GitLab Response to the Consultation on the EU Cloud and AI Development Act

GitLab welcomes the European Commission's objective of strengthening Europe's digital sovereignty, resilience, security, and capacity to develop and deploy cloud and AI technologies. We share those objectives, and appreciate the opportunity to offer a technical perspective as the Cloud and AI Development Act (CADA) proposal advances through the legislative process.

GitLab is the intelligent orchestration platform for software teams and their AI agents. It brings together planning, source code management, continuous integration and delivery, security testing, compliance management, deployment, and AI-powered capabilities across the software development lifecycle. This enables a broad range of entities, including public sector and highly regulated organisations, to develop, secure, and maintain software in-house, including software supporting public services and other mission-critical activities.

GitLab's product and commercial architecture gives customers choice within their own enterprise guardrails. Customers can select deployment models, cloud environments, infrastructure arrangements, and AI models according to their security, compliance, latency, operational, and sovereignty requirements. GitLab's AI architecture is model agnostic by design. GitLab does not train generative AI models, and all GitLab AI model sub-processors are restricted from using model input and output to train models. GitLab Duo and Duo Agent Platform (DAP) use a diverse set of models and model hosts rather than being tied to a single provider, and customers can connect approved third-party models, including models developed in the Union such as Mistral's, or their own customer-hosted models, through an AI Gateway. The AI Gateway is a standalone service that provides a common interface between GitLab Duo and Duo Agent Platform and approved AI models, handling authentication, model selection, request routing, and the return of model responses.

GitLab also brings a deep open-source perspective, as well as a long-standing presence in the Union itself. GitLab began as an open-source project in 2011, led by Ukrainian and Dutch co-founders, and our open-source software is published under the MIT License. As an all-remote company, a substantial share of our workforce, including engineering, security, and open-source community roles, is based across Member States, and we maintain a commercial and legal presence within the Union alongside a long-standing base of European public and private sector customers. Through our Developer Relations and open-source programmes, we support projects, contributors, communities, and industry initiatives across Europe.

In our experience, sovereignty outcomes benefit from measures such as meaningful customer control, portability, interoperability, provider diversity, and risk-based technical and organisational safeguards. We offer our comments in that constructive spirit, with the aim of helping CADA's obligations track as closely as possible to the risks they are designed to address, and supporting the ability of European public authorities to access independent, secure, innovative, and technically appropriate tools for developing software.

I. Distinguish SaaS from IaaS and create a pathway for single-tenant services

For services within CADA's cloud-service scope, the framework would benefit from more clearly distinguishing among IaaS, PaaS, multi-tenant SaaS, and single-tenant SaaS. These categories describe different functions and different control relationships.

In simple terms, IaaS provides the basic building blocks, such as computing capacity, storage, and networks. PaaS provides a managed environment in which customers can build or run applications, and SaaS provides a finished application that customers use as a service. A multi-tenant SaaS service supports multiple customers in a shared technical environment, with logical controls separating each customer's data and activity. A single-tenant SaaS service provides a dedicated environment for one customer, while the provider operates and maintains the application.

These distinctions matter because they determine who controls which risks. An infrastructure provider, a platform provider, and an application-layer provider built on top of them each have important responsibilities, but they do not control the same systems or risks, and a single set of requirements applied uniformly across all three will not track those differences. Requirements designed for physical infrastructure, such as data centre audits, broad establishment requirements, or blanket localisation rules, may not translate as precisely to application-layer services if applied without regard to the risks those services actually control.

CADA's own structure supports this distinction. Annex II states expressly that software within the meaning of the Cyber Resilience Act[1] falls within the scope of Annexes II and III, and that hardware falls outside it. Yet Annex III's audit criteria A and B (evidence for establishment and location) are built around physical premises, hardware, and facility operations. That asymmetry suggests the assurance regime was calibrated with infrastructure providers principally in view, notwithstanding that its stated scope is software. Establishing distinct treatment for application-layer services would give effect to a boundary the proposal has already drawn, rather than introduce a new one.

