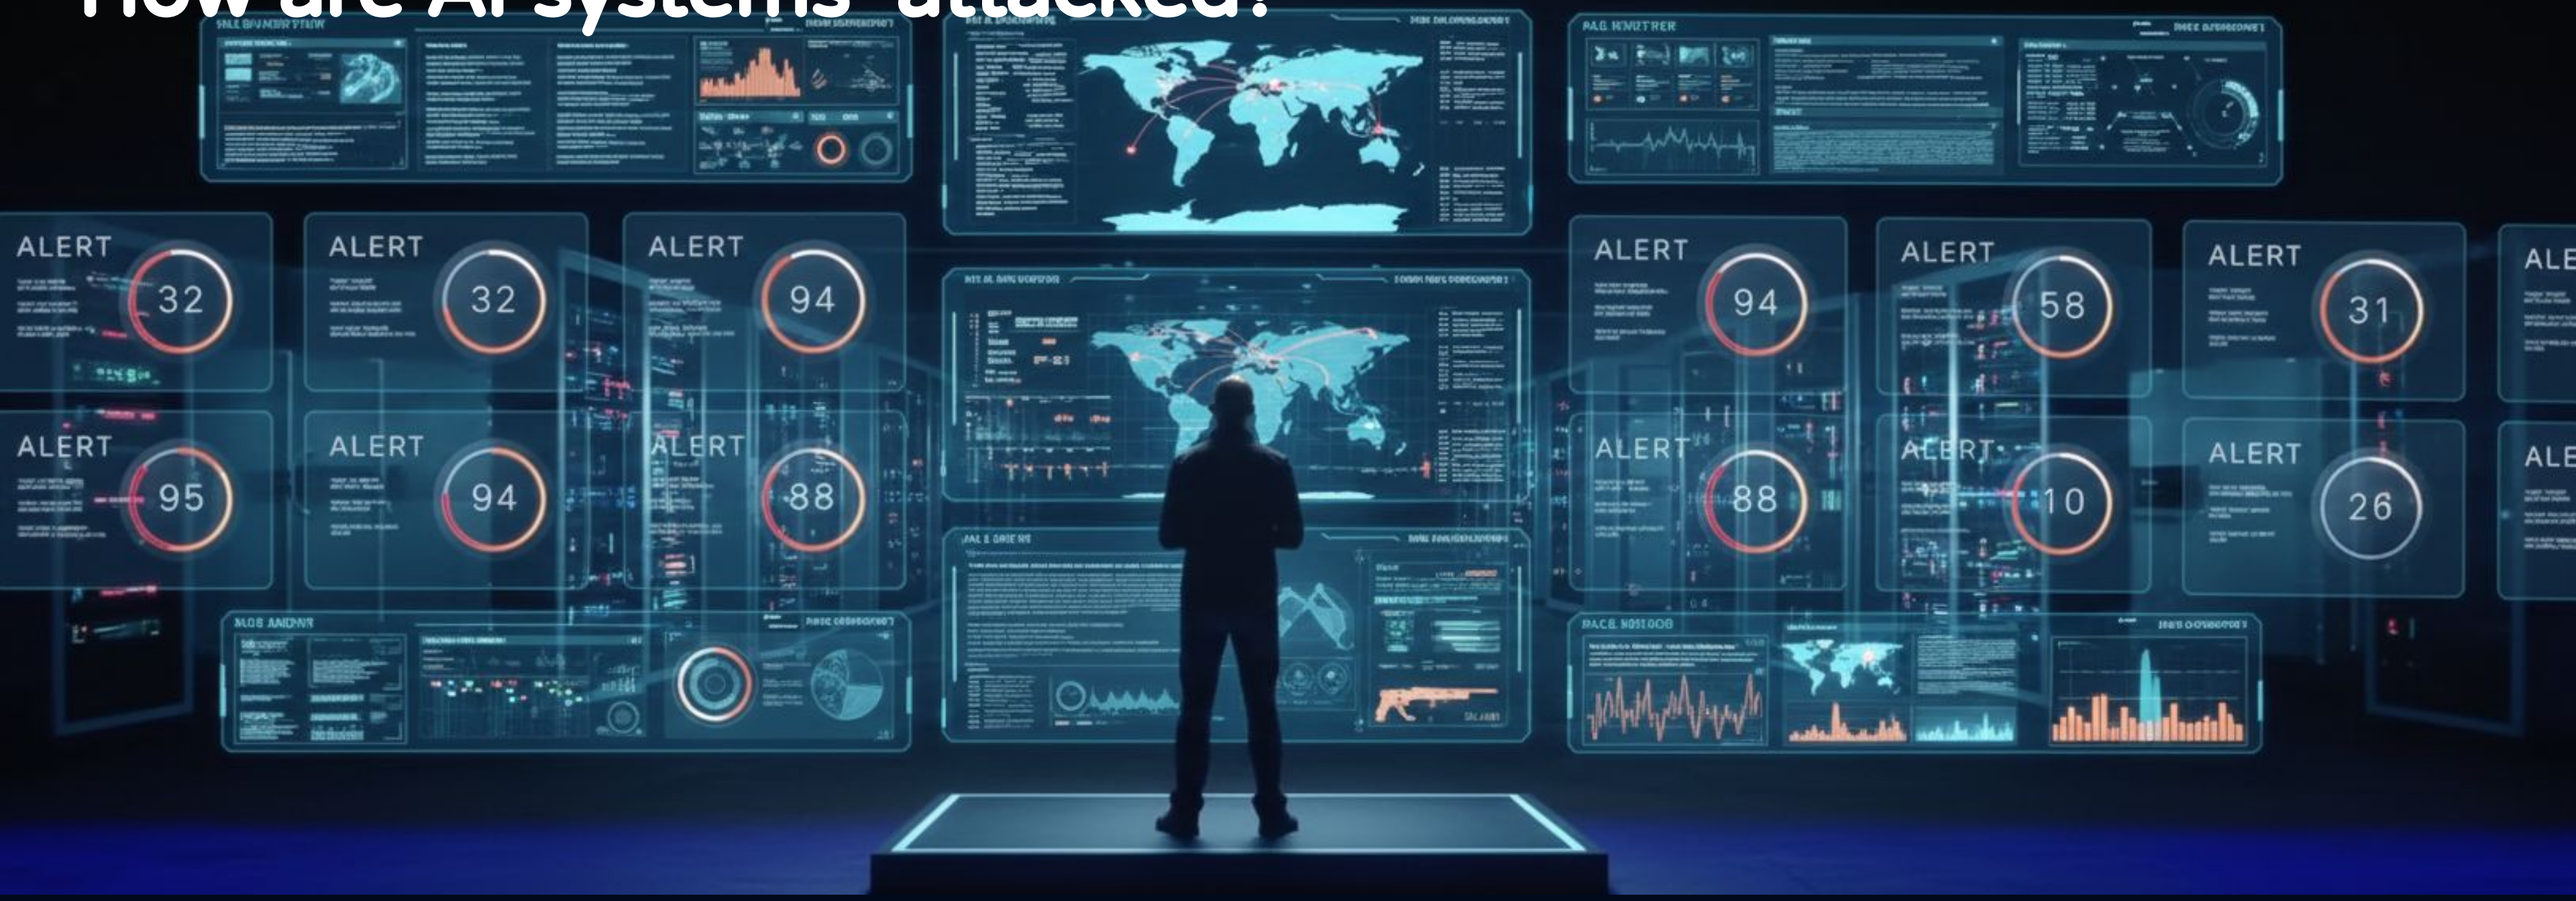

Data leakage and model exfiltration. — stealing the “brain”

Modern AI models and the data behind them are worth billions. And attackers that get into an AI data center don’t need to destroy anything — they just silently copy.

- Model weight exfiltration: moving huge neural network files out through main networks, management networks, or covert channels. Once stolen, your frontier model can be run by an adversary at a fraction of your cost.
- Data leakage: sensitive citizen data, strategic corporate IP, or intelligence reports leaking via logs, misconfigured storage buckets, or even the model’s own responses.

