

TECHNOLOGY, MEDIA & TELECOMMUNICATIONS

Model AI Governance Framework for Agentic AI Updated with Case Studies and Best Practices

Executive Summary

On 20 May 2026, the Infocomm Media Development Authority announced that the Model AI Governance Framework for Agentic AI ("**Agentic MGF**") has been updated to include: (i) more than ten case studies of real-world agentic deployments that provide practical illustrations on how organisations can operationalise different dimensions of the framework's recommendations to meet their needs; and (ii) new best practices in key areas such as multi-agent systems, managing risks from third-party agents, and guarding against automation bias.

For background, the Agentic MGF was first launched in January 2026 as a first-of-its-kind framework for reliable and safe agentic Artificial Intelligence ("**AI**") deployment. It builds upon the governance foundations of the Model Governance Framework for AI and provides guidance to organisations on how to deploy agents responsibly, recommending technical and non-technical measures to mitigate risks while emphasising that humans are ultimately accountable. For more information, please refer to our February 2026 Legal Update titled "[Supporting the Safe Adoption of Agentic AI – Singapore Launches New Model Governance Framework for Agentic AI](#)".

The additions in the [updated Agentic MGF \(version 1.5\)](#), some of which were contributed by Rajah & Tann Singapore's Head of Technology, Media & Telecommunications, [Rajesh Sreenivasan](#), are intended to help more organisations with agentic AI by providing them with examples of how their counterparts have operationalised the Agentic MGF recommendations for their systems, so that they can do the same.

This Update explores some of the key case studies and best practices presented.

Assessing and Bounding the Risks

The Agentic MGF provides guidance that organisations planning to use agentic AI should: (i) determine suitable use cases for agent deployment by considering agent-specific factors that can affect risk; and (ii) make design choices to bound the risks upfront by limiting an agent's access to tools and systems, and establishing a robust identity and permissions framework.

The updated Agentic MGF demonstrates this principle aptly through the case study of Dayos, an enterprise AI automation company headquartered in Singapore with operations in the United States of America. Dayos built an AI-powered ticketing agent, which handles every internal Information Technology request that comes in, and either resolves it automatically or routes it to a human. Dayos uses tiered risk levels to guide and bound the actions taken by its agent. Every type of ticket is assessed for severity of impact, reversibility, and feasibility of human oversight. Pursuant to the assessment, the agent would have a different autonomy level. For example:

1. **For Tier 1 actions** (low severity, fully reversible) such as password resets, this is fully automated by the agent, but its actions are audited biweekly.
2. **For Tier 2 actions** (moderate severity, partially reversible), the agent diagnoses and proposes the fixes, but can only act with a human's approval.
3. **For Tier 3 actions** (high severity, limited reversibility) such as permissions modifications, the agent cannot act in relation to these.

Enabling Meaningful Human Accountability

Challenges arise in ensuring meaningful human accountability for agentic AI because agent actions emerge dynamically and adaptively from interactions (instead of fixed logic). To address these challenges, the Agentic MGF encourages organisations to consider: (i) a clear allocation of responsibilities within and outside the organisation by establishing chains of accountability across the agent value chain and lifecycle, while emphasising adaptive governance; and (ii) measures to enable human oversight, such as requiring human approval at significant checkpoints, and establishing human audits paired with automated monitoring.

Building on this guidance, the updated Agentic MGF identifies the case study of Tencent, a multinational technology company headquartered in the People's Republic of China, responsible for gaming, messaging application WeChat and AI models such as Hy. CodeBuddy is an agentic AI coding system developed by Tencent Cloud and used by its engineers. It can autonomously plan, write, test and deploy code through natural language instructions, with access to filesystems, terminal commands, external Application Programming Interfaces and Model Context Protocol ("MCP") tools. To allow for meaningful human oversight without overly fatiguing the user, it employs a mix of preset secure defaults and configurable permissions. For example:

1. **Defining significant checkpoints for human intervention:** CodeBuddy defines by default which actions need human approval, such as editing files, running shell commands, making network requests, or using external tools.
2. **Enabling humans to effectively evaluate approval requests:** Complex commands are explained in plain English, aiding the user in making informed decisions.
3. **Complementing with automated monitoring:** As an additional safety net, continuous and automated real-time monitoring takes place, and suspicious commands that are identified will still require human approval, even if the command has been pre-approved previously.

