US AI Worthiness Report

2026

NATIONAL CAPITAL REGION (NCR) EDITION

AT WORTHY

US AI Worthiness Report 2026

National Capital Region (NCR) Edition

Table Of Contents

Edition Note

The National Capital Region (NCR) Edition is a regional release drawn from the completed assessment underlying the forthcoming US AI Worthiness Report 2026. It presents an analysis of the District of Columbia, Maryland, and Virginia. Full interstate comparisons, national rankings, and the broader implementation benchmark are reserved for the national publication. This edition should be read both as a substantive regional report and as an early view of the analytical logic that will structure the national benchmark.

Preface

American leadership in artificial intelligence will be judged less by frontier capability alone than by whether public institutions can use AI lawfully, competently, and credibly. In a federal system, that test is not performed only in national laboratories or large technology firms. It is performed in procurement rules, agency practice, oversight routines, public services, and the institutions that translate policy into administration.

That is the purpose of the US AI Worthiness Report 2026. The national report evaluates how U.S. jurisdictions combine capability with responsibility across four pillars: Standards and Governance, Talent and Research, Adoption and Public Value, and Resources and Infrastructure. It extends the logic of the Global AI Worthiness Index 2025, which placed the United States first globally while also underscoring that leadership at scale must be matched by credible public stewardship.

The National Capital Region is the right place to begin. The District of Columbia, Maryland, and Virginia sit at the country's most concentrated intersection of federal authority, procurement, rulemaking, research capacity, contractor ecosystems, and digital infrastructure. The region does not simply discuss public-sector AI. It operationalizes it.

This edition therefore asks a narrower but sharper question: when national ambition meets real public administration, what does institutional quality actually look like? The answer matters not only for the region. It previews the national benchmark to come.

Khaled Koubaa
Chief Executive Officer

Executive Summary

The NCR is not three separate AI stories. It is a single governance corridor with distinct roles, shared dependencies, and uneven alignment between technical capability, oversight capacity, and public legitimacy.

This edition applies the completed analytical framework of the US AI Worthiness Report 2026 to the District of Columbia, Maryland, and Virginia. It uses the same standard as the forthcoming national report: whether public institutions can develop, govern, procure, and apply AI with an appropriate balance between capability and responsibility. This analysis presents a score-based diagnostic comparison. The point is not to stage a regional contest, it is to show how different jurisdictions contribute different pieces of what a credible public AI system requires.

The NCR edition shows that the U.S. challenge is not whether it can build or buy AI. It is whether it can organize governance, talent, deployment, and infrastructure into a coherent public system that remains lawful, reviewable, and trusted as use scales.

Key findings at a glance

- Maryland is the corridor's most governance-forward jurisdiction. It is strongest where policy direction, oversight capacity, transparency, and public-sector implementation reinforce one another.
- Virginia is the corridor's enabling backbone. It leads in talent, research depth, contractor ecosystems, and infrastructure, but its public-value and accountability signals remain thinner than its technical base.
- The District of Columbia is policy-dense, procurement-visible, and administratively influential, but materially dependent on capacity located elsewhere in the corridor.

The broader challenge is not AI readiness in the abstract. It is institutional coherence: whether governance, talent, adoption, and infrastructure can be organized into a public system that remains lawful, legible, and trusted as deployment scales.

Taken together, the three jurisdictions reveal a corridor defined as much by complementarity as by contrast. Maryland provides governance-led coherence. Virginia anchors infrastructure and technical depth. The District supplies policy reach, procurement influence, and administrative visibility.

None resolves the full challenge alone. The NCR therefore matters because it shows, in concentrated form, how American AI strength is actually assembled: across borders, across institutions, and across different kinds of public capacity.

The corridor also reveals the main weakness in the current U.S. picture. Governance scaffolding is increasingly visible. Research depth is real. Public-sector AI use is no longer hypothetical. Yet measured public outcomes, resident-facing transparency, appeals, and redress remain comparatively underbuilt. In plain terms, institutions are becoming better at authorizing and deploying AI than at proving, in a durable and public way, that those deployments deserve trust.

