

TECHNOLOGY, MEDIA & TELECOMMUNICATIONS

IMDA Issues Guidelines for Generative AI Chatbots to Ensure Transparency, Accountability and Trust

Executive Summary

Generative artificial intelligence ("AI") has become increasingly interwoven into consumers' daily lives. Consumers need clear and accessible information to use these tools responsibly and confidently. However, such information may be too technical, fragmented across various sources (e.g. terms of service, privacy notices and blog posts), or may be over-simplified for marketing purposes. This leaves consumers uncertain about the tool's limitations, negatively affecting their exercise of prudence. At the ecosystem level, such information asymmetry reduces trust, especially when minors and other vulnerable users are concerned.

As such, on 20 July 2026, the Infocomm Media Development Authority ("IMDA") published its [Transparency Guidelines for Generative AI Chatbots](#) ("**Guidelines**"). The Guidelines, which are voluntary, seek to narrow this information gap for generative AI chatbots as a starting point. They set out how chatbot deployers can provide meaningful transparency to users through a chatbot information card ("**Card**"). The Card functions as a single consolidated reference point, similar to a medical label, where consumers can find essential information about the chatbot. The Guidelines are organised around three key principles of the relevance, accessibility, and timeliness of the information in question.

This Update summarises the scope of the Guidelines, the key principles, and the practical implications for various stakeholders.

Scope of the Guidelines

Application Layer

The Guidelines focus on transparency at the chatbot-application layer, rather than the underlying models layer. This is for two reasons. First, most users experience generative AI through such applications, giving them the widest reach. Second, there are growing societal concerns surrounding the use of such chatbots, including exposing children to inappropriate content, endangering the safety of vulnerable users, and data privacy issues.

External Chatbots

The Guidelines address generative AI chatbots that are external-facing (i.e. those used by customers or the general public), rather than chatbots deployed solely for internal organisational use. External-facing chatbots can fall into one or more of the following categories:

1. **General-purpose chatbots:** These chatbots are designed for open-domain interactions with users that span a wide range of tasks.
2. **Companion chatbots:** These chatbots are designed to facilitate social interaction and emotional engagement with target users.
3. **Domain-specific enterprise chatbots:** These chatbots are specifically designed for particular fields or tasks.

A chatbot may span more than one category. Further, newer chatbots such as agentic assistants may not fit neatly within any category. Deployers should therefore consider the following risk and context factors when determining which of the risks and disclosures in the Guidelines are most relevant to their chatbots:

1. **Usage:** How consequential is the use of the chatbot? For example, does it only answer the user's queries by providing information, or does it go further to take actions on behalf of the user?
2. **Domain:** What field does the chatbot operate in? For example, is this a general-use public chatbot, or one designed for users in a specialised field such as healthcare or finance?
3. **User profile:** Who are the chatbot's main users? For example, is it accessible to minors?

Chatbot Deployers and Upstream Parties

The Guidelines state that chatbot deployers bear primary responsibility for providing transparency to end-users and conducting pre-deployment risk assessments and mitigating measures. However, upstream parties such as model providers should also support responsible deployment by providing deployers with relevant, up-to-date information on the underlying models, enabling them to conduct risk assessments and mitigations and ensure meaningful transparency for end-users.

Key Principles in the Guidelines

Minimum Baseline of Meaningful Transparency

The Guidelines are organised around three key principles: relevance, accessibility, and timeliness. While voluntary, chatbot deployers are strongly encouraged to adopt at least the minimum described as a baseline of meaningful transparency for their users.

Relevance

The Card should contain all the essential information that consumers need to know to use the chatbot responsibly and with confidence.

1. **Identify key information:** The Card should address the following four questions ("**Four Questions**") that users care about most:
 - **What the chatbot can and cannot do:** Including what the chatbot can or cannot be used for, its capabilities, limitations and prohibitions (e.g. age restrictions, out-of-scope uses).

