



Regulatory Sandbox Playbook

Introduction

No country globally has solved the conundrum of how to regulate Artificial Intelligence (AI) effectively, especially in high-consequence sectors like healthcare, given the pace at which new models are being developed and deployed. Health AI regulation is, at its core, an exercise in balancing two goals that are both legitimate and frequently in tension: the goal of getting beneficial technology to patients and healthcare providers quickly, and the goal of ensuring that technology is safe and effective before and when it reaches them. AI use cases in health that promise to improve administrative processes and reduce inefficiencies, solve access to care issues, address staffing challenges, improve the quality of health care, transcend and challenge all aspects of traditional regulatory structures and approaches.

Globally, regulators are recognizing that prior regulatory pathways need evolving, and most are in a period of taking stock of what needs to be done. In the US, in their words, “the FDA’s traditional paradigm of medical device regulation was not designed for adaptive artificial intelligence and machine learning technologies”. In June 2026, the European Parliament overwhelmingly voted to officially delay the EU AI Act compliance deadlines to Dec 2 2027 and beyond. The MHRA in the UK established the National Commission into the Regulation of AI in Healthcare as an expert, non-statutory advisory body to review current regulations, reporting recommendations to the MHRA in September 2026.



Louis Brandeis, US Supreme Court Justice, 1916 - 1939

Louis Brandeis, US Supreme Court Justice, famously wrote in 1932 that “it is one of the happy incidents of the federal system that a single courageous State may, if its citizens choose, serve as a laboratory; and try novel social and economic experiments without risk to the rest of the country”.

US states have been especially active in AI oversight, serving in some cases as laboratories for regulatory experimentation. But the federalist system creates a tension by vesting authority to regulate AI as a medical device with federal regulators like the FDA, while regulation of the practice of medicine, which AI arguably does, remains reserved to the states. Although some states have acted to slow uses which are viewed as potential avenues of patient harm (for example, Illinois HB 1806 effectively bans AI from providing direct mental health and psychotherapy services), other states are taking a different approach which represents evidence informed policy making. Regulatory Sandboxes, defined by the OECD as “a controlled environment, established to understand the opportunities and risks associated with specific innovations and to develop an appropriate regulatory environment to accommodate them effectively”, have emerged as a key mechanism: structured, time-limited environments to test regulatory approaches, develop or deploy novel technologies, generate evidence, and iteratively refine oversight as the technology evolves.

We want to draw attention that exit pathways remain the least developed dimension of the sandbox landscape reviewed here. Most sandbox programs offer no clear route from sandbox participation to permanent use or deployment as the underlying regulatory framework from which mitigation was needed remains in place, leaving participants to exit into much the same regulatory uncertainty they entered from. We return to this in more detail in Section 3.

Audience

Scope and Audience

The Coalition for Health AI (CHAI) is a non-profit membership organization that serves as connective tissue between healthcare, technology, government, academia, and patients to advance responsible AI in health through frameworks, best practice guidance, and tooling. Our members span the breadth of the health AI ecosystem and are all navigating an increasingly fragmented and fast-moving state regulatory landscape. With the federal regulatory vacuum prompting states to enact their own AI legislation and offer sandbox programs, the prospect of a 50-state patchwork creates real friction for organizations trying to deploy responsible AI at scale.

CHAI, in our role as trusted convener, wants to take stock of what healthcare related sandbox frameworks exist across federal and state contexts. This playbook centers on the US state and federal landscape, drawing on a small set of international examples where they illuminate graduation and exit design specifically. We have also primarily focused on AI Regulatory Sandboxes relevant to healthcare, i.e. general AI Regulatory Sandboxes, or health-specific. It offers a light-touch framework of design choices, and taxonomies existing and proposed approaches. Subsequent playbooks will document more in-depth qualitative and quantitative considerations for policymakers grappling with how to create a structured space for health AI to be tested and eventually deployed safely.

Acknowledgments

This playbook was developed by CHAI's Policy Subgroup on Regulatory Sandboxes, a working group of clinicians, academics, lawyers, former federal agency directors, and health AI leaders convened to examine the emerging landscape of state regulatory sandboxes for health AI. We are grateful to the following individuals for their expertise, review, and substantive contributions to this work:

David Blumenthal, Mandy Cohen, Carmel Shachar, Randi Seigel, Neena Haider, Quinn Shean, Katherine Mulready, Frank Meyers, and Eric Fish.

