

STRATEGIC
FUTURES
INSTITUTE

Governing in an Era of Technology

Shaping Strategic Futures for Society

WHITE PAPER NO. 1

Strategic Futures Institute · November 2026

Executive Summary

> **Thesis.** Technology does not determine how societies are governed. The choices governments make about technology — what to build, what to buy, what to permit, and what to refuse — determine it. Governments that make those choices early, transparently, and with public trust as the explicit objective will govern well in the digital age. Governments that drift into technological change will be governed by it.

Every era of rapid technological change has rewritten the relationship between the state and its citizens. The printing press decentralized knowledge and made public opinion a political force. The telegraph centralized control and turned colonial administration into imperial command. Early computing made mass administration possible while teaching a durable lesson: governments adopt technology first to cut internal costs, and only later — if ever — to serve citizens better. The digital era is repeating each of these patterns at greater speed and scale.

This paper surveys the technologies now reshaping governance — digital service delivery, artificial intelligence, data analytics, connected infrastructure, distributed ledgers, and advanced networks — and the challenges that accompany them: privacy and cybersecurity risk, algorithmic bias, the digital divide, misinformation, and regulatory lag. Its central finding is that these challenges share one common denominator. Each, left unaddressed, spends down the same asset: public trust. Trust is not a by-product of good digital governance. It is the precondition for it, and the correct design objective for every system a government deploys.

The paper closes with seven recommendations for the United States, each assigned to the institutions positioned to act on it:

- **Give AI accountability an inventory with consequences** — complete, public agency AI use-case inventories, audited, with named officials answerable for each high-risk system.
- **Make the NIST AI Risk Management Framework the federal baseline** for agency AI, in procurement as well as in-house development.
- **Build standing foresight capacity** — expand technology assessment for Congress and require horizon-scanning in agency strategic plans.

- **Regulate iteratively** — principles-based rules, regulatory sandboxes, and sunset clauses for fast-moving technologies, replacing the cycle of lag and overcorrection.
- **Finish the security migration** — zero-trust architecture on a deadline, funded legacy-system retirement, and supply-chain screening for critical networks.
- **Measure citizen experience, not internal savings** — publish service-performance data and hold agencies to it, correcting the public sector's oldest technology habit.
- **Treat digital literacy as civic infrastructure** — sustained funding through libraries and schools, aimed at the information environment citizens actually inhabit.

The Strategic Futures Institute publishes this paper as its founding statement of method: history before speculation, evidence before recommendation, and the interests of citizens — not the convenience of institutions — as the test of every proposal.

I. Introduction: Governance in the Digital Age

Technology now permeates nearly every facet of public life, and it has changed what citizens expect from their governments: services as seamless as the ones they use everywhere else, decisions informed by data rather than habit, and institutions that can explain themselves. Meeting those expectations requires more than adopting new tools. It requires reimagining how government functions — how it delivers services, makes decisions, and accounts for itself to the people it serves.

Technology governance is the steering of technological development — across the state, business, and civil society — so that it aligns with societal objectives. Properly understood, it is not primarily a technical discipline. It is the practice of implementing shared values and protecting human rights through principles, codes of practice, and regulation, with democratic participation and technology assessment built into the process rather than bolted on.[1] The OECD describes the same imperative at the level of the state: governments must shape technological trajectories proactively, because trajectories left to chance diverge from the public interest.[2]

One distinction matters throughout this paper. *E-government* is the delivery of public services by digital means. *E-governance* is broader: it includes citizens' ability to access public data, monitor government activity, and participate in decisions.[3] Governments routinely achieve the first and neglect the second, and then misread the resulting public skepticism as resistance to technology. It is usually resistance to being processed rather than represented.

Well-designed digital governance does three things at once. It builds trust, by making decisions about data, security, and system changes transparent and subject to oversight. It manages risk, by addressing cybersecurity threats and compliance failures before they become public failures. And it enables innovation, because clear rules about what is permitted let agencies take calculated risks instead of avoiding all risk — the usual bureaucratic equilibrium.

