

THE GLOBAL COMPLIANCE **ARBITRAGE** PLAYBOOK

One evidence spine.
Three execution planes.
No imaginary passport.

Operational Edition 1.0



OPERATIONAL PLAYBOOKS

Engineering Framework & Operational Architecture

THE GLOBAL COMPLIANCE ARBITRAGE PLAYBOOK

One evidence spine. Three execution planes.
No imaginary passport.

THE THESIS. There is no globally compliant AI control stack. There can be a globally reusable evidence system that assembles local controls without rebuilding the product from scratch.

Operational Edition 1.0

Cross-border agentic AI across the European Union, the United States and China.

Evidence cut-off August 29, 2026.



Contents

Contents	2
Document control.....	4
How to use this playbook.....	4
Reading conventions.....	4
Executive brief	5
Five decisions for the next 30 days.....	6
The contrarian position	6
1. The thesis. Compliance does not travel	7
The economic equation	7
2. What changed by august 2026	8
The August 2026 change ledger.....	9
What the ledger invalidates	9
What non-compliance costs	10
3. Europe. A product regime with a moving clock	12
Europe's hidden advantage and hidden tax.....	13
What is already fully applicable	13
EU hard stops.....	14
4. United states. Conduct, sectors and courts.....	15
The US operating map	16
The fairness asymmetry	16
5. China. The service gate is operational	17
China scope and launch triage	18
The data-localization myth.....	18
China hard stops	19
6. The regulatory dictionary	20
False friends that create expensive controls.....	21
Human in the loop	21
Algorithm filing	21
Model validation	21
Data sovereignty	21
7. Architecture. One evidence spine, three execution planes	23
Architecture principles	23
The minimum viable evidence schema.....	24
The minimum viable separation.....	25
8. The policy resolver.....	26
9. Control portability and rewrite economics	28

How to price the rewrite.....	29
10. The evidence factory.....	30
11. Human oversight is not one control.....	32
The AEPD Rule of 2 and the four-eyes principle	33
12. Data, identity and memory boundaries	34
Memory is a regulated subsystem.....	34
Cross-border route decision	35
13. The agentic risk multiplier	36
From model risk to decision risk	36
Decision-risk worksheet.....	37
14. Launch gates and hard stops.....	38
Minimum launch record	38
Hard-stop catalog	39
Exception discipline	39
15. Operating model and decision rights.....	41
The regulatory change function	41
Source hierarchy	42
16. The 90-day implementation roadmap	43
Day-90 acceptance criteria	43
17. Metrics, testing and control drift	45
Test design and evidence retention.....	46
18. Commercial scenarios and no-go cases	48
Scenario logic.....	48
Board scenario table.....	49
The perspective most companies miss.....	49
Capital allocation questions.....	49
19. Ten-question deployment matrix.....	50
Scoring	50
Deployment worksheet.....	51
Exception worksheet	52
20. Primary-source register.....	53
Verification notes.....	54
21. Method, limitations and disclaimer	55
Method	55
Limitations	55
Disclaimer.....	55
Closing position	55

Document control

Field	Value
Document	The Global Compliance Arbitrage Playbook
Edition	Operational Edition 1.0
Evidence cut-off	August 29, 2026
Primary audience	Boards, general counsel, chief risk officers, product leaders, AI platform teams, and regional executives
Decision scope	Commercial AI and agentic systems deployed across the EU, the US and China
Method	Primary-source legal and regulatory review, control decomposition and RefleXio operational synthesis
Status	Strategic playbook. Not legal advice.
Publisher	RefleXio Intelligence
Series	RefleXio Intelligence, Strategic Dossier No. 06
Version	Operational Edition 1.0
Prepared by	RefleXio Intelligence editorial desk

Table 1. Document control.

How to use this playbook

This is not a catalog of rules. It is a deployment system. Read the executive thesis first, use the jurisdiction chapters to classify the problem, then move directly to the operating model, launch gates, and 90-day build plan. The appendices are designed for workshops and evidence reviews.

The text distinguishes law, official guidance, and RefleXio analysis. Dates, thresholds, and filing counts are tied to primary sources whenever those sources publish them. Qualitative portability ratings and recommended operating thresholds are explicitly labeled as RefleXio assessments. Where a rule is unsettled, the playbook says so. Regulatory uncertainty is a control input, not a rounding error.

