



ALLIANCE AGAINST ISLAMOPHOBIA (AUSTRALIA)

Submission by the Alliance Against Islamophobia (Australia) on Modernising and Harmonising Classification Standards

Addressing Vilification, Dehumanisation and Harm in Australia's Classification Framework

Public Submission

25 April 2026

TABLE OF CONTENTS

Introduction	3
Executive Summary	4
Part I – Core Reform Issues and AAI’s Substantive Position	5
1. <i>The Current Framework Does Not Adequately Address Vilification Harms</i>	5
2. <i>Support for the Harm-Informed Direction in the Report, With an Important Expansion</i>	6
3. <i>Adoption of a Clear Concept of Dehumanising Material</i>	7
4. <i>Support for Clearer Content Descriptors and Greater Transparency</i>	7
5. <i>Cultural and Community Expertise Must Be Built Into the System</i>	8
6. <i>Online and Streaming Environments Must Be Addressed</i>	8
7. <i>Caution on Self-Classification and Automation</i>	9
8. <i>Complaint and Review Mechanisms Should Be Strengthened</i>	9
Part II – AAI Response to Selected Reform Concepts	10
9. <i>AAI’s Position on Selected Reform Concepts</i>	10
10. <i>AAI Response to Selected Reform Concepts</i>	11
C02 – Overhaul of the classifiable element “Themes”	11
C03 – Improved reliance on evidence over morality	11
C07 – Clarify definitions	12
C09 – Classification Advisory Council	12
C18 – Classification website: transparent decisions	12
C20 – Element icons	13
Conclusion	14

Introduction

The Alliance Against Islamophobia (AAI) welcomes the opportunity to make this submission on the proposed modernisation and harmonisation of Australia's classification standards.

AAI supports a classification framework that helps people make informed choices, protects children and vulnerable communities from harm, and remains consistent with Australia's multicultural and anti-racist values.

We also support a framework that is practical, transparent, and capable of operating effectively across contemporary media environments, including streaming services, digital platforms, and other online channels.

AAI's interest in this consultation arises from long-standing concerns about the circulation in Australia of films and related media content that depict Muslims through dehumanising, conspiratorial, or exclusionary narratives.

These concerns arise in part from a discernible pattern in some recent Bollywood, wider Indian cinema, and related media content that has depicted Muslims through dehumanising, conspiratorial, exclusionary, or threat-based narratives. AAI's concern is not with Indian cinema as a whole, but with this subset of content and the harms it can generate.

These issues are further complicated by the fact that some of the relevant narratives are culturally coded and communicated through specific historical, religious, linguistic, and political references.

In the Australian context, those constructs may not always be readily understood within existing classification settings, particularly where cultural and linguistic barriers make it harder to identify dehumanising or vilifying meaning and its likely impact on diaspora communities.

In AAI's experience, and as reflected in AAI's earlier submissions, such content can contribute to hostility, fear, and social exclusion in diaspora settings, including Australia.

Current classification settings do not adequately address vilification, discrimination, or the public-health and social-cohesion harms that can flow from content that normalises hatred towards protected groups.

We therefore support many of the reform directions identified in the Social Research Centre and Mettlesome report. However, we respectfully submit that the present reform process should go further by making vilification, dehumanisation, and incitement of hatred explicit classification considerations rather than merely matters of consumer information.

Executive Summary

AAI supports reform of the classification framework, but submits that reform should do more than modernise language and improve usability. It should also address a substantive policy gap.

In particular:

1. Vilification, discrimination, dehumanisation, and incitement of hatred should be made express classification considerations.
2. The revised “Themes” element should not stop at “discrimination”, but should expressly encompass dehumanising and vilifying portrayals of protected groups.
3. Classification should remain evidence-based and harm-informed, but harms to dignity, belonging, equality, and community safety must be recognised as real harms.
4. A community-led Classification Advisory Council should include representatives from communities with lived experience of racism, religious hatred, and cultural harm.
5. Greater transparency is needed through clearer descriptors, stronger public reasons, QR-linked title information, and accessible explanations of decisions.
6. Online, streaming, and platform-distributed content should not be allowed to evade effective oversight.
7. Any move towards self-classification or automated classification must be accompanied by strong safeguards, auditing, appeals, and the capacity for prompt correction.
8. Complaint and review pathways should be simplified and affordable, including fee waivers where matters raise public-interest harms to vulnerable communities.
9. The revised classification standards should operate coherently with relevant state and territory racial and religious vilification protections, so that harmful material is not assessed under disconnected frameworks.