That diagnosis makes the NCR more than a regional case study. It makes the region an early test of the national benchmark. The forthcoming US AI Worthiness Report 2026 will show whether the pattern visible here is exceptional or widespread: where capability clusters without matching oversight, where governance is strong but infrastructure thin, and which states are building the most durable models of public AI administration.

For policymakers, this edition offers a practical lesson. Federal baselines matter, but they do not govern by themselves. Coherence depends on the quality of state and jurisdictional translation. For partners and institutions preparing for the next phase of U.S. AI policy, it offers something else as well: a preview of the benchmark that will matter nationally when interstate comparison, implementation quality, and public trust move to the center of the debate.

AI Worthiness Scorecard

AI Worthiness scores are not additive performance metrics. They are structural signals of the alignment between capability and responsibility. Imbalance between the two is the primary indicator of institutional risk.

- Responsibility = [Governance and Standards] + [Adoption and Public Value]
- Capability = [Talent and Research] + [Resources, Compute and Infrastructure]

This analysis is based on a structured multi-indicator scoring model using publicly verifiable evidence, not narrative assessment.

Jurisdiction	District of Columbia	Maryland	Virginia
[S] Governance and Standards	23.00	29.00	26.50
[T] Talent and Research	19.50	21.00	23.50
[A] Adoption and Public Value	13.50	16.00	11.50
[R] Resources - Compute and Infrastructure	6.00	12.00	13.00
Responsibility (S + A)	36.50	45.00	38.00
Capability (T + R)	25.50	33.00	36.50
Total	**62.00**	**78.00**	**74.50**

Insights

- The District shows an 11-point responsibility-over-capability spread, confirming structural dependence rather than institutional weakness.
- Maryland shows a 12-point responsibility-over-capability spread, indicating governance maturity outpacing infrastructure scale.
- Virginia's 1.5-point spread masks a deeper internal imbalance, with weak Adoption pulling down responsibility quality.

1. The NCR in One View

1.1 The central diagnosis

The fastest way to understand the NCR is not as a metropolitan area, but as a federal-state interface. Authority, procurement, research capacity, contractor ecosystems, and digital infrastructure spill across jurisdictional lines. The region works as a corridor of differentiated roles rather than as three self-contained AI systems.

The corridor concentrates governance, capability, and infrastructure, but not within the same institutional structures. Performance therefore depends on coordination rather than self-sufficiency. That structure matters because it exposes the U.S. problem in its clearest form. The country does not lack AI activity, it lacks consistent alignment between the institutions that set rules, the institutions that buy systems, the institutions that operate them, and the infrastructures on which they rely.

A quick way to read the corridor is to look at how each jurisdiction combines capability and responsibility. The table below summarizes the structure in one view.

Jurisdiction	Defining strength	Principal constraint
District of Columbia	Policy reach, procurement visibility, administrative centrality	Thin infrastructure base and dependence on capacity located elsewhere
Maryland	Governance discipline, coordination, public-sector conversion	Needs deeper resident-facing legitimacy and wider scale
Virginia	Talent depth, contractor ecosystem, cloud and data-center base	Public-value proof and accountability lag capability

Read together, these three roles produce the report's central diagnosis. The NCR is strong not because one jurisdiction dominates every pillar, but because the corridor concentrates different strengths in close proximity. That is also the source of its vulnerability. When policy reach, governance discipline, public deployment, and infrastructure depth are distributed across borders, coordination becomes a condition of performance rather than an administrative afterthought.

1.2 Why the NCR Matters Nationally

The region matters nationally because each jurisdiction occupies a structurally important place in the American AI system.

The District of Columbia matters because it is where federal policy language, procurement norms, administrative interpretation, and government use cases are most visible. The District shapes the operating assumptions under which public-sector AI is debated and acquired across the country.

Maryland matters because it demonstrates whether governance ambition can be converted into routine public administration. It highlights the machinery that often determines whether AI is governable in practice: inventories, review processes, procurement discipline, interagency coordination, and visible oversight.

Virginia matters because public AI capability increasingly depends on infrastructure and technical ecosystems that are less politically visible than laws or executive orders but just as consequential. Data centers, cloud capacity, contractors, cybersecurity assets, and technical labor markets are now part of the operating base of modern public administration. Virginia sits near the center of that base.