- **How reliable and safe the chatbot is:** Including the common risks posed by such chatbots (e.g. content safety risks, emotional safety risks, age-related safety risks), the mitigatory safeguards that are put in place to manage those risks, the residual risks that users should still take note of regardless, and the recommended precautions that users should adopt for each risk (e.g. fact-checking, limited sharing, parental controls).
 - **How user data will be used and protected:** Including what data is collected, who has access to the data, whether the data is utilised to train models, and the user control options available.
 - **How users can report issues:** Including the right channels for reporting, the types of issues that can be raised, and the acknowledgements and follow-up actions that may be taken.
2. **Provide adequate disclosure:** At a minimum, the Card should address the Four Questions above, and there should be at least one substantive disclosure for each area. A substantive disclosure is a concrete, specific statement that a user can act on. For example, a statement that the *"chatbot can produce factual errors and that its responses should be verified before being relied on for important decisions"* is preferable to a generic statement that *"we take safety seriously"*. Beyond this, deployers can decide what and how much to disclose.
3. **Calibrate extent of disclosure:** The following considerations should be factored in when calibrating what information to publicly disclose:
- **Ensuring proportionality:** Deployers should strike the balance between the usefulness of the information for end-users (in making informed decisions and/or taking meaningful action) against the effort that would be required to obtain or produce the same.
 - **Protecting proprietary information:** Deployers do not have to publish proprietary details related to their chatbots. However, where their safety measures showcase a competitive advantage, they should disclose the protections in place and the effectiveness of the same, without having to provide implementation details.
 - **Safeguarding security:** Deployers should provide purposeful and calibrated disclosure to ensure that users receive sufficient information while withholding critical details that may compromise security or expose the application to exploitation or malicious attacks.
 - **Disclosing only reasonably available information:** Deployers do not have to attest to upstream information that is not reasonably available to them or that they cannot verify, such as the model provider's internal safety testing.

Accessibility

The information in the Card should be presented in a manner that is easy to understand, navigate and find. The Guidelines set out recommended design considerations which deployers can adapt to suit their application interface, platforms and users.

1. **Easy to understand:** Deployers should strike the balance between: (i) reducing the technicality of their statements; and (ii) ensuring that these statements provide sufficiently meaningful detail to enable users to make informed decisions. This can be achieved by being specific and substantive, and using plain language.
2. **Easy to navigate:** The Card should be presented in a manner that is easily digestible. This can be achieved by:
 - offering layered disclosure (e.g. summary information presented upfront with a drop-down menu for further clicking);
 - structuring the content for readability (e.g. using clear headers, bullet points, bolded key terms, and short segmented sections); and

- using visualisations.
3. **Easy to find:** The Card should be easily discoverable, by:
- serving as a reliable one-stop shop or reference point for comprehensive key information;
 - being surfaced at onboarding and reachable from within the chatbot interface; and
 - being standardised across different modalities (e.g. website and mobile application).

At a minimum, upon first use of the chatbot, deployers should present a high-level safety statement with a clearly identifiable link to the Card, and thereafter, the Card should be easily reachable from within the chatbot. Beyond this, deployers can decide how best to design and display this to users.

4. **Consolidated Card:** Deployers who offer a family of closely related chatbot experiences may publish a single, consolidated Card for the family rather than separate Cards for each variant. This can also be applied for chatbots that are deployed across multiple platforms or embedded within different applications. The consolidated Card should document:
- characteristics common to the family (e.g. shared safety measures or data practices);
 - key differences between variants (e.g. features, target user groups or safety settings); and
 - ways for users to identify the variant that they are using (e.g. through product names).

Timeliness

The Card should be published by the time the chatbot is made available to users and kept updated as it evolves with information on the latest capabilities, risks and safety policies. The recommendations in the Guidelines on publishing and updating may be adapted to each deployer's context, provided publication remains timely and updates keep pace with developments affecting users.

1. **At deployment:**
- Deployers should publish their Cards by the time their chatbots are made available to external users, including in beta or pilot form.
 - For phased rollouts, deployers can publish the information that is known at launch and update this as capabilities and risks settle.
 - For chatbots that were publicly accessible before the Guidelines were issued, deployers are encouraged to issue the corresponding Cards as soon as reasonably practicable.
2. **Ongoing updates:**
- The Card should be updated within a reasonable time whenever a new development meaningfully affects users through a change in the chatbot's capabilities or safety profile, whether the development arises from changes to the chatbot or from external factors.
 - Even when no such change has occurred, deployers are encouraged to periodically review their Cards to ensure that the information is kept up to date.
 - The Guidelines provide examples of: (i) potential trigger points when updates should be considered; and (ii) changes that may not require such updates.
3. **Versioning:** Deployers are encouraged to clearly version their Cards to help users to verify that they are reading up-to-date information, including stating: (i) the date the Card was most recently revised; and (ii) the version of the chatbot to which the Card relates, where the chatbot uses version identifiers.

Key Insights

The Guidelines are a welcome step towards ensuring greater accountability, trust and responsibility in AI use. When properly implemented, **consumers** will be empowered to make more informed decisions about their use of chatbots, enabling greater end-user responsibility and ownership. **Deployers** will be able to demonstrate their commitments to safety and reliability, which the public can use as benchmarks to measure compliance. Deployers are encouraged to view the Guidelines as an opportunity to set standards of excellence and enhance public trust, rather than as inhibitions on their operations.

