You Can't Secure What You Can't Control: A Hierarchy of AI Controls

Borrow Maslow's hierarchy of needs and apply it to technical controls. There's an order to it: until you can actually get your hands on the programs running in your environment — control what executes and what it can reach — you can't meaningfully do security, compliance, or threat defense on top. Control is the base of the pyramid; everything else depends on it.

Control comes before compliance

Compliance is detail-oriented by nature: lists of controls, each one accounted for. Think of OSHA accounting for every safety measure in a factory, or a pilot working through pre-flight checklists — nothing is left to memory. But you can only build and enforce a checklist for things you can actually see and manage in real time. If you can't assess what AI tools are doing *as it happens*, and can't control them, you can't produce a checklist, can't enforce it, and therefore can't credibly claim compliance.

Control comes before stopping attackers

Attackers live in the gaps of the checklist — they make programs behave outside the expected rules. If there's no checklist at all because you have no visibility or control, there's nothing for them to work around; the whole environment is open field. Control is what turns "anything goes" into a defined, defensible surface.

What "control" means for AI specifically

For generative AI, control is the ability to:

- **See** which AI programs are executing across your endpoints, in real time.
- **Stop** the ones you haven't sanctioned — at the kernel layer, by policy.
- **Scope** the ones you allow to specific systems, data, and groups.

Only once that base is in place can the higher layers stand up: real compliance checklists you can enforce, and defense against attackers (and data exposure) that assumes a known, governed environment.

The sequence is the point. Visibility and control first; compliance and threat defense follow. FortressAI is built to establish that base layer for AI in your environment.

Adapted from a conversation with Cyber Crucible founder Dennis Underwood.

Revision #1
Created 2026-11-24 04:49:01 UTC by Dennis Underwood
Updated 2026-11-24 04:49:01 UTC by Dennis Underwood