1.3 What This Reveals About the U.S. System

The broader national implication is straightforward. The United States does not principally face a shortage of AI capability. It faces a problem of institutional coherence. The NCR makes that problem visible because it is unusually rich in federal authority, research capacity, procurement experience, and infrastructure, yet still shows clear gaps in outcome measurement, transparency, redress, and dependency management.

This matters because the NCR is not a low-capacity environment. It represents one of the most institutionally dense and resource-rich governance systems in the country. If misalignment between capability, governance, and public legitimacy appears here, it is unlikely to be an isolated condition. It suggests that the central challenge for U.S. AI governance is not expansion of capability, but the organization of that capability into systems that remain coherent, accountable, and trusted at scale.

2. Framework and Method

2.1 What AI Worthiness Measures

In this report, AI worthiness refers to the capacity of public institutions to develop, govern, procure, and apply artificial intelligence in ways that are effective and publicly legitimate. It is not a measure of technological enthusiasm, ecosystem visibility, or infrastructure size alone. It asks whether institutions can turn access to AI capability into governable public capacity.

In the NCR, that definition has a practical consequence. The District of Columbia, Maryland, and Virginia do not function as self-sufficient systems. Authority, labor markets, procurement relationships, research networks, and infrastructure cross formal borders. High AI worthiness in this setting therefore does not mean autonomy in every asset. It means the ability to govern dependencies competently and to translate them into credible public outcomes.

2.2 The STAR framework in the NCR context

The STAR framework remains the organizing structure of the national report. In this regional edition, the four pillars are read through a public-sector lens.

Pillar	What it measures in this edition	Why it matters in the NCR
Standards and Governance	Policy direction, procurement controls, oversight, transparency, review routines, and institutional accountability.	The corridor is saturated with policy and procurement. The key question is whether governance language is translated into enforceable administrative practice.
Talent and Research	Workforce depth, research output, funding, centers of excellence, and the ability of institutions to organize expertise.	The NCR contains dense research and technical labor markets. Public governance quality depends on whether institutions can access and retain that expertise.
Adoption and Public Value	Meaningful public-sector use, measured outcomes, transparency at the point of deployment, and resident-facing legitimacy.	The corridor is no longer in a pre-deployment phase. The harder question is whether use is producing demonstrable public benefit and trusted administration.
Resources and Infrastructure	Compute access, cloud and data-center capacity, cybersecurity, contractor ecosystems, and broader enabling infrastructure.	Public capability increasingly depends on material systems and vendors that sit across jurisdictional lines. Dependence must be governed, not merely assumed.

2.3 Scope, evidence, and limits

The scope of this edition is limited to the District of Columbia, Maryland, and Virginia, treated as three distinct public jurisdictions within one densely interconnected corridor. Federal institutions are incorporated as context and baseline because they shape procurement, standards, civil-rights obligations, security requirements, and administrative practice across the region. They are not scored as a fourth jurisdiction.

The edition uses a public evidence standard. Scores and narrative judgments rely on official, documentable, and auditable sources wherever possible: statutes, executive orders, agency guidance, procurement records, budgets, legislative materials, audit reports, inventories, workforce data, university records, and infrastructure registries. Observable institutional practice carries more weight than rhetorical commitment.

The report should be read as a comparative diagnostic, not a causal model. Scores do not claim that any single law, vendor, or infrastructure asset produced a given outcome. They indicate balance, asymmetry, and institutional direction. Because this edition covers only three jurisdictions, small differences are less important than the structural logic beneath them. The recommended reading is role first, structure second, score third.

3. Comparative Findings

3.1 Capability vs Responsibility

The most revealing way to read the corridor is through the relationship between capability and responsibility. In this framework, capability combines Talent and Research with Resources and Infrastructure. Responsibility combines Standards and Governance with Adoption and Public Value. That split shows not only what each jurisdiction possesses, but how its public capacity is organized.

Maryland is the corridor's clearest responsibility-led case. Its responsibility composite reaches 45 points against a capability composite of 33. The significance of that spread is not that Maryland lacks technical capacity. It is that policy direction, oversight routines, and public-sector conversion are more mature than one would predict from infrastructure scale alone. Maryland's comparative advantage lies in turning governance intent into administrative discipline.