The views expressed in this playbook reflect the collective input of the working group and do not necessarily represent the official positions of the individual contributors or their affiliated organizations.





Lead Author

Lucy Orr-Ewing

Director of Strategy and Policy for CHAI

lucy@chai.org

The state of the nation: American landscape



US State Activity (see Appendix for detail)

Enacted: At the time of writing (June 2026), seven states have enacted legislation creating sandboxes or commission to consider establishing a sandbox open to healthcare-related AI applications. Several state bills (Utah SB 149, Texas HB 149, Kansas HB 2291, Connecticut SB 5) set out sector-agnostic sandboxes broad enough to cover healthcare AI. Wyoming SF0156 enacted a sandbox in 2019 that is healthcare-specific and covering digital innovations more broadly (the Medical Digital Innovation Sandbox Act), though there is no record of any participant in that sandbox to date. Arizona HB 2434’s sandbox began as a fintech vehicle in 2018 and later in 2022 the state passed an expansion bill, HB 2731, to include all regulated applications, including healthcare AI. Delaware HJR 7 (enacted in 2025) required the state’s AI Commission work with the Secretary of State to create a regulatory sandbox framework for the testing of agentic AI technologies, with a written report due to the legislature, Governor, and other stakeholders by January 2, 2026.

Died: Four states (Idaho, Iowa, New Hampshire, and Virginia) introduced bills that would have created an AI sandbox in the 2026 legislative session, but all of those bills died. Idaho and Iowa introduced bills ([HB 945](#) and [HSB 766](#), respectively) based on the Cicero Institute’s [AI Medical Services Act](#) . Both bills create a new category of licensed AI provider (Artificial Intelligence Augmented and Autonomous Service Provider), establish a Board of Autonomous Medical Practice, create tiered licensing based on AI autonomy and clinical function, include provisional “sandbox” licenses, contemplate insurance reimbursement, and incorporate extensive oversight and auditing requirements. Oklahoma HB 1916 died in committee.

In progress: Over the past two years, several bills have been introduced following the model of Kansas HB 2291 (enacted in July 2025, over a governor’s veto), including progress (Ohio HB 176, Pennsylvania SB 245, and Rhode Island HB 7383). Pennsylvania SB 939 was amended in Senate Committee in February 2026.

US Federal Activity

At the federal legislative level, SB 2750, “SANDBOX Act” was introduced by Senator Ted Cruz (TX) in September 2025 and would have created a federal sandbox program enabling regulatory relief from FDA requirements; however the bill has not progressed, and likely won’t.

At the federal regulatory level, sandbox-style activity is emerging but still nascent. The FDA announced in December 2025 their “TEMPO” (Technology-Enabled Meaningful Patient Outcomes) pilot program, linked to CMMI’s A.C.C.E.S.S Initiative, which operates as a mini-regulatory relief sandbox by allowing participating s manufacturers of unauthorized digital health devices to request enforcement discretion while collecting real-world performance data if they are participating in the ACCESS program. The FDA is currently actively evaluating submissions for the pilot, and announced their first participant on Wednesday July 27th, Dexcom’s glucose monitoring software. To date, the FDA has not addressed state sandbox programs that may allow tools that otherwise would require FDA approval to operate under such program; this is a departure from historical norms. Accordingly, states designing s, and sandbox design should anticipate the need to adapt if the FDA changes its posture.

Relatedly, the Advanced Research Projects Agency for Health (ARPA-H) - a government agency within Health and Human Services (HHS) that provides funding for high-risk, high-reward biomedical research - currently has numerous grants in process, where they follow a similar theme of allowing technology otherwise requiring some regulatory approval or license the opportunity for limited deployment subject to frequent reporting, with the goal of evidence-informed policy development. For example their, Agentic AI-Enabled CardioVascular CAre TransfORMation (ADVOCATE) program seeks to encourage the development of autonomous AI to connect to patient records, assist in scheduling appointments, provide diet and physical therapy recommendations, and write and modify prescriptions, among other services.

