August 24, 2026

BSA RECOMMENDATIONS ON DRAFT NOTICE TO IMPLEMENT ARTICLE 32 OF AI BASIC ACT

Submitted Electronically to the Ministry of Science and ICT (MSIT)

The Business Software Alliance (**BSA**)[1] would like to provide our recommendations to the MSIT on the draft Notice on Methods for Implementing the AI Safety Assurance Obligation (**Draft Notice**), which implements Article 32[2] of the Act on the Development of Artificial Intelligence and Establishment of Trust (**AI Basic Act**).

BSA is the leading advocate for the global software industry before governments and in the international marketplace. Our members are at the forefront of providing AI-enabled products and services, as well as tools used by others in the development of AI systems and applications.

BSA has engaged with MSIT throughout the development of the AI Basic Act and its subordinate legislation.[3] We have also collaborated with MSIT in ongoing efforts to refine the AI Basic Act so that it can better meet Korea's AI ambitions. Given said ongoing efforts to refine the AI Basic Act, we urge MSIT to defer the Draft Notice until that work concludes. Our remaining recommendations pertain to various aspects of the draft Notice which require greater specificity.

Summary of BSA's Recommendations

1. **Defer the Draft Notice until the ongoing review of the AI Basic Act is concluded.** Issuing binding rules now would pre-empt the AI Basic Act Improvement Working Group on questions squarely within its remit, add a further instrument to an already complicated regulatory landscape, and require businesses to prepare for obligations that MSIT may soon revise.

[1] BSA's members include: Adobe, Alteryx, Amadeus, Amazon Web Services, Asana, Atlassian, Autodesk, Avalara, Bentley Systems, Box, Cisco, Cloudflare, Cohere, Cohesity, Dassault Systemes, Databricks, Datadog, Docusign, Dropbox, Elastic, EY, Graphisoft, HubSpot, IBM, Kyndryl, MathWorks, Microsoft, Notion, Okta, OpenAI, Oracle, PagerDuty, Palo Alto Networks, PTC, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Trend Micro, TriNet, Veeam, Workday, Zendesk, and Zoom Communications Inc.

[2] Article 32 (Obligation to Ensure AI Safety) requires an AI business operator to identify, assess, and mitigate risk across the entire AI lifecycle, and to establish a risk management system to monitor and respond to AI-related safety incidents, where the following conditions as set by AI Basic Act Enforcement Decree are met: 1) the cumulative compute used to train the AI is at or above a threshold of 10^26 FLOPs; 2) the AI constructed and operated using the most advanced technology; and 3) the AI pose a risk of widespread and significant impact on human life, physical safety and fundamental rights. Article 32(3) specifically delegates the implementation methods to a ministerial notice, which is the subject of this consultation.

[3] See: BSA Recommendations on Updated Draft Enforcement Decree to the AI Basic Act, December 2025, https://www.bsa.org/policy-filings/korea-bsa-recommendations-on-updated-draft-enforcement-decree-to-the-ai-basic-act; BSA Recommendations on Korea AI Basic Act Draft Subordinate Legislation and Guidelines, October 2025, https://www.bsa.org/policy-filings/korea-bsa-recommendations-on-korea-ai-basic-act-draft-subordinate-legislation-and-guidelines; BSA Comments on Korea AI Basic Act, March 2025, https://www.bsa.org/policy-filings/korea-bsa-comments-on-korea-ai-basic-act.

2. **Make the scope of application more specific.** MSIT should define "substantial modification" and "substantial contribution", confirm that the Article 32 obligations attach to AI models rather than AI systems, and allocate obligations according to each entity's role in the AI ecosystem.
3. **Narrow the definitions underpinning incident reporting and revise the reporting timelines.** MSIT should define "risk" and "AI-related safety incident" by reference to specified harms that meet a high threshold of severity, extend the deadline for the initial incident report, clarify which entity bears the reporting obligation, require the final report as soon as practicable, and confirm that operators may rely on their existing incident response frameworks.

Defer the Draft Notice until ongoing review of AI Basic Act is concluded

The Draft Notice will establish binding rules to implement Article 32 of the AI Basic Act. Article 32 of the AI Basic Act governs AI which fulfill *all* of the following conditions: 1) cumulatively trained above the compute threshold set by the AI Basic Act Enforcement Decree (currently 10^26 FLOPs); 2) "constructed and operated" using the "most advanced AI technology"; and 3) pose a risk of "widespread and significant impact on human life, physical safety and fundamental rights" (**in-scope AI**). It requires relevant AI Business Operators to identify, assess, and mitigate risks across the in-scope AI's lifecycle, to establish a risk management system for AI-related safety incidents, and to submit the results to the Minister of Science and ICT (collectively, the **Article 32 obligations**).