Reading conventions

Label	Meaning
MUST	A binding legal or regulatory obligation within the stated scope.
SHOULD	An official expectation, supervisory principle or strong operating practice.
REFLEXIO GATE	A recommended hard stop or decision threshold. It is not itself law.
WATCH	A proposal, draft, litigation front or policy signal that could alter the control graph.

Table 2. Reading conventions.

Executive brief

BOARD POSITION. Approve a cross-border AI launch only when the company can identify the accountable entity, resolve the applicable control graph, prove the control state, and reverse the deployment without contaminating another region.

The first version of this playbook was directionally right and operationally incomplete. It correctly argued that Europe, the United States and China do not share a compliance language. It then assumed that a jurisdiction-neutral agent core could sit behind three governance envelopes. That is useful as a sketch. It is unsafe as an operating model. Data routes, user identity, decision logic, tool authority, model choice, and even the definition of the regulated object can differ by market. The core is not always neutral merely because the architecture diagram says so.

The upgraded thesis is sharper. Compliance itself is local. Evidence can be portable. The strategic asset is an invariant evidence spine that records what the system is, which data and models it used, what it did, what tests it passed, who approved it, and how it changed. A policy resolver then assembles regional controls and artifacts from that spine. Execution is separated where law, risk or economics demand separation. This is compliance arbitrage in the legitimate sense. It is lawful control reuse without regulatory evasion. It reduces duplicated engineering without pretending that incompatible obligations have magically become interoperable.

Three changes make this urgent. The EU AI Omnibus, Regulation (EU) 2026/1744, entered into force on July 27, 2026 and moved the main high-risk deadlines to December 2, 2027 for Annex III uses and August 2, 2028 for AI embedded in Annex I products. The United States is simultaneously pushing federal preemption, retaining sector enforcement and allowing major state duties to mature. China has continued to operationalize filings and, from July 15, 2026, added a targeted regime under the Interim Measures for the Administration of Artificial Intelligence Anthropomorphic Interaction Services (the anthropomorphic interaction measures). The result is not convergence. It is synchronized divergence.

Board question	Required answer before approval
What is the regulated object?	A mapped system, use, decision, service, and legal entity. Never just a model name.
What is reusable?	Evidence primitives, tests, lineage, event records, and change history.
What is local?	Scope decision, notices, rights, filings, data routes, identity rules, and accountable approvals.
What stops the launch?	Unknown scope, unresolved sensitive-data flow, untested human escalation, missing artifact or stale evidence.
What expires the approval?	Material model, data, tool, purpose, geography or legal change, plus specified incidents.

Table 3. Executive brief.

Source note. European Commission, 2026a, 2026b. Federal Reserve Board, 2026. Colorado General Assembly, 2026. CAC, 2026a, 2026b.

Five decisions for the next 30 days

This playbook is designed to produce decisions, not ceremonial alignment. The next 30 days should end with an operating register, named owners, and a list of launches that must pause. A policy committee without an executable control graph is simply a meeting that has learned the word agentic.

Decision	Executive owner	Evidence of completion
Adopt the evidence-spine architecture	CTO and General Counsel	Approved target architecture and funded backlog
Create one cross-border system inventory	Chief Risk Officer	Register linking systems, uses, entities, data, models, and markets
Name jurisdiction control owners	Regional CEOs	Accountable owner for every obligation and exception
Approve hard-stop conditions	Board Risk Committee	Launch policy with no override by commercial teams alone
Select two pilot journeys	Chief Product Officer	One high-impact decision flow and one public-facing agent flow

Table 4. Five decisions for the next 30 days.

The contrarian position

The strictest-wins strategy sounds prudent and often creates the wrong product. It imports restrictions without importing the legal objective, while still failing local duties that use different evidence and different beneficiaries. A universal human approval button does not satisfy every oversight rule. A European technical file does not answer a US adverse-action inquiry. A Chinese filing number does not prove a system is safe. More compliance is not the same as the right compliance, and the gap between them is where budgets die.

The right global standard is not the harshest jurisdiction. It is the highest-quality evidence system combined with the smallest justified set of local execution differences. That architecture is harder to explain in a keynote and easier to defend in a regulator meeting. This is usually a good trade.

REFLEXIO RULE. Globalize facts. Localize legal conclusions. Automate the assembly. Keep the approval revocable.