Definitions, goals and considerations

Sandboxes are designed to provide limited regulatory relief from enforcement of laws and regulation that are identified by participants as possible barriers to full implementation of new products or practices. Which rules require relief depends heavily on the state’s legal and regulatory framework and on the new tool’s use case. Rather than working from a fixed statutory list, states and applicants should begin by identifying the regulatory bodies with authority over the proposed activity — at the state level: often its the medical board, pharmacy board, nursing board, insurance department, health department, and attorney general, and federally, if there may be a software as a medical device implication, the FDA and engage them early to determine jointly what relief is actually required. Common categories in need of limited relief include: licensing and scope of practice, the corporate practice of medicine or other professions, supervision and delegation rules, telehealth rules, state privacy statutes, insurance and utilization review law, and consumer protection laws.

Sandboxes allow both innovators and regulators to create parameters in which novel technologies can go to market with bounds. However, differences in the statutory authorizations reveal differing objectives in furtherance of this goal. The OECD and Utah definitions treat sandboxes primarily as “policy learning tools”, being that the point is for regulators to understand a technology well enough to build appropriate rules around it. Some states (Pennsylvania, Ohio, Texas) position their sandboxes for the benefit of innovators who require waiving of regulatory parameters in order to gain limited market access to test a technology . This connotes slightly different purposes for the sandbox (one being for the regulators, one being for the companies).

Three goals are emerging from existing state programmes:

- 1

Accelerate compliant market access, i.e. creating a ‘fast-track’ for innovation to generate evidence and reach consumers, including patients, sooner given the states’ regulatory barriers to deployment
- 2

Learn in a controlled environment to inform regulatory frameworks, i.e. the output will produce evidence to understand the safety, effectiveness, benefits and risks of this technology, that informs oversight frameworks and potentially rulemaking
- 3

Shape market dynamics toward public goals, i.e. directing innovation toward problems the market has not solved on its own, such as rural access or care affordability

These goals are not mutually exclusive, but they require slightly different designs - as we are seeing play out in proposed and enacted bills - and they reflect different weightings of the core tension at the heart of sandbox policy: the balance between enabling innovation and protecting patients. A programme oriented toward goal (1) will prioritize speed and volume of participants; goal (2) demands reporting and a regulatory hypothesis to test; goal (3) requires explicit eligibility criteria tied to public benefit. States that have not articulated which archetype they are pursuing risk building a programme that is too permissive to generate learning, too burdensome to attract applicants, or too generic to move the needle on the problems they actually care about. Harmonizing key design principles, which is an outcome of this document, would assist innovators avoid patchworks but also creates commonality so that ultimate regulatory policy avoids further patchwork, avoids jurisdictional infighting, etc.

Before any organization is accepted to the program, states should establish: (a) a pre-specified regulatory hypothesis or target outcome against which they are measuring the participant, (b) a baseline against which participant performance will be measured and how they will be measured, (c) what reporting and data is required, to whom and when and (d) a defined threshold that constitutes “success” for graduation or rulemaking purposes.

Then, whichever goals a state wishes to pursue, the design will require clear answers to three core questions - with more granular design choices nested.

Admission and Eligibility

What are the challenges trying to be solved and what are the tools that solve those challenges? What evidence is needed to demonstrate those tools may be able to address those challenges? What state agencies are involved in determining eligibility?

Operation and Oversight

What are the requirements of participation? I.e. how is performance evaluated, by whom, and against what performance standard? What reporting obligations apply, and at what frequency? What non-waivable consumer protections remain are key to to protect patients - existing or new? What must be reported to the public?

Exit pathways and graduation

what happens when the sandbox program ends? I.e. How long does the sandbox run, and is that duration appropriate to the risk and evidence needs of health AI specifically? Is there a meaningful exit - a path to full market authorization, licensure, or scope-of-practice change? Does completion of one state’s sandbox carry any reciprocal weight in another state, or does every company start over?

Taxonomy and Analysis

Admission & Eligibility: what technologies can participate?

