RKS Global

April 2026 · rks.global

Russia's Draft AI Law: Why It Threatens Human Rights

Table of Contents

1. Introduction

Artificial intelligence (AI) technologies are rapidly transforming the economy and public administration, and are also being actively deployed in law enforcement and by the military. Consequently, many countries face the conflicting challenge of, on the one hand, protecting domestic developers from the expansion of foreign tech giants, and, on the other, exporting their own technologies and regulatory approaches. For a long time, Russia refrained from comprehensive regulatory legal regulation in this area, adopting a wait-and-see approach and giving preference to sectoral regulation and soft law mechanisms.

On 18 March 2026, the draft Federal Law "On the Foundations of State Regulation of the Areas of Application of Artificial Intelligence Technologies in the Russian Federation" was made available for public consultation.

Experts at RKS Global, who systematically monitor AI regulatory developments in Eurasian countries (see the report “Human vs. Model: How Governments Use AI for Censorship and Surveillance”), have prepared an overview of the draft law, focusing on its implications for human rights and the digital technology industry.

2. General Overview and Structure of the Draft Law

The draft law contains twenty-one articles and covers a broad range of issues, ranging from terminology to the obligations of individual entities and the competence of government authorities, as well as international cooperation. At the same time, the text is quite inconsistent with some provisions being described superficially, while others are given excessive details. For example, Article 20 provides detailed regulation of energy benefits for data processing centres,including preferential grid connection, reduced tariffs, and exemption from connection fees. This falls partially outside the stated subject matter of the draft law.

A very important and alarming symptom from the perspective of fundamental rights is that more than twenty key issues have been delegated to secondary legislation of governmental bodies including the Government of the Russian Federation, the Federal Service for Technical and Export Control of Russia (FSTEC), the Federal Security Service of the Russian Federation (FSB), and other authorities. In addition, new powers have been granted to specific government authorities. For example, the President of the Russian Federation is now exclusively empowered to determine the specifics of AI deployment in defence, national security, public safety, and law enforcement.

From a logical and legislative drafting perspective, the document remains insufficiently developed and internally contradictory. For example, the text introduces terms that are not defined in the terminology section or appear only once; certain conceptual provisions are included but not elaborated upon; and the rights and obligations of entities conflict within the lifecycle stages of AI systems. It may be assumed that different parts of the document were drafted by several stakeholders at different times, and subsequently compiled under significant time pressure. Drivers behind this urgency may include both the upcoming State Duma elections and the emerging regulatory lag behind other Eurasian Economic Community member states.

3. Scope and Exceptions

Part 4 of Article 1 explicitly excludes from the scope of the law the use of AI for the purposes of defence, state security, law enforcement, counter-terrorism protection, and prevention of natural and man-made emergencies. While the exclusion of defence is, unfortunately, a fairly common trend among legislators worldwide, the exclusion of security and law enforcement is extremely alarming. As early as 2023, in the case of Glukhin v. Russia (no. 11519/20, 4 July 2023), the ECtHR pointed to a regulatory vacuum and de facto arbitrariness in the use of facial recognition technologies by government authorities in public spaces. This vacuum has since never been addressed through regulation, while federal and regional authorities have deployed a significantly larger number of automated surveillance systems in both the physical and digital domains.

Part 3 also states that the law shall apply to natural and legal persons engaged in the development, application, and deployment of AI technologies on the territory of the Russian Federation. However, the law further includes sections defining the powers of government authorities and the principles of international cooperation.

If adopted, the law is expected to enter into force on 1 September 2027.

4. Terminology

The definition of AI (para. 1, Art. 3) reproduces the wording used in the National AI Development Strategy of 2019 and. is based on the criterion of "imitation of human cognitive functions." This approach differs significantly from those adopted at the international level and in most jurisdictions.

Tying the definition to the "imitation of human cognitive functions" does not provide clarity and creates a risk of excluding technological solutions that are not directly based on machine learning methods or neural network architectures.