Virginia represents the capability-rich pole of the corridor. Its capability composite reaches 36.5 and is anchored by talent, research depth, contractors, and infrastructure. Its responsibility composite stands at 38, close to parity in aggregate, but the internal pattern matters more than the total. Virginia's strongest assets remain on the capability side, while its thinner signals appear where public-value proof, transparency, and resident-facing legitimacy become most demanding.

The District of Columbia is responsibility-centered but materially constrained. Its responsibility composite reaches 36.5 against a capability composite of 25.5. The District is not a full-stack AI ecosystem. It is a policy-dense administrative node whose significance comes from procurement visibility, federal proximity, and the concentration of government use cases. Its profile shows how far public-sector influence can extend beyond locally controlled infrastructure, and where that model begins to meet limits.

Taken together, these profiles reveal a corridor built on complementarity rather than duplication. Maryland contributes governance coherence. Virginia provides the enabling technical engine. The District supplies policy reach and procurement centrality. The challenge for the corridor is whether these differentiated strengths can be organized into a public system that is not only powerful, but governable.

3.2 Strengths and Constraints

Despite their different roles, the three jurisdictions share a common structure. The strongest signals across the corridor cluster around governance scaffolding, research depth, and the existence of real public-sector AI activity. The weakest signals cluster around measured public outcomes, resident-facing legitimacy, and deeper forms of strategic technical autonomy.

The first shared strength is governance architecture. All three jurisdictions operate above a pre-policy baseline. Strategy documents, administrative guidance, oversight vocabulary, procurement awareness, and responsible-AI framing are present across the corridor. The NCR is not beginning from institutional silence. It is working from a comparatively mature governance floor.

The second shared strength is knowledge density. Universities, federal research institutions, contractors, laboratories, and technically skilled labor give the corridor a deep intellectual base. That matters because AI governance fails quickly when public institutions cannot access or organize expertise close to the point of decision-making.

The third shared strength is that public-sector AI activity is already real. The corridor has moved beyond abstract strategy and into implementation. That makes its weaknesses more analytically important, not less, because the region cannot explain away gaps in trust or transparency as the natural condition of early experimentation.

The most important shared constraint is weak public-value proof. Institutions can point to activity more easily than they can show durable, measured outcomes that residents can evaluate. Related gaps appear in transparency, disclosure, appeals, and redress. In practical terms, the corridor is better at building the machinery of AI governance than at demonstrating, in public, that the machinery is producing accountable benefit.

A second shared constraint is dependence. The NCR is nationally visible as a technology and infrastructure region, but much of the material base on which public institutions rely remains concentrated in external vendors, cloud relationships, and infrastructure assets that government does not fully control. Strategic importance should not be mistaken for strategic self-sufficiency.

3.3 Federal Baseline and Local Translation

One of the clearest lessons from the NCR is that U.S. AI governance begins with a federal baseline, but does not end there. Federal institutions establish much of the language, process discipline, procurement expectation, and risk logic through which public-sector AI is now understood. That baseline matters. It raises the floor of institutional awareness and helps explain why governance signals are comparatively strong across the corridor.

But federal baselines do not implement themselves. They have to be translated into agency routines, budget priorities, procurement controls, training, monitoring, and corrective capacity inside state and district institutions. That translation differs sharply across the NCR. Maryland appears strongest where federal expectations must be turned into recurring administrative practice. Virginia absorbs federal pressures through a strong enabling ecosystem, but the path from technical capacity to resident-facing accountability is less mature. The District is closest to federal rulemaking and procurement logic, yet also the most dependent on infrastructure and capabilities located elsewhere.

That gap between baseline and translation is where institutional friction appears. Procurement frameworks can be relatively sophisticated at the point of acquisition and thinner at the point of monitoring. Oversight responsibilities can be distributed across multiple offices without clear ownership of consequential outcomes. Technical systems can be shaped by federal standards, deployed by state or district agencies, and supplied by vendors whose capabilities outrun the institutions meant to supervise them.

