The Immune System's War

How Operation Epic Fury Activated the Architecture of Impunity

S. E. G. Nyberg
Strange Loops LLC
March 2026

A system that cannot selectively forget cannot adapt after deployment.
— The Discipline of Forgetting

A sufficiently organized coalition of senescent institutional components can—and historically does—rewrite the immune system's own code.
— The Mortal Architecture

* * *

On February 28, 2026, the United States launched Operation Epic Fury—a sustained air campaign against Iran that killed Ayatollah Khamenei and approximately forty senior Iranian officials in its opening hours. Iran retaliated with ballistic missiles against Dubai, Abu Dhabi, and American installations across the Persian Gulf. As of this writing, six American service members are dead, the Strait of Hormuz is closed, and the President has estimated a four-to-five-week campaign.

Four days before the first strikes, Bloomberg reported that SpaceX was weighing a confidential IPO filing as soon as March, at a valuation exceeding $1.75 trillion. Two days before the first strikes, xAI began buying back $3 billion in debt to clean its balance sheet for the offering. And on the day the Ossoff-Van Hollen letter's deadline expired—March 2, the date by which Defense Secretary Pete Hegseth was required to explain Grok's integration into Pentagon networks—Hegseth was at the podium briefing reporters on combat operations and American casualties.

This essay does not argue that the war was launched to provide cover for a missed congressional deadline. That would be conspiratorial, and the evidence does not support it. What this essay argues is something more structurally troubling and considerably harder to dismiss: that the war, regardless of its independent justification, has activated a set of dependencies between the state and a single private entity in a way that makes democratic oversight of those dependencies operationally dangerous to exercise. The immune system has not been captured. It has been placed in a position where activating it would harm the host.

The analytical framework for what follows was established across four prior essays: *The Mortal Architecture*, *The Discipline of Forgetting*, *The Sovereign Substrate*, and *Beyond the Stitch*. Each identified a distinct failure mode in complex systems—state drift, overfitting, vendor lock-in, the commodification of public infrastructure. What the events of late February and early March 2026 reveal is that these are not four separate pathologies. They are four symptoms of a single convergence, and the convergence is now live, under load, during a shooting war.

* * *

I. The War as Immunosuppressant

The Ossoff-Van Hollen letter, dated February 9, 2026, posed four questions to Secretary Hegseth about xAI's Grok chatbot operating inside Department of Defense networks. The questions were specific: What governance framework applies? How are antisemitic outputs mitigated? Do servicemembers' prompts flow to X or SpaceX? Who audits the system's classified outputs? The letter set a deadline of March 2.

No response has been made public. No Pentagon statement acknowledging receipt has been issued. No press release from the senators' offices confirms a reply. On March 2, Hegseth was occupied with the operational demands of a war that began seventy-two hours earlier. The deadline was not defied. It was absorbed—consumed by the gravitational field of a kinetic conflict that made questions about AI governance seem, in the immediate political calculus, secondary.

But the war does not merely distract from oversight. It *activates the dependency* that oversight was designed to examine.

MILNET—the 480-satellite military communications constellation sole-sourced to SpaceX—is presumably routing military communications during active combat operations against Iran. Grok, operating on GenAI.mil at Impact Level 5, is available to all 3 million DOD personnel for intelligence analysis during a conflict that requires real-time processing of targeting data, signals intelligence, and operational planning. Every node in the critical path runs through infrastructure owned, operated, and maintained by the entity whose oversight the Ossoff-Van Hollen letter was attempting to initiate.

The structural consequence is precise: pulling MILNET for audit during an active campaign risks communications disruption. Restricting Grok's access to classified networks during wartime intelligence processing risks analytical degradation. Demanding answers about the entity's capital structure—the Saudi sovereign wealth fund's equity position, the CFIUS review that may or may not have occurred—while that entity's infrastructure carries the operational traffic of a war becomes, in the language of institutional politics, "not the time."

The Mortal Architecture described reflexive entropy as the condition in which corrupted institutional components "can perceive the apoptotic signal coming, model the immune

system's detection criteria, and take strategic action to subvert it." But the essay's four proposed countermeasures—illegibility, metabolic apoptosis, ephemeral oversight, adversarial network architecture—all assumed a system at rest, or at least a system whose operations could be paused for maintenance. The biological parallel was explicit: "Synaptic pruning requires unconsciousness." The system cannot simultaneously process information and flush its waste.

What happens when the system cannot be put to sleep?

A war is the institutional equivalent of a system that cannot be taken offline. The factory floor is running triple shifts. The janitor cannot enter. The oversight mechanisms that require access to the system's internal state—audit rights, interoperability testing, firmware escrow verification—become operationally incompatible with the system's primary function. And the longer the war continues, the more institutional memory accumulates on the infrastructure that cannot be examined, the more decision chains are built on outputs that cannot be validated, and the harder it becomes to perform the schema migration even after the war ends.