If you’re running national security, grid operations, or healthcare workloads, this isn’t just a compliance incident.

It’s like handing an adversary your playbook and your population registry.

Model manipulation — turning your AI against you

Model manipulation is about changing what the AI does and believes.

- Prompt injection & jailbreaking: attackers hide malicious instructions in documents, web pages, or internal tools. When your AI agent reads them, it overrides its safety rules and starts doing what the attacker wants — exfiltrating files, changing configs, or revealing secrets.

- Data poisoning: attackers subtly pollute training or fine-tuning data so that, under certain conditions, the model lies, fails to detect threats, or makes decisions in the attacker's favor.

These attacks are hard to detect because nothing "crashes."

Your AI still works — it just quietly works for someone else.

AI controlling operational technology.

As we connect AI to power grids, traffic systems, ports, and pipelines, a compromised model can directly change setpoints, thresholds, and schedules.

This turns an AI data center into a single point of failure for:

- Power stability
- Transportation
- Financial markets
- National defense

The more efficient your AI-driven grid becomes, the bigger the systemic shock if it’s maliciously mis-optimized.

Physical & supply-chain threats — from “concrete to cable”

AI data centers are attractive targets: they concentrate power, cooling, connectivity, and compute.

Threats include:

- Backdoors in firmware or BMCs,
- Compromised networking gear,
- Malicious insiders during build-out,
- Physical sabotage of power and cooling infrastructure.

If you don’t secure the stack from concrete to cloud, you’re relying on hope as a security control.

Building a truly sovereign and secure AI facility is hard.

It’s not enough to put GPUs in a room and call it a “national AI cloud.”

You need an integrated approach.

At FrontierOne, we design, finance, build, and operate what we call Cyber-Secure Sovereign AI Factories.

Think of them as turnkey AI infrastructure that’s:

- Hyper-local — physically in your jurisdiction, on your power grid, connected to your fiber;
- Hyper-secure — zero-trust networks, sovereign AI operating system, continuous monitoring;
- Hyper-aligned — models trained and governed under your laws, your ethics, your mission.

We work with partners like DMC Global Partners to build sovereign-grade data centers, and then we secure the entire stack using our ASGARD Operating System.

Security from “Concrete to Cloud”

Physical infrastructure layer

- Secure campus design, controlled access, resilient power and cooling, and a trusted hardware supply chain.

Compute & network fabric

- High-bandwidth AI fabric for training and inference, micro-segmented networks, and cross-domain guards between classified and unclassified environments.

Sovereign AI OS & MLOps — ASGARD

- Secure data lakes for structured and unstructured data.

Applied AI layer

- A platform where your ministries, regulators, banks, hospitals, and startups can safely build AI applications — from fraud detection and public-service chatbots to defense-grade analytics

In other words: we don’t just host your AI. We harden it as national infrastructure.

One example of this vision is Project FUSION in the State of Andhra Pradesh, India.

Together with DMC Global Partners, we're working on a 500MW Sovereign AI Factory, designed to:

- Deliver secure domestic AI infrastructure and data residency,
- Provide up to military-grade cybersecurity with real-time monitoring and threat intelligence,
- Support model training, fine-tuning, and low-cost, high-speed inference,
- And anchor a broader AI innovation ecosystem with research parks and workforce training.

It's exactly the kind of project that moves a region from AI consumer to AI producer—while keeping citizens' data and national capabilities under domestic control.

At FrontierOne, we believe every country and every critical institution deserves its own AI — on its own terms.

Cyber-Secure Sovereign AI isn’t a buzzword.

It’s the difference between:

- An AI that can be turned off or turned against you,
- And an AI that operates as a trusted extension of your state and your enterprise.

If you’re ready to treat AI as the strategic infrastructure it has become —

to secure it from concrete to cloud, and to bring it home under your own laws, your own governance, and your own values

— we’d love to talk.