The sections that follow proceed in four steps. Section II draws the lessons of earlier technological eras, which are more instructive than is usually admitted. Section III surveys the technologies now reshaping governance, including a candid account of where the claims made for them outrun the evidence. Section IV examines the challenges, and identifies the single vulnerability that connects them. Section V states what should be done, by whom.

II. Lessons from History

Technological revolutions in governance are not new, and neither are the mistakes governments make during them. Four episodes repay study.

The printing press: information escapes control

The printing press was the first mass-distribution mechanism for ideas, and its political consequences followed from a single fact: it made suppression impractical. Before print, censorship could be nearly total — burn the manuscripts, silence the author. After print, banning a book reliably increased demand for it. Print did not cause the Renaissance, but it accelerated the recovery and spread of classical knowledge, cut the price of a book from the cost of a house toward a month's salary, and gave scientists something they had never had: identical copies, so that effort went into new findings rather than into correcting copyists' errors.[4]

Politically, print created the precondition for public opinion. The works of Locke, Voltaire, and Rousseau reached a broad literate public, and argument from reason began to compete with argument from authority. Thomas Paine's *Common Sense* sold an estimated 100,000 copies within months of publication in 1776 — in a colonial population of roughly 2.5 million, a reach no American publication has matched proportionally since. The lesson is not that information technology liberates. It is that once the cost of distributing ideas collapses, control of the information environment is no longer available as an instrument of rule, and governments that continue to rely on it fail.

The telegraph: communication escapes distance

The telegraph decoupled communication from transportation for the first time: information could move faster than any messenger. Its clearest governance lesson comes from British India. Before the telegraph, correspondence between London and India took weeks or months, and distance imposed autonomy — local officials necessarily decided most questions themselves. The telegraph abolished that autonomy. During the rebellion of 1857, telegraph lines allowed the British to coordinate troop movements and responses in something close to real time, and the technology converted a loosely supervised colonial administration into something closer to direct imperial command.[5] Cable stations became geopolitical assets, controlled and taxed as such. The same technology that made real-time journalism and modern diplomacy possible also demonstrated that a communications network is an instrument of central power over every periphery connected to it.

Early computing: efficiency first, citizens later

Public administration computerized in the 1960s around mainframes processing taxation, social insurance, and other mass transactions. The gains were real: administrations absorbed growing workloads without proportional staff growth, and routine work was automated at scale. But the research record of the period is blunt about what did not happen. Studies found "disappointing" results in new or better services for the public; the focus was internal rationalization, with little attention to the range or quality of what citizens received.[6] Computerization also split the workforce's experience — knowledge workers were freed for more interesting work while data-entry staff got jobs that were more fragmented and pressurized — and it raised, early, the privacy questions that arise when personal records can be linked across databases.

This is the most directly relevant precedent for the present moment, because the pattern has never been repealed. Government technology programs are still justified, funded, and evaluated primarily on internal cost savings. The citizen's experience is measured, if at all, as an afterthought.

Radio: the public sphere, and its capture

Radio created a mass public sphere: political education, criticism of policy, and citizen participation at national scale. It also demonstrated, within a decade, every pathology of a captured medium — commercial pressure, political co-optation, and organized disinformation campaigns run through a technology the public had initially trusted.[7] The parallel to the platform era requires no elaboration.

What the record shows

Two conclusions survive all four cases. First, technology is neither centralizing nor decentralizing in itself. Print decentralized knowledge; the telegraph centralized command; computing did both in sequence. The direction is set by existing power structures and by deliberate design choices made during adoption — which means it is a legitimate and necessary object of policy, not a force of nature to which policy can only react. Second, every new information technology expands the public sphere and simultaneously degrades it, by admitting manipulation at the same scale as participation. Societies do not get the first without developing defenses against the second. The digital era is these two dynamics running concurrently, at machine speed.