1. The thesis. Compliance does not travel

Executives often ask whether an EU-compliant AI system can be deployed in the United States or China. The question contains its own error. Systems are not compliant in the abstract. A particular provider or deployer uses a defined system for a stated purpose, with specific data, users, tools, and effects, under an identified legal entity. Change any of those facts and the answer can change.

Europe largely regulates products, systems, providers, and deployers through a horizontal risk framework layered on existing product safety and fundamental-rights law. The United States regulates conduct through sector statutes, state laws, consumer-protection authority, examination, and litigation. China regulates public-facing services through an interlocking stack of content, algorithm, cybersecurity, personal-information, and data rules, with filing, and security-assessment mechanisms for specified services. The nouns overlap. The legal objects do not.

This produces three forms of non-portability. Semantic non-portability occurs when the same label means different duties. Artifact non-portability occurs when one regulator does not accept another regulator's proof. Architectural non-portability occurs when the technical design needed in one market conflicts with another market. The third category is the expensive one. It is also the category most often discovered after the sales team has announced a launch date.

Failure mode	Typical symptom	Cost signal
Semantic	Teams say risk management but mean different processes	Repeated meetings and contradictory policies
Artifact	A technical file cannot answer an examiner or filing request	Manual evidence reconstruction
Architectural	Data, identity or tool boundaries cannot be separated	Regional rebuild or market withdrawal

Table 5. The thesis. Compliance does not travel.

The economic equation


The relevant business case is not compliance cost alone. It is the expected cost of duplicate controls, delayed revenue, operational friction, regulatory exposure, and forced rework. A control that costs more today may be cheaper if its evidence can be reused and its policy can be changed without touching the product core. Reversibility has option value. Most business cases omit it because spreadsheets are easily impressed by launch dates.

UNIQUE THESIS. The company with the best cross-border evidence compiler can move faster than the company with the thinnest compliance team, even when both use the same model.

2. What changed by august 2026

Three jurisdictions now run on three different clocks

Current law and official policy position as of 29 August 2026



A single global launch date is a fiction. The change queue is the product.

Figure 1. The three regulatory clocks. Current law and official policy positions through August 29, 2026.

The most important update is not that one region became stricter. It is that all three regions changed the timing and shape of their gates. Europe delayed key high-risk duties while activating transparency obligations. The United States intensified its policy conflict between a national preemption agenda and state-level implementation. China continued to add service-specific rules and operational evidence of filing at scale.

The EU delay should not be mistaken for repeal. Article 50 transparency duties under Regulation (EU) 2024/1689 became applicable on August 2, 2026. The AI Omnibus, Regulation (EU) 2026/1744, changed the timing for high-risk systems and reduced some administrative burden, but it did not turn Europe into a voluntary regime. Companies now have more time to build the right system, and zero excuse for building the wrong one.

The US federal picture is equally easy to misread. The December 11, 2025 executive order created a federal policy position against burdensome state fragmentation and directed the Attorney General to establish an AI Litigation Task Force (White House, 2025b, 2025c; DOJ, 2026). That did not erase state law. California's finalized automated decision-making technology (ADMT) regulations, 11 CCR §§ 7200 et seq., took effect on January 1, 2026, with ADMT compliance required by January 1, 2027. Colorado repealed and reenacted its 2024 AI law, SB 24-205, replacing it with SB 26-189, ch. 131, 2026 Colo. Sess. Laws, codified at C.R.S. § 6-1-1701 et seq., effective for key duties on January 1, 2027. Sector law, consumer protection, and antitrust remain very real.

China's 868 filed generative AI services and 530 registered applications or functions as of April 30, 2026 show a live administrative system, not a PowerPoint concept. The anthropomorphic interaction measures are also narrower than many summaries suggest. They target continuing emotional interaction that simulates natural-person characteristics and exclude ordinary customer service, question answering, work assistance, education, and research when no continuing emotional interaction is involved. Scope discipline matters on both sides of the Pacific.

Source note. European Commission, 2026a. White House, 2025b. DOJ, 2026. CPPA, 2025. Colorado General Assembly, 2026. CAC, 2026a, 2026b.