To address this, the term "AI technologies" has also been introduced. It defines a list of technological solutions: technologies based on AI, including computer vision, natural language processing, speech recognition and synthesis, intelligent decision support, and advanced AI methods. However, enshrining a list of technologies in law is not an optimal solution, which is why most legislators worldwide strive to make the text technology-neutral, supplementing it with recommendations, technical standards, codes of practice, and other frameworks.

In addition, the text contains definitions of AI models and large foundation AI models. As regards the latter, the National Strategy set a minimum threshold of 1 billion parameters (approximately the level of GPT-2 1.5B or modern small models). The current text states that this threshold will be further determined by the Government. It should be noted that this distinction has no practical significance, as the text does not differentiate

the requirements for developers of ordinary AI models and those of large foundation models.

5. General Principles of Regulation

A risk-based approach has been declared as one of the regulatory principles, likely as a nod to international trends and the EU approach; however, this concept is not developed in the rest of the draft law: the text contains neither a prohibition of any practices nor a classification of systems into groups according to the risks posed by their development and deployment.

The principle of technological independence envisages the development of "sovereign and national AI models." At the same time, the development and deployment of AI technologies must be carried out on the basis of traditional Russian spiritual and moral values:

"life, dignity, human rights and freedoms, patriotism, civic-mindedness, service to the Fatherland and responsibility for its fate, lofty moral ideals, a strong family, creative labour, the priority of the spiritual over the material, humanism, mercy, justice, collectivism, mutual assistance and mutual respect, historical memory and continuity of generations, the unity of the peoples of Russia."

To oversee the implementation of these provisions, the Government of the Russian Federation *"shall designate an authorised body (expert organisation) responsible for developing criteria for the compliance of AI models, systems, and services with traditional Russian spiritual and moral values, and shall determine the procedure for its activities."*
These requirements cannot be translated into specific metrics and technical standards. This allows the authorities to apply these norms arbitrarily and equip them with an unrestricted instrument for controlling market participants. Moreover, such requirements create the risk of transforming AI models from neutral sources of information into channels for disseminating state ideology. Combined with the advanced persuasion capabilities of modern models, such AI systems could become a specialised instrument for propaganda and the manipulation of public opinion.

6. Sovereign, National, and Trusted AI Models

Article 7 elaborates on the concepts of "sovereign and national" large foundation AI models. It should be noted that both terms are not defined separately in the text, but are used only in conjunction.

For a model to be recognised as "sovereign and national," three conditions must be met:

1. All stages of AI model development and training are carried out on the territory of the Russian Federation.
2. All stages of development, training, and operation of AI models are carried out by citizens of the Russian Federation and Russian legal entities.
3. AI model training use datasets that are compiled on the territory of the Russian Federation by citizens of the Russian Federation and Russian legal entities.

However, the specific criteria for meeting these requirements have also been delegated to further rulemaking by the Government.

Notably, the law provides few positive consequences for a model being recognised as "sovereign and national." Article 7 contains only an abstract reference to ensuring favourable conditions for development, deployment, and application.

Article 8 sets out the requirements for the "trusted AI models," which may exclusively be used in governmental information systems (GIS) and critical information infrastructure (CII) facilities:

1. confirmed compliance with the applicable security and quality requirements established by FSTEC, FSB, and other authorities and state corporations;
2. data processing carried out exclusively on the territory of the Russian Federation;

3. confirmed compliance with the applicable security and quality requirements established by FSTEC, FSB, and other authorities and state corporations.

Models meeting the requirements are included in a register of the "trusted models." The procedure for maintaining the register and the requirements for compliance confirmations are also to be determined by the Government.

Thus, these provisions may be used by the authorities to pressure businesses if recognition of a model as "sovereign and national" or "trusted" becomes an important barrier to its operation and access to funding. At the same time, the vague wording will allow the creation of a simplified regime for large state-affiliated AI developers.

7. Legal Entities and Their Obligations

The draft law identifies a number of entities and their corresponding rights and obligations, thereby attempting to reflect the complexity of interactions within AI supply chains.