Establish a distinct SaaS treatment and a pathway for services hosted on recognised infrastructure

We therefore recommend that CADA establish a distinct SaaS service model definition, separate from IaaS and PaaS, with a correspondingly tiered obligation set in Annex II focused on matters within the SaaS provider's control. These include contractual and organisational safeguards, identity and privileged access controls, encryption and key management, data processing arrangements, service continuity, incident response, auditability, and the provider's actual ability to access customer content or the infrastructure storing it. Beyond rendering the obligations more proportionate, such changes would have a pro-competitive impact, better enabling meaningful compliance by small and medium-sized software providers that do not maintain their own physical or cloud infrastructure.

We would also recommend that Article 17, together with Annex II, establish a pathway for services hosted on recognised infrastructure. Under it, where a provider's service runs on infrastructure already recognised at or above the relevant Union assurance level, the provider would demonstrate compliance in its own right only for the criteria attributable to the layer it operates (as described in the paragraph above), and would reference the existing infrastructure recognition rather than reproduce it. Giving effect to this would require Annex II's criteria to be attributable to the layer that controls them.

The tenancy model should be a factor in determining which criteria are attributable to the provider's layer. In a multi-tenant service, customers share a common technical environment and separation is achieved through logical controls the provider designs and operates, so those controls are properly part of what the provider must demonstrate. In a single-tenant deployment, each customer's service runs in a dedicated instance, typically in a customer-selected region, with its own data stores and, where offered, customer-managed encryption keys and private network connectivity. Because there is no data plane shared across customers, the assurance question narrows to the integrity of the underlying infrastructure, which the infrastructure provider has already demonstrated, and the provider's own operational access and tooling through which that dedicated instance is managed, which is verifiable directly. Single-tenant offerings are therefore the clearest case for the pathway, though the underlying principle – that a provider should demonstrate the controls it operates and inherit those established at layers it does not – applies across service models.

II. Define metadata and telemetry functionally

Distinguish customer content from information about that content

Building on the layer distinctions noted above, Annex II should not treat customer content, metadata, and telemetry as a single category across all service models.

[1] Regulation (EU) 2024/2847

A useful starting point is to distinguish between the material itself and information about that material. Customer content may include source code, documents, configurations, or other information that a customer stores in the service. Metadata may describe that content or its use, such as the identity of a user, the time of an action, the location of a service, or the relationship between a project and a workflow. Telemetry is generally operational information about how a service is functioning, including service availability, error rates, processing times, pipeline status, scan results, or aggregated security signals.

These categories can overlap, and some metadata or telemetry may be sensitive. Information that identifies a customer, reveals the existence of a sensitive project, or permits access to customer-stored content should receive appropriate protection. But it does not follow that every operational signal should be treated as customer content or be subject to the same localisation requirement.

Annex III (audit criterion C) defines a customer as the contracting public sector body, not an individual natural person. Identification of a customer, for this purpose, should refer to the contracting organisation and its stored content, not to whether an individual end user is named or otherwise identifiable within routine operational metadata, such as the identity of the user who triggered a particular action. Personal data relating to identifiable end users remains fully subject to the General Data Protection Regulation, including its rules on international transfers, and importing that analysis into CADA's criteria would broaden a definition that can already extend to any data derived from use of the service.

Blanket treatment impairs security and resilience

For an application-layer service, operational telemetry helps security and reliability teams understand whether a service is operating normally. A provider may need to know that a workflow is failing in a particular region, that a vulnerability-scanning service is returning an unusual pattern of errors, or that an attack is generating abnormal API traffic. This information can be necessary to investigate incidents, manage vulnerabilities, maintain service continuity, and shift operations across regions when appropriate.

Treating all such information as customer content could make it harder to detect threats and maintain resilience. It could also prevent security teams from working together across regions when a distributed response is necessary.

Apply a functional test with predictable waiver criteria

We recommend that CADA define in-scope metadata and telemetry through a functional test, namely whether the data enables identification of a customer or permits access to that customer's stored content. The description of customer data in Annex III (audit criterion C) should be amended accordingly. Aggregated telemetry or operational telemetry that does not satisfy that functional test and is used solely for security, reliability, and availability purposes should not be subject to a blanket localisation requirement, provided it is minimised, used for defined purposes, encrypted where appropriate, access controlled, and auditable.