Part I – Core Reform Issues and AAI’s Substantive Position

This Part explains AAI’s broader policy position in relation to the present consultation. It identifies the principal shortcomings in the current framework and the substantive reforms AAI considers necessary to ensure that the classification system is modern, transparent, harm-informed, and responsive to vilification and dehumanising content.

Part II then addresses selected reform concepts from the consultation materials more directly.

1. The Current Framework Does Not Adequately Address Vilification Harms

AAI’s central concern is that the present framework does not squarely address content that dehumanises or vilifies people on the basis of religion, race, ethnicity, or other protected characteristics.

The issue is not whether controversial political or historical material should be censored. It is whether Australia’s classification system should continue to treat content depicting protected communities as inherently violent, deceitful, polluting, primitive, or civilisational threats as though those harms are outside classification concern.

In AAI’s experience, this gap means that content capable of worsening discrimination, fear, humiliation, and social exclusion can still reach Australian audiences without adequate classification reasoning, without meaningful warnings, and without a clear mechanism for early intervention.

AAI submits that this is not a minor omission. It is a structural gap.

For example, content may repeatedly depict a protected group through threat-based, conspiratorial, dehumanising, or culturally coded patterns, yet still pass through the current framework without any explicit treatment of vilification or dehumanisation as a classification issue, particularly where cultural and linguistic barriers make the underlying meaning and likely diaspora impact less readily visible.

2. Support for the Harm-Informed Direction in the Report, With an Important Expansion

AAI supports the report's shift towards clearer, evidence-based and harm-informed classification, and agrees that the framework should move away from narrow moralism. We also support the proposal to improve how "Themes" are understood and communicated.

However, the current proposal should be strengthened in two ways.

First, the proposed inclusion of "discrimination" within Themes is welcome, but too narrow on its own. The framework should expressly recognise:

- vilification;
- dehumanisation;
- incitement or encouragement of hatred against protected groups; and
- coded or conspiratorial narratives that portray protected groups as existential threats.

Second, the revised guidelines should make clear that these harms are relevant both to classification outcome and to consumer advice.

In other words, harmful portrayals should not merely trigger better labelling; in serious cases, they should affect the classification itself.

These are not abstract concerns, but harms consistent with public-health and social-cohesion research on dehumanising and vilifying portrayals, including where transnationally circulating narratives echo early-stage genocidal propaganda and normalise genocidal framing about Muslims within diaspora settings.

That concern is reflected in a discernible pattern in some recent Indian films and related media content, including *The Kerala Story*, *The Kashmir Files* and *Chhaava*, which have been associated with dehumanising portrayals of Muslims, communal polarisation, and broader climates of hostility.

3. Adoption of a Clear Concept of Dehumanising Material

AAI recommends that the revised framework adopt a clear and workable concept of dehumanising material.

A useful approach is to recognise content as dehumanising where an ordinary person would conclude that it portrays a class of persons identified by a protected characteristic as lacking qualities intrinsic to humans, or as inherently dangerous, polluting, inferior, deceitful, or fit for exclusion from the community.

This kind of formulation would help decision-makers distinguish between:

- criticism of governments, political movements, or states, which must remain protected; and
- content that targets people as a group in a way that strips them of equal dignity and moral worth.

That distinction is essential if the system is to remain both rights-respecting and effective.

4. Support for Clearer Content Descriptors and Greater Transparency

AAI strongly supports proposals aimed at making classification information more meaningful and more accessible to the public.