The August 2026 change ledger

Change	Status	Operational consequence
EU AI Omnibus, Regulation (EU) 2026/1744, entered into force	Binding from Jul 27, 2026	Rebaseline the roadmap. Annex III high-risk duties now apply Dec 2, 2027 and Annex I embedded AI Aug 2, 2028.
EU Article 50 transparency	Applicable from Aug 2, 2026	Mark machine-readable synthetic output where required and label specified deepfakes and AI-generated text.
US federal preemption campaign	Executive policy and litigation capability	Do not assume state duties are void. Track challenges and preserve state-level policy routing.
Federal SR 26-2	Final supervisory guidance, Apr 17, 2026	Replace SR 11-7 references in bank model governance. Tailor rigor to materiality and model exposure.
Colorado SB 26-189	Signed May 14, 2026	Rebuild Colorado mapping around covered ADMT, notices, records, data correction, and meaningful human review.
California ADMT regulations, 11 CCR §§ 7200 et seq.	Effective Jan 1, 2026	Prepare notices, access and opt-out capability for covered significant decisions by Jan 1, 2027.
The anthropomorphic interaction measures	Effective Jul 15, 2026	Run a narrow scope test for continuing emotional interaction. Do not apply blindly to ordinary work assistants.
China filing volume	868 services and 530 apps by Apr 30, 2026	Treat filing data, model identifiers and display obligations as product operations.
China large PI processor rules (CAC, 2026c)	Draft dated Aug 7, 2026	Watch only. Do not encode draft obligations as binding law.

Table 6. The August 2026 change ledger.

What the ledger invalidates

- An August 2, 2026 default date for all EU high-risk obligations is outdated.
- SR 11-7, Guidance on Model Risk Management (April 4, 2011), is no longer the current federal banking model-risk guidance. SR 26-2 supersedes it for Federal Reserve supervised institutions. Confirm the position of the OCC and FDIC separately.
- The original Colorado AI Act is no longer the final operating text. SB 26-189, ch. 131, 2026 Colo. Sess. Laws, repealed and reenacted the framework codified at C.R.S. § 6-1-1701 et seq. Enforcement is exclusively by the Attorney General as a deceptive trade practice, with no private right of action.

- China filing is not necessarily a pre-launch ten-day gate. The algorithm-recommendation rule, Article 24 of the Provisions on the Administration of Algorithmic Recommendations for Internet Information Services, requires specified providers to file within ten working days from the date services are provided, while other security-assessment and generative AI pathways can affect launch readiness.
- The anthropomorphic interaction measures are not a general chatbot law. The emotional-continuity scope test matters.

CONTROL ACTION. Create a dated legal baseline in the control graph. Never allow a policy engine to cite a superseded rule without a version flag.

What non-compliance costs

Every hard stop in this playbook has a quantified downside. The first edition left it unstated; this edition prices it. The ceilings below are the statutory maxima, not expected outcomes. They are included because a board cannot weigh a control investment against an unquantified exposure, and because the asymmetry between regimes is itself a design input. The EU prices prohibited practices above everything else, China attaches personal liability to responsible individuals, and the US state regimes price per violation, which scales with the number of affected consumers rather than with turnover.

Regime	Ceiling	Basis
EU AI Act, prohibited practices	EUR 35,000,000 or 7% of total worldwide annual turnover, whichever is higher	Regulation (EU) 2024/1689, Art. 99(3)
EU AI Act, other operator duties, including Art. 50	EUR 15,000,000 or 3%, whichever is higher	Regulation (EU) 2024/1689, Art. 99(4)
EU AI Act, incorrect or misleading information to authorities	EUR 7,500,000 or 1%, whichever is higher	Regulation (EU) 2024/1689, Art. 99(5)
EU AI Act, SMEs and start-ups	The lower of the amount or the percentage, inverting the general rule	Regulation (EU) 2024/1689, Art. 99(6)
EU AI Act, GPAI model providers	EUR 15,000,000 or 3%, whichever is higher	Regulation (EU) 2024/1689, Art. 101(1)
GDPR (General Data Protection Regulation), upper tier, including principles, rights and transfers	EUR 20,000,000 or 4%, whichever is higher	Regulation (EU) 2016/679, Art. 83(5)
GDPR, lower tier, controller and processor duties	EUR 10,000,000 or 2%, whichever is higher	Regulation (EU) 2016/679, Art. 83(4)
China PIPL (Personal Information Protection Law), serious circumstances — entity	Up to RMB 50,000,000 or up to 5% of prior-year turnover, at the regulator's election	PIPL Art. 66, second paragraph
China PIPL (Personal Information Protection Law), serious circumstances — individual liability	RMB 100,000 to 1,000,000 fine on responsible individuals, plus possible disqualification from office	PIPL Art. 66, second paragraph

