

THE
SENTINEL
ON
NATIONAL INTELLIGENCE,
WRITTEN IN THE YEAR 2026,
BY
MR. FOXWELL.

PREFATORY REMARKS.

The central idea is the common surface. Most agency work will remain agency work: collection, investigation, military support, technical exploitation, diplomatic analysis, and source handling that do not seek community status or cross-domain use. The common surface appears when an intelligence object is reused across agencies, supported by common funds, presented as a community judgment, routed across lawful domains, powered by a shared model, or recorded as an exception to a common rule.

The operator is named functionally rather than branded permanently. Today the lawful home is the national intelligence leadership created by Congress and housed in ODNI. Future statutes may adjust the office, title, or internal chart. The theory does not depend on one proper noun. It depends on a public office with authority to recognize common status, deny common recognition, require repair, report avoidance, and leave records for inspectors and Congress. This is a theory of public control over common intelligence, not a proposal to freeze an org chart in amber.

The papers claim no classified knowledge. They reason from public law, public institutional structure, official materials, and visible incentives. They distinguish public-law claims from design hypotheses and open risks. The central question is not whether every workflow can be redesigned from the outside. It is whether a republic can require lawful conditions at the points where secret agency work becomes shared national judgment.

THE SENTINEL.

NO. I.

BY RIVER FOXWELL.

Introduction.

The United States does not lack intelligence capacity. It lacks a settled theory of commonness. Agencies collect, analyze, investigate, warn, protect, and operate under different laws and missions. That specialization is necessary. Yet the most consequential intelligence objects increasingly cross the walls that created them: commercial datasets, model outputs, cross-domain warnings, community assessments, private-sector telemetry, and presidential briefings built from many streams. Once an object crosses those walls, possession is no longer a sufficient theory of authority.

The common surface is the part of intelligence where agency work becomes shared government power. It is not every secret, every database, or every operation. It is the surface on which separate institutions ask the rest of the state to rely on their work as common: use this data, trust this assessment, follow this route, certify this tool, accept this exception, spend this money, or remember this decision. That request for shared reliance is the moment public law should become sharper.

The older drafts tried to describe this through a protocol layer. The phrase is useful only if it is not allowed to lie. Law is not hardware. Agencies can resist, relabel, delay, or move work toward a stream where common authority has less grip. A router cannot continue as normal while ignoring the network. An intelligence agency can continue doing many things while avoiding community recognition. The design must therefore speak less like engineering and more like government: recognition, eligibility, refusal, consequence, record, and review.

The common surface has four governing questions. First, what has entered the surface? Second, what lawful conditions travel with it? Third, who may contest those conditions? Fourth, what remains when the decision is later examined? These questions can be answered through records, budgets, adversarial review, exposure limits, and routing decisions; they should not be dressed as twenty inventions. The unity is conceptual, not bureaucratic: common reliance creates public obligations.

This compression is important. A data record, an analytic label, a route, a privacy limit, and a black-box entry all look related because each asks the same public-law question: what must be true before this secret object can be used by more than one institution in the name of the whole? But they are not the same mechanism. A budget condition changes incentives. An adversarial review tests judgment. An exposure limit constrains how visible a person becomes to the state. A routing decision names ownership across law. A durable record gives future oversight something other than memory to inspect.

The authority to operate the surface should be held by the lawful national intelligence leadership, under a charter that defines common status and forbidden zones. Today that means ODNI and the DNI, because current law already places national intelligence integration, NIP

budget functions, information-sharing duties, and anti-duplication responsibilities there. But the theory is deliberately functional. If the internal title changes, the principle survives: the office that gives national recognition to common intelligence must also impose lawful conditions on that recognition.

The surface is not a throne. It cannot command every collector, analyst, investigator, or commander. It cannot make an agency virtuous, a President restrained, or Congress vigilant. It can deny community status to an unmarked product, deny common-system admission to an ungoverned dataset, require a routed seam to name ownership, make a common warning carry dissent, make high-risk reuse of sensitive information spend justification, and preserve a record when an exception is used. That is not total power. It is the power to say what counts as common.