Implementing Technical Controls and Processes

The Agentic MGF also urges organisations to consider additional controls during key implementation stages:

1. **Design and development stage:** New and tailored controls should be implemented, and least-privilege access to tools and data should be enforced.
2. **Pre-deployment stage:** It is important to test for new dimensions at different levels and across varied datasets to capture the full spectrum of agent behaviour.
3. **Deployment stage:** Agents should be gradually rolled out and thereafter monitored continuously, on a real-time basis, in production.

For guidance on the real-world implementation of this principle and best practice ideas, users can refer to GovTech in the updated Agentic MGF. GovTech is a Singapore statutory board that develops digital government services and drives public sector transformation. It rolled out agentic coding assistants within the government, with a phased approach to rollout to allow for incremental monitoring of risks while preparing controls for new features. In the first phase, the tool was limited to GovTech's employees and low-risk systems only, and external tools (MCP servers) were not allowed. This approach presented a number of advantages. For example:

1. **Limiting damage:** The potential damage was limited in the event that anything went wrong.
2. **Building safeguards:** During this time, GovTech was able to: (i) build the necessary safeguards for a wider rollout, such as central logging, monitoring, and a framework to safely connect to approved external tools; and (ii) test the system against potential attacks to ensure the effectiveness of its guardrails.
3. **Improving rollout:** Lessons from the first phase helped to improve the rollout, including by fixing early technical issues, reducing the cognitive load to be borne by human approvers, and making the overall setup easier to adopt.

Enabling End-User Responsibility

Finally, human accountability extends to the end users who use and rely on agents as well, who should be provided with sufficient information to promote trust and enable responsible use. Thus, the Agentic MGF guides organisations to consider: (i) transparency, that is, informing users of the agents' capabilities and the contact points for escalation of issues if the agents malfunction; and (ii) educating users on the proper use and oversight of agents, the potential loss of trade craft as agents take over more functions, and the need for sufficient training to ensure that humans retain their core skills.

The updated Agentic MGF illustrates this well through the case study of Workday, a global enterprise AI platform for managing people, money and AI agents. Workday implemented AI agents to streamline its internal financial and human resources operations, including: (i) Recruiter Agent, which supports screening and evaluation of candidates for job roles; and (ii) Conversational Scheduling Agent, which manages the full interview, orientation, and event session scheduling lifecycle. To enable end-user responsibility, Workday informed users of:

1. **The agent's identity and range of actions:** Workday's interface ensured that users who interacted with its AI agents through platforms like Slack or Microsoft Teams were informed upfront that they were engaging with an AI-powered tool.

2. **Agent's reasoning:** Workday ensured that the agent's reasoning was visible to users so that they were equipped with relevant information to review and use the results in a responsible manner.

Key Insights

The updated Agentic MGF represents a significant step forward for organisations seeking to deploy agentic AI safely and responsibly. Drawing from successful real-world implementations offered by other organisations, the framework is now further enhanced with actionable, practical guidance that can be adapted to a variety of industry contexts. Businesses can leverage these case studies to benchmark their own AI governance practices, and tailor the best practices identified to their own specific regulatory environments and operational realities.

Critically, early engagement with the Agentic MGF and its updated case studies and best practices is strongly encouraged. With Singapore set to drive AI adoption across businesses of all sizes in key industries as part of the National AI Strategy, organisations that proactively adopt robust governance frameworks will be better positioned to manage emerging risks, build trust with users and regulators, and capitalise on the competitive advantages that responsible AI deployment can offer.

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 [Regional Technology, Media & Telecommunications Practice](#) for more information.

Contacts

TECHNOLOGY, MEDIA & TELECOMMUNICATIONS

Rajesh Sreenivasan

HEAD

D +65 6232 0751
rajesh@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