Risk-to-patient stratification

Most enacted bills do not mandate a risk framework in statute. A foundational element emerging across sandboxes, however, is a requirement, imposed by the authorizing entity, that participants provide evidence of suitable risk mitigation, which may include clinical performance standards or other vehicles for redress should harm occur. Ohio's proposed bill explicitly leaves this to be created by rule rather than specifying it up front, and several other states in their proposed bills delegate risk assessment to the administering agency rather than stratifying by use case (e.g., autonomous vs. human-in-the-loop) in the bill text itself. Arizona's bill has roots in a 2018 FinTech Sandbox bill, and is the only state in this review that manages risk through consumer-count and transaction-size limits. Wyoming stratifies risk product-by-product: sandbox participants must post a surety bond - a financial guarantee that compensates patients if the AI tool causes harm during the pilot - and the bond amount, along with monitoring intensity, scales to each product's individual "risk profile," set by the Department of Health, rather than applying uniform rules across the sandbox. Utah takes a different route to the same end, using a risk-proportionate application standard rather than a financial guarantee.

Participant requirements

This is where the bills diverge most sharply in prescriptiveness. Pennsylvania, Kansas and Ohio are the most demanding upfront, both requiring explicit criminal conviction bars, articulated consumer benefit, risk disclosure, and a failure plan. Kansas is the only state that has a mandatory denial, not agency discretion, if applicants don't comply. Texas asks for a detailed system description, a benefit assessment covering consumer, privacy, and public safety impact, and an adverse-consequence mitigation plan. Utah asks for evidence of financial resources, technical expertise, and mitigation planning, but frames it as risk-proportionate rather than a fixed checklist. Arizona relies wholly on disclosure, leaving the Attorney General full discretion, while also requiring applicants to attest to their cybersecurity precautions and how they protect patient data. Wyoming goes further than any other state, requiring background checks on every individual substantially involved in development or management, on top of its bond requirement.

Who has the authority to make what decisions? Who decides?

Three governance patterns show up: dual-track separation of operator and advisory functions (Texas); and agency-embedded administration, sitting inside Commerce in Utah and inside Health in Wyoming. Pennsylvania's proposed bill would have taken a fourth approach, housing the sandbox in the Governor's Office rather than a regulatory agency. Kansas introduces a new Regulatory Relief Division inside the Attorney General's office, with Arizona leaving all discretion to its Attorney General, with no advisory board. Utah leans instead on the Office of AI Policy to draw up an appropriate participation agreement (a negotiated contract between the state office and each sandbox entrant) to do the work a board would otherwise do in terms of setting reporting frequency, mitigation requirements, etc.

Taxonomy and Analysis

Operation and Oversight: what are the requirements of participation?

Performance evaluation

Who evaluates performance and against what performance standard is perhaps the least statutorily developed design question across current bills, and one where health AI sandboxes may have something to learn from adjacent oversight models - including the role of independent monitoring bodies (e.g., data and safety monitoring boards in clinical trials) with authority to pause or terminate testing on safety grounds.

Reporting and non-waivable protections

Texas (enacted), Ohio (proposed), and Pennsylvania (proposed) all require quarterly reporting plus immediate notification of consumer harm. Kansas sets the fastest incident-reporting clock of any state reviewed: three business days to notify of an incident causing consumer harm, with full details due within seven. Utah again leaves reporting frequency to the Office of AI Policy's participation agreement with the vendor rather than fixing it in statute, though they offer a minimum requirement that participants must report any incident causing consumer harm, privacy breach, or unauthorized data use immediately. Routine reporting cadence and content is left entirely to the Office's discretion in drafting the agreement itself. Arizona's reporting requirement is wholly discretionary, with the statute providing only that the Attorney General "may" establish periodic reporting requirements, and no statutory floor at all. Wyoming's obligations are the most detailed: five years of record-keeping post-sandbox, with unannounced inspection rights. On protections, every bill preserves a non-waivable floor regardless of sandbox status; the specific prohibitions vary but consistently reach manipulation, unlawful discrimination, and biometric misuse. Wyoming adds two protections none of the others have: a consumer protection bond and mandatory consumer disclosure before a product is made available. Kansas categorically excludes licensing requirements from the waiver power altogether, i.e. the sandbox can waive or suspend underlying rules and regulations, but never a license or licensing requirement itself.

Duration

Most state sandbox terms last 12 to 36 months, with extension provisions in most bills. Duration of the sandbox varies depending on the goals of the effort.

Taxonomy and Analysis

Exit pathways and graduation: what happens when the sandbox program ends?

