

# Memo: Flexible Hardware-Enabled Governors (flex-HEGs)

**Successful AI governance requires compliance verification.** As AI advances, the potential for catastrophic misuse or accidental deployment of dangerous capabilities increases. This poses novel and as-of-yet unsolved governance challenges. Whether governments (or coalitions of governments) can effectively impose rules requiring safety mechanisms (of whatever kind) depends on their ability to assure that no credible efforts would be able to flout such rules. Without such an assurance, the dominant considerations would remain shaped by concerns about who can release more powerful capabilities sooner. After-the-fact penalties for rule violations, as is commonly done in law, may not on its own be a viable approach if the far-reaching consequences of an AI catastrophe could render any such enforcement action moot. If these assumptions are true, it is crucial to develop technological mechanisms which undergird the deployment and monitoring of agreed-upon rules and safety mechanisms in a way that affords high confidence among all strategic players that *no* player has the technical option of unilaterally violating any such rules.

**We believe that such technological mechanisms are possible**, and that, with a consolidated R&D effort, we would be able to demonstrate and document **compelling prototypes by the end of calendar year 2025**. The goal is to lower the barriers to adoption of such a technology—both for the hardware firms involved in producing it, and for the governmental or intergovernmental processes involved in requiring its use in some contexts.

**Regulatory capabilities this may enable**, if implemented in all high-performance AI accelerators, include:

- Limiting the size of training runs in terms of total FLOP (and/or total FLOP/s)
- Limiting the size of datasets that can be used in a training run
- Adjusting these limits over time
- Requiring that code be reviewed by a regulator before running (including only if above a certain size of compute)
- Requiring the possession of a non-expired licence to run computations of a certain size range (without requiring the computation itself to be reviewed by a regulator before running)
- Requiring a standardized evals protocol to be incorporated into the computation graph of training
- Requiring weights to be encrypted in such a way that inference with them is only possible on other flex-HEGs-compliant hardware, subject to another flexible policy for inference, including potential “speed limits” and “kill switches”
- If geolocation implemented: agreeing on differing limits or policies depending on jurisdiction

A further benefit of a technological solution as envisioned here is that it would enable significantly more targeted and reliably governance mechanisms than, for example, export controls alone. But this solution is synergistic: export controls can be shaped to provide an early incentive for hardware companies to actually use flex-HEGs. Also, if an agreement can be enforced by flex-HEGs, it would **not** be necessary to require a centralised “chip registry”, unilateral “kill switch”, nor even geolocation or human inspections, which could reduce strategic concerns with AI governance proposals. Demonstrating flex-HEGs’ technological viability would be of substantial

strategic and economic interest, bolstering the belief that the implementation of AI governance is feasible.

**We envision a technology stack enabling flexible on-chip verification of governance compliance.** The central capability we wish to deliver is for a cryptographically certified firmware layer to be loaded onto a high-performance computing device, which can enable on-chip checking for assured compliance with a flexible set of rules resulting from a multilateral decision-making process. Such rules might impose, for example, limitations on the type and size of computations a chip is permitted to run, compliance with a regulatory licensing scheme, etc. There are two critical desiderata for such a set-up, which together set it apart from other compute governance [proposals](#):

- (1) First, it is essential that there be **no technically viable means to bypass the check**, including by highly sophisticated physical tampering, without rendering the computing device permanently inoperable.
- (2) Second, the mechanism must **enable maximum flexibility for the decision-makers without retooling the hardware**. To achieve this, the on-device compliance processor should be capable of making its compliance determinations using general-purpose computation with inputs from various kinds of information about the code and its context.

**A potential approach** for designing a technology stack with the stated desiderata consists of four functional layers to be implemented through several technological mechanisms.<sup>12</sup>

- **Compliance processor:** an auxiliary computing device for determining compliance, integrated into a securely assembled chassis, with sufficient hardware capabilities and software libraries to securely and flexibly implement and update the desired safety mechanisms, e.g. through:
  - [Static analysis of compute graphs](#) (e.g. maximum FLOP count) and [inputs](#) (e.g. cryptographically verified timestamp)
  - [Multilateral cryptographic compliance-firmware](#) renewed with periodic [broadcast updates](#), e.g. using multisignature contracts
- **Compliance inputs:** peripheral connections to the compliance processor which supply it with the context which might be needed to make compliance determinations under a broad variety of potential regulatory scenarios, e.g. through:
  - [Compliance-compute interlock](#), e.g. executable instructions must pass through the compliance processor in order to be loaded into the primary compute processor
  - Dynamic [performance counters](#) and/or [runtime trace analysis](#)
  - [Incorruptible inputs from other compliance processors](#), e.g. authenticated with [continuous group key agreement \(CGKA\)](#)
  - [Real-time clock](#) which is covered by tamper-detection mechanisms, and [resistance to any clock-specific tampering attacks](#)
  - Optionally: [provable geolocation](#), e.g. with time-of-flight measurements from new space-based assets, or maybe [Iridium STL](#)
- **Tamper-detection and on-device enforcement:** to detect any attempts of physical tempering and disable the chip if any such attempts were detected, e.g. via:
  - Hard-wired chassis-intrusion detection mechanism, e.g. [FIPS 140-style active mesh](#)

---

<sup>1</sup> We are indicating the [technological readiness level](#) of these different mechanisms with the following colour code: [red](#) for TRL 2, [orange](#) for TRL 3, [yellow](#) for TRL 4, [purple](#) is TRL 5, and [blue](#) is TRL 6. The desired readiness level for all of these by the end of such an effort is TLR 7, and above.

<sup>2</sup> What remains out of scope for this specific proposal is solutions to secure the assembly chain, e.g. through fab detection and inspection, and tracing components while in transit between manufacturing facilities.

- RF-based chassis-intrusion detection mechanism, e.g. [anti-tamper radio](#)
- Permanent self-disabling mechanism, e.g. [nanothermite](#)
- **Offline power source:** to ensure the functional integrity of the entire stack, it is essential that the above mechanisms do not rely on external power sources, to be achieved e.g. through:
  - Power supplied during shipment (analogous to refrigerated shipping)
  - Batteries
  - Fail-safes in the event of power loss, e.g. a circuit inhibiting the ignition of the nanothermite itself relying on energy, such that the absence of power automatically activates the self-destructive mechanism

**Call to action:**

- If you're a funder and interested in accelerating work in this area, feel free to reuse ideas and language used in this memo e.g. in your own funding calls.
- Teams interested in directly pursuing work in this area may work on one or several of the following types of projects:
  - (i) 'on-paper' system design and analysis,
  - (ii) design of one or more of the components,
  - (iii) prototyping and testing implementations of one or more of the components,
  - (iv) system integration and demonstration of components,
  - (v) complete design, implementation and demonstration of the full stack.
- You can send us a 1-2 page expression of interest outlining the team composition, anticipated timelines and a rough budget.
- Note that this memo is **not** a funding call. However, we anticipate one or several funders to open such calls in the near future, and would be happy to keep you informed about these.