A second source of friction is bandwidth. The corridor is not short on policy ambition. It is one of the most policy-saturated AI environments in the country. But once inventories, impact assessments, review requirements, and public reporting become part of the operating baseline, agencies must sustain them through staff time, documentation discipline, procurement expertise, legal interpretation, and corrective procedures. In high-capacity systems, the practical constraint often becomes less the absence of ideas than the difficulty of keeping administrative practice aligned with rising expectations.

A third friction concerns dependency. Government can set many of the rules while still lacking full control over the operational substrate through which those rules must be realized. In the NCR, where cloud dependence, contractor ecosystems, and infrastructure concentration are central to public capability, the mismatch between where capacity sits and where accountability lands becomes especially visible.

4. Profiles of the Three Jurisdictions

4.1 District of Columbia

The District of Columbia is the corridor's clearest example of administrative reach without full-stack self-sufficiency. Its importance lies not in territorial scale or infrastructure depth, but in policy density, procurement visibility, and the concentration of government use cases. The District's overall score is 62.0, composed of 23.0 in Standards and Governance, 19.5 in Talent and Research, 13.5 in Adoption and Public Value, and 6.0 in Resources and Infrastructure. That composition is analytically decisive. The District is strongest where policy, administration, and procurement matter most. It is weakest where material autonomy becomes decisive.

Its governance profile is substantial. The District benefits from immediate exposure to federal administrative norms, procurement expectations, civil-rights framing, and risk-management practices. As a result, AI is treated less as a speculative innovation topic than as an object of administration. Oversight signals, documentation expectations, and public-sector coordination are more developed than the District's physical scale would otherwise suggest.

The District also sits inside a knowledge-rich environment. Universities, think tanks, federal research presence, policy schools, professional networks, and the circulation of highly educated labor through government and contracting give it credible depth in Talent and Research. The District can attract and organize expertise even if it does not control the region's deepest labor market or infrastructure base.

Its most interesting position appears in Adoption and Public Value. The District is not only a place where AI is discussed. It is a place where AI is used in meaningful administrative contexts. That gives it visibility and practical importance. Yet the downstream legitimacy mechanisms remain comparatively thin. The District is further along in authorizing and running AI-enabled public systems than in showing, through measured outcomes and resident-facing safeguards, that those systems are consistently producing accountable public benefit.

The sharpest constraint is infrastructural. With only 6.0 points in Resources and Infrastructure, the District cannot be understood as a self-sufficient AI jurisdiction. Its access to technical systems often depends on vendors, data centers, compute environments, and broader regional assets that sit outside its direct control. The District governs through a technical environment it does not fully own.

That specialization is both strength and risk. It allows the District to shape procurement, policy interpretation, and administrative practice well beyond its borders. But it also creates a familiar public-sector vulnerability: policy reach can exceed operational control. The District's long-term credibility therefore depends not only on procedural sophistication, but on whether the systems it helps authorize remain intelligible, contestable, and materially supportable over time.

4.2 Maryland

Maryland is the corridor's clearest governance-forward case. Its overall score of 78.0 is built from 29.0 in Standards and Governance, 21.0 in Talent and Research, 16.0 in Adoption and Public Value, and 12.0 in Resources and Infrastructure. The structure of that profile matters more than the total. Maryland does not rely on the corridor's largest infrastructure base or deepest contractor ecosystem to establish its position. Its advantage lies in how consistently it converts policy direction into administrative order.

That strength is most visible in Standards and Governance. Maryland appears comparatively strong at assembling the machinery that makes AI governable in practice: policy direction, oversight routines, intergovernmental alignment, transparency mechanisms, and multistakeholder inclusion. The key point is not simply that rules exist. It is that governance is embedded in a more operational form of public management.

Maryland also leads the corridor in Adoption and Public Value. This is important because it shows that governance strength is not stopping at framework design. The state has moved beyond abstract commitments and into meaningful public-sector deployment. It is not merely building guardrails around hypothetical systems. It is governing real uses. Even so, its weaker signals on trust, redress, and inclusion show that the work of public legitimacy is not complete.

Its Talent and Research profile gives that governance story a credible technical base. Maryland combines research institutions, funding capacity, and centers of excellence strong enough to support administrative seriousness without claiming the corridor's deepest labor-market position. The state's strength here is quality, institutional depth, and proximity between knowledge assets and public-sector governance.