The central claim is therefore smaller than earlier rhetoric and stronger because it is smaller. Secret power cannot be made safe by fragmentation alone, and it cannot be made legitimate by usefulness alone. It becomes governable when the objects that cross agency walls enter a common surface with conditions the state can later prove it obeyed.

NO. II.

BY RIVER FOXWELL.

The 2004 Test.

The most dangerous precedent for this proposal is the last major reform. The Intelligence Reform and Terrorism Prevention Act of 2004 created the Director of National Intelligence and ODNI after a public diagnosis of diffuse responsibility and failed integration. The reform gave the new office real authority. It did not merely create a symbol. Public law gives the DNI access to national intelligence, important National Intelligence Program budget responsibilities, information-sharing roles, authority over priorities and standards, and a duty to eliminate unnecessary duplication.

Yet the settlement preserved the departmental architecture beneath the national title. Money, personnel, missions, operational loyalty, professional identity, and congressional patrons remained inside host departments. The result was not simple powerlessness. It was something more ambiguous: national responsibility attached to an office that still had to operate through institutions with their own chains of command and sources of leverage. That is the *de jure/de facto* problem this series cannot evade.

A critic can state the objection brutally. If the DNI already had budget authority over the National Intelligence Program, access to intelligence, and integration duties, why would a common-surface regime do better? If agencies could resist the last reform while formally accepting it, why would they not resist this one? If the real mass of power sits in people, daily money, lawyers, commanders, field offices, and departmental patrons, why should recognition at the common surface matter?

The answer begins by refusing overclaim. The common-surface regime does not master every intelligence activity. If an agency can keep work truly local, fund it locally, use it locally, and never ask others to treat it as a community asset or judgment, the regime may not bind the work

directly. That limit is not a footnote. It is part of the theory. The proposal governs the points where intelligence becomes common in fact; it does not pretend that all intelligence is always common.

This makes the project narrower than an intelligence constitution for the entire state. It is a regime for transitions: local-to-common data, agency-to-community analysis, operational-to-national warning, private telemetry-to-government use, model output-to-policy judgment, department claim-to-cross-domain route, secret exception-to-reviewable memory. Those transitions are not minor wiring. They are where many of the most consequential failures occur: confidence inflates, privacy travels poorly, dissent disappears, responsibility fragments, and no one later knows who turned agency work into national judgment.

The difference from 2004 is not that the new office receives a grander name. Names have disappointed before. The difference is conditional recognition. The law should not only tell a center to coordinate. It should define what cannot receive common status without satisfying common conditions. A product that lacks lineage and dissent treatment can remain an agency product, but it cannot be sold as the community's voice. A dataset without provenance and use rules can remain inside a lawful local mission, but it cannot enter shared use. A cross-domain crisis may proceed locally for a time, but it cannot become the government's integrated route without a named owner and recorded boundaries.

Conditional recognition will still be evaded. An agency may relabel the work as local, use departmental funds, or avoid asking for community status. The answer is not to pretend the law can prevent every evasion. The answer is to make avoidance itself an object. When an agency asserts local-only status against evidence of multi-agency reliance, shared funding, cross-domain use, presidential consumption, or material risk to rights, the assertion should be reported to inspectors and appropriate committees. That notice may not compel compliance. It changes the political cost of avoidance and gives Congress a specific fight rather than a mist of coordination failures.

This is why the National Intelligence Program and the Military Intelligence Program must be treated honestly. NIP levers are real where they reach. MIP and military operational streams are not simply extensions of the DNI's hand. A common-surface rule can bind defense-funded intelligence only when Congress, the President, and the Secretary of Defense decide that a defense activity has crossed into national common use. Without that political decision, the common-surface regime becomes mostly a notice, audit, and status system at the military edge. That is a weakness. Stating it is stronger than hiding it.