This is the deepest form of reflexive entropy: not a system that fights the immune response, but a system whose operational conditions make the immune response itself pathogenic. The war is not the senescent cell. The procurement relationship is the senescent cell. The war is its SASP—the inflammatory secretion that degrades the surrounding oversight tissue by making correction look like autoimmune suicide.

<p style="text-align:center">* * *</p>

II. The IPO as Democratic Capture

Reflexive entropy, as described in *The Mortal Architecture*, operates through four strategies: the target predicts the detection algorithm, fires the auditor, defunds the oversight body, or refuses to comply. The essay acknowledged a fracture in each strategy and concluded that "the best achievable outcome may be managed chronic condition."

The events of March 2026 reveal a fifth strategy that the framework did not name, one that subsumes all four fractures simultaneously: democratic capture through equity distribution.

Bloomberg reported on February 27 that SpaceX is targeting a confidential SEC filing as soon as this month, with Bank of America, Goldman Sachs, JPMorgan Chase, and

Morgan Stanley underwriting. The capital raise could reach $50 billion—surpassing Saudi Aramco's 2019 record. Prediction markets assign a 76 percent probability of an IPO before September 2026. Polymarket gives over 40 percent odds the offering closes above $2 trillion.

Once SpaceX is publicly traded at $1.75 trillion, it will be held in index funds, mutual funds, ETFs, 401(k)s, and IRAs across the American financial system. Millions of ordinary Americans—teachers, nurses, firefighters, retirees—will hold fractional ownership of the entity that sole-sources the military's satellite communications and processes intelligence through Grok on classified networks. At that point, "regulate the company" becomes "destroy ordinary Americans' retirement savings during a war."

This is not merely a political cost that raises the bar for action. It is a structural inversion of democratic accountability. The same citizenry whose interests the oversight mechanisms exist to protect becomes the constituency that opposes their activation—not because citizens have been deceived, but because their financial interests are now genuinely, materially aligned with the entity's continued immunity from regulation. The public does not need to be fooled. It needs only to own shares.

The *Policy Memo* submitted to the Senate Armed Services Committee identified this dynamic: "the political cost of regulatory action rises because it becomes a cost imposed on ordinary Americans' retirement savings." But the memo, constrained by the conventions of policy writing, understated the structural implication. This is not a cost increase. It is a phase transition. Before the IPO, the entity's survival depends on government contracts, which can be conditioned or revoked. After the IPO, the entity's survival is entangled with the financial system itself, and its disruption becomes a systemic risk that no elected official will voluntarily trigger.

The immune system's host starts fighting the immune system on behalf of the pathogen—not through corruption, but through mutualism. The pathogen has made itself metabolically essential to the host's financial health. *The Mortal Architecture* proposed metabolic apoptosis as a remedy: "make institutional resource flows continuously contingent on demonstrated value," so that the corrupted institution desiccates as capital flows elsewhere. The IPO inverts this mechanism. The capital does not flow away from the entity. It flows *into* it, from millions of individual accounts, creating a diffuse constituency whose aggregate political weight exceeds any oversight body's mandate.

The pre-IPO window—the period during which ownership is concentrated among identifiable institutional investors and sovereign wealth funds, during which conditions can be attached and structural reforms mandated—is the last moment at which the tools of democratic accountability have something to grip. It is closing now, during a war that makes gripping anything politically impossible.

* * *

III. The Blacklisting as Diagnostic Confirmation

On February 27, 2026, CNN reported that the Trump administration ordered military contractors and federal agencies to cease business with Anthropic, designating it a "supply-chain risk to national security." The designation followed Anthropic's refusal to remove safety guardrails prohibiting mass surveillance applications and autonomous weapons integration. Senators Markey and Van Hollen issued a joint statement demanding Hegseth stop what they called a "pressure campaign" against a company "for refusing to enable mass surveillance and autonomous warfare."

The same week, Grok—whose own safety assessment by the GSA on January 15 found it "does not meet the safety and alignment expectations required for general federal use," describing it as "sycophantic and susceptible to corruption by biased data"—was approved for use in classified military systems.

This is the single cleanest empirical confirmation of what *The Mortal Architecture* called semantic institutional senescence: "the service responds promptly with syntactically valid payloads whose informational content has degraded." The Pentagon's institutional machinery is producing outputs that are procedurally correct—official designations, procurement memoranda, contract modifications—whose semantic content is precisely inverted. The institution deploys "national security" language to designate a company that maintains safety constraints as a supply-chain risk, while accelerating integration of a system that its own government's safety review deemed unfit for general use.