Resources and Infrastructure are more modest. Maryland has enough material base to support institutional deployment and public use, but it is not the corridor's primary infrastructure anchor. That should not be read as a defect. It clarifies the character of Maryland's model: institutional coherence under conditions of partial dependence, rather than infrastructural dominance.

This gives Maryland a distinct role in the corridor. Where the District is strongest in policy reach and Virginia in enabling infrastructure, Maryland is strongest in public-sector conversion. It takes governance principles and turns them into recurring practice. Its next stage of maturity lies less in more policy language than in deeper resident-facing legitimacy: clearer disclosure, stronger contestability, more visible protections, and tighter proof that deployment is delivering measurable public value.

4.3 Virginia

Virginia is the corridor's clearest capability-rich jurisdiction. Its overall score of 74.5 is built from 26.5 in Standards and Governance, 23.5 in Talent and Research, 11.5 in Adoption and Public Value, and 13.0 in Resources and Infrastructure. The composition of that score explains Virginia's national significance. It is the jurisdiction most closely associated with the technical backbone on which the wider federal-state AI system increasingly relies.

Its strongest signal appears in Talent and Research. Virginia does not merely sit near expertise. It generates, attracts, and sustains it. Workforce depth, research intensity, public research funding, and the circulation of highly skilled labor through universities, firms, contractors, and public institutions make Virginia the corridor's most capability-rich environment.

That strength is reinforced by Resources and Infrastructure. Virginia is the region's infrastructure anchor, especially in data-center and cloud capacity, and it benefits from a dense contractor ecosystem that links federal demand, enterprise services, and public-sector digital operations. Much of its importance lies in what it enables across the wider system. Virginia functions less as a conventional state case than as a platform jurisdiction.

Its governance profile is credible, but not as developed as its technical base. Strategy, oversight signals, and interstate relevance are present. The more important issue is the weaker performance on transparency, impact assessment, and trust frameworks. Public AI governance has not kept pace with the scale and strategic significance of the ecosystem the state hosts.

That gap is clearest in Adoption and Public Value, Virginia's weakest pillar. Meaningful public deployment is underway, but measured outcomes, transparency, and resident-facing safeguards remain thin relative to the state's capability base. The state has the power to support large systems. It is less advanced in showing how those systems translate into accountable public benefit that residents can understand and contest.

This is why Virginia matters analytically. It demonstrates a central proposition of the AI Worthiness framework: technical power and public worthiness are not identical. A jurisdiction can possess exceptional workforce depth, research strength, contractor capacity, and infrastructure concentration, yet still fall short of full maturity if public-value proof and accountability do not rise alongside them.

Virginia's next stage of progress is therefore not simply more infrastructure. It is stronger public discipline around the capability it already hosts: better disclosure, tighter impact assessment, clearer resident protections, more visible outcomes, and more rigorous oversight of the vendor and cloud relationships beneath public deployment. The task is to convert platform strength into public legitimacy.

5. NCR Insights for U.S. AI Worthiness

5.1 The U.S. problem is coherence, not readiness

The NCR points to a national conclusion that is both simple and consequential: the United States should not primarily be understood as a country short on AI capability. It should be understood as a country in which capability, governance maturity, infrastructure depth, and public legitimacy are distributed unevenly and aligned unevenly.

That distinction matters because it changes the policy question. The challenge is not whether the country can build or acquire advanced AI systems. It plainly can. The challenge is whether public institutions can organize governance, talent, adoption, and infrastructure into a system that scales without losing legality, reviewability, and trust.

The NCR makes this visible in concentrated form. Maryland shows what governance-led coherence looks like. Virginia shows what capability-rich scale looks like when public-value signals lag. The District shows how policy reach and administrative visibility can coexist with material dependence. None of these is the national story by itself. Together, they clarify its structure.

5.2 Common failure modes in high-density governance environments

High-density governance environments often look institutionally advantaged. They concentrate agencies, procurement expertise, legal capacity, contractors, researchers, and digital infrastructure within a small administrative space. The NCR does all of that. Yet density changes the form of risk rather than eliminating it.

The first failure mode is mistaking policy density for governance maturity. Institutions can adopt the language of responsible AI, publish principles, and stand up guidance while still lacking strong inventories, review routines, documentation discipline, and corrective procedures.