This is the least developed part of the current bill landscape for health-related applications. Fintech sandboxes typically separate the sandbox from the license: a firm tests execution of an activity that already has a defined path to licensure, using the sandbox to rehearse compliance with requirements that exist. A healthcare AI sandbox testing autonomous diagnosis or triage has no such target to rehearse toward because there is no existing licensure framework for the activity itself.

None of the enacted bills reviewed specify or appear to have a clear graduation pathway to full market authorization, licensure, or scope-of-practice change; participants largely exit the sandbox into the same regulatory uncertainty they entered from. The proposed bills in Pennsylvania and Ohio create explicit licensing pathways: participants must seek state licensure to continue post-sandbox. Texas and New Hampshire offer legal protection during testing but leave participants to navigate applicable state law themselves. Utah explicitly states participation creates no property right, with findings feeding rulemaking rather than licensing. Wyoming requires formal wind-down within 60 days of sandbox conclusion unless the participant has obtained legal authority to continue. Practically speaking, these various exit pathways do not grapple with the absence of an existing state framework for semi-autonomous or autonomous clinical applications to operate.

The Cicero Institute's AI Medical Services Act, a model bill framework, introduced in Iowa and Idaho, proposes a solution to the graduation problem. Under this Model Act, the applicant has an exit from the sandbox into the market: successful completion against published, human-benchmarked safety data converts a provisional license automatically into a permanent license as an AI Augmented & Autonomous Service Provider (AAASP), with a purpose-built Board of Autonomous Medical Practice holding full licensing, auditing, and disciplinary authority over participants. The Model Bill also propose reciprocity between states with similar frameworks upon application; this is a more developed proposal for reciprocity than the language seen in other bills, like Kansas and Connecticut.

It is not surprising that exit pathways are not developed yet in the United States, as exit pathways are underdeveloped for sandboxes globally. The best example comes from Saudi Arabia's SFDA sandbox that produced a clean graduation route: an rPPG smartphone-vitals app, submitted by a Canadian developer and run jointly with the Ministry of Health's Seha Virtual Hospital, completed SFDA sandbox testing and transitioned directly into the Authority's Innovative Medical Devices pathway, including a dedicated in-country clinical trial, and reached marketing authorization.

The UK's MHRA AI Airlock cohort has produced early results but no completed graduations to full authorization; Automedica's actual exit was an acquisition, not regulatory clearance, and post-Airlock testing still found its LLM vulnerable to unsafe outputs. Singapore's AI-SaMD sandbox has yet to produce a market launch, though regulators are now proposing stronger controls and broader eligibility; a related initiative, Singapore's Global AI Assurance Sandbox, has been more useful as a stress-testing exercise, surfacing real safety gaps (premature medication recommendations, dosing errors, prompt-injection failures) across several tested tools. None of this amounts to a proven graduation model, but it suggests sandboxes may do more to surface risk than to confer market authorization.

Conclusion

In under two years, the regulatory sandbox has moved from a fintech borrowing to one of the few instruments states have to govern health AI while federal authority remains unsettled. That momentum is genuine and, notably, bipartisan: the appeal of sandboxes cuts across party lines. But the analysis above shows that a sandbox is only as useful as its design, and design varies widely across the few states that have acted it. In order to make sandboxes a more long term solution to address regulatory barriers, there needs to be an exist for companies that successfully participate and demonstrate safety, quality and care improvement.

Three conclusions follow from the landscape as it stands in July 2026:

First, exit is the hard part, and most states have not solved it. Fintech sandboxes work because they rehearse compliance with a licensure path that already exists. Autonomous diagnosis and triage have no such destination, so participants graduate into the same uncertainty they entered from. Only the Cicero Model Act attempts a purpose-built answer with its AAASP license and dedicated board. A sandbox without a defined exit is helpful for learning but is not going to give innovators the confidence to scale or drive longer term adoption and deployment of AI in healthcare. .

Second, coherence beats comprehensiveness. A program built to accelerate market access, one built to generate regulatory evidence, and one built to steer innovation toward unsolved public problems require different designs. States that have not named which of these they are pursuing risk a program too permissive to teach them anything, too burdensome to attract anyone, or too generic to move the problems they claim to care about.

Third, coordination among states should be part of the design and thinking - as we have learned from the patchwork of laws that have made scaling in digital health challenging . Every additional state that builds in isolation raises the cost of the patchwork that a sandbox regime is supposed to relieve.