The 2004 test therefore does not defeat the project. It disciplines it. The reform must not promise a center that commands all intelligence. It must promise a rule for the moments when intelligence asks to become common. If those moments are rare, the project shrinks. If modern intelligence increasingly depends on shared data, shared models, community assessments, private-sector surfaces, and cross-domain warning, the surface is large enough to matter. The factual question belongs in pilots, audits, and congressional fights, not in metaphor.

NO. III.

BY RIVER FOXWELL.

What Enters.

The hardest question is not who operates the common surface. It is how much of intelligence ever touches it. A theory that governs only the connective tissue while the organs do the decisive work elsewhere may be elegant and marginal. A theory that pulls everything into the surface becomes centralization under another name. The boundary must be argued, not assumed.

Agency-specific work remains outside the surface when it is conducted under one agency's lawful mission, funded and governed within that mission, not presented as a community judgment, not reused by multiple elements, not entered into a shared system, and not likely to affect rights through cross-agency use. The point is not to create loopholes. It is to preserve the legitimacy of specialized craft. A Navy intelligence stream supporting fleet operations, an FBI investigative lead used inside a case, or a CIA source report held inside a compartment may be lawful and important without becoming common.

Work enters the surface when the government asks other institutions to rely on it as more than local. The first threshold is shared reliance: more than one element uses the same dataset, model, feed, system, or archive for national intelligence purposes. The second is community voice: an assessment is presented as a community product rather than an agency position. The third is cross-domain routing: a threat requires law enforcement, intelligence, military, diplomatic, private-sector, and civil-liberties equities to move in relation to one another. The fourth is rights exposure: information about U.S. persons becomes queryable, linkable, or reusable beyond the purpose that first justified it. The fifth is public consequence: a warning, attribution, sanctions package, use-of-force option, or presidential decision relies on combined secret material.

These thresholds are not clean lines in every case. That is why the boundary needs a contest procedure. An agency asserting local-only status should be able to say why: source protection, operational command, legal restriction, investigative integrity, mission specificity, speed, or lack of shared reliance. The common operator should be able to answer why the threshold is crossed. If unresolved, the dispute should not vanish into meetings. It should become a reviewable boundary claim.

The boundary claim is a second kind of device in this manuscript. It is not a wrapper attached to data. It is an adversarial institutional act. One side says: this remains local. The other says: this has become common. The record states the reasons. Inspectors can later compare the claim to actual use. Congress can ask whether repeated local-only claims are hiding common reliance. The design does not abolish evasion; it gives evasion a shape when it is discovered.

The claim may still be too narrow. Many important decisions happen inside agencies before a common threshold is reached. An investigation may frame the facts before an assessment becomes community-wide. A military intelligence stream may shape options before national recognition occurs. A classified tool may discipline analysts before anyone calls it shared infrastructure. The common-surface regime cannot honestly say it governs all of that. It can say that when those

upstream decisions are later used as common justification, the system must disclose lineage, dependence, uncertainty, and unresolved dissent at the point of national reliance.

That point of national reliance is the answer to the charge that the surface is only wiring. In the machine-age intelligence system, wiring is not incidental. The wiring determines which data travels, which model summarizes, which assumption propagates, which route loops, which warning reaches the President, which dissent disappears, and which record exists for later law. The surface may be a fraction of activity by volume. It may be a majority of consequence by effect.

A pilot should test that claim. Count how many high-consequence judgments depend on common data. Count how many cross-domain routes recur. Count how often community products rely on shared vendor feeds or shared models. Count how often local-only claims later become community reliance. If the common surface is thin, the regime should stay thin. If the surface is thick, the state should stop pretending that agency ownership is enough.

The boundary of the claim is therefore an admission and a bet. The admission is that most intelligence activity will not be governed directly by this proposal. The bet is that the most republic-facing failures increasingly occur when agency work becomes common without adequate law attached to its commonness. The honest reformer does not bury that bet. He designs the pilot that can prove or disprove it.

NO. IV.

BY RIVER FOXWELL.

The Defection Case.