As a general observation, it is premature to implement further binding rules pertaining to the AI Basic Act at this time. In March 2026, MSIT officially launched the AI Basic Act Improvement Working Group (**Working Group**), which brought together experts spanning academic institutions, industry associations, and civil society organisations.[4] In its press release, MSIT noted that in the course of initial consultations on the AI Basic Act and its subordinate legislation, it received a "wide range of opinions" which "required further consultation and consensus among various stakeholders or additional revision of the AI Basic Act and its Enforcement Decree". As a participant in the Working Group, BSA is grateful for the opportunity to work with MSIT and other relevant stakeholders to "identify areas requiring institutional improvement in the AI Basic Act", including Article 32.

In the circumstances, we are concerned that establishing binding rules for Article 32, while the Working Group process to review the AI Basic Act is ongoing, could result in the following unwanted consequences:

- **It would require businesses to prepare for obligations which MSIT may soon revise, resulting in wasted resources.** Businesses would need to prepare documentation, risk assessment structures, and incident reporting processes on basis of requirements that may be superseded. Businesses would then need to adjust their established processes

[4] MSIT Press Release: "*AI Basic Act Improvement Working Group Launched with Participation of Industry, Academia, and Civil Society*", March 2026, https://www.msit.go.kr/eng/bbs/view.do%3Bjsessionid%3DZzEuiEqQ8uA3vkPVkC_0eL9cvcgmFuF8rR_P8RfN.AP_msit_1?bbsSeqNo=42&mId=4&mPid=2&nttSeqNo=1238&sCode=eng

once more, thus resulting in wasted effort, increased costs and further regulatory confusion.

- **The Draft Notice adds another layer to an already complicated regulatory landscape.** Article 32 is currently addressed across the AI Basic Act, the Enforcement Decree, and non-binding guidelines. The Draft Notice adds yet another without stating how it interacts with MSIT's existing guidelines on AI safety obligations. Adding more pieces of subordinate legislation will increase compliance complexity, while undermining legal certainty.
- **It would pre-empt the Working Group on questions that fall squarely within its remit.** The Working Group is expected to consider issues related to Article 32, such as whether a cumulative compute threshold is still an appropriate criterion, whether the obligation should attach to AI models or to AI systems, and which parties should bear the obligation. The Draft Notice makes assumptions on those issues before the Working Group has reached its own conclusions, essentially undermining the Working Group process.

Recommendations: Defer finalizing the Draft Notice until the Working Group has concluded its review process and MSIT has settled the resulting amendments to the AI Basic Act, and consult again on a revised draft that reflects them.

Make the scope of application more specific

The Article 32 obligations attach to: 1) the developer of an in-scope AI; and 2) to any AI Business Operator that turns an AI into an in-scope AI through a "substantial modification". As previously highlighted, any AI that is trained above the cumulative compute threshold of 10^26 FLOPs will be considered in-scope.

As presently drafted, the terms that determine scope are undefined, the object of the obligations is unclear, and the allocation of those obligations does not reflect the roles that AI developers and AI deployers actually play.

- **Key terms, such as "substantial modification" and "substantial contribution", are undefined.** Under Article 3(1)2 of the Draft Notice, an AI Business Operator that substantially modifies an AI so that it becomes an in-scope AI assumes the full set of Article 32 obligations. The Draft Notice does not specify what makes a modification substantial, but Article 3(2) of the Draft Notice notes that "any training-related computation, including data pre-processing and fine-tuning" would count towards calculating the AI's cumulative compute. This is further complicated by a subsequent exclusion clause specifying that "computation that made no substantial contribution to the actual improvement of the AI's capabilities… may be excluded from the cumulative compute". The lack of specificity regarding what constitutes a "substantial contribution" is concerning. Fine-tuning, continued training, and adapting a model for a particular application are routine activities, and businesses need to clearly understand which activities would carry the potential consequence of triggering Article 32 obligations.
- **The Draft Notice (and the AI Basic Act more broadly) does not clearly specify whether the Article 32 obligations attach to AI *models* or to AI *systems*.** BSA urges MSIT to recognise the distinction between AI *models* (which create underlying technology that may be used for a broad range of different applications) and AI *systems* (which combine one or more AI models with other elements to create an AI application for a particular use). For

example, one large language AI model may be integrated into hundreds of different AI systems, such as translation apps or AI-powered chatbots. The model is the underlying technology, while the system puts it to a particular use. Given the focus on training compute, the Article 32 obligations should clearly attach to AI models, rather than AI systems.