Practical next steps for a state designing or amending a sandbox:

Start by naming the goal(s), what types of applicants it hopes to attract, and writing down the regulatory hypothesis(es) that are being tested, the baseline metrics or problem, and the threshold that defines success, all before the first applicant is admitted. Define key reporting metrics to measure success and failure. A sandbox that cannot say in advance what it is trying to learn will not learn it.

Design backward from the exit. Decide what a participant graduates into, whether full authorization, a provisional-to-permanent licensure path, or evidence that feeds rulemaking, and build the entrance to match that destination.

Preserve a non-waivable protection floor. Consumer disclosure, prohibitions on manipulation and unlawful discrimination, and clear incident-reporting timelines should not be waivable under the sandbox.

Work with other states to build for reciprocity from the first draft. Adopt shared design principles and portable evidence standards now, so that a company completing one state's program carries weight in the next rather than starting over.

For CHAI and the broader ecosystem, the opportunity is to supply the connective tissue the statutes lack: a common vocabulary of design choices, a shared bar for what graduation-quality evidence looks like, and a reciprocity framework states can adopt off the shelf. Harmonizing the design layer now is the most direct way to keep fifty separate experiments from hardening into fifty incompatible regimes.

Appendix

Failed

Oklahoma HB 1916 (2025)

● Failed [Link To Bill](#)

Introduced by Rep. Alonso-Sandoval, this bill establishes a nine-member AI Council (appointed by the Governor, House Speaker, and Senate President Pro Tempore) to enforce the act, run an AI Regulatory Sandbox Program for testing innovative systems. It would create a comprehensive state AI governance framework, and create a tiered AI risk-classification system (unacceptable/high/limited/minimal), while requiring risk assessments, human oversight, audits, and individual notification rights for high-risk AI in areas like health care, law enforcement, and finance. It was set to take effect November 1, 2025, but died in committee in December 2025.

New Hampshire HB 1725 (2025)

● Failed [Link To Bill](#)

This bill establishes a 7-member AI Council to advise the state and a 36-month regulatory sandbox for testing new systems (with no waiver of the core prohibitions). It also requires clear disclosure when consumers interact with AI (including in health care), bans AI developed with intent to unlawfully discriminate (disparate impact alone isn't enough). This bill died in committee in February 2026.

Iowa HSB 766 (2026)

● Failed [Link To Bill](#)

Iowa's HSB 766 creates a new licensed category - the AI-Augmented and Autonomous Service Provider (AAASP) — regulated by a new Board of Autonomous Medical Practice, with three license classes (state-only, FDA-reciprocity, and non-diagnostic support) and four autonomy modifiers (voluntary, advisory, supervised, fully autonomous). Every AAASP starts under a provisional license good for up to two years, functioning as a de facto regulatory sandbox: the board can impose geographic limits, patient volume caps, scope restrictions, and phased human-supervision requirements to allow controlled, evidence-gathering deployment, and provisional licensees get limited noneconomic damages liability as long as they're in substantial compliance with disclosure rules (unless there's gross negligence or willful misconduct). Conversion to a full license requires the board to find that published safety benchmarks have been met, though licensees can request expedited conversion by demonstrating they already exceed them.

The bill also builds in cross-jurisdictional recognition — a “sandbox reciprocity state” with a substantially similar testing environment gets streamlined reciprocal licensure, and the board must align its own benchmarks with federal (FDA Class B) standards and other states' frameworks “to the maximum extent practicable.” Notwithstanding any sandbox-stage restrictions, high-need patients (e.g., those in provider-shortage areas, with terminal or severe chronic conditions, or unable to access a human clinician in time) must still be allowed access with informed consent and physician attestation. Unlike Idaho's companion bill, Iowa doesn't use the term “sandbox” outright, but the provisional-licensure period serves the same function.

Idaho HB 945 (2026)

● Failed [Link To Bill](#)

Idaho HB 945 follows the same Model Bill template as Iowa HSB 766, one developed by Cicero Institute.