Annex II's own text supports a more granular approach at lower assurance levels. Union assurance level 4's criterion 4.1(c) already limits its customer data obligation to data that, following a risk assessment, is identified as sensitive, a qualifier that is absent from the equivalent criterion at Levels 1 through 3 (1.1(c), 2.1(c) and 3.1(c)). Extending that same risk-based qualifier, or the functional test proposed above, to those levels would more closely align the framework's treatment of metadata and telemetry with an approach the Regulation itself already adopts at its highest tier.

Localisation waivers should also use clear and harmonised criteria. Providers and public authorities should be able to determine the applicable requirements before entering into a procurement or designing a service architecture.

Criterion (f) at Union assurance levels 2 to 4 provides that data generated by using the service “are not transferred outside the Union in any case,” without the exemption that criterion (c) provides for a

public sector body to determine otherwise. As the European Data Protection Supervisor (EDPS) has noted (Opinion 14/2026, par. 23), the two should be reconciled. We recommend clarifying that criterion (f) restricts the use of such data to train or fine-tune AI systems operated by a third country or a third-country entity, and that transfers are governed by criterion (c) and, for personal data, by Chapter V of the General Data Protection Regulation.

III. Recognise equivalent assurance and tie obligations to actual access and risk

GitLab shares CADA's core security objective of ensuring that the cloud and AI services underpinning essential public functions meet a consistent, verifiable standard of security and operational resilience across the single market. The recommendations in this section are provided in furtherance of that objective.

Leverage existing assurances and avoid duplication

To avoid duplication, Article 20 should expressly recognise existing, independently assessed application-layer certifications and attestations, including SOC 2 Type II, ISO/IEC 27001, and CSA STAR, where their scope covers equivalent controls.

A layered assurance model would also help avoid asking different providers to demonstrate the same protection multiple times. The infrastructure provider may already have independently assessed controls for physical security, hardware, and the underlying computing environment. The SaaS provider may separately demonstrate controls for application access, software changes, customer data handling, incident response, and privileged personnel. Requiring both parties to repeat the other's assessment would add cost without necessarily improving security.

Tie personnel and establishment requirements to actual access

Establishment and personnel requirements would track the underlying risk most precisely if tied to actual access and operational responsibility rather than to a provider's global headcount. At Union assurance level 2 the personnel requirement in 2.1(b) should apply only to the roles that hold access to customer data or to the infrastructure on which it is stored. Personnel without that access should not be captured merely because they are employed within the same corporate group.

Annex III already contemplates this approach. Criterion B directs the auditing organisation to assess whether personnel store, transmit, access or process customer data, whether they enable administrative access to or visibility into it, and whether their compromise could disrupt the service. Those are functional criteria tied to access and impact. Annex II paragraph 2.1(b), by contrast, states a categorical locational requirement, and unlike paragraph 2.1(c) on customer data it carries no provision for a contracting authority to determine otherwise. Aligning 2.1(b) with the assessment Annex III already sets out would give effect to the risk-based logic present in the Regulation's own audit criteria.

The treatment of technical and operational support would also benefit from greater proportionality between levels. Paragraph 1.1(d) expressly permits support to be performed outside the Union where legal, technical and organisational measures ensure traceability, security and governance. Paragraph 2.1(h) permits none. A framework that moves from conditional permission to absolute prohibition between adjacent assurance levels leaves no intermediate option for workloads that warrant heightened but not maximal assurance. Level 2 should provide a pathway for time-bound assistance by specialist personnel where operationally necessary, subject to the safeguards already articulated in paragraph 1.1(d). Consistent with the treatment of subcontractors generally, the provision of such assistance should not itself cause the supporting entity to be treated as a subcontractor subject to the full establishment regime.

Calibrate assurance to the workload rather than the buyer

Article 29 implementing guidance should require a data classification step in the procurement risk assessment methodology. Assurance should track the sensitivity of the data processed by the specific tool and workload, not simply the classification of the contracting authority or sector.

A software development platform used by a defence ministry for ordinary source code management, testing, and delivery workflows may process materially different information from the ministry's classified records or operational systems. The fact that the same authority uses both systems should not, by itself, determine the assurance level for each service. As the EDPS has recommended (Opinion 14/2026, par. 52), that step can draw on the data protection impact assessments public bodies already carry out under Article 35 GDPR, a well-established mechanism supported by EDPB-endorsed guidance.