III. The Present Wave

Digital government, where it has worked

The successes of digital government share a pattern: they are cases where the state redesigned a service around the citizen rather than digitizing a form. Estonia built its digital state on a secure data-exchange layer (X-Road) and a legal guarantee that citizens can see who has accessed their records; nearly all public services are available online, and the foundation is legal and institutional, not merely technical.[8] Brazil's PIX instant-payment system, launched by the central bank in 2020, brought tens of millions of previously unbanked people into the formal financial system within a few years. Ukraine's Diia application maintained digital identity and access to public services under wartime conditions — a stress test no peacetime system designer would have dared specify. The British government has estimated the potential savings from full digital transformation of its public services in the tens of billions of pounds annually, though such estimates should be read as upper bounds rather than forecasts. In each case the decisive variable was political will joined to technical competence, sustained over years.

Artificial intelligence

AI is the technology on which the governance stakes are now highest, and the one where public discussion is least disciplined. Three distinct things are commonly conflated. *Predictive systems* — forecasting demand, flagging anomalies, triaging caseloads — are mature, genuinely useful, and already widespread in administration. *Generative systems*, built on large foundation models, are transforming document work, citizen-facing communication, and software production inside government, with real productivity gains and real risks of fluent, confident error. *Agentic systems,* which chain model outputs to tools and take multi-step actions with limited supervision, are arriving in administrative workflows now and raise the accountability questions of the first two categories in a sharper form: when a system acts, someone must be answerable for the act.

Two structural facts deserve more attention than they receive. First, the capability frontier is concentrated: a small number of firms control the largest models, the specialized chips they require, and increasingly the energy supply behind both — which makes compute itself an object of governance, through export controls, procurement policy, and competition policy. Second, the public is not persuaded. Survey evidence finds majorities in both the United States (53 percent) and the United Kingdom (64 percent) expecting AI to increase unemployment, and low trust in both companies and governments to manage the technology well.[9] The Congressional Budget Office's assessment is more measured — significant effects on productivity and the federal budget over time, with wide uncertainty bands.[10] Policy made

inside that gap between institutional enthusiasm and public skepticism will be fragile, and analysts do the debate no favors by narrowing the gap artificially in either direction.[11]

Data analytics

The shift from anecdote-driven to data-driven administration is real and largely beneficial: traffic systems tuned on live data, public-health interventions targeted by evidence, disaster response coordinated in real time, tax and benefits fraud detected at a scale manual review cannot reach. The governing constraint is no longer analytical capability. It is the quality, provenance, and lawful basis of the data — and the discipline to notice when a correlation has been promoted to a cause because a dashboard displayed it confidently.

Connected infrastructure and advanced networks

Sensor networks and the Internet of Things give cities continuous measurement of traffic, air quality, energy use, and infrastructure condition; credible estimates put the potential quality-of-life gains from well-implemented smart-city systems at 10–30 percent across safety, health, and time-use indicators. Fifth-generation networks, and the 6G systems expected around 2030, extend this connectivity with far lower latency and far greater device density. The Department of Homeland Security's horizon-scanning work states the other half of the ledger plainly: hyper-connectivity expands the attack surface, deepens supply-chain exposure, and makes mission-critical services dependent on networks that must themselves be defended.[12] Connectivity is capability and vulnerability in the same package, priced together.

Distributed ledgers, without the promotion

Blockchain systems offer tamper-evident, append-only records, and they have found legitimate public-sector uses in narrow settings — land-title pilots, credential verification, audit trails. The promotional claims made for them in government have generally outrun the evidence. Most public records do not need a distributed ledger; they need a competent database, access controls, and an audit log. Proposals for blockchain-based voting in particular collide with a broad technical consensus that internet voting, with or without a ledger, cannot currently meet the security requirements of public elections. A serious governance agenda distinguishes the narrow real uses from the decade of conference presentations.