 [See Here For The Cicero Institute's AI Clinical Services Act](#)

In progress

Ohio HB 176 (2025)

● Pending [Link To Bill](#)

Ohio HB 176 creates a general-purpose “universal regulatory sandbox” (not AI-specific) letting businesses test new products or services in the state without otherwise-required licenses or regulatory compliance. It sets up a Regulatory Relief Division and an advisory committee (business reps, regulators, and legislators) to review applications, consult the relevant state agency on consumer-safety risks, and recommend approval or denial to the governor, who has final say. Approved participants get a written agreement allowing a 24-month test (extendable up to 5 times) exempt from the specific waived laws — though taxes and consumers' right to restitution can never be waived, and participants stay liable for civil harm and crimes. They must disclose to consumers that the offering is an unlicensed test, keep records, report incidents/complaints quarterly, and can be removed at any time. The bill also creates a public web portal for residents to suggest regulations to cut, reported to the legislature quarterly. This passed the house in March 2026.

Rhode Island HB 7383 (2026)

● Pending [Link To Bill](#)

This bill is also a general-purpose sandbox (not AI-specific) covering products, services, and financial/insurance offerings involving any innovation. It creates a Regulatory Relief Office (ORR) within the Commerce Corporation, led by a director appointed by the Secretary of Commerce, plus an 11-member advisory committee (6 business reps, 3 regulator reps, 1 senator, 1 house member) to review applications and recommend admission or denial.

Pennsylvania SB 245 (2025)

● Pending [Link To Bill](#)

Pennsylvania SB 245 follows the same model as Ohio HB 176 and RI H 7383 in creating a Regulatory Relief Office in the Governor's Office and an advisory committee (2 director-appointed, 3 Senate-appointed, 3 House-appointed — including one senator and one representative) to review applications from businesses wanting to test products, services, or financial/insurance offerings free from specified state laws.

Pennsylvania SB 939 (2025)

● Pending [Link To Bill](#)

This bill creates an Artificial Intelligence, Data Center and Emerging Technology Regulatory Sandbox Program, letting approved participants test innovative AI/ data-center/emerging-tech products or services in the state for 12 months (extendable up to 12 more) without otherwise-required state licenses or authorizations. Applicants must show jurisdiction, expertise, a testing plan, and consumer-protection safeguards (liability coverage, financial reserves); the office reviews applications (in consultation with relevant agencies), can deny/approve with a 60-day cure period, and must publish granted waivers publicly. Participants aren't shielded from civil or criminal liability, must give consumers clear disclosures that the product is experimental/unlicensed, must keep records and file quarterly reports, and can be terminated from the program at the office's discretion at any time. This was amended in Feb 2026.

Appendix

Enacted

Wyoming SF0156 (2019)

● Passed [Link To Bill](#)

Wyoming SF0156 created the Medical Digital Innovation Sandbox Act, letting Wyoming-based companies apply to the Department of Health for time-limited waivers (up to 24 months, extendable by 12) from licensing requirements across 19 healthcare professions (nurses, physicians, PTs, psychologists, EMS, etc.), so they can test innovative medical digital products/services on consumers before full regulatory authorization exists. Applicants pay a \$500 fee, may face background checks, and must be approved within 90 days based on the product's risk profile, consumer-protection plan, and personnel expertise. Participants must give consumers written disclosures about the product's experimental/temporary status, post a consumer-protection bond (min. \$10,000), keep 5 years of records, and remain fully subject to criminal law and civil liability throughout — the sandbox waives specific regulatory hurdles, not liability. The department can revoke authorization for violations or consumer harm, and must report annually to the legislature. It took effect January 1, 2020, and is widely cited as an early template for state-level health AI/digital-health regulatory sandboxes.

Arizona HB 2434 (2018) / HB 2731 (2022)

● Passed [Link To Bill](#)

Arizona HB 2434 established a regulatory sandbox for FinTech since 2018, but HB 2731 expanded its scope to non-financial sectors.

It is run by the Attorney General (not a subject agency), letting companies test innovative products without normal licensure. Applicants need an in-state presence, are reviewed within 90 days, and the AG has sole discretion to admit (non-appealable denial). Approved participants get a 24-month testing period, subject to caps (max 10,000 AZ consumers; transaction limits for lenders/money transmitters) and sector-specific rule carve-outs for insurance/investment products. Consumers must get bilingual disclosures that the product is an unendorsed, temporary test. Participants must keep records, report breaches/failures, and can get a 1-year extension to pursue permanent licensing. The AG can remove violators (non-appealably) and may sign reciprocity deals with other states. Funded by application fees into the state's consumer fraud fund; the program sunsets July 1, 2028.