Sectoral regulators can refer to the Guidelines when implementing meaningful transparency for generative AI applications in their respective domains. **At an ecosystem level**, the Guidelines standardise how information is shared publicly and enact a consistent disclosure framework. This empowers consumers to compare applications and locate relevant information efficiently, rather than navigating disparate and inconsistent sources. In turn, a trusted environment is built around meaningful transparency, enabling the public to engage responsibly and confidently with generative AI.

For deployers of newer chatbots such as agentic assistants, it will be useful to augment the reading of the Guidelines with a review of IMDA's May 2026 discussion paper titled ["Legal Responsibility for AI Agents"](#). For more information, please refer to our June 2026 Legal Update titled ["IMDA Issues Discussion Paper on Allocation of Legal Responsibility for AI Agents"](#).

If you have any queries on the above, please reach out to our team set out on this page.

For regional Technology, Media & Telecommunications matters, please see Rajah & Tann Asia's [Technology, Media & Telecommunications Practice](#) for more information.

Contacts

TECHNOLOGY, MEDIA & TELECOMMUNICATIONS

Rajesh Sreenivasan

HEAD

D +65 6232 0751
rajesh.sreenivasan@rajahtann.com

Steve Tan

DEPUTY HEAD

D +65 6232 0786
steve.tan@rajahtann.com

Benjamin Cheong 张汶权

DEPUTY HEAD

D +65 6232 0738
benjamin.cheong@rajahtann.com

Please feel free to also contact Knowledge Management at RTApublications@rajahtann.com.

Regional Contacts

Cambodia

Rajah & Tann Sok & Heng Law Office

T +855 23 963 112 | +855 23 963 113
kh.rajahtannasia.com

China

Rajah & Tann Singapore LLP Representative Offices

Shanghai Representative Office

T +86 21 6120 8818
F +86 21 6120 8820

Shenzhen Representative Office

T +86 755 8898 0230
cn.rajahtannasia.com

Indonesia

Assegaf Hamzah & Partners

Jakarta Office

T +62 21 2555 7800
F +62 21 2555 7899

Surabaya Office

T +62 31 5116 4550
F +62 31 5116 4560
www.ahp.co.id

Lao PDR

Rajah & Tann (Laos) Co., Ltd.

T +856 21 454 239
la.rajahtannasia.com

Malaysia

Christopher & Lee Ong

T +603 2273 1919
F +603 2273 8310
www.christopherleeong.com

Myanmar

Rajah & Tann Myanmar Company Limited

T +951 9253750
mm.rajahtannasia.com

Philippines

Gatmaytan Yap Patacsil Gutierrez & Protacio (C&G Law)

T +632 8248 5250
www.cagatlaw.com

Singapore

Rajah & Tann Singapore LLP

T +65 6535 3600
sg.rajahtannasia.com

Thailand

Rajah & Tann (Thailand) Limited

T +66 2656 1991
th.rajahtannasia.com

Vietnam

Rajah & Tann LCT Lawyers

Ho Chi Minh City Office

T +84 28 3821 2673 | +84 28 3521 2832

Hanoi Office

T +84 24 3267 6127 | +84 24 3267 6128
vn.rajahtannasia.com

Rajah & Tann Asia is a network of legal practices based in Asia.

Member firms are independently constituted and regulated in accordance with relevant local legal requirements. Services provided by a member firm are governed by the terms of engagement between the member firm and the client.

This Update is solely intended to provide general information and does not provide any advice or create any relationship, whether legally binding or otherwise. Rajah & Tann Asia and its member firms do not accept, and fully disclaim, responsibility for any loss or damage which may result from accessing or relying on this Update.

Our Regional Presence



Rajah & Tann Singapore LLP is one of the largest full-service law firms in Singapore, providing high quality advice to an impressive list of clients. We place strong emphasis on promptness, accessibility and reliability in dealing with clients. At the same time, the firm strives towards a practical yet creative approach in dealing with business and commercial problems. As the Singapore member firm of the Lex Mundi Network, we are able to offer access to excellent legal expertise in more than 100 countries.

Rajah & Tann Singapore LLP is part of Rajah & Tann Asia, a network of local law firms in Cambodia, China, Indonesia, Lao PDR, Malaysia, Myanmar, the Philippines, Singapore, Thailand and Vietnam. Our Asian network also includes regional desks focused on Brunei, Japan and South Asia.

The contents of this Update are owned by Rajah & Tann Singapore LLP and subject to copyright protection under the laws of Singapore and, through international treaties, other countries. No part of this Update may be reproduced, licensed, sold, published, transmitted, modified, adapted, publicly displayed, broadcast (including storage in any medium by electronic means whether or not transiently for any purpose save as permitted herein) without the prior written permission of Rajah & Tann Singapore LLP.

Please note also that whilst the information in this Update is correct to the best of our knowledge and belief at the time of writing, it is only intended to provide a general