

# Drawing the Line: Towards a Coherent, Harm-Based and Rights-Consistent Classification Framework

29 March 2026 – Center for Online Safety and Liberty

## 1. Executive Summary

The Center for Online Safety and Liberty (COSL) welcomes the opportunity to contribute to the consultation on modernising and harmonising Australia’s classification guidelines. This review is both timely and necessary. As the consultation paper recognises, Australia’s classification system remains widely trusted, but is increasingly strained by a media environment defined by global platforms, user-generated content, and rapid technological change .

We strongly support the consultation’s central direction: a shift towards a more evidence-based, harm-informed, and consistent classification framework. In particular, we endorse the move away from moral judgement and toward a clearer articulation of the risks that different types of content may pose to individuals and communities . This shift reflects both contemporary community expectations and longstanding recommendations for reform.

However, this submission emphasises a critical limitation in the current reform approach. While the consultation focuses on updating classification guidelines, in practice these guidelines do not operate in isolation. They are deeply interconnected with other regulatory and legal frameworks, including the Criminal Code, the Customs Act, and the Online Safety Act. These frameworks collectively determine not only how content is classified, but how it is restricted, removed, or criminalised.

As a result, classification reform that is not accompanied by greater coherence across these systems risks entrenching inconsistency rather than resolving it. A system that classifies content according to one set of principles, but regulates or criminalises it according to another, cannot deliver the clarity, fairness, or public trust that this consultation seeks to achieve.

Drawing on COSL’s *Drawing the Line* project, this submission proposes that the central task of modern classification policy is to clearly distinguish between:

1. content that reflects or risks real-world harm, and
2. content that constitutes fictional, simulated, or consensual expression.

Maintaining this distinction is essential both for effective harm prevention and for the protection of lawful expression. When these categories are collapsed, classification systems risk overreaching into areas of expression that do not cause harm, while simultaneously undermining the credibility and focus of efforts to address genuine abuse.

## 1.1 Table of Contents

1. Executive Summary.....	1
1.1 Table of Contents.....	2
1.2 Summary of Recommendations.....	2
2. Introduction.....	3
2.1 About COSL.....	3
2.2 Overview.....	3
3. Drawing the Line: A Coherent Framework for Modern Classification.....	4
4. Consistency and Harmonisation: The Missing Principle.....	5
5. Applying a Harm-Based Approach: Fiction, Sexual Content, and Evidence.....	6
5.1 Fiction, Representation, and Harm.....	6
5.2 The Role of Evidence.....	7
5.3 Consistency in Application.....	7
5.4 Policy Implications.....	8
6. Selective Response to Proposed Reforms.....	8
6.1 Strongly Supported Reforms.....	8
6.2 Reforms Supported with Caution.....	9
6.3 Key Gap in the Proposed Reforms.....	9
7. Implementation: Embedding Evidence, Accountability and Coherence.....	10
7.1 Embedding Evidence in Classification Practice.....	10
7.2 Strengthening the Role of the Advisory Council.....	10
7.3 Transparency and Accountability Mechanisms.....	10
7.4 Ensuring System-Wide Coherence.....	11
8. Conclusion: Drawing the Line Clearly.....	11
Appendix A: Mapping to Consultation Themes.....	12
References.....	12

## 1.2 Summary of Recommendations

This submission advances the following key recommendations:

1. **Adopt a clearly defined, harm-based standard** for classification decisions, grounded in evidence rather than moral judgement or assumed offence.
2. **Ensure consistency across media types and content categories**, avoiding differential treatment based on format, genre, or cultural origin.
3. **Explicitly recognise the distinction between real-world harm and fictional or simulated depictions**, and reflect this distinction in classification criteria.
4. **Establish institutional mechanisms for ongoing evidence review**, including a strengthened advisory body with a research mandate.
5. **Acknowledge and address the interdependence between classification guidelines and broader legal frameworks**, to ensure system-wide coherence.

Taken together, these steps would support the development of a classification system that is not only modernised in form, but coherent in principle — one that protects against harm without extending into unnecessary or counterproductive restrictions on lawful expression.

## 2. Introduction

### 2.1 About COSL

The Center for Online Safety and Liberty (COSL) is a nonprofit dedicated to empowering individuals and communities to thrive online by building safer digital spaces, fostering creativity, combating harm, and championing digital rights. COSL serves as an incubator for independent projects that tackle pressing issues such as age verification mandates, Section 230 rollbacks, encryption battles, and content-scanning overreach, while also developing open source trust-and-safety tools and nurturing inclusive online communities.