The smart paradox

Running through all of these technologies is a single tension. The mechanisms that produce efficiency — automation, large-scale data collection, automated decision-making — are the same mechanisms that produce the characteristic harms: opacity, discrimination at scale, privacy loss. This is not a balance to be struck once, by choosing a point between efficiency and ethics. It is a design condition, present in every system, that has to be managed

continuously. Governments that treat ethics review as a stage gate before deployment discover this. Governments that build it into design, procurement, and operation have a chance of staying ahead of it.

IV. The Challenges, and the One That Connects Them

Privacy and cybersecurity

Government holds the most sensitive data in society and defends it, too often, with the oldest systems. Three developments have worsened the position. Re-identification research has shown that supposedly anonymized datasets can frequently be re-linked to individuals, so a privacy technique that was adequate in 2005 is a liability now. Connected devices and advanced networks multiply entry points faster than agencies retire them. And supply chains for digital infrastructure are global and opaque, admitting compromised components at tiers no procurement officer ever sees. The countermeasures are known — security-by-design, zero-trust architecture, privacy-enhancing technologies such as differential privacy, regular resilience testing — and are unevenly applied, mostly for budgetary and organizational rather than technical reasons.[13]

Algorithmic bias and accountability

Automated systems reproduce the biases of their training data and their design choices, and then apply them with a consistency no human bureaucracy could match — in criminal justice, healthcare, hiring, and benefits administration. The technical sources are well understood: unrepresentative data, proxy variables that encode protected characteristics, objective functions that optimize the measurable at the expense of the relevant. The governance gap is different in kind: in most deployments, no identified official is answerable for a system's outcomes, affected citizens have no clear route to contest a decision, and the system itself cannot explain one. Frameworks for trustworthy AI converge on the same requirements — fairness, transparency, explainability, documented accountability[14] — and the AI use-case inventory, maintained publicly, is the basic instrument for making them real.[15]

The digital divide

Every service moved online is a service withdrawn from citizens who lack connectivity, devices, skills, or trust in the channel. The divide tracks income, age, education, disability, and rurality, and poorly designed digital services widen it independently of access: a portal that requires a confident, literate, broadband-connected user excludes by design. Digital inclusion is therefore not a social program adjacent to digital government. It is a condition of its legitimacy.

Misinformation and the information environment

Generative AI has industrialized the production of persuasive false content — text, image, audio, and video — while distribution networks deliver it faster than institutions can respond. The effect is not primarily that citizens believe specific falsehoods. It is the corrosion

of the background assumption that shared facts exist, which democratic decision-making requires.[16] The historical record of Section II applies: the answer developed for earlier media was never censorship, which fails, but a combination of literacy, trusted institutions, and platforms whose incentives are not aligned against the public sphere. None of the three is currently in good repair.

Regulatory lag

Legislative and regulatory cycles run in years; the technologies they govern now change materially in months. The predictable result alternates between two failures: long periods of under-governance in which harms accumulate unpriced, followed by reactive rules written against the last crisis. The OECD and others have named the alternative — agile, anticipatory governance, in which regulators monitor development continuously, regulate principles rather than specific implementations, and revise iteratively.[2] It is easy to recommend and hard to institutionalize, which is why Section V assigns it to specific bodies with specific instruments.

The trust deficit

These challenges are usually managed as separate portfolios, by separate offices, against separate metrics. That is a mistake, because their costs are paid from a single account. A data breach, a biased benefits decision, a service that excludes the poor, a viral falsehood unanswered — each is recorded by the public not as an isolated incident but as evidence about the same question: whether government is competent and on their side. Trust is the reserve currency of digital governance. Every initiative draws on it; well-governed initiatives replenish it; and when the reserve is spent, even beneficial programs fail for want of adoption and consent. The practical consequence is that trust cannot be a communications objective pursued after the engineering is done. It has to be a design requirement, specified like any other: transparency that citizens can actually use, accountability with names attached, recourse that actually works.