- **Third, the Draft Notice conflates the distinct roles of AI *developers* and AI *deployers*.**[5] We have consistently recommended that compliance obligations should be allocated according to each entity's role in the AI ecosystem, and each entity should be accountable only for risks within its knowledge and control. As previously mentioned, any AI Business Operator that brings an out-of-scope AI within scope through a substantial modification assumes the full set of Article 32 obligations. The difficulty is that Article 32 requires risk to be identified, assessed, and mitigated across the entire AI lifecycle, including the development and training stages in which a deployer took no part. A deployer that fine-tunes a model does not acquire access to the base model's training data, architecture, evaluation results, or known limitations, and it is therefore not in a position to assess risks embedded during the original training process. The result is that the party with the least information about how an AI was built would bear the full burden of documenting it, while the developer of the base model would bear none under Article 32, because that model falls below the threshold on its own. Where a modification does bring an AI within scope, the modifying entity's obligations should be confined to the risks arising from its own modification, and MSIT should not expect it to account for training-stage decisions it did not make and cannot inspect. Likewise, the downstream risks that arise from subsequent modifications should not automatically fall on the developer of the base model, given that those risks arise from the modifying entity's own choices of post-training techniques, datasets, and end uses, which sit outside the developer's knowledge and control..

Recommendations: MSIT should define "substantial modification" in Article 3(1)2 of the Draft Notice and specify what constitutes a "substantial contribution" under Article 3(2) of the Draft Notice. MSIT should also confirm that the Article 32 obligations attach to AI models rather than AI systems, and allocate obligations appropriately depending on whether a business is an AI developer or deployer.

Narrow definitions of "risk" and of "AI-related safety incident" and extend the reporting timeline

The Draft Notice creates an additional incident reporting obligation for AI Business Operators. This is set out in the following provisions:

- **Article 2(1) of the Draft Notice** defines "risk" as "the likelihood that a person's life, body, or fundamental rights will be infringed, together with the severity of the potential harm".
- **Article 2(4) of the Draft Notice** defines an "AI-related safety incident" as "a case in which a risk is actually realized, or a serious harm to public safety arises, due to a malfunction of an AI or a similar cause".

[5] AI developers are entities that design, code or produce an AI system. AI deployers are entities that use an AI system that is either developed internally or by a third party. An entity can be both a developer and deployer.

- **Article 8(3) of the Draft Notice** then requires a relevant AI Business Operator to report the AI-related safety incident to MSIT in three stages: 1) an incident-occurrence report within 24 hours of becoming aware of the incident; 2) an initial-response report within 7 days of that first report; and 3) a report on incident-handling results within 15 days of it. The final report must set out the cause of the incident, the scale of harm, follow-up measures including the removal of residual risk, and a plan to prevent recurrence.

It is clear that the AI-related safety incidents that would trigger the reporting obligation must meet a high threshold of severity relating to life, physical safety, fundamental rights or public safety. On principle, BSA agrees with setting a high threshold. However, we raise the following points for MSIT's consideration:

- **The definitions of "risk" and "AI-related safety incident" are not specific enough**. For example, it is unclear if an infringement of a person's "body" requires physical harm (and if so, the severity of such harm). Likewise, given the broad scope of what are considered "fundamental rights" under the Constitution of the Republic of Korea,[6] it is not clear what incidents would constitute an infringement that is severe enough to trigger the reporting obligation. The reporting trigger should be based on objective and reasonably ascertainable criteria that allow businesses to distinguish reportable safety incidents from minor AI errors and malfunctions (e.g., inaccurate output). In the absence of specifics, it is difficult for businesses to determine if an event is reportable. This may result in over-reporting, which consumes time and resources that would be better directed at resolving the incident, while requiring MSIT to identify the incidents that genuinely threaten public safety from among a large volume of notifications that do not. Furthermore, requiring incident reporting for all risks that materialize, regardless of their severity or impact, would lead to comparatively minor issues such as inaccurate output being reported as an incident. This may then incentivize AI Business Operators to narrow the range of risks they assess in the first place to avoid reporting obligations, which would be counterproductive to MSIT's policy objectives. In the circumstances, it is critical for MSIT to define specific harms that meet a high threshold of severity for incident reporting, such as death, serious physical injury, or property damage that meets a significant threshold for financial loss.[7]
- **The reporting process is too vague.** The process for reporting AI safety incidents creates significant practical challenges, including because the requirement hinges on discovering an AI-related safety incident, which is vague. Companies often must quickly and thoroughly