Regime	Ceiling	Basis
China PIPL, ordinary violations after refusal to correct	Up to RMB 1,000,000, plus RMB 10,000 to 100,000 on responsible individuals	PIPL Art. 66, first paragraph
California CCPA	USD 2,663 per violation and USD 7,988 per intentional violation or violation involving a consumer under 16, as CPI-adjusted from January 1, 2025	Cal. Civ. Code §§ 1798.155, 1798.199.90
Colorado SB 26-189	Up to USD 20,000 per violation under the Consumer Protection Act, and up to USD 50,000 where committed against an elderly person	C.R.S. § 6-1-112

Table 7. Statutory penalty ceilings by regime.

Source note. Regulation (EU) 2024/1689, Arts. 99 and 101. Regulation (EU) 2016/679, Art. 83. PIPL, Art. 66. Cal. Civ. Code §§ 1798.155 and 1798.199.90, as CPI-adjusted by the CCPA effective January 1, 2025. C.R.S. § 6-1-112. Figures are statutory maxima and are not legal advice.

3. Europe. A product regime with a moving clock

Europe remains the most structured horizontal AI regime of the three markets, but the structure is layered rather than singular. The AI Act sits beside GDPR, product-safety rules, sector law, labor rules, and national enforcement. Agentic systems add a practical complication because a single orchestration layer can contain several AI systems, tools, and non-AI components. The inventory must separate the regulated system from the surrounding service.

For an operator, the critical EU sequence is scope, role, risk class, obligation set, evidence package, and post-market monitoring. Role is not a formality. Provider, deployer, importer, distributor, and product manufacturer obligations differ. Substantial modification can also move responsibility. The company that fine-tunes, wraps or repurposes a third-party system may inherit more than a vendor invoice.

The AI Omnibus buys time for high-risk implementation, not for inventory. The most expensive work is factual. Organizations need a defensible description of intended purpose, decision impact, data flows, human oversight, model, and tool dependencies, logs, incident thresholds, and change history. Those facts are needed across jurisdictions even when the legal conclusions differ. That is why the evidence spine should be built before the final local dossier.

EU control domain	Minimum operating evidence	Portability
Role and intended purpose	Entity map, system boundary, intended use, prohibited use, and change owner	High
Risk management	Lifecycle risk register, acceptance criteria, residual-risk approvals	High evidence, local conclusion
Data governance	Lineage, quality tests, relevance, representativeness, and access controls	High evidence, medium control
Logging	Event schema, retention, actor, and model identity, time source, and integrity controls	High
Human oversight	Competence, information, timing, authority, intervention, and rehearsal records	Medium
Transparency	User notice, synthetic-content marking and label logic	Medium
Conformity and monitoring	Technical documentation, evaluation, approval, incident, and drift records	Low artifact, high source evidence

Table 8. Europe. A product regime with a moving clock.

Source note. Regulation (EU) 2024/1689. European Commission, 2026a, 2026c. Regulation (EU) 2016/679. AEPD, 2026.

Europe's hidden advantage and hidden tax

Europe's advantage is not simply that it regulates early. Its documentation, logging, data-governance, and post-market disciplines can create high-quality evidence primitives that are useful elsewhere. A well-built European system often begins with better facts than a lightly governed US deployment. The hidden advantage is evidence quality.

The hidden tax appears when teams confuse that evidence with legal portability. An EU technical file does not clear a US adverse-action inquiry, a bank examiner's validation test, or a discovery request. Evidence quality is not evidence portability. EU human oversight focuses on effective intervention and fundamental-rights protection in covered contexts. A US bank examiner may ask whether a model is conceptually sound, independently validated, and monitored in proportion to exposure. A Chinese service regulator may focus on filing, content safety, identity, labels, user dependency or local records. Similar control families, different proof.

Europe also creates a timing trap. The delay to 2027 and 2028 can encourage a pause just as Article 50 and existing GDPR duties require active work. The correct response is sequencing. Build the inventory, evidence spine, transparency controls, and agentic data safeguards now. Defer only the components whose legal deadline actually moved. A deadline is not a product strategy, although it has launched several consulting practices.