We support:

- clearer sub-categories within Themes;
- icons for high-concern content areas;
- stronger explanations of why a classification was given;
- QR-code or similar quick-access links to fuller title information;
- improved title listings and a more usable classification website; and
- public education to improve understanding of what classifications do and do not mean.

For AAI, these measures are particularly important because affected communities often report confusion and distrust when content that appears plainly hostile or degrading to them is classified without any visible acknowledgement of that issue.

If the system is to remain trusted, people must be able to see not only the outcome, but the reasoning.

5. Cultural and Community Expertise Must Be Built Into the System

AAI supports the proposal for a community-led Classification Advisory Council.

That body should include representatives with lived experience and subject-matter expertise in:

- racism and religious vilification;
- anti-Muslim hatred and other forms of group-based dehumanisation;
- multicultural community safety;
- child and adolescent wellbeing;
- trauma-informed practice; and
- digital and platform-based hate propagation.

The Advisory Council should not be symbolic. It should have a real role in informing guideline development, reviewing emerging harms, and providing advice on contested or evolving forms of content.

The Council should also operate in a way that supports coherence with adjacent regulatory and safety settings, including online-content and communications regulators, so that emerging harms, including transnationally circulating extremist narratives associated with far-right Hindu extremist groups, are not addressed in institutional silos.

AAI also supports multilingual and accessible classification information so that culturally and linguistically diverse communities can meaningfully use the system.

6. Online and Streaming Environments Must Be Addressed

AAI supports harmonisation across media and agrees that fragmentation in the current system creates confusion and leaves obvious gaps.

For communities affected by transnational hate narratives, one of the most serious problems is that harmful content now travels easily across cinemas, streaming platforms, social media, short-form video, online news environments, and networked sharing systems. A framework that is modern in form but still narrow in jurisdiction will not meet the problem.

AAI therefore supports stronger coherence between the classification scheme and the online environment. Content should not avoid scrutiny simply because it reaches audiences through a different technological route.

7. Caution on Self-Classification and Automation

AAI recognises the case for efficiency, including classify-once arrangements, classification updates where content changes, and better digital tools.

However, AAI urges caution about any broad move towards self-classification as the default, or an expanded reliance on automation, where the content in question may involve racism, vilification, discrimination, or coded dehumanising narratives.

Such material is often subtle, contextual, and culturally coded. It can be missed by systems designed primarily for obvious content elements such as violence, sex, or language.

If self-classification and automation are expanded, there should be:

- robust auditing and quality assurance;
- strong independent oversight;
- community-triggered review pathways;
- transparent correction mechanisms; and
- special escalation processes for high-risk or contested content.

8. Complaint and Review Mechanisms Should Be Strengthened

AAI supports more accessible complaint, review, and reconsideration processes.

Where a community organisation raises a credible public-interest concern that content vilifies or dehumanises a protected group, the pathway for review should be simple, timely, and affordable. Fee waivers should be available where there is a clear public-interest rationale or risk of harm to vulnerable communities.

A preventive system is preferable to one that leaves affected communities to pursue legal complaints only after harm has already occurred.

Part II – AAI Response to Selected Reform Concepts

9. AAI's Position on Selected Reform Concepts

AAI's broad position on the report's reform directions is as follows:

Supported, with strong support

- clearer, harm-informed descriptors;
- overhaul of Themes;
- improved definitions;
- Classification Advisory Council;
- transparent reasons and website improvements;
- element icons;
- quick links to fuller classification information;
- public education.

Supported in principle, but requiring safeguards

- single-touch classification;
- updates to classifications;
- greater use of digital tools;
- expanded automated classification tools;
- any move towards self-classification.

Neutral or secondary from AAI's perspective

- PG13 and MA16+ reforms;
- alignment of publication labels with other media;
- other usability changes not directly related to vilification harms.

These latter issues may improve clarity, but they should not displace the more urgent need to address discriminatory and dehumanising content.

10. AAI Response to Selected Reform Concepts

AAI considers it useful to respond expressly to several of the reform concepts identified in the consultation materials.