A design tested only on cooperation is a brochure. Consider a hypothetical ransomware attack on a regional water utility. Pumps and billing systems fail, a private vendor holds the logs, local officials fear contamination, and early indicators resemble a criminal ransomware crew. Some infrastructure also appears in foreign-intelligence holdings associated with state-tolerated cyber activity. The case is a seam: local emergency, private infrastructure, criminal investigation, network defense, foreign warning, possible sanctions, and potential presidential interest.

The cooperative version is easy. The cyber route opens, the integration lead is named, the vendor feed enters shared use with provenance and restrictions, privacy exposure is limited, analysts compare criminal and state-tolerated hypotheses, and the record preserves emergency decisions. That case proves little. The real test begins when actors defect.

Suppose the FBI office labels the matter criminal-only and resists community status because it wants to protect the investigation, preserve prosecutorial options, and avoid intelligence handling rules that may complicate evidence. Some of that resistance is legitimate. Criminal process has integrity of its own. But the incident is already being briefed to national officials, using foreign-intelligence comparisons, with victim information relevant to public warning and possible sanctions. The local-only label is now contestable because the work is being used beyond the case.

Under the common-surface regime, the national operator cannot seize the prosecution. It can challenge the status claim. The FBI may continue the investigation as FBI work. It may not simultaneously treat the assessment as local-only while allowing it to anchor a community cyber

warning without lineage, dissent, purpose, and privacy treatment. If the Bureau refuses community status, the refusal is recorded with grounds: evidentiary integrity, investigative sensitivity, source protection, or urgency. The community product either excludes the contested claim, marks it as FBI-specific, or escalates the dispute. The result is imperfect. It slows integration. It also prevents a hidden investigative preference from quietly becoming the national view.

Suppose next that a Defense element says related intelligence is operational and therefore outside the common surface. The claim may be correct. Military intelligence in support of force protection is not automatically national common infrastructure. But if the same stream supports a presidential attribution judgment, a cross-domain route, or a warning to civilian water utilities, the defense-local claim changes character. The common operator can require a boundary notice: which portion remains operational, which portion is being used for national reliance, what can be shared, and what cannot. If Defense declines, the record moves to the DNI, the Secretary of Defense, and the committees with jurisdiction. That is not total compulsion. It is the law admitting where the real fight is.

Now suppose both defections succeed in the short term. The FBI proceeds criminally, Defense holds its stream, the vendor negotiates narrow terms, and the common product reaches senior officials later than it should. This is the point a serious proposal must not evade. The regime may not catch the defection in time. It may only create evidence: a local-only assertion, an operational-boundary assertion, an unresolved route, a delayed warning, an exception that expired, and a public top-line next year showing repeated avoidance in water-sector cyber incidents.

The floor must be said plainly. In the worst case, the common-surface regime becomes a machine for generating audit trails no one reads. Congress may ignore the evidence. Inspectors may lack resources. Officials may learn to satisfy the forms while preserving the old behavior underneath them. The design cannot compel courage from oversight, judgment from committees, or honesty from departments that have already chosen evasion. Its payoff is contingent on actors it can inform but not command. At its floor, it produces paperwork.

That is a low floor, but not zero. Paperwork is not accountability; it is the minimum material from which accountability can later be made. A record ignored in one year can be read by another committee, sampled by an inspector, preserved for a commission, or used to prove that avoidance was a pattern rather than an anecdote. Without the regime, a later inquiry asks whether agencies coordinated sufficiently. That question is soft enough for everyone to answer yes and no at once. With the regime, the inquiry asks who asserted local-only status, who accepted it, what comparison was withheld or marked, who allowed operational boundaries to block civilian warning, and whether the same avoidance pattern recurred. The design does not guarantee accountability. It prevents failure from having no artifact.

The water case also tests privacy. Vendor telemetry collected for network defense may include customer identifiers, employee access records, or logs from unrelated systems. Intelligence analysts want to reuse it for attribution. The repurposing decision should be narrow: the new purpose, the fields exposed, the roles permitted, the expiration, and whether derived products inherit restrictions. In a crisis, some access may occur first and be repaired later. That repair does not undo the exposure. The design cannot buy back visibility already created. It can only prevent emergency visibility from becoming permanent infrastructure.