What is already fully applicable

A playbook dated August 2026 that discusses only the deferred high-risk clock omits the two parts of the AI Act that are already fully binding. The Article 5 prohibited practices have applied since February 2, 2025, and the general-purpose AI model obligations since August 2, 2025. Neither was deferred by the Omnibus. An organization can be in breach today while still inside the high-risk grace period.

Obligation	Applicable since	Status
Chapters I and II, including the Article 5 prohibited practices	February 2, 2025	Fully applicable. Not affected by the Omnibus delay. Certain new Article 5 points added by Regulation (EU) 2026/1744 apply from December 2, 2026.
Chapter V general-purpose AI model obligations, plus Chapters VII and XII and Article 78	August 2, 2025	Fully applicable. Article 101 penalties for GPAI providers apply from August 2, 2026.
Article 50 transparency duties	August 2, 2026	Fully applicable.
Chapter III Sections 1 to 3 for Annex III high-risk uses	December 2, 2027	Deferred by Regulation (EU) 2026/1744.
Chapter III Sections 1 to 3 for AI embedded in Annex I products	August 2, 2028	Deferred by Regulation (EU) 2026/1744.

Table 9. The EU AI Act clock in full.

Source note. Regulation (EU) 2024/1689, Article 113, as amended by Regulation (EU) 2026/1744.

EU hard stops

Hard stop	Why it blocks launch
The provider or deployer role is unresolved	Obligations and technical-documentation ownership cannot be assigned.
Intended purpose is unstable	Risk class, testing, and user instructions cannot be defended.
Human oversight exists only as a button	The operator may lack time, competence, information or authority to intervene.
Agent logs cannot reconstruct tool action	Monitoring, incident analysis, and accountability are structurally weak.
Personal-data memory cannot honor rights	GDPR access, erasure, minimization, and retention controls may fail.
Article 50 marking or labels are missing where applicable	The obligation is already applicable from Aug 2, 2026.

Table 10. EU hard stops.

EU PLAY. Use the extra high-risk implementation time to build reusable evidence primitives. Do not spend it writing a larger policy manual.

4. United states. Conduct, sectors and courts


The United States has no single horizontal AI statute equivalent to the EU AI Act. That does not make it unregulated. It makes the control graph more dependent on the use, sector, state, representation, and harm. The same model can be a marketing claim, a credit decision component, a medical-necessity tool, an employment screener or an antitrust mechanism. Each context activates a different liability path.

Federal policy since January 2025 has emphasized innovation and the removal of perceived barriers. The December 2025 executive order and January 2026 DOJ memorandum added a formal preemption campaign against selected state AI laws. In March 2026 the White House published a national legislative framework favoring a uniform federal approach. These are material signals, but they are not a blanket immunity. State laws remain operative unless preempted, enjoined, repealed or otherwise displaced.

At the same time, federal sector and conduct controls remain practical. The FTC has used its unfairness and deception authority against unsafe facial-recognition deployment and deceptive AI claims. The SEC has brought AI-washing cases. Creditors still need specific and accurate adverse-action reasons under ECOA, 15 U.S.C. § 1691(d)(3), and Regulation B, 12 C.F.R. § 1002.9(b)(2), which requires the statement of reasons to be specific and to identify the principal reasons. The Federal Reserve issued SR 26-2 to replace SR 11-7, emphasizing model materiality, exposure, validation, monitoring, governance, and proportionality. DOJ's RealPage enforcement shows that algorithmic coordination can become antitrust evidence rather than innovation theater.

The same system meets three different gates

These regimes overlap, but they do not recognize one another's artifacts



Passing one gate is evidence. It is not a passport.

Figure 2. Product conformity, conduct liability, and service filing are three different regulatory gates.

Source note. White House, 2025a, 2025b, 2026. DOJ, 2026. FTC, 2023, 2024. SEC, 2024, 2025. CFPB, 2022, 2023. Federal Reserve Board, 2026. DOJ, 2025.