Entity	Obligations
Model Developer Product type: AI Model	Eliminate bias; Notify of prohibited use cases; Document the architecture and limitations of the model; Conduct risk modelling and establish the maintenance procedure.
System Operator Product type: AI System	Prepare a safe operation manual (including prohibition of manipulation and exploitation of vulnerabilities); Test the system for unlawful use scenarios; Inform users of the purpose and limitations; Ensure maintenance; immediately suspend operation in

	the event of a threat of harm; Maintain an incident log and appoint persons responsible for security.
Service Owner Product type: AI Service	Establish access rules prohibiting unlawful use; Take measures against abuse; Inform users about interaction with AI (except in obvious cases); For audiences exceeding 500,000 users per day, comply with information dissemination organisers (ORI) requirements.
User	Comply with the access rules for the service; Use AI for lawful purposes and not circumvent built-in safety mechanisms.

Despite the intention to reflect the multi-tiered nature of supply chains, "model – system – service – user," these obligations are not comprehensively defined. For example, the model developer is required only to model risks, with no reference to the typology and sources of such risks or to risk management requirements.

Also this article sets out, what is, in essence, the only prohibited practice, which appears nowhere else and applies solely to system operators: *"the inadmissibility of using the system for manipulating behaviour and exploiting human vulnerabilities."* Part 1 of Article 10 also introduces a new entity – the developer of an AI system – which appears nowhere else in the text and is absent from the terminological apparatus.

From a human rights perspective, the provision imposing the obligations on information dissemination organisers (ORI) – within the meaning of Article 10.1 of Federal Law No. 149-FZ of 27 July 2006 "On Information, Information Technologies, and Information Protection" – on owners of AI services with a daily audience exceeding 500,000 users located on the

territory of the Russian Federation is a cause for concern. This entails an obligation to store data on users' interactions with the service, including the content of their messages, within the Russian Federation, and to ensure its availability to authorised bodies through a mechanism providing round-the-clock remote access for an authorised FSB unit to the ORI's information system. This also potentially opens up broad opportunities for the implementation of systems that will automate the detection of undesirable content in users' interactions with AI systems.

8. Specific Human Rights

Article 9 of the draft law is intended to introduce new rights for citizens in relation to AI technologies:

1. to be informed about the use of AI systems in the sale of goods/services and about autonomous decision-making affecting their rights;
2. to refuse interaction with AI systems (in cases determined by the Russian Government);
3. to pre-trial appeal of decisions by government authorities and state-owned organisations made using AI systems;
4. to compensation for harm caused by unlawful use of AI systems (in accordance with the procedure established by the Civil Code of the Russian Federation).

However, these rights are formulated declaratively, essentially repeating provisions of general legislation, and do not provide effective human rights protection mechanisms. The provisions on informing about autonomous decision-making are worded in a manner that allows an extremely narrow interpretation: mere notification of the fact without explaining the logic behind a decision that has affected a person's rights.

The possibility of refusing interaction with AI systems is limited to cases to be subsequently determined by the Government. It may be assumed that this will apply to only essential services. The right to compensation for harm is provided only for unlawful use of AI systems under the procedure established by the Civil Code of the Russian Federation. In effect, no new right is created; the provision merely references general law. Nor does it

provide for the right to compensation for harm caused by a defective product during lawful use of an AI system. The very title of the article is also noteworthy, as it effectively extends these rights only to Russian citizens, disregarding foreigners and stateless persons.

9. Liability

Article 11 of the draft law allocates liability among entities within supply chains in proportion to each entity's degree of fault, but provides no further specifics, merely referring to general legislation.

The article also effectively relieves model developers, system operators, and service owners of liability. It arises only where they *"knowingly knew or should have known of the possibility of obtaining an unlawful result through the use of their products, unless otherwise established as a result of investigative actions."* It should be noted that the article does not distinguish between different grounds for legal liability. One might assume that the authors intended only criminal and administrative liability, since the caveat regarding investigative actions is inapplicable to civil liability. However, Part 7 of the article states that criminal and administrative liability arises in accordance with the procedure established by general legislation.