V. Recommendations

The following recommendations are addressed to the United States federal government, with state governments as partners where noted. They are ordered by leverage, not by ease.

1. Give AI accountability an inventory with consequences. Federal agencies are already directed to maintain inventories of their AI use cases. The inventories should be completed, published in usable form, and audited — and for every system that affects rights, benefits, or safety, a named senior official should be recorded as answerable for its outcomes, with documented routes for citizens to contest decisions and obtain human review. The Office of Management and Budget should report compliance annually, by agency, in public. An inventory that no one is punished for neglecting is a suggestion.

2. Make the NIST AI Risk Management Framework the federal baseline. The National Institute of Standards and Technology's AI Risk Management Framework is the strongest available common language for AI risk in the United States.[17] It should be the default standard in federal AI procurement and in-house development alike — required in solicitations, evidenced in contract deliverables, and incorporated by reference into agency policy — so that a vendor selling to government sells to one standard, and agencies inherit a tested framework instead of improvising their own.

3. Build standing foresight capacity. Congress abolished its Office of Technology Assessment in 1995, at the beginning of the internet era, and has been legislating about technology with borrowed expertise since. The Government Accountability Office's Science, Technology Assessment, and Analytics team should be expanded toward a full restoration of that capacity, and executive agencies should be required to include a technology horizon-scan — what is coming, on what timeline, with what consequences for their mission — in their strategic plans. Foresight that exists only in speeches does not exist.

4. Regulate iteratively. For fast-moving technologies, Congress and the agencies should default to principles-based statutes implemented through rules that carry sunset dates and scheduled reviews; structured regulatory sandboxes, in which new services operate under supervision with real consumers and real safeguards, on the model several states have already run in financial services; and safe harbors that expire. The aim is a regulatory system whose revision cycle is shorter than the technology's, replacing the present alternation of neglect and overcorrection.

5. Finish the security migration, on a deadline. Zero-trust architecture should move from strategy documents to funded milestones, agency by agency, reported publicly. The Technology Modernization Fund and agency budgets should prioritize retiring the legacy systems that generate a disproportionate share of federal cyber risk, and supply-chain screening for critical network components should be treated as a standing cost of hyper-connectivity rather than an emergency measure.[12] The standard is simple: no mission-critical service should depend on a system the government would be embarrassed to describe.

6. Measure citizen experience, not internal savings. The oldest habit in government technology — justifying systems by internal cost reduction — should be inverted in how programs are evaluated. Agencies should publish, for their highest-volume services, the measures citizens experience: completion rates, time required, error and abandonment rates, and satisfaction, disaggregated so that exclusion is visible. Funding reviews should weigh those measures at least as heavily as cost per transaction. What gets measured gets managed; for sixty years government has measured the wrong thing.

7. Treat digital and media literacy as civic infrastructure. The information environment cannot be governed by takedown, and the historical record says attempts to do so backfire. The durable investment is in citizens: sustained federal support, through the Institute of Museum and Library Services and state and local partners, for digital and media literacy delivered where people already are — libraries, schools, and community institutions — including the specific skill the present moment demands, which is evaluating synthetic media. This is the slowest recommendation in this paper and, on the evidence of every previous media transition, the one with the highest return.

VI. Conclusion

The future of governance in the digital era is not predetermined, and it is not primarily a technology question. The printing press did not decide that public opinion would matter; the telegraph did not decide that empires would tighten their grip; the mainframe did not decide that efficiency would outrank service. People in institutions decided those things, mostly by the accumulation of small choices about design, procurement, and oversight — which is good news, because it means the choices remain available.

The United States enters this period with real assets: a deep research establishment, the firms at the capability frontier, standards institutions of the first rank, and a public that, however skeptical, still expects its government to be capable of governing. What it lacks is not knowledge but institutional habit — the habit of anticipating rather than reacting, of designing for trust rather than auditing for it afterward, and of measuring success by what citizens experience rather than what agencies save.