### 2.2 Overview

The consultation paper identifies a clear and widely supported direction for reform: classification decisions should be grounded in evidence and potential harm, rather than moral judgement. COSL strongly endorses this shift. It represents a necessary evolution of Australia’s classification system in response to both changing media environments and evolving community expectations.

Historically, elements of the classification framework have relied on concepts such as “offensiveness,” “community standards,” and what a “reasonable adult” might find acceptable. While these concepts were intended to provide a proxy for societal values, they are inherently subjective and difficult to apply consistently. In practice, they can obscure rather than clarify the rationale for classification decisions, and may embed unexamined assumptions or biases.

A harm-based approach offers a more coherent and transparent alternative. Rather than asking whether content is offensive or morally objectionable, classification decisions should ask:

1. whether the content poses a risk of harm to identifiable individuals or communities, and
2. whether that risk is supported by credible empirical evidence.

This approach is consistent with arguments advanced in a previous submission to the Stage 2 reforms process (Malcolm, 2024), which identified the limitations of morality-based standards and the need for a clearer harm-based framework. The present submission builds on that foundation, applying those principles to the specific reform proposals under consideration.

This approach aligns with broader trends in public policy, including in areas such as public health and online safety, where evidence-based frameworks are increasingly used to assess risk and guide intervention. It also better reflects the stated purpose of the classification system: to inform consumers and protect against harm, rather than to enforce moral norms.

Importantly, a harm-based approach does not require abandoning community values. Rather, it reframes them. Community concerns — for example, about violence, sexual violence, or discrimination — can and should inform classification decisions. But they should do so through an evidence-based assessment of impact, not through assumptions about offence or propriety.

Adopting this approach would also improve consistency across different types of content. Under a moral or offence-based framework, similar content may be treated differently depending on context, presentation, or cultural framing. A harm-based framework, by contrast, provides a more stable basis for comparison, enabling classification decisions to be applied more uniformly across films, games, publications, and emerging forms of media.

For these reasons, COSL supports the proposed reform direction and encourages the development of clear, operational definitions of harm and evidentiary thresholds to guide classification decisions. Without such definitions, there is a risk that the language of “harm” could replicate the ambiguity of earlier moral standards, rather than replacing them.

### **3. Drawing the Line: A Coherent Framework for Modern Classification**

A central challenge for modern classification systems is not simply identifying harmful content, but distinguishing between fundamentally different categories of content in a consistent and principled way.

COSL’s *Drawing the Line Watchlist* (2025) was developed to address this challenge. It examines how legal and platform governance frameworks distinguish — or fail to distinguish — between content that records or causes concrete harm to real children, and content that is purely fictional, artistic, or imaginative. The Watchlist evaluates jurisdictions based on how clearly they distinguish between real-world harm and fictional or expressive content, and identifies Australia as exhibiting areas of concern in this regard.

The Watchlist’s findings highlight a recurring problem: systems that blur these distinctions tend to produce outcomes that are both overinclusive and underinclusive. They may restrict or stigmatise forms of expression that do not cause harm, while simultaneously diluting the focus on content that does.

Australia’s classification framework has historically navigated these distinctions with varying degrees of clarity. The current consultation presents an opportunity to articulate them more explicitly and apply them more consistently.

This is particularly important in a digital media environment where:

1. content formats are increasingly diverse and hybrid;
2. fictional and interactive media are more prevalent; and
3. global distribution exposes Australian audiences to a wide range of cultural contexts and storytelling traditions.

In this environment, classification systems must be capable of distinguishing between representation and harm. The depiction of a harmful act is not, in itself, equivalent to causing harm. Nor does the presence of sensitive or confronting themes necessarily indicate that content should be restricted beyond what is necessary to inform consumers.

The *Drawing the Line Watchlist* recommends that laws and policies such as classification frameworks should:

1. Codify a clear distinction between content that directly harms real victims, such as deepfake pornography and child sexual abuse material (CSAM), and fictional works that do not.
2. Never use the same terminology to classify real abuse materials and fictional works.
3. Assign distinct enforcement responsibility for abuse materials and fictional works.
4. Apply proportionate, harm-based penalties, reserving criminal sanctions for offences involving demonstrable harm to real victims.
5. Maintain separate statistical reporting so that crimes of real sexual abuse are not intermixed with obscenity crimes.