The US operating map

Layer	Question	Evidence that matters
Federal consumer protection	Was the claim or practice deceptive or unfair?	Marketing substantiation, testing, foreseeable harm, monitoring, and remediation
Banking and credit	Was the model governed, explainable and lawful for the decision?	Inventory, materiality, validation, reasons, monitoring, and adverse-action records
Securities	Were statements about AI accurate and supportable?	Product capability evidence, governance records and disclosure review
Antitrust	Did the system coordinate, share sensitive data or reduce competition?	Data provenance, competitor inputs, objective function, recommendations, and override behavior
California	Is covered ADMT used for a significant decision as defined at 11 CCR § 7001(ddd)?	Pre-use notice, access, opt-out workflow, risk assessment, and data map
Colorado	Does covered ADMT materially influence a consequential decision?	Technical documentation, interaction notice, adverse-outcome explanation, records, and human review
Employment and local law	Does an automated tool affect applicants or workers?	Bias audit where required, notice, consent, retention, and decision records

Table 11. The US operating map.

The fairness asymmetry

One global fairness test is not enough. Europe can require data-governance and fundamental-rights analysis without reproducing the US doctrine, evidentiary posture or protected-class data environment. US statutes and state rules can require specific discrimination analysis, notices, reasons or review. China's public-facing generative AI measures prohibit specified discriminatory content and impose content-governance duties, but their beneficiary and enforcement logic differ.

The operational answer is a test catalog with policy-controlled data access. Do not export protected-attribute datasets or test outputs merely to keep the dashboard uniform. Store the method, test identifier, population definition, result, confidence, limitation, and approval in the evidence spine. Keep sensitive test data and legal conclusions in the appropriate regional plane.

US PLAY. Treat federal deregulation rhetoric and state implementation as simultaneous facts. The control graph needs both a litigation watch and live state policies.

5. China. The service gate is operational

China's AI governance is best understood as a service stack rather than a single AI law. The Algorithm Recommendation Provisions, Deep Synthesis Provisions, Interim Generative AI Measures, Personal Information Protection Law, Data Security Law, revised Cybersecurity Law, Regulations on Network Data Security Management and sector rules can all matter. Scope depends on service type, audience, content function, data, and regulatory attributes such as public-opinion or social-mobilization capability.

The filing system has visible scale. CAC reported 868 generative AI services filed and 530 applications or functions registered by April 30, 2026. The underlying algorithm filing regime is also precise about timing. Providers with public-opinion or social-mobilization attributes must file within ten working days from the date services are provided, under Article 24 of the Algorithm Recommendation Provisions. That is different from saying every AI service needs a filing number before launch. A mature playbook keeps filing, security assessment and go-live readiness separate.

The 2026 anthropomorphic interaction measures add a targeted layer. They apply to public services in China that use AI to provide continuing emotional interaction simulating a natural person's characteristics, thought patterns and communication style. The official explanation excludes ordinary customer service, knowledge question answering, work assistance, learning, education, and scientific research without continuing emotional interaction. Covered services face duties related to safety, sensitive interaction data, dependency risks, minors, disclosures, use-duration reminders, and human intervention in specified crisis conditions. A broad chatbot label would overstate scope and waste controls. A narrow scope memo with product evidence is better.

The portable asset is evidence, not compliance




Figure 3. The portable asset is a shared evidence spine, while compliance decisions remain local.

Source note. CAC, 2022a, 2022b, 2023, 2024a, 2024b, 2026a, 2026b. National People's Congress, 2021a, 2021b, 2025. State Council, 2024.

China scope and launch triage

Question	If yes	If no
Is the service offered to the public in China?	Continue public-service scope analysis.	Assess enterprise, research and sector rules separately.
Does it generate text, image, audio or video?	Test the Interim Generative AI Measures and deep-synthesis labeling.	Map other algorithm, data and sector rules.
Does it have public-opinion or social-mobilization attributes?	Assess filing and security-assessment duties.	Document why the attribute is absent and monitor change.
Does it provide continuing emotional interaction?	Test the anthropomorphic interaction measures.	Do not import that regime merely because the interface is conversational.
Does it publish or distribute deep-synthesis content?	Apply identity, labeling, content, and logging duties within scope.	Retain evidence of the product boundary.
Does personal or important data leave China?	Apply the current cross-border pathway and thresholds.	Keep the route controlled and prove localization where promised.

Table 12. China scope and launch triage.