Building that habit is the work this Institute exists to support. The technologies will keep arriving on their own schedule. Whether they arrive in a society equipped to govern them is a decision, and it is being made now, in the ordinary course of business, by everyone who signs a procurement, drafts a rule, or ships a system. This paper's argument reduces to one sentence: make those decisions as if public trust were the asset at stake — because it is.

About the Strategic Futures Institute

The Strategic Futures Institute is a nonprofit research and education organization based in Spanish Fort, Alabama. Its mission is to educate and inform policymakers and decision makers on governing and leading in a world of rapidly advancing technology, through nonpartisan research, analysis, and public education.

The Institute is a nonpartisan organization. It does not support or oppose candidates for public office, and it does not engage in lobbying. Its federal tax-exempt status under Section 501(c)(3) of the Internal Revenue Code is pending before the Internal Revenue Service.

White Paper No. 1 · November 2026 · © 2026 Strategic Futures Institute. This paper may be reproduced for educational purposes with attribution.

Sources

1. Technology governance. Wikipedia, drawing on the literature of technology governance and the OECD framework. en.wikipedia.org/wiki/Technology_governance

2. OECD. Technology governance; and An Overview of National AI Strategies and Policies. oecd.org/en/topics/technology-governance.html

3. Global Development Research Center. Defining E-Governance. gdrc.org/u-gov/egov-01.html

4. History.com. 7 Ways the Printing Press Changed the World. history.com/articles/printing-press-renaissance

5. Introduction to Communication and Media Studies, ch. 8.2, Impacts of the Telegraph. rotel.pressbooks.pub/comm-media-studies

6. The Computerization of Public Administration (supplement, 1988). aei.pitt.edu/94300; and Oxford Academic, The Public Sector in the Digital Age, chapter on the initial impact of computing.

7. The role of radio as the public sphere for public political education in the digital era. Cogent Social Sciences (Taylor & Francis), 2023. tandfonline.com/doi/full/10.1080/23311886.2023.2239627

8. Finally 100% Digital: Estonia's 30-Year Journey from the USSR to e-Estonia. ComplexDiscovery. complexdiscovery.com

9. Brookings Institution. What the public thinks about AI and the implications for governance. brookings.edu

10. Congressional Budget Office. Artificial Intelligence and Its Potential Effects on the Economy and the Federal Budget. cbo.gov/publication/61147

11. Brookings Institution. The coming AI backlash will shape future regulation. brookings.edu

12. U.S. Department of Homeland Security. 5G/6G Horizon Scanning Report (2022); and Science & Technology Directorate, 5G/6G. dhs.gov

13. Federation of American Scientists. Increasing Responsible Data Sharing Capacity throughout Government. fas.org; and CyberScoop, Privacy-boosting tech could prevent breaches, data misuse with government aid.

14. IBM. What Is Algorithmic Bias? ibm.com/think/topics/algorithmic-bias; and The Greenlining Institute, Algorithmic Bias Explained (2021).

15. National Telecommunications and Information Administration. Artificial Intelligence Accountability Policy Report. ntia.gov; and TechPolicy Press, AI Accountability Starts with Government Transparency.

16. Carnegie Endowment for International Peace. Can Democracy Survive the Disruptive Power of AI? (2024). carnegieendowment.org; and New America, Why Americans Crave Fake News.

17. National Institute of Standards and Technology. AI Risk Management Framework. nist.gov/itl/ai-risk-management-framework

18. United Nations. The Role of e-Governance in Bridging the Digital Divide. UN Chronicle. un.org; and UN DESA, E-Government Survey, Chapter 5: The Future of Digital Government.

19. University of Bristol, Smart Internet Lab. From 5G to 6G Governance. bristol.ac.uk

20. SSTI. TBED101: Distinguishing critical and emerging technologies in policy. ssti.org