The case ends without triumph. The utility is restored. Attribution remains probabilistic. The criminal case is slower because some intelligence cannot enter evidence. Defense refuses part of

the boundary claim. A year later, inspectors find that local-only and operational labels were used in multiple critical-infrastructure incidents. Congress receives a pattern. The question becomes political: will it amend the charter, condition funds, demand joint rules, or tolerate the escape? That is how public control actually appears - not as a perfect device, but as a fight made visible enough to have.

NO. V.

BY RIVER FOXWELL.

Evidence Without Echo.

Intelligence often fails by hearing one thing many times. A claim begins in a vendor feed, moves through a model summary, appears in an agency report, is cited by another assessment, and returns to the President as apparent convergence. The system now believes it has several witnesses. It may have one witness wearing several uniforms.

The remedy is evidence lineage, not certainty theater. A community product should not merely assign confidence. It should describe the ancestry of agreement: which judgments share data, which share models, which share source chains, which share prior products, which rest on distinct methods, and which dependencies remain unknown. This is not a demand for perfect genealogy. It is a defense against the most basic error in a data-saturated state: counting echoes as independent confirmation.

Lineage has levels because tempo varies. In crisis, the product may carry a minimal profile: independent, partly dependent, common upstream likely, or unknown. That is enough to warn without lying about confidence. In ordinary tempo, lineage should identify major source classes, vendor feeds, partner streams, models used, prior products relied upon, and known contamination risks. After major decisions, a deeper audit should repair what the crisis profile could not know. The label is not final truth. It is a promise to update the ancestry of judgment.

This is where adversarial review matters. A lineage claim should be examined by people not invested in the product's prestige: tradecraft officers, source specialists, data stewards, model evaluators, and analysts from outside the main line. Their task is not to choose the conclusion. It is to test whether apparent agreement rests on distinct foundations. This is a different kind of mechanism from passports and logs. It is contest, not annotation.

Dissent changes under this view. A dissent based on separate evidence challenges the evidentiary base. A dissent based on the same evidence challenges inference. A dissent based on a model-retrieval flaw challenges the evidence universe itself. All three matter differently. The consumer should not receive dissent as undifferentiated friction. He should know what kind of disagreement exists and what would change it.

Artificial intelligence makes lineage both harder and more necessary. A model can retrieve, rank, summarize, translate, cluster, and draft before a human knows what universe of evidence has been created. Two teams may appear independent while relying on the same embedding space, vendor corpus, model-generated summary, or retrieval index. Agreement between them is not worthless, but it is not clean convergence. It is model-mediated agreement and should be treated as such.

The machine also creates a new excuse. The analyst says the tool surfaced the pattern. The manager says the workflow selected the package. The policymaker says the community agreed. Responsibility can dissolve across human, model, vendor, and database. The answer is not to ban useful opacity. Some tools will remain partly opaque. The answer is role accountability: who used the tool, what the tool did, which output was accepted, what was rejected, what was not independently verified, and who signed the final judgment despite those limits.

Evidence lineage has costs. It can slow production. It can expose uncomfortable dependence on vendors or partners. It can make polished consensus look less impressive. It can also be gamed. Agencies may overstate distinctiveness; vendors may obscure common upstream sources; model providers may resist disclosure. Those risks do not make lineage unnecessary. They mean lineage itself must be sampled, challenged, and corrected after the fact.

The deepest value is intellectual humility enforced by workflow. A republic should not require intelligence to be certain before it acts. It should require intelligence to know whether its uncertainty is independent, inherited, machine-shaped, politically pressured, or simply unknown. Confidence is not the same as truth. Lineage tells the government what kind of confidence it has bought.

NO. VI.

BY RIVER FOXWELL.

The Citizen in the Substrate.