The data-localization myth

China does not impose a single blanket localization rule for all data, although localization does apply to specified categories and operators. The legal pathways differ for critical information infrastructure operators, important data, personal-information volumes, sensitive personal information, and specified exemptions. The March 2024 Provisions on Promoting and Regulating Cross-Border Data Flows narrowed and clarified several obligations. Counted cumulatively from January 1 of the current year, and for handlers that are not critical information infrastructure operators, personal information of 100,000 individuals or fewer is exempt from all three mechanisms (Article 5(4)); between 100,000 and 1,000,000 individuals, or sensitive personal information of fewer than 10,000 individuals, requires a standard contract or certification (Article 8); and 1,000,000 individuals or more, sensitive personal information of 10,000 or more individuals, or any export of important data, requires a CAC security assessment (Article 7(2)). Article 5 adds further non-volume exemptions for contract necessity, cross-border human-resources management and emergency protection of life or health. CAC has reported that the March 2024 provisions reduced the volume of security-assessment and standard-contract filings, although the published figures should be checked against the current CAC release before use in a filing forecast.

The correct architecture is not data never leaves. It is data leaves only through a classified, approved and observable route. A product promise that all data stays local is stricter than the law and becomes its own enforceable representation. Marketing promises become enforceable architecture the moment they are published. Compliance owns that veto.

China hard stops

- The China provider, filing applicant or accountable operating entity is unresolved.
- The service cannot display the applicable model or filing information where required.
- Content labeling and moderation controls cannot distinguish public output from internal enterprise workflows.
- Real-identity duties are assumed from PIPL rather than mapped to the applicable service rule.
- Cross-border data is unclassified or routed through a global memory store by default.

6. The regulatory dictionary

Cross-border failures often begin as vocabulary failures. Teams hear the same phrase and assume the same control objective. The dictionary below translates the operating object, beneficiary, and artifact, not merely the word. That is the difference between a glossary and a control system.

Term	EU meaning	US meaning	China meaning
Risk management	Lifecycle process tied to system risk and rights	Sector-specific model, conduct or enterprise risk	Classified safety, content, data, and service governance
Conformity	Market-entry and product-assurance logic	No general analogue. Validation or lawful use may matter	No general analogue. Filing and assessment are different objects
Human oversight	Effective intervention, competence and authority	Qualified review, appeal, validation or decision responsibility by context	Content review, user protection, crisis handoff or sector supervision
Transparency	Notices, instructions and synthetic-content duties	Specific reasons, consumer notice or truthful claims by context	AI disclosure, labels, filing display, and content-governance signals
Data governance	Quality, relevance, representativeness, privacy, and rights	Purpose, accuracy, fairness, security, and sector requirements	Personal information, important data, content, and cross-border controls
Accountable actor	Provider, deployer and supply-chain roles	Developer, deployer, creditor, employer, adviser or other regulated actor	Service provider, technical supporter, platform, and sector entity

Table 13. The regulatory dictionary.

A governance resolver converts context into controls




Figure 4. The policy resolver turns deployment context into an executable control bundle and evidence package.

Source note. RefleXio Intelligence analysis. The resolver logic is a RefleXio construct, not a regulatory artifact.

False friends that create expensive controls

Human in the loop

A person clicking approve is not proof of human oversight. Effective oversight requires information, competence, time, independence, authority, and a technically reliable way to intervene. In some US contexts, the human is the legally responsible decision maker. In China, human intervention may be part of content governance or crisis handling. The same interface can fail all three regimes in three different ways.

Algorithm filing

A filing is neither a certification nor a safety seal. It fixes information for supervision and accountability within a specified scope. Treating it as an approval badge overstates its meaning. Failing to operationalize the filing number, model mapping and change process understates its effect.

Model validation

US banking validation is not EU conformity assessment. SR 26-2 calls for a risk-based model-risk program proportionate to materiality and exposure, including effective challenge, validation, monitoring, and governance. An EU technical file can provide source evidence, but the institutional test and accountable reviewer are different.

Data sovereignty

Data residency is one control, not sovereignty. Identity, encryption keys, support access, model telemetry, prompt, and memory retention, vendor sub-processors, legal entity, and operational control

also matter. A local data center connected to a global admin plane may be geographically local and operationally foreign. Maps are patient. Access logs are less sentimental.

If a team says	Ask instead
We are EU compliant	For which system, role, purpose, risk class, entity, and date?
The model is approved in China	Which service, filing or assessment pathway and which displayed identifier?
A human reviews it	What information, authority, time, and reversal mechanism does the human have?
The US is deregulated	Which sector, state, conduct rule, representation, and litigation exposure were removed?
Data stays local	Which data, logs, keys, support channels, memories, and backups are included?

Table 14. Data sovereignty.