Utah SB 149 (2024)

● Passed [Link To Bill](#)

Utah SB 149 (2024), the Artificial Intelligence Policy Act, creates an Office of AI Policy (under the Dept. of Commerce) running an AI Learning Laboratory Program which acts as a Regulatory Sandbox. Accepted applicants can sign a regulatory mitigation agreement (temporary reduced fines/cure periods/restitution terms) for a specific AI technology, capped at 12 months + one 12-month extension. Eligibility requires technical/financial capacity, a favorable risk-benefit case, and a monitoring plan; participants stay bound by unwaived laws, must self-report harms/breaches, and can be removed anytime with no property right created — state disclaims liability/endorsement.

Kansas HB 2291 (2025)

● Passed [Link To Bill](#)

Kansas HB 2291 creates a general-purpose regulatory sandbox (not AI-specific) in the Attorney General's office — modeled on Ohio's HB 176. A Regulatory Relief Division and 11-member advisory committee review applications, consult affected agencies on consumer risk, and recommend approval to the director, who has final say. Approved participants get a 24-month test (extendable once, up to 12 more months) exempt from waived rules, though taxes, licensing, and restitution rights can't be waived, and participants stay liable for crimes and civil harm. Kansas-resident consumers must get “unlicensed test” disclosures; participants report incidents/complaints and can be removed anytime. Applications stay confidential, the committee can meet in executive session, and the office runs a public suggestion portal with regular legislative reporting.

Texas HB 149 (2025)

● Passed [Link To Bill](#)

The Texas Responsible AI Governance Act (TRAIGA) creates a regulatory sandbox (Ch. 553), run by the Department of Information Resources (not the AG), letting companies test AI systems without normal licensing/registration. Approval requires DIR plus any relevant licensing agency sign-off; applications need a system description, a consumer/privacy/safety benefit assessment, a risk-mitigation plan, and proof of federal AI-law compliance. Sandbox participants get up to 36 months (extendable for good cause) with real legal protection — the AG and state agencies are barred from prosecuting or sanctioning violations of the specific waived laws/regs during testing — but core AI-protection duties (Ch. 552: disclosure, no social scoring, no discrimination, etc.) can't be waived and remain enforceable. Participants owe quarterly reports to DIR (performance, risk mitigation, consumer feedback), can be recommended for removal by the new Texas AI Council or an agency for undue risk or unwaived violations, and DIR reports annually to the legislature. The 7-member AI Council (gubernatorial/lit. gov/speaker appointees) oversees policy study and sandbox recommendations but has no binding rulemaking power. Effective January 1, 2026.

Delaware HJR 7 (2025)

● Passed [Link To Bill](#)

This is a joint resolution, not a statute, so it creates no actual sandbox yet. It directs the state's existing AI Commission (created 2024) to work with the Secretary of State to design a regulatory sandbox framework specifically for testing agentic AI (autonomous, decision-making AI systems, as distinct from generative AI), with particular attention to its use in corporate governance and potential conflicts with Delaware's General Corporation Law. The Commission and Secretary of State must deliver a written report of findings and recommendations by January 2, 2026 to the legislature's leadership, the Controller General, Legislative Services, and the Governor. No application process, waivers, participant protections, or enforcement mechanics are established — it's a study-and-recommend mandate, presumably intended to inform future sandbox legislation.

Connecticut Public Act 26-15 / SB 5

● Passed [Link To Bill](#)

This Act (effective July 1, 2027) is a planning mandate, not an operating sandbox — similar in spirit to Delaware's resolution but embedded in statute. It directs the Commissioner of Economic and Community Development, in consultation with the Banking, Administrative Services, Public Health, and Insurance Commissioners, to develop a plan for an AI regulatory sandbox that would let applicants test innovative products/services under reduced licensing and regulatory requirements. The plan's stated purpose is to make Connecticut competitive for AI development/deployment. Notably, the commissioner is also directed to reach out to other states' existing AI sandboxes to assess the feasibility of a reciprocal multistate sandbox program — an interoperability angle not present in the other bills you've shared. By January 1, 2028, the commissioner must submit legislative recommendations to the Governor and the relevant legislative committees (commerce, banking, insurance, public health) to actually implement the plan.