The common surface becomes most dangerous at the substrate: datasets, metadata, identity fields, query logs, model outputs, access roles, retention clocks, and derived products. These are the materials that make intelligence governable. They are also the materials that could make the state more capable of seeing a person from many angles at once. A civil-liberties critique need not exaggerate. Governed concentration is still concentration.

The answer cannot be that scattered systems are safer. They may be. They may also hide abuse. A person whose information sits in several systems under incompatible rules is not necessarily protected; he may simply be harder for anyone to defend because no one can reconstruct how the data moved, who queried it, or why it was reused. The real question is comparative: does common governance reduce exposure while increasing reconstructability, or does it merely make the state more competent at search?

That question is empirical. The design should therefore treat privacy as a budget of exposure, not only as a principle. Sensitive categories should have limits on who may query them, how long linked records remain available, how many roles can see raw identity fields, how often high-risk queries are sampled, and when dormant access expires. The state should spend justification when it makes a person more visible to itself. If exposure expands in an emergency, the expansion should expire or be ratified. If it becomes routine, public-law oversight should know that routine has formed.

Purpose matters because silent migration is the ordinary path of abuse. Data collected or bought for one mission should not automatically become available for another merely because it sits in the same environment. A new use should require a new purpose, a lawful basis, proportional

review, and downstream limits. In crisis, access may come first. But repair must follow, and the repair should be visible to inspectors. Emergency should not become a memory hole where purpose binding dies.

Transformation matters because modern data rarely stays raw. It is linked, masked, scored, summarized, translated, clustered, and fed into models. A privacy rule that attaches only to the original file will vanish as soon as the file becomes a derived artifact. Restrictions should travel downstream unless law changes them. If a model output owes its existence to sensitive inputs, it should inherit enough of their obligations to prevent laundering privacy risk through transformation.

Query accountability may be more important than collection accountability. Old data becomes newly powerful when searched for a new person, pattern, or purpose. Queries involving sensitive U.S. person information should record role, purpose, authority, search class, and review status. Repeated unusual patterns should be visible to privacy officers and inspectors. The question is not only what the government has collected. It is how often the government causes a person to become visible inside its systems.

Exposure limits will be contested. Operators will say that a tight budget slows warning. Privacy officers will say that a loose budget legitimizes surveillance by spreadsheet. Both may be right in different cases. This is one of the collisions the design does not fully buy back. A pilot should measure whether common governance actually reduces exposure: fewer raw-identity views, fewer dormant accounts, fewer repurposings without review, fewer untraceable queries, more deletions when purpose expires. If those numbers do not improve, the common substrate has failed the citizen even if it has pleased the overseer.

The citizen does not need a romantic promise that secrecy will become harmless. He needs a legal architecture in which every expanded visibility has a reason, a role, a duration, and a record. The state may know dangerous things when law permits. It should not be allowed to forget how often it made a person knowable to itself.

NO. VII.

BY RIVER FOXWELL.

The Human System.

No design survives the people who can route around it. Intelligence is carried by professionals with missions, habits, pride, caution, ambition, exhaustion, loyalty, and fear. They are not villains for protecting their offices. They are not angels because they serve in secret. They are institutional actors under pressure. A public design that ignores that fact will become a form everyone learns to satisfy while continuing the old behavior beneath it.

The first human pressure is possession. Information means relevance, money, access to senior consumers, and proof that an office matters. A manager who guards a source or feed may be protecting quality and security. He may also be protecting his program's place in the ecosystem. The motives can merge. The point of the common surface is not to denounce possession. It is to change the moment when possession becomes reliance by the whole. At that moment, the holder becomes a steward with duties rather than a proprietor with preferences.

The second pressure is conformity. Analysts do not merely read evidence. They work inside climates. Some conclusions travel easily; others create work, anger, or isolation. Dissent can be correct and still be costly. The solution is not another speech about courage. Courage is easier when the institution gives it a channel. A marked dissent, a lineage challenge, or a model-process objection gives the analyst a lawful form for disagreement. It does not make him brave. It reduces the penalty for saying the system may be wrong.