Moreover, these entities may also avoid liability *"if they took exhaustive measures to prevent such a result and complied with the requirements of the legislation of the Russian Federation."*

Taken together, these provisions create a regulatory framework that opens broad possibilities for arbitrary enforcement. In some cases, state-loyal companies may evade liability, while in others, restrictions and demonstrative penalties may be imposed.

The draft law also provides for the possibility of recourse claims by system operators against model developers where the operator proves that *"the harm arose exclusively as a result of defects in the product that existed at the time of its transfer to the operator and could not have been identified through the exercise of all reasonable diligence provided for by the operating conditions and technical documentation."* On the one hand, this

provision disregards the right of another entity – the service owner – to bring such a claim; on the other, it imposes a disproportionate burden of proof on the system operator, who will have no practical means of substantiating their position when using proprietary models. At the same time, there are no additional procedural guarantees for the disclosure of information to downstream entities in the supply chain.

10. Labelling of Synthetic Content

All obligations for explicit and implicit labelling of synthetic content are assigned to the final link in the supply chain – the service owner. International approaches, by contrast, proceed from the need to introduce multi-layered labelling applicable at both the system and AI model levels.

Several provisions have been implemented in line with China's "Interim Measures for the Management of Deep Synthesis Systems" (2022): the obligation to verify labelling is assigned to digital platform operators, and labelling may be excluded from generated content, provided a contract with the user is concluded. However, while this measure is less burdensome for businesses, it provides less transparency than in China, where, in addition to the contract, logs must be retained for at least six months.

11. Intellectual Property

The draft law proceeds from the assumption that generated content is not protectable by default, except where it constitutes *"exclusively original works meeting the protectability criteria established by civil legislation."* At the same time, these provisions directly contradict the logic of the remaining parts of Article 13, which delegate the transfer of intellectual property rights to agreements between the model developer, service owner, and user. The contradiction lies in the fact that, on the one hand, the article does not recognise the protectability of synthetic content; yet, on the other hand, it obliges entities to agree on the distribution of rights over outputs that may not constitute intellectual property at all.

The article also follows the logic of correlating the rights and obligations of different entities within the supply chain. However, as with other provisions, there is an imbalance: on the one hand, model developers must ensure that models are created using materials that do not infringe the intellectual property rights of third parties; on the other hand, the contract between the service owner and the user must include *"information on the lawfulness of the origin of the material used for training the system."*

In most cases, the service owner will not simultaneously act as the model developer or system operator and will not possess the necessary information on the composition of the datasets used to create the model. At the same time, the draft law contains no corresponding requirements for the disclosure of information about datasets.

12. Government Authorities

In addition to the relationships between end-users and commercial entities, the draft law also confers authority on various federal and regional government bodies.

The President of the Russian Federation will approve the National AI Development Strategy (as he has done previously) and determine *"the specifics of the application of AI technologies in the areas of defence, national security, public safety, and law enforcement."* These provisions pose a danger to human rights, as they effectively remove public safety and law enforcement from the legislature's exclusive competence. Henceforth, the application of AI systems in these areas may not require the adoption of any law. In the long term, this provision becomes particularly dangerous in the context of the rapid development of technologies deployed in the area of public safety and law enforcement: police robotics, predictive policing systems, speech and emotion recognition technologies in places of mass gathering, behavioural analysis through body-worn cameras of law enforcement officers, and automated social media monitoring. Each of these technologies poses significant risks to human rights, such as privacy, freedom of assembly, freedom of expression, and the prohibition of discrimination.

The main regulatory and administrative powers are vested in the Government of the Russian Federation and a separate competent authority to be designated subsequently. These functions are currently performed by the Ministry of Digital Development, Communications and Mass Media of the Russian Federation.

Certain powers have also been assigned to regional government authorities; however, given the narrow scope of their competence as defined by the Constitution of the Russian Federation, it may be assumed that their activities will largely focus on inter-regional cooperation and the implementation of federal policy. In terms of regulatory acts, one should expect direct replication of federal regulations.