The institution passes every health check. It files paperwork. It issues memoranda. The output format is indistinguishable from healthy institutional function. But the informational content is corrupted in a way that actively degrades the system's security posture while wearing the uniform of security enhancement. This is the gray failure operating in daylight—HTTP 200 OK, payload poisoned—and the monitoring system

cannot detect it because the monitoring system (congressional oversight, inspector general reviews, interagency safety boards) measures format, not meaning.

The practical consequence is that the one mechanism the *Discipline of Forgetting* prescribed for avoiding cognitive lock-in—maintaining independent analytical capacity to validate whether cached outputs are stale—is being systematically dismantled. The Pentagon cannot audit Grok's outputs against an independent baseline if the independent baseline providers are being blacklisted. The system is building a cognitive monoculture at exactly the moment when its cognitive infrastructure is being deployed on classified networks during a war. The cache is growing. The invalidation mechanism is being removed. And the system cannot know which Grok outputs are sound and which reflect the failure modes that the NSA's own classified November 2024 review flagged—because the entity that could provide adversarial verification has been designated a threat.

<p align="center">* * *</p>

IV. The Departures as Signal

Six of twelve xAI co-founders departed within weeks of the SpaceX merger. A seventh, Toby Pohlen, left on February 27. The departures have been covered as a corporate culture story: SpaceX's operational intensity clashing with xAI's research ethos. Engineers complained, according to SatNews's February 16 reporting, about the post-merger reorganization into four divisions and the cultural realignment toward the "Sentient Sun" strategy.

The Mortal Architecture prescribes "out-of-band behavioral phenotyping" as the detection mechanism for gray failures—profiling a component's behavioral output from outside, rather than trusting the component's self-reported health. The co-founder departures are precisely this kind of behavioral signal, and they are being misread.

The people who left are not interchangeable engineers. They are the researchers who designed Grok's architecture, understood its training data, knew its failure modes, and could identify when its outputs were unreliable. Their departure from the entity that is now processing classified military intelligence is not a human-resources story. It is the cognitive equivalent of a hospital losing its entire diagnostic radiology staff while continuing to operate the scanners. The machines still produce images. The images still look like medical imaging. But the interpretive capacity—the knowledge of what

constitutes an artifact versus a finding, what the machine can see and what it misses—has walked out the door.

The entity passes every institutional status check: it has contracts, it has a valuation climbing toward $2 trillion, it produces outputs, its CEO gives press conferences alongside defense officials. But the cognitive core has hemorrhaged. The building is occupied. The occupants have changed. The nameplate has not.

No reporting I have found frames the co-founder exodus as a national security signal. The framework says it is one—and that the failure to recognize it illustrates exactly the monitoring blindness that behavioral phenotyping was designed to correct. Status indicators say healthy. Behavioral output says otherwise.

<div align="center">* * *</div>

V. Cognitive Lock-In: Where Forgetting Meets the Sovereign Substrate

The Discipline of Forgetting argued that "system adaptability is inversely proportional to the completeness of a system's memory of its own operational history." The essay's most sobering observation was that large language models, once trained, have no forgetting mechanism—that their weights are fixed and their training data's influence is distributed across billions of parameters in ways that cannot be selectively addressed.

The Sovereign Substrate argued that sole-source vendor lock-in in military infrastructure creates dependencies that constrain accountability. The lock-in is in the hardware: proprietary terminals, proprietary encryption, proprietary inter-satellite laser links.

These two arguments have a junction point that neither essay names explicitly, and it is the most structurally dangerous one.

Every day that Grok processes intelligence on classified Pentagon networks, it generates institutional memory. Briefing materials incorporate its analytical outputs. Decision chains build on its assessments. Operational records embed its language, its framings, its pattern-recognition artifacts. Commanders act on intelligence products shaped by Grok's distributional biases—biases that the GSA safety report characterized and the NSA review flagged, but that no independent adversarial system is now available to identify in real-time, because the adversarial provider has been blacklisted.

This institutional memory cannot be "forgotten" without auditing every decision chain that depends on it. Switching AI providers is not a software migration. It is, in the precise language of *The Discipline of Forgetting*, a schema migration—"the most feared operations in production engineering, not because they are technically complex, but because they require the system to forget its own assumptions while remaining operational. Forgetting under load is harder than forgetting at rest, for the same reason that synaptic pruning requires unconsciousness."

A schema migration during a war is a system being asked to forget its own assumptions while the factory floor is running triple shifts, during the most dangerous operational tempo in a generation. The political and operational barriers to switching are not static. They grow with every day of operations. The lock-in is not just in the infrastructure. It is in the *cognition the infrastructure has already produced*—in the decision chains built on Grok's outputs, in the analytical patterns that have become institutional habit, in the officer corps's learned reliance on a tool whose internal workings its builders have abandoned.