The third pressure is program loyalty. People defend what they build, sometimes for good reasons. A system, model, source network, or analytic shop can carry years of expertise. Reform that treats every defense of program as selfish will produce resistance and deserve some of it. The better design redirects prestige. If the public honor lies in making the routed problem whole, preserving lineage, reducing exposure, and sustaining lawful memory, then pride can serve the common function without pretending that pride has disappeared.

The fourth pressure is secrecy itself. Secrecy narrows audience, protects mistakes from embarrassment, and makes insider language sound like proof. It can also protect lives and sources. The danger is not secrecy alone but unreviewable secrecy: custom hidden as necessity, delay hidden as caution, avoidance hidden as expertise. Durable records and inspector access change behavior because they create the possibility that today's informal habit will become tomorrow's pattern. That possibility is a cultural force.

The fifth pressure is contractor and vendor dependence. Modern intelligence increasingly touches private data, private clouds, private telemetry, private models, and private expertise. Vendors have their own incentives: profit, liability, opacity, repeat business, and protection of proprietary methods. The common surface should not assume that every dependency is governmental. A vendor feed can become an upstream sovereign if no one can see how many judgments depend on it. This is why lineage and exposure are human-system rules as much as technical ones; they reveal dependencies people may prefer not to name.

The human system can defeat every proposal in this series. Staff can fill fields carelessly. Agencies can classify around review. Managers can transform dissent into nuisance. Inspectors can be underfunded. Congress can accept briefings without pressing. A serious theory should not claim more than it can command. It can arrange incentives so evasion costs more, stewardship carries status, dissent has a channel, and repeated avoidance becomes visible. That is not virtue by statute. It is friction placed in the path of predictable vice.

The irreducible cases remain. Some officials will conform without being pressured. Some will game metrics. Some will confuse loyalty to office with loyalty to country. Some will betray. The design does not eliminate these failures. It tries to keep them from becoming the hidden operating system of the intelligence community. Good institutions do not make human beings safe. They make human failure harder to hide and easier to correct before it becomes tradition.

NO. VIII.

BY RIVER FOXWELL.

Conclusion.

The path into law should begin with a charter for common intelligence functions, not with a grand reorganization. Congress should define common status, assign the operating responsibility to the lawful national intelligence leadership, preserve forbidden zones, establish pilots, and require records. The office title matters less than the function: someone with public authority must be able to recognize common status, deny it, require repair, record refusal, and report avoidance. If no one can sign those acts, the theory remains prose.

The charter should include both entry and exclusion. Entry covers shared acquisition, shared datasets, community analytic products, cross-domain routes, high-risk AI workflows, enterprise services, and secret exceptions that affect common functions. Exclusion protects analytic conclusions, lawful tradecraft, prosecution, military command, diplomacy, and genuinely local systems. A statute that names only entry becomes a jurisdictional appetite. A statute that names only exclusion becomes another coordination memo. Both lists are needed.

Enforcement should be layered because no single lever reaches the whole intelligence system. NIP funding conditions matter where common national-intelligence assets depend on NIP money, transfer, reprogramming, or enterprise approval. Community product status matters for analysis. Common-system admission matters for shared data. Model certification matters for machine workflows. Boundary notices matter where an agency claims local-only or operational status against evidence of common reliance. Inspector reporting matters when money cannot reach. Congressional notification matters when departments route around the common surface. The layered design is less elegant than one master switch. It is more honest.

The military boundary requires special candor. The DNI participates in Military Intelligence Program development and provides guidance, but the MIP is not simply the NIP under another label. Military operational intelligence has its own constitutional and command logic. The common-surface charter should therefore create a defense-common rule through defense and intelligence legislation: when military-funded intelligence is used as national common reliance, community status, shared model input, cross-domain civilian warning, or presidential attribution outside the operational chain, common-surface conditions attach to that use. If the activity remains purely operational, the charter should not pretend to govern it. If departments abuse the operational label to avoid common duties, the remedy is notification, audit, and appropriations conflict, not fictional omnipotence.