[6] Per Articles 10 to 37 of the Constitution of the Republic of Korea, fundamental rights encompass a wide range of matters that extend well beyond life and physical safety to include equality (Article 11), freedom of occupation (Article 15), privacy (Article 17), freedom of expression and protection of honour (Article 21), and the right of property (Article 23)

[7] The thresholds for incidents in other jurisdictions are highly specific. For reference:

- Article 3(49) of the EU AI Act defines a "serious incident" as an incident or malfunctioning of an AI system that directly or indirectly leads to any of the following: (a) the death of a person, or serious harm to a person's health; (b) a serious and irreversible disruption of the management or operation of critical infrastructure; (c) the infringement of obligations under Union law intended to protect fundamental rights; or (d) serious harm to property or the environment.
- California's Transparency in Frontier Artificial Intelligence Act takes a similar approach, defining a "critical safety incident" as any of the following: (a) unauthorized access to, modification of, or exfiltration of the model weights of a foundation model that results in death, bodily injury, or damage to, or loss of, property; (b) harm resulting from the materialization of a catastrophic risk; (c) loss of control of a foundation model causing death or bodily injury; or (d) a foundation model that uses deceptive techniques against the developer to subvert the developer's controls or monitoring, outside the context of an evaluation designed to elicit this behavior, in a manner that demonstrates materially increased catastrophic risk.

investigate whether an incident has occurred, and it is unclear whether the reporting requirement is triggered when a company first suspects, reasonably believes, or confirms that an AI-related safety incident has occurred. This is compounded by the imprecision in the underlying definition, which leaves operators uncertain not only when the clock starts but whether it starts at all. In addition, the Draft Notice does not clarify which actor across the AI value chain bears the reporting obligation. A developer may not be aware of an incident arising during deployment unless the deployer communicates it, and an entity should not be responsible for failing to report an incident it could not reasonably have known about.

- **The reporting timeline is too short.** The 15-day timeline for submitting a final report requires companies to identify the cause of the incident, quantify the harm, remove residual risk, and produce a recurrence prevention plan. Such detailed analysis is impracticable in a 15-day window. The reporting requirements further obligate companies to divert resources from addressing the safety of their systems to fulfilling short-fuse reporting requirements, as an incident is unfolding. Instead, a final report should be submitted as soon as practicable, without creating a fixed time limit. The 24-hour deadline for the initial report is also out of step with comparable regimes (e.g., the EU AI Act and California's Transparency in Frontier AI Act), which have considerably longer reporting timelines.[8]
- **There is no alignment with existing reporting frameworks.** Where the Draft Notice requires an AI Business Operator to address residual risk or prevent recurrence following an incident, MSIT should expressly permit operators to satisfy the requirement by aligning with their existing incident response frameworks, such as those based on the NIST Cybersecurity Framework or ISO/IEC 27035. Many AI safety incidents, particularly those involving security vulnerabilities or malicious exploitation, already fall within the scope of an operator's cybersecurity incident response program, and a duplicative AI-specific process would add compliance burden without improving safety outcomes.

Recommendations: MSIT should:

- Define "risk" and "AI-related safety incident" by reference to specified harms that meet a high threshold of severity, such as death, serious physical injury, or major property damage.
- Make clear that the reporting obligation is triggered only once an AI Business Operator has confirmed that a reportable incident has occurred.
- Extend the deadline for the initial report to align with reporting timelines in other comparable regimes, such as the EU AI Act and California's Transparency in Frontier AI Act;
- Clarify which entity bears the reporting obligation so that no operator is responsible for reporting an incident it could not reasonably have known about.
- Require the final report to be submitted as soon as practicable, rather than within a fixed period that may not allow a meaningful root cause analysis to be completed.

[8] Under Article 73 of the EU AI Act, serious incidents involving high-risk AI systems must be reported within 15 days, shortened to 10 days where a death may have been caused and 2 days for a widespread infringement or a serious and irreversible disruption of critical infrastructure. Under California's Transparency in Frontier AI Act, a critical safety incident must be reported within 15 days of discovery, or within 24 hours only where it poses an imminent risk of death or serious physical injury.

- Confirm that operators may satisfy residual-risk and recurrence-prevention requirements by aligning with their existing incident response frameworks.

Conclusion

We thank MSIT for the opportunity to comment on the Draft Notice. We hope our recommendations are useful as MSIT continues its work on implementing the AI Basic Act. BSA values the opportunity to participate in the AI Basic Act Improvement Working Group, and we look forward to continuing that work with MSIT and other stakeholders. Please do not hesitate to contact me if you have any questions regarding this submission or if I can be of further assistance.

Yours sincerely,

Tham Shen Hong
Director, Policy – APAC