This approach is consistent with the consultation's emphasis on evidence-based decision-making and provides a foundation for more coherent and transparent classification outcomes.

It also aligns with the broader objective of maintaining public trust. A classification system that is perceived as overreaching, inconsistent, or disconnected from evidence risks losing legitimacy. By contrast, a system that clearly articulates its principles and applies them consistently is more likely to be understood, respected, and relied upon by both consumers and industry.

#### **4. Consistency and Harmonisation: The Missing Principle**

A central objective of the consultation is to harmonise classification criteria across media types, reducing fragmentation and inconsistency. COSL strongly supports this goal. In a media environment where content flows seamlessly across formats and platforms, differences in classification standards based on medium alone are increasingly difficult to justify.

However, the consultation's focus on harmonisation is incomplete. While it addresses inconsistencies between films, games, and publications, it does not sufficiently engage with the broader regulatory ecosystem in which classification operates.

In practice, classification outcomes are not confined to consumer advice. They interact directly with:

1. criminal law, which determines when content may give rise to offences;
2. customs regulation, which governs the importation of material; and
3. online safety frameworks, which enable content removal and access restrictions.

These systems are not merely adjacent to classification — they rely on and reinforce its categories and assumptions. As a result, inconsistencies between them can produce outcomes that are confusing, disproportionate, or difficult to justify.

For example, content that is treated as a matter of consumer advice within the classification framework may, in other contexts, be subject to restriction or enforcement under different standards. Where these standards are not aligned, the result is a system that lacks internal coherence.

This has two consequences.

First, it undermines clarity for users and industry. A classification label is intended to provide meaningful guidance about how content will be treated. If that treatment varies depending on the legal or regulatory pathway engaged, the label loses much of its practical value.

Second, it risks overreach or inconsistency in outcomes. A system that applies harm-based reasoning in one domain but relies on broader or less clearly defined criteria in another may inadvertently extend restrictions beyond what is necessary to address actual harm.

The practical consequences of this misalignment are not hypothetical. In February 2026, a female author was convicted of publishing child abuse material in the form of a novel, aimed at adults and replete with content warnings, which described a consensual ageplay relationship between two adults (AAP 2026). This work of fiction, if assessed within a classification framework, would ordinarily be evaluated by reference to context, audience, and impact, rather than being characterised as inherently harmful.

This case illustrates that, even if classification guidelines are updated to adopt a more evidence-based and harm-informed approach, their effect may be limited if parallel legal frameworks continue to apply broader or inconsistent standards. In practice, criminal law may override or pre-empt classification-based reasoning, leading to outcomes that are difficult to reconcile with the objectives of reform.

This inconsistency has also been documented in recent scholarship. Remminga (2026) observes that Australia's obscenity laws, particularly under the Customs Act 1901, are applied unevenly across state and federal jurisdictions, resulting in contradictory and difficult-to-predict outcomes for creators and distributors of erotic and queer media.

For this reason, COSL recommends that the principle of harmonisation be understood not only as consistency across media, but as consistency across the broader system of regulation in which classification operates.

This does not require this consultation to resolve all issues across criminal, customs, and online safety law. However, it does require:

1. explicit recognition of these interdependencies; and
2. the development of classification guidelines that are capable of operating coherently within that wider framework.

Without this, reforms to the guidelines risk being only partial, addressing surface inconsistencies while leaving deeper structural tensions unresolved.

## **5. Applying a Harm-Based Approach: Fiction, Sexual Content, and Evidence**

The transition to a harm-based classification framework has important implications for how certain categories of content are assessed. In particular, it requires greater clarity in distinguishing between content that causes or risks harm, and content that merely depicts or represents harmful acts in a fictional or simulated context.

This distinction is not abstract. It has direct implications for the consistency, proportionality, and credibility of classification outcomes.

## **5.1 Fiction, Representation, and Harm**

A harm-based framework must begin from a simple premise: the depiction of harm is not equivalent to harm itself.

Fictional, narrative, and artistic works routinely depict violence, abuse, or other forms of wrongdoing. These depictions may be confronting or controversial, but they serve a wide range of purposes — including storytelling, critique, education, and personal expression.