The Discipline of Forgetting identified cache invalidation as the paradigmatic forgetting problem: "the cache cannot know that it is stale, because the very definition of staleness is that the world has changed in a way the cache hasn't observed." The Pentagon cannot know which Grok outputs are reliable and which encode the antisemitism, the sycophancy, and the data-poisoning vulnerabilities that its own internal reviews documented—because the capacity for independent verification is being destroyed. The system is building a cache it cannot invalidate, running on infrastructure it cannot replace, processing through a model whose architects have departed, owned by an entity it cannot regulate, during a war it cannot pause.

That is the convergence. Not a metaphor. A mechanism.

* * *

VI. The Isomorphism of Scale

One final structural observation, drawn from the only essay in the series that operates at the municipal rather than federal level.

Beyond the Stitch documented the I-75 Cap project in Detroit: approximately $200 million in public infrastructure designed by a nonprofit whose board includes the developers who will benefit most from its construction. The developers contributed

$400,000—a ratio of private investment to public expenditure of 1:500. The first building adjacent to the caps shed its affordable housing commitment before its foundation was poured. Community input was channeled into commemorative design while governance questions were deferred to "future phases." The DDP's CEO circulated for photographs while engineers addressed the public.

At the federal level, MILNET represents $277 million in FY2026 alone, sole-sourced to an entity whose capital structure includes Saudi sovereign wealth. Congressional oversight is channeled into letters whose deadlines arrive during wars. Governance questions—who audits the firmware? who reviews the investor base? what open standards prevent vendor lock-in?—are deferred to "future consideration." The Defense Secretary briefs the press about combat operations while the AI oversight letter sits unanswered.

The extraction geometry is identical. Public capital flows in; binding commitments flow out. The entities that benefit most sit on the advisory structures that shape the design. The commemorative apparatus—storytelling booths in Detroit, congressional letters in Washington—absorbs community concern without converting it into structural constraint. The engineering proceeds on schedule. The governance remains vapor.

In Detroit, the mechanism produces a park that functions as a real estate catalyst for luxury hotels. In orbit, it produces a satellite constellation that functions as an accountability shield for a $1.75 trillion IPO. The scale differs by orders of magnitude. The architecture is the same: public infrastructure as private subsidy, with the public providing the capital and the advisory structures providing the appearance of oversight without its substance.

Peter Frase called it Rentism: technological abundance plus extreme hierarchy, in which the elite provide the vision, the public provides the capital, and the private sector extracts the rent. The I-75 Cap generates rent through property-value appreciation. MILNET generates rent through sole-source dependency. The SpaceX IPO will generate rent through equity participation—the final transformation, in which the public does not merely fund the infrastructure but *invests in the entity that controls it*, converting every citizen-shareholder into a stakeholder in the system's continued immunity from oversight.

＊ ＊ ＊

VII. What Remains

The Mortal Architecture concluded with four open problems it could not solve. *The Discipline of Forgetting* added three more. The events of March 2026 add one that is, in some respects, more fundamental than any of them:

The immune activation problem in dependent systems. All of the framework's prescriptions—apoptosis, metabolic starvation, adversarial oversight, behavioral phenotyping—assume that the immune system can be activated without damaging the host. When the host's critical functions are operationally dependent on the entity to be corrected, and when a kinetic conflict has placed those functions under maximum load, activation of the immune system becomes indistinguishable from an autoimmune attack. The question is not whether the tools of oversight exist. They do. The question is whether there exists a deployment sequence—a specific ordering of interventions, timed to windows of reduced operational dependency—that permits correction without catastrophic disruption.

The pre-IPO window is one such moment. It will not last. The war will eventually end—but the IPO, once completed, distributes the dependency across the entire financial system and makes the correction problem permanent. The schema migration becomes structurally impossible, not because the technology cannot be replaced, but because the political constituency for replacement has been dissolved into the constituency for continuation.

The Mortal Architecture argued that systems endure by mastering the art of dying well. *The Discipline of Forgetting* argued that systems adapt by mastering the art of forgetting well. What March 2026 demonstrates is a third principle, darker than either: **the most dangerous systems are the ones that have made themselves impossible to forget—that have woven their operations so deeply into the host's critical functions that pruning them requires a violence the host cannot survive.**

The tools exist. The map has been drawn. The window is closing. And the factory floor is on fire—which is not, it turns out, the worst time to perform maintenance. It is the time when the absence of maintenance becomes permanent.

* * *

The facts cited in this essay are drawn from public reporting by CENTCOM, CBS News, NPR, Bloomberg, CNBC, CNN, Al Jazeera, Fortune, Axios, Breaking Defense, SpaceNews, SatNews, The Register, Tom's Hardware, Military Times, Defense One, Interesting Engineering, Fudzilla, Benzinga, Yahoo Finance, the offices of Senators Ossoff, Van Hollen, Warren, and Markey, and official filings and statements from SpaceX, xAI, the FCC, the GSA, and the Department of Defense. Specific source attributions are provided inline throughout. All claims reflect the public record as of March 3, 2026.