Reconsider the private-sector delegation in Article 31

We would encourage the co-legislators to reconsider paragraph 3 of Article 31, the only provision in that Article capable of producing binding obligations, and which empowers the Commission to adopt delegated acts specifying both the need for an impact assessment and the mitigation measures private entities shall take. Entities in sectors of high criticality already conduct risk assessments under Directive (EU) 2022/2555, and financial entities do so under Regulation (EU) 2022/2554. A delegated power to layer a further assessment regime on those entities risks duplication without a corresponding gain in resilience, and would extend the assurance framework's practical effects well beyond the public procurement context for which it was designed.

If paragraph 3 is retained, any delegated act adopted under it should attach to the specific critical function or system that creates the relevant risk rather than to the entity as a whole, and should not extend to general-purpose tools used for activities unrelated to that function.

IV. Treat open-source risk consistently and invest in the ecosystem

CADA should avoid treating open-source software as inherently higher risk, or imposing anti-tampering requirements solely because a component is open source. Open-source software's visibility can support review and reuse, but its actual risk still depends on governance, maintenance, provenance, and vulnerability management, the same factors that determine risk for proprietary software. This position is further supported by the EDPS (Opinion 14/2026, par. 88), which observes that open-source solutions can enhance transparency and auditability, because source code can be inspected and malicious functionality can be more easily identified.

Requirements should therefore be triggered by demonstrable risk factors, such as third-country control, unaccountable governance, an abandoned codebase, inadequate vulnerability disclosure practices, or uncertain provenance. A well-governed and actively maintained open-source project should not be treated like software of unknown origin or uncertain maintenance status. Under Article 13(5) of the Cyber Resilience Act, a manufacturer must exercise due diligence over the third-party components it integrates, including open-source components. CADA should apply the same principle: the entity that selects and integrates a component is responsible for managing its risk, whatever the component’s licence.

A corresponding safeguard against these risk factors is ensuring that the components European public administrations depend on are actively maintained. The Commission has itself recognised open source as a lever for sovereignty, competitiveness, and security in the EU Open Source Strategy, several of whose measures are advanced through CADA, and Chapter V rightly encourages public bodies to prefer open-source components assessed against objective criteria. Realising that ambition requires sustained investment in maintenance, security, and skills, not procurement preference alone. The Commission's own 2021 study valued open-source software's contribution to EU GDP at €65–95 billion annually, yet the components generating that value remain underfunded relative to their systemic importance. Existing national programmes, and ongoing work towards an EU-level equivalent, offer models worth building on. GitLab's open-source programmes

and Developer Relations engagement across European projects, communities, and industry organisations reflect the practical contributions that support Europe's digital resilience.

Conclusion

GitLab supports the Commission's ambition to strengthen Europe's digital sovereignty and ensure that public authorities can rely on secure and resilient technology. The pursuit of those objectives would be advanced were CADA to distinguish SaaS from IaaS and PaaS in Annex II; recognise infrastructure assurance already established at the relevant level, so that a provider demonstrates only the criteria attributable to the layer it operates; define in-scope metadata and telemetry functionally, with waiver criteria knowable before a procurement begins; clarify that criterion (f) restricts the use of data for AI training rather than operating as a blanket transfer prohibition; recognise equivalent application-layer certifications; tie personnel and establishment requirements at Union assurance level 2 to actual access, with a proportionate route for time-bound specialist assistance; calibrate assurance to the sensitivity of the specific workload rather than the sector of the contracting authority; reconsider the delegated power in Article 31(3); and assess open-source components on governance and maintenance rather than licence type, alongside sustained investment in the components public administrations rely on.

These changes would strengthen sovereignty by making it more technically precise and more useful in practice, giving public authorities the ability to develop secure software in-house, choose architectures that fit their requirements, and maintain meaningful control over their data, infrastructure, models, workflows, and security policies. GitLab looks forward to supporting the Commission's efforts as a long-standing technology partner to European public authorities, and welcomes the opportunity for continued engagement as this proposal develops.