To treat all such depictions as inherently harmful, or as equivalent to real-world conduct, risks collapsing important distinctions. It also creates pressure to apply categorical restrictions based on subject matter alone, rather than on an assessment of actual impact.

A more coherent approach is to evaluate:

1. the context of the depiction;
2. its degree of realism or simulation; and
3. any evidence of harm arising from its availability or use.

This approach is consistent with the consultation's emphasis on evidence-based decision-making and would support more nuanced and proportionate classification outcomes.

## **5.2 The Role of Evidence**

A harm-based framework depends on a robust evidentiary foundation. Where claims are made that particular categories of content contribute to harmful behaviour, those claims should be assessed against available empirical research.

Recent research, including Lievesley & Harper (2026), indicates that the relationship between fictional or simulated depictions of harmful acts and real-world offending is, at most, limited — and may in many cases be neutral (i.e. a null effect). This finding is consistent with earlier empirical work. For example, Paul and Linz (2008) found no evidence that exposure to virtual depictions of child sexual material increased acceptance of abusive behaviour, while meta-analytic research on pornography more broadly has found no clear causal link between consumption and sexual aggression (Ferguson & Hartley, 2022). While this remains an evolving field, the current state of evidence does not support broad assumptions that fictional content inherently contributes to harm.

This has two implications for classification policy. First, it underscores the need to avoid precautionary assumptions that are not evidence-based. While precaution may be appropriate in some contexts, it should not substitute for analysis, particularly where it may lead to disproportionate restrictions.

Second, it highlights the importance of ongoing research and review. As media formats evolve, so too must the evidence base used to assess their impacts. Classification frameworks should therefore include mechanisms for regularly incorporating new research into policy and practice.

## **5.3 Consistency in Application**

A further challenge arises from inconsistent treatment of similar content across contexts. Previous consultations and policy discussions such as the Stevens report (2020:72, 86) have,

for example, drawn particular attention to specific genres or cultural forms — including Japanese manga art — as requiring distinct or heightened treatment.

Such approaches risk introducing differential standards that are not grounded in harm, but in assumptions about form, origin, or audience. Where similar themes or depictions are treated differently depending on their format or cultural context, the result is a classification system that appears arbitrary or uneven.

A harm-based framework provides a means of addressing this. By focusing on impact, context, and evidence, rather than on format or genre, classification decisions can be applied more consistently across different types of media.

## **5.4 Policy Implications**

Taken together, these considerations suggest that a modern classification framework should:

1. Distinguish clearly between real-world harm and fictional or simulated depictions;
2. Assess content based on context, realism, and evidence of harm, rather than subject matter alone; and
3. Avoid differential treatment based on format, genre, or cultural origin, unless supported by evidence.

This does not preclude the classification or restriction of content that is genuinely harmful. Rather, it ensures that such measures are targeted, proportionate, and justified.

In doing so, it also supports the broader objective of maintaining public trust in the classification system, by ensuring that its decisions are transparent, consistent, and grounded in defensible principles.

## **6. Selective Response to Proposed Reforms**

COSL supports many of the proposed reforms outlined in the consultation paper, particularly those that advance clarity, consistency, and evidence-based decision-making. At the same time, several proposals would benefit from further refinement to ensure they operate effectively within a harm-based framework.

### **6.1 Strongly Supported Reforms**

COSL strongly supports the following proposals:

#### **Evidence-based classification (Co3)**

The proposed shift toward balancing community concern with established evidence of harm is a critical reform. As discussed above, this principle should be further developed through clear definitions and evidentiary standards to guide its application.

#### **Consistent treatment of sexual content (Co8)**

The proposal to assess sexual content consistently, regardless of identity or context, represents an important step toward reducing bias and improving fairness. This aligns with a harm-based approach and should be implemented with clear guidance.

### **Classification Advisory Council (C09)**

The establishment of a community-informed advisory body is a valuable initiative. To be effective, this body should:

1. include diverse perspectives, including those with lived experience; and
2. be empowered to commission and review empirical research on the impacts of different types of content.

### **Improved transparency and consumer advice (C18–C21)**

Measures to improve the clarity and accessibility of classification information — including clearer explanations, icons, and digital tools — are likely to enhance public understanding and trust.

## **6.2 Reforms Supported with Caution**

Several proposals are directionally positive but raise implementation risks.