13. International Cooperation

Articles 15–18 define the vectors of international cooperation, and are predominantly declarative and programmatic. Overall, the emphasis is on Russia's "leading role," the formation of a "unified space of trusted AI technologies," and the promotion of Russian approaches and standards. Article 17 concerns so-called "cross-border AI technologies." It contains provisions that create the legal basis for further restricting or prohibiting the use of foreign AI technologies on the territory of the Russian Federation: the operation of cross-border AI technologies may be prohibited or restricted in cases established by Russian law.

The vague wording, combined with Article 7 (sovereign and national models), outlines the contours of a digital isolationism regime in the field of AI, mirroring that in the broader sphere of the Internet and digital technologies. In the future, this could result in the prohibition of popular services and models, such as ChatGPT, Claude, Llama, Mistral, and others. In turn, this will significantly affect Russian businesses, which rely heavily on foreign systems and models to build their products. Therefore, businesses will be forced either to circumvent restrictions to continue using foreign services via VPNs or intermediary foreign entities, or use Chinese alternatives. Companies that rely on participation in government projects or

procurement will be compelled to use only software from Russian developers.

14. Computing Infrastructure

Article 20 provides detailed regulation of the development and operation of data centres (supercomputers) for AI purposes, including specific economic incentives such as preferential grid connection, exemption from connection fees, reduced electricity tariffs, tax benefits, and budget financing.

The level of detail in these provisions is at odds with the rest of the document. These provisions would be more appropriate in a separate sector-specific regulation or a state support programme. Their inclusion in the AI Act creates the impression that the interests of a particular industry group are being lobbied for.

15. Conclusions

1. The draft law declares a risk-based approach, but in practice does not implement it in any way. There is no classification of systems by risk levels, no prohibited practices are defined and no differentiated obligations are established. The actual backbone of regulation is not risk management, but access control – through the register of trusted models and the regime of sovereign and national models.
2. The areas of defence, national and public security, and law enforcement have been removed from the scope of the law and transferred to the exclusive competence of the President of the Russian Federation. This means that the application of AI systems precisely in those areas where the risks to human rights are the highest – facial recognition, predictive policing, and automated surveillance – remains in a regulatory vacuum, as identified by the ECtHR as early as 2023.

3. The introduction of spiritual and moral values as criteria for evaluating AI models is unprecedented in global practice. These requirements cannot be translated into technical standards and business processes, and, combined with the vague admission criteria for the register, create a mechanism of market segmentation in which access is determined not by the quality of the model but by the developer's loyalty.

4. Citizens' rights are formulated declaratively and are not backed by effective protection mechanisms . The right to information has been reduced to being merely notified of the fact of AI deployment, without disclosing the logic of decision-making. The right to refuse interaction with AI systems is contingent on a Government decision and will most likely be limited to a narrow range of cases. Compensation for harm is tied to the "unlawfulness" of use and refers to the general provisions of the Russian Civil Code.

5. The liability regime is structured in favour of developers and operators: the formulation "knowingly knew or should have known," combined with the possibility of exemption upon compliance with "exhaustive measures," creates the conditions for selective enforcement. The information asymmetry within supply chains is not offset by procedural guarantees for information disclosure.

6. More than twenty key issues have been delegated to the Russian Government, FSTEC, FSB, and other authorities, which dilutes the role of the federal law in the system of regulatory legal regulation. For market participants, this means that it is impossible to assess their obligations until secondary legislation is adopted. It also opens broad possibilities for arbitrary interpretation of the law's provisions by enforcement authorities, creating a threat of human rights violations.

7. The linkage of Articles 7, 8, and 17 creates a three-tiered regime for restricting the domestic market: preferences for domestic developers (sovereign and national models), admission through state certification (register of trusted models), and a legal basis for prohibiting foreign technologies (cross-border restrictions). Together with the declarative nature of the articles on international cooperation, this indicates that the draft law is not oriented towards integration into the global regulatory environment but rather towards the establishment of a

closed national AI perimeter, consistent with the broader sovereign internet framework.

8. Overall, the draft law in its current version represents less an instrument for the protection of citizens' rights and AI risk management, and more a regulatory infrastructure for state control over the AI market and its participants.