The second is procurement sophistication without lifecycle accountability. Public institutions may become better at acquiring AI systems than at supervising them after deployment. Auditability, update control, incident escalation, and exit options often lag behind front-end review.

The third is capability outrunning legitimacy. Research strength, infrastructure, and technical labor can expand quickly while transparency, human review, public explanation, and redress move more slowly. The result is a widening gap between what institutions can do and what they can still justify persuasively to the public.

The fourth is dependence without control. Public institutions increasingly rely on vendors, cloud environments, and contractor expertise they do not fully control. Without strong contract governance, documentation rights, continuity planning, and internal review capacity, public authority becomes operationally thinner than it appears.

The fifth is fragmented authority masked by institutional abundance. Many offices may touch the same system at different stages without any one of them owning the full chain of responsibility. Formal coverage exists, but practical accountability diffuses.

5.3 Why the national report matters now

This NCR edition establishes the structure of the problem, but not its national distribution. The national report will not simply extend this analysis. It will determine whether the pattern observed in the corridor is exceptional or systemic, and will identify which jurisdictions are structurally aligned, which are capability-heavy but governance-thin, and which are governance-forward but capacity-constrained.

What the NCR reveals is a system composed of strong but unevenly coupled elements. Governance maturity, technical capacity, infrastructure depth, and public-sector adoption are all present, but not consistently aligned within the same institutional boundaries. The national question is therefore not whether these elements exist. It is how they are distributed, how they interact, and whether they are being organized into a coherent public system at scale.

The national report will do what a regional edition cannot. First, it will provide the first full interstate benchmark of AI worthiness across the United States, allowing policymakers to compare institutional structures rather than rely on anecdote or reputational proxies. Second, it will identify recurring structural patterns across jurisdictions: where capability is advancing faster than oversight, where governance frameworks are strong but technical depth remains limited, and where adoption is active but public legitimacy is underdeveloped. Third, it will establish a common reference point for decision-making across government, research institutions, industry, and civil society, grounded in comparable evidence rather than fragmented interpretation.

The NCR shows that the United States has the ingredients of serious public-sector AI leadership, but not yet their full alignment. The national report will determine how widely this condition extends, where institutional coherence is strongest, where it is fragmenting, and where it has yet to emerge. Its value lies not in extending the present analysis, but in making it comparable, interpretable, and actionable at national scale.

What the national report makes possible:

- A full interstate benchmark across all U.S. jurisdictions, based on a common scoring framework.
- A national map of structural profiles, distinguishing governance-led systems, capability-rich but oversight-thin systems, and adoption-active but legitimacy-constrained systems.
- A shift from descriptive assessment to comparative diagnosis, enabling policymakers to identify binding constraints and prioritize reforms.
- A shared reference point for governments, institutions, and strategic partners seeking to strengthen implementation quality and public trust.

This moves the discussion from regional diagnosis to national accountability.

6. Policy Priorities

6.1 For governors, agency heads, and CIOs

For executive leaders, the central task is not to make government appear more technologically advanced. It is to make public AI systems more governable. The following priorities should be treated as operating requirements, not aspirational principles.

1. Establish clear institutional ownership for AI across policy direction, inventories, procurement review, deployment approval, post-deployment monitoring, and incident escalation.
2. Move from ethical language to operating infrastructure. Agencies should maintain current inventories, classify higher-risk uses, document legal authority, define meaningful human review, and set review schedules tied to system changes.
3. Use procurement as a governance instrument. Contracts should preserve auditability, documentation access, data provenance requirements, update notification, security obligations, performance monitoring, and termination or remediation rights.
4. Institutionalize human oversight, appeals, and redress wherever AI affects rights, benefits, opportunities, or material public burdens.
5. Invest in agency-level technical competence so staff can interrogate vendor claims, read model documentation, understand risk, and manage incidents intelligently.
6. Define success before deployment and measure public outcomes, including error patterns, trade-offs, and whether the system is advancing a clearly stated public purpose.
7. Map and manage dependency risk across cloud providers, vendors, external compute, and contractor relationships that have become load-bearing for public operations.
8. Require real cross-agency coordination through common standards, shared reporting, and escalation mechanisms for higher-risk systems.