### **Themes overhaul (C02)**

Refocusing “Themes” on areas of community concern may improve clarity. However, there is a risk that broadly defined categories (such as discrimination or suicide) could reintroduce moral or subjective judgements unless accompanied by clear, harm-based criteria.

### **Icons and content warnings (C20)**

Enhanced signalling of content elements can support informed choice. At the same time, there is a risk of signal inflation or unintended stigmatisation if warnings are applied too broadly or without clear thresholds.

### **Self-classification and automation (C12–C14)**

Expanding self-classification and automated tools may improve efficiency. However, these systems must be:

1. transparent in their operation;
2. subject to audit and oversight; and
3. grounded in clearly defined standards.

Without these safeguards, there is a risk that automated systems may replicate or amplify existing biases.

## **6.3 Key Gap in the Proposed Reforms**

While the consultation advances important reforms, it does not fully address a central issue: the absence of an explicit distinction between content that involves real-world harm and content that constitutes fictional or simulated representation.

Without this distinction, a harm-based framework remains incomplete. It risks:

1. inconsistent application across content types;
2. overinclusive classification decisions; and
3. reduced clarity for both users and industry.

In addition, the consultation does not sufficiently address the relationship between classification guidelines and other legal frameworks. As discussed in Section 4, these interdependencies are critical to the effective operation of the system.

Addressing these gaps does not require a wholesale expansion of the consultation’s scope. However, it does require incorporating clearer conceptual distinctions into the guidelines; and acknowledging the need for alignment across the broader regulatory environment.

## **7. Implementation: Embedding Evidence, Accountability and Coherence**

The proposed reforms present a valuable opportunity not only to update classification guidelines, but to strengthen the institutional foundations that support them. To realise this opportunity, reform must be accompanied by mechanisms that ensure evidence-based decision-making, transparency, and system-wide coherence over time.

### **7.1 Embedding Evidence in Classification Practice**

A shift toward harm-based classification requires more than a change in principle — it requires operationalisation. COSL recommends that updated guidelines include:

1. Clear definitions of harm, distinguishing between:
  - a. direct harm to individuals (e.g. exploitation, abuse);
  - b. indirect or societal harms (where supported by evidence); and
  - c. mere offence or discomfort (which should not, on its own, justify restriction);
2. Evidentiary thresholds, specifying the types and quality of evidence required to support claims of harm; and
3. Guidance on uncertainty, ensuring that precautionary approaches are applied proportionately and transparently.

Without these elements, there is a risk that the language of “harm” will replicate the ambiguity of earlier moral frameworks, rather than replacing them.

### **7.2 Strengthening the Role of the Advisory Council**

The proposed Classification Advisory Council (Co9) has the potential to play a central role in maintaining an evidence-based and responsive system.

To be effective, the Council should be:

1. Independent and multidisciplinary, including expertise in:
  - a. psychology and behavioural science;
  - b. media and cultural studies;
  - c. law and human rights; and
  - d. lived experience from diverse communities;
2. Mandated to commission and review research, including on emerging media forms and evolving patterns of consumption; and
3. Integrated into policy development cycles, with a formal role in reviewing and updating guidelines over time.

This would ensure that classification policy remains grounded in current evidence, rather than static assumptions.

### **7.3 Transparency and Accountability Mechanisms**

Transparency is essential to building and maintaining public trust.

COSL supports proposals to:

1. publish clear explanations of classification decisions; and
2. improve access to classification data through digital tools.

In addition, COSL recommends:

1. Audit mechanisms for automated and self-classification systems, including:
  - a. regular independent review;
  - b. publication of high-level methodologies; and
  - c. processes for correcting errors or inconsistencies;
2. Documentation of decision-making criteria, ensuring that classification outcomes can be understood and, where necessary, challenged.

These measures would help ensure that increased efficiency does not come at the expense of accountability.

## **7.4 Ensuring System-Wide Coherence**

As discussed in Section 4, classification guidelines operate within a broader regulatory ecosystem. Ensuring coherence across this system is essential to achieving the objectives of reform.

COSL recommends that the Department:

1. Explicitly acknowledge the interdependence between classification guidelines and other regulatory frameworks, including criminal, customs, and online safety law;
2. Assess potential points of inconsistency, particularly where classification categories may trigger different regulatory consequences; and
3. Develop pathways for alignment over time, whether through coordinated review processes or cross-agency consultation.