Pilots should be narrow enough to fail. Commercially available information is one pilot because public policy already recognizes the need for documentation, safeguards, and attention to repeat access. Cyber routing is another because it crosses the public-private, criminal, foreign-intelligence, military, diplomatic, and civil-liberties lines by nature. AI-assisted community products are a third because machine influence is arriving before law has learned to name it. Each pilot should test actual claims: whether lineage improves judgment, whether exposure limits reduce privacy risk, whether routes reduce loops, whether records help inspectors, and whether agencies avoid the surface in practice.

Memory should be built before scandal. Directives, boundary claims, avoidance notices, emergency exceptions, repurposings, exposure expansions, route disputes, product-status changes, machine-use approvals, presidential directions touching analytic integrity, and inspector findings should enter a classified civic record. The record should be tiered to protect sources and methods. It should still preserve action, authority, objection, expiration, and oversight notice. Without such memory, every later investigation becomes argument with smoke.

Public reporting should be low enough to protect operations and real enough to matter. Counts and categories can reveal whether the system exists: common-admission decisions, refusals, boundary claims, high-risk queries, exposure expansions, expired exceptions, sustained dissents, model approvals, route openings, unresolved loops, and inspector audits. These top-lines will not prove virtue. They give citizens and committees a vocabulary for asking whether secret governance is being used or merely advertised.

The floor of the whole project is the same. Records may become archives no one opens; top-lines may become numbers no one interprets; sunsets may become votes no one wants to cast. If that happens, Sentinel 2.0 will have built a better memory than the status quo and still failed politically. It cannot force the republic to use the instruments it gives it. But this is a comparative argument. The status quo's floor is not freedom; it is often amnesia. An unread trail can still be recovered. An uncreated trail cannot.

Exit is not a decorative safeguard. Every new power should sunset unless evidence supports renewal. Congress should ask whether common governance reduced duplicative purchase, clarified evidence lineage, protected dissent, lowered privacy exposure, shortened routing loops, improved inspection, or merely produced paperwork. If a pilot works, expand it. If it fails, narrow or end it. If an agency avoids the surface, decide whether the avoidance is legitimate specialization or proof that the charter is too weak. Permanent power should not arise from the inertia of a pilot that no one wanted to revisit.

This final paper leaves the central tension alive. The state needs secrecy to know danger. The republic needs law to know what secrecy has done. Sentinel 2.0 does not claim to reconcile those needs once and for all. It proposes a discipline for the common surface: recognize when intelligence has become shared national judgment, attach lawful conditions to that status, measure the exposure created, preserve the dissent and lineage behind it, record the exceptions, and force temporary powers to justify their survival. That is how secret power remains useful without being allowed to become private government.

PUBLIC SOURCE NOTE.

These papers rest on public materials rather than classified access. Public-law anchors include Title 50 authorities concerning the Director of National Intelligence and the Intelligence Community Inspector General; public ODNI materials on the members of the Intelligence Community, the National Intelligence Program and Military Intelligence Program, analytic objectivity, commercially available information, artificial-intelligence principles, civil liberties and privacy, and the National Counterterrorism Center; and the public reform record after September 11, including the Intelligence Reform and Terrorism Prevention Act of 2004.

The manuscript distinguishes three kinds of claims. Public-law claims describe authorities, offices, budget categories, standards, official missions, and oversight structures visible in public sources. Design claims propose mechanisms that would require statutory drafting, executive implementation, technical pilots, and inspector review. Risk claims identify unresolved tensions: speed against lineage, source protection against memory, privacy against integration, AI usefulness against explainability, NIP leverage against MIP escape, and operator authority against capture or evasion.

The document does not use public sources as a substitute for classified experience. It uses public sources to define the legitimate field of public argument: what law has already assigned, what public reports have already diagnosed, and what a constitutional system may demand at the points where secret agency work becomes common national judgment. Operational details would have to be tested by authorized pilots before any permanent regime could claim success.