C02 - Overhaul of the classifiable element “Themes”

AAI supports reform of the Themes element. However, the proposed inclusion of “discrimination” should be expanded so that the framework also expressly captures vilification, dehumanisation, incitement or encouragement of hatred, and coded conspiratorial portrayals of protected groups as existential threats.

For AAI, this is not a semantic issue. A classification framework that refers only to “discrimination” risks understating the seriousness of media content that depicts protected communities as inherently violent, deceitful, polluting, primitive, disloyal, or unworthy of equal belonging.

Where such portrayals are present, the system should be able to identify them clearly and respond to them meaningfully.

C03 - Improved reliance on evidence over morality

AAI strongly supports a move away from moralistic or culturally narrow approaches to classification. A harm-informed, evidence-based framework is preferable.

However, the evidence base should not be confined to physical or immediately visible harms alone.

The framework should recognise evidence of broader harms associated with dehumanising and vilifying portrayals, including humiliation, fear, social exclusion, reinforcement of prejudice, erosion of belonging, increased hostility towards targeted communities, and damage to social cohesion.

In that sense, C03 should be understood as supporting a serious account of cultural, psychological, and public-health harms, not excluding them.

C07 - Clarify definitions

AAI strongly supports clearer and better-informed definitions.

This concept is central to the present consultation. One of the key difficulties under the current framework is that harmful portrayals of protected groups are often not assessed through clear concepts that decision-makers can apply consistently.

The revised guidelines should therefore define, with sufficient clarity, matters such as vilification, dehumanisation, incitement of hatred, and discriminatory or exclusionary portrayals of protected groups.

AAI also submits that the guidelines should draw a careful distinction between criticism of governments, political ideologies, or states, which must remain protected, and content that attacks people as a class in a way that strips them of equal human dignity.

C09 - Classification Advisory Council

AAI supports the creation of a Classification Advisory Council.

That Council should include community members with lived experience and subject-matter expertise relevant to racism, religious vilification, anti-Muslim hatred, multicultural community safety, trauma-informed practice, and transnationally circulating hate narratives.

It should also be structured so that consultation with affected communities is real and ongoing, not occasional or symbolic.

A properly constituted Council would improve legitimacy, responsiveness, and trust in the classification system.

C18 - Classification website: transparent decisions

AAI strongly supports greater transparency in classification decisions.

Where communities experience content as plainly hostile, degrading, or dehumanising, yet the public record offers little or no visible explanation of how those concerns were considered, distrust naturally follows. Clearer public reasons, or accessible summaries of decision-making logic, would strengthen confidence in the system and help affected communities understand why a title received a particular classification.

Transparency is especially important in contested matters, or where harmful portrayals are subtle, coded, or context-dependent.

C20 - Element icons

AAI supports the introduction of element icons and more visible content indicators.

Icons are not a substitute for stronger classification standards, but they are an important consumer-information tool. They may be especially useful where content includes sexual violence, suicide, discriminatory themes, or other issues of significant community concern. If implemented well, icons would improve clarity, support informed viewing choices, and make the system more accessible and intuitive.

AAI submits, however, that icons relating to discrimination or high-concern thematic harms should be designed carefully so that they communicate genuine risk without trivialising the seriousness of vilifying or dehumanising content.

Conclusion

AAI welcomes the present consultation and supports reform that makes Australia's classification framework more modern, transparent, and useful.

But modernisation should not be understood only as simplification, digitisation, or administrative efficiency. It should also mean that the framework can recognise and respond to the real harms caused by content that vilifies, dehumanises, and normalises hatred towards protected groups.

AAI therefore urges the Government to ensure that the revised classification standards:

- expressly address vilification, discrimination, dehumanisation, and incitement of hatred;
- embed lived community expertise into the ongoing framework;
- improve public transparency and accessibility; and
- preserve strong independent safeguards, especially where automation or self-classification is expanded.

A classification system that reflects Australia's contemporary values should help people make informed choices, protect children, respect freedom of expression, and recognise that the dignity and safety of marginalised communities are central, not peripheral, to the public interest.