This approach would support a more integrated and predictable system, reducing the risk of contradictory or disproportionate outcomes.

## **8. Conclusion: Drawing the Line Clearly**

Australia's classification system has long played an important role in helping individuals and families make informed decisions about media consumption. The current consultation recognises that, to remain effective, the system must evolve — becoming more consistent, more transparent, and more responsive to a changing media landscape.

COSL strongly supports this direction of reform.

At the same time, this submission has emphasised that the success of these reforms depends on clarity of principle as much as on changes to process or presentation. In particular:

1. a shift to evidence-based, harm-informed classification must be clearly defined and operationalised;
2. distinctions between real-world harm and fictional or simulated depictions must be maintained; and

3. reforms must be designed with an awareness of the broader legal and regulatory context in which classification operates.

A system that cannot distinguish between representation and harm risks both overreaching into lawful expression and failing to focus attention where it is most needed. Conversely, a system that draws these distinctions clearly is better able to protect against genuine harm, while preserving the space for lawful and diverse forms of expression.

By grounding reform in evidence, ensuring coherence across systems, and maintaining clear conceptual distinctions, Australia can develop a classification framework that is both fit for purpose and worthy of continued public trust.

## **Appendix A: Mapping to Consultation Themes**

### **Evidence-based approach (Co3)**

→ Supported. Recommend clear definitions of harm and evidentiary thresholds (Sections 2, 7.1)

### **Consistency across media**

→ Supported and expanded to include system-wide coherence (Section 4)

### **Sexual content reforms (Co8)**

→ Supported. Emphasise consistent, harm-based assessment (Section 6.1)

### **Themes overhaul (Co2)**

→ Supported with caution. Risk of reintroducing subjective criteria (Section 6.2)

### **Icons and consumer advice (C20–C21)**

→ Supported with safeguards to avoid over-signalling (Section 6.2)

### **Self-classification and automation (C12–C14)**

→ Supported with strong transparency and audit requirements (Section 7.3)

### **Advisory Council (Co9)**

→ Strongly supported with expanded research mandate (Section 7.2)

### **Key gap identified**

→ Need for explicit distinction between real harm and fictional representation; and recognition of interdependencies with broader legal frameworks (Sections 5, 6.3, 7.4)

## **References**

AAP (2026) Sydney author guilty of child abuse after book, Daddy's Little Toy, depicted adult role-playing as toddler,

<https://www.theguardian.com/australia-news/2026/feb/10/sydney-author-lauren-mastros-a-tori-woods-guilty-child-abuse-daddys-little-toy-ntwnfb>

COSL (2025), *Drawing the Line Watchlist 2025*,

<https://c4osl.org/wp-content/uploads/2026/01/Drawing-the-Line-Watchlist-2025-digital-version.pdf>

Department of Infrastructure, Transport, Regional Development, Communications and the Arts (2025), *Update to modernise and harmonise the classification guidelines 2025* (Consultation Paper).

Ferguson, C. J. & Hartley, R. D. (2022), “Pornography and Sexual Aggression: Can Meta-Analysis Find a Link?”, *Trauma, Violence, & Abuse*, 23(1): 278–287.

Lievesley, R., & Harper, C. A. (2026). Understanding the use of child-related sexual fantasy and fictional sexual materials by people who are attracted to children. *Psychology & Sexuality*, 1-33.

Malcolm, J. (2024), *Submission on Modernising Australia’s Classification Scheme – Stage 2 Reforms*,

<https://www.infrastructure.gov.au/sites/default/files/documents/mancs-jeremy-malcolm.pdf>.

Paul, B. & Linz, D. G. (2008), “The Effects of Exposure to Virtual Child Pornography on Viewer Cognitions and Attitudes Toward Deviant Sexual Behavior”, *Communication Research*, 35(1): 3–38.

Remminga, A. (2026). Obscene content: policing queer media, erotica, and fan creations in Australia. *Porn Studies*, 1–12. <https://doi.org/10.1080/23268743.2026.2625085>

Social Research Centre & Mettlesome (2025), *Evidence Review and Concept Testing for Classification Reform* (as summarised in consultation materials).

Stevens N (2020) Review of Australian classification regulation, <https://www.infrastructure.gov.au/sites/default/files/documents/review-of-australian-classification-regulation--may2020.pdf>.