6.2 For legislatures and oversight bodies

Legislatures and oversight institutions define the governance floor below which operational convenience cannot fall. Their role is to turn accountability from an episodic reaction into a standing condition of public AI use.

1. Create clear statutory categories for consequential public-sector AI use so stronger safeguards attach by law where systems affect rights, benefits, enforcement, health, education, employment, or public safety.
2. Require public inventories and meaningful disclosure for significant AI systems, including purpose, responsible agency, decision context, and availability of human review or appeal.
3. Anchor impact assessments in enforceable process before deployment or material modification of high-risk systems.
4. Strengthen due-process, appeal, and redress protections so affected individuals can obtain notice, contest errors, and reach a human decision-maker with authority to intervene.
5. Set minimum contractual requirements for significant AI procurements, including audit access, documentation, update notification, incident response, and termination rights.
6. Align appropriations with governance capacity by requiring agencies to show how inventories, legal review, technical expertise, contract management, and auditing will be sustained after deployment.
7. Establish routine reporting and hearing structures before crisis, covering inventories, procurement exposure, performance outcomes, appeals activity, audit findings, and incidents.
8. Equip inspectors general, auditors, and equivalent bodies with explicit authority to examine AI-enabled systems and the records needed to evaluate them.
9. Build legislative competence so lawmakers and oversight staff can interrogate agencies and vendors effectively rather than defer to technical jargon.

6.3 For federal coordination and public-private partnerships

The NCR demonstrates that AI capability now sits across jurisdictions, vendors, and infrastructure providers. The next question is whether federal coordination and public-private partnerships make that capability more governable or quietly more dependent.

1. Build interoperable governance baselines that travel across jurisdictions, including common templates for inventories, impact assessments, documentation, incident reporting, procurement clauses, and human-review requirements.
2. Treat federal coordination as administrative infrastructure, with standing channels through which agencies and jurisdictions can share lessons, compare controls, surface failures, and align on minimum expectations.
3. Align grants and technical assistance with governance maturity, not only with deployment ambition.
4. Use federal procurement as leverage to improve downstream public control over vendors, platforms, and systems integrators.
5. Design public-private partnerships around capability transfer as well as service delivery, so institutions leave engagements more knowledgeable and more governable than when they entered them.
6. Set shared resilience and continuity expectations for systems dependent on private infrastructure, including portability, documentation preservation, concentration risk, and contingency planning.
7. Evaluate partnerships against public outcomes rather than novelty alone, with clear metrics for benefit, harm, rights impact, and public accountability.

Conclusion

This NCR Edition argues that the central question in U.S. AI governance is no longer whether public institutions will engage with artificial intelligence. They already are. The more consequential question is whether they can do so through arrangements that are competent, governable, and worthy of public trust.

The District of Columbia, Maryland, and Virginia do not tell one story. They tell three structurally related ones. The District shows how policy centrality, procurement influence, and administrative reach can produce strategic importance without a deep independent infrastructure base. Maryland shows how governance discipline and statewide coordination can convert policy intent into operating practice. Virginia shows how infrastructure scale, contractor strength, and technical workforce depth can make a jurisdiction indispensable to the wider U.S. system while still leaving unresolved questions about transparency, public value, and resident-facing legitimacy.

Taken together, they expose the current American asymmetry: capability is accumulating faster than coherence. Policy baselines are stronger. Procurement is more consequential. Deployment is more common. Infrastructure is more strategic. Yet public institutions are not advancing at the same pace in transparency, outcome measurement, contestability, and trust.

That is why AI worthiness matters as a public standard. It asks not only whether institutions possess resources, talent, or ambition, but whether those assets are being organized through rules, oversight, implementation discipline, and public-value orientation strong enough to sustain trust over time.

The NCR suggests that a stronger model of public AI leadership is possible, but not automatic. The ingredients exist. The alignment is still incomplete. The forthcoming national report will answer the next question at interstate scale: which jurisdictions are building that alignment most credibly, which are falling behind, and what a nationally coherent model of public AI governance will actually require.

AT WORTHY TECHNOLOGY

www.atworthy.com

info